

AI, WHERE IT HAS BEEN AND WHERE IT IS GOING

Arthur L. Samuel

Computer Science Department, Stanford University

A recent article on Expert Computer Systems carried a banner headline that said, and I quote, "Artificial Intelligence is no longer science theory", end of quote. This statement carries the clear implication that the science of Artificial Intelligence is now a closed book and there is no need for further research. I am sure that the writer of this headline had no such thoughts but I fear that there is a growing tendency to believe the implication. The main thrust of my talk today will be to deny this implication and to urge that more rather than less attention be given to AI research.

It is with a profound feeling of personal inadequacy that I address you today on this subject. While it is true that I have worked in Artificial Intelligence, most of my work was done a long time ago when the total amount of knowledge in this field was much less than it is today. It may well be that my concerns are, in part, colored by my lack of complete familiarity with the work that is now being done. I agreed to talk only because it seemed to me that I should be able to view the field of AI research from a perspective that is denied a newer worker in the field. I can not hope to say much about AI that you do not already know but perhaps the retelling of old ideas will stimulate your thinking.

In sharing my perspective with you, it seemed that it might be worthwhile for me to review some of the earlier history of AI and to recall what our thoughts were at some earlier time as to where AI research was going. By comparing where AI is today with our earlier expectations we might be able to get a clearer view of what still remains to be done. Then, perhaps, we might make some new predictions as to where AI will be at some future time, say for the year 2000.

Recalling where AI has been is relatively easy for me to do, since it has been my good fortune to have been involved with the modern digital computer and with AI almost from the start. Recalling earlier expectations and earlier events comes rather close to reminiscing and I hate to start reminiscing as this is supposed to be a mark of age.

You know, of course, that there are three ways to tell if a person is getting old. In the first place, the older person likes to reminisce. I must plead guilty on this count but then I have always had this failing. Secondly, old people tend to forget things. And then there is a third way—but I have forgotten what it is.

From one point of view, AI started in 1834 or shortly thereafter, when Charles Babbage suggested the possibility of having his Analytical Engine play chess. It was not until the 1940's, with the emergence of the modern digital computer, that we saw a renewal of interest in Artificial Intelligence. There were, of course, numerous attempts at making automata of various sorts but most of the time the aim was to fool the public rather than to get machines to exhibit behavior which if done by humans would be assumed to involve the use of intelligence. So we can date the emergence of AI as a discipline to some time in the mid 1940's.

Through a fortuitous combination of circumstances, I became interested in AI at its very beginning, not in 1834, I hasten to say, but in 1947, when the availability or soon to be realized availability of the modern digital computer made AI something more than idle speculation.

Actually, my involvement in computing began even earlier, in 1924 at MIT when Vannevar Bush got me interested in his Differential Analyser. When I first was introduced to this device, it consisted of a single stage using a modified watt-hour meter as the integrating device, although Bush was even then planning to use a Kelvin Integrating Disk for a second stage.

One of my first assignments was to see what could be done in the way of using the differential analyser to solve the non-linear equations in exterior ballistics, which, at the time were solved by the elaborate hand process of numerical integration using log tables and the old-fashioned desk calculator. Trying to solving these equations to the desired accuracy on the Differential Analyser, which was then capable of perhaps 5% accuracy, seemed quite hopeless to me. I felt that this problem required the accuracy of a digital solution and that maybe we should be working on a digital computer. Yes, digital computers were being talked about in 1924 although the distinction between computers and calculators was not then as clear as it later became. Bush, with his usual enthusiasm, soon talked me out of any such revolutionary ideas and set me to work, with a few hints as to how the ballistic problem might be solved.

The solution method, which Bush suggested and which I implemented, was to write a set of difference equations to account for the departure of an actual trajectory from a simple parabolic path, and then to solve these difference equations on the Differential Analyser. By this method,

we were able to achieve an accuracy of something better than 1%, not the desired accuracy, but near enough to hold out hope that the desired accuracy could ultimately be obtained.

This solution method was still in use during the war years, on the very much improved Differential Analysers that were, by then, available. Interestingly enough, it was the eventual dissatisfaction with this solution method that led Herman Goldstine to press for government support of the digital computer work at the University of Pennsylvania. Contrary to a popular misconception, the digital computer appeared on the scene too late to have any effect on the war although there were a few rather advanced electronic calculators that did do useful work, primarily in the field of cryptography.

In retrospect, I wonder what would have happened had I been successful in deflecting Bush from his interest in the Differential Analyser or alternatively, had we been less successful in applying the differential analyser to the exterior ballistic problem. Just think what would have happened, if Bush, with his keen intellect and driving enthusiasm, had started to work on digital computers in 1924.

A brief note in the Fifty-Year-Ago column of a recent issue of the Scientific American will serve to point out the way matters stood in 1933. Let me quote:

"The fundamental ideas of behavioristic and Gestalt psychologies justify attempts to construct and develop machines of a new type—machines that think. Such machines are entirely different from integrators, tide calculators and the like to which the term 'thinking machine' is sometimes applied. Unlike the latter they are not designed to perform with mathematical regularity but can 'learn' to vary their actions under certain conditions."

The original writer of this note was, of course, thinking about some of the early work that was then being done on self-organizing systems and not about digital computers as such but the argument over the relative merits of these two approaches was already taking shape in 1933.

Yes, we were talking about 'thinking machines' in 1933 and even about the use of electronic devices to solve problems in Artificial Intelligence although the term Artificial Intelligence had not yet been coined and the digital computer was still only a dream. I am sure that our ideas at the time were hazy in the extreme. One is apt to remember what one should have thought at an earlier time and not what one actually did think. On the other hand, there is the opposing tendency to forget the long gradual evolution that usually lies behind any new development such, for example, as the development of the modern digital computer. I can remember talking with others at the Bell Laboratories about the need for better computing facilities and about the feasibility of using the long-life vacuum tubes, that I had worked on, for the construction of some ill defined machine that would help me in solving some of my problems.

The war years put a stop to my idle speculation and it was not until 1946 after I had moved to the University of Illinois that I was again able to think about electronic aids to calculations and specifically about digital computers. My

old obsession with solutions of space charge problems with complex boundary conditions still persisted and I began to agitate for a digital computer.

Perhaps, it will give you a flavor of conditions in the late 1940's if I tell you something about our problems at Illinois during this period. After a lot of delays in getting started we eventually got an appropriation of \$110,000, a lot of money in those days, to build a computer. I was busy running an Electron Device Laboratory but I did have a lot of graduate students tumbling all over themselves trying to find suitable thesis subjects so I put several of them to work on new computer components, building a cathode ray storage tube, developing a magnetic core storage scheme, developing some special non-sequential adder and multiplier circuits etc.

We were overambitious, and it soon became evident that we did not have the amount of money that it then took to build the kind of computer that we wanted. The only solution seemed to be that of putting something together that was not ideal. The hope was that we could use this interim computer to do something spectacular that would induce the University authorities to give us some more money.

Since I had been making noises about unconventional things that one could do with a Computer, someone, I think it was Taube, suggested that I was the logical person to undertake the programming task. My first thought was that I should fulfill Babbage's old dream of having it play chess. I was disheartened when I found that Claude Shannon had apparently beaten me to the punch, but then, when I had arranged to meet Claude, I learned that he was much less far along with the problem that newspaper reports had led me to believe. My talk with Claude did awaken me to the difficulties that chess presented.

Checkers was my next choice. It happened that a world champion checker match was to be held in the neighboring town, Kankakee, the next spring and it seemed quite reasonable at the time for us to put together some sort of a computer in a few months and for me to write a checker program during this same period of time that could challenge and beat the new world checker champion at the conclusion of this match. This would give us the publicity that we needed. How naive can one be?

So I started to work, writing a checker program for a machine that did not exist, using an instruction set that I was forced to create as I needed it, writing directly in octal, and assigning fixed memory locations for my variables, since the idea of an assembly language and of an assembler had not yet been invented. The situation was ideal to highlight the need for a symbolic notation and why I did not invent one, I will never know.

Finding myself thwarted in my desires to devote all of my time to computer activities by the teaching load that I had to carry, I decided to leave University life in 1949 and go with IBM. I had written a checker program of sorts but the computer was still a long way from completion.

I had, by this time, become a confirmed hacker so I wanted to get my checker program to run on an IBM

machine. The program, of course, had to be completely rewritten, first for an interim computer called the Defense Calculator and then for the 701 as soon as the instruction set for this machine had begun to take shape. This first 701 version was still written in octal and I was able to get it to run on the very first experimental model of the 701. Late, after the idea of an assembly program had been conceived and implemented, I was faced with the difficult task of converting a running machine-language program from its octal form into symbolic notation. This led me to write a "Disassembly" program, which was to do yeoman duties in converting the library of system and application programs, that had by then been written, from octal to symbolic form.

IBM, in those days did not take kindly to one of their engineers wasting company time playing checkers, even if it was against a machine, and so most of my checker work had to be done on my own time. I had dressed my efforts up with a certain amount of respectability by adding a learning feature to the program but even then it was the use of the program as a computer test vehicle which could be adjusted to run continuously that got me the machine time that I needed to try out my experimental learning routines.

Later, during the 704 days, I can remember a period when there were often several 704's on the test floor at the same time with test crews working two shifts a day. The test crews would arrange to let me have all the time I wanted from 11 PM to 7 AM, sometimes on as many as four machines at the same time, in return for reports on the machine behavior.

Several events occurred along about this time that make it possible for me to recall the exact time sequence of AI development during these early years. Two of these events had a significant effect on the course of future events.

The first event was the Dartmouth Summer Research Project on Artificial Intelligence that John McCarthy ran during the summer of 1956. The name Artificial Intelligence was, of course, coined by John in the process of naming this project. Only a very few people were involved, but never again would it be possible for such a small a group to have such a profound influence on the course of AI research. The interchange of ideas that occurred at Dartmouth did a great deal to establish AI as a viable subject for continuing research.

The second event was the publication in 1963 of the book "Computers and Thought" edited by Feigenbaum and Feldman. This book reprinted twenty papers by twenty eight authors that the editors thought would summarize or at least would be typical of the state of the art in some ten aspects of the general field of artificial intelligence. Most of you know this book, I am sure, but I do want to call your attention to two rather significant features.

In the first place, most of the papers discussed were what might be called "toy" problems, that is problems that were specifically chosen because they required the development of new AI principles but were, never-the-less, sufficiently well defined that one could measure one's progress toward meeting some desired goal. A second feature was the sharp division of the papers into two distinct categories,

the first set that tried to look at the way that people solved the chosen problem and ape this action by machine, and the second group that considered the problem more or less in the abstract and tried to invent machine methods that were tied to the unique characteristics of the computer. This book was to become the bible of the early workers in the field.

It may be instructive to compare a recent publication, "The Handbook of Artificial Intelligence", for which Feigenbaum was again one of the authors, with this earlier book. The difference in treatment, and the difference in size, attest to the fact that AI is growing up.

The third and fourth events, not of earth shaking significance, but important to me as mnemonic aids, were the publication by me of two papers, the first one contained my assessment of the state of AI in 1962, and the second made predictions as to where computers would be in 1984.

The first paper was called "Artificial Intelligence: A Frontier of Automation". I am tempted to quote from this paper at some length since many things that I said then could equally well be said today, but I will forbear. The paper ends with a statement that "We are still in the game-playing stage, but, as we progress in our understanding, it seems reasonable to assume that these newer techniques will be applied to real life with increasing frequency and that the effort devoted to games and other toy problems will decrease." I will have more to say about toy problems later because I have come to attach greater significance to them than I did in 1962.

As a minor aside, I revealed my own personal bias between the two approaches to AI problems by using the phrase "Bird-Watching" to head the section describing the how-people-solve-the-problem approach and by using the phrase "Back to Aerodynamics" to head the section describing the second approach. Here again, there has been a softening of my disdain for inquires as to how people handle problems and in particular as to the way in which they learn. But more on this later.

I cannot refrain from quoting another excerpt from this paper to wit: "Progress is being made in machine learning and ...we will be able to devise better machines or even to program existing ones so that they can outperform man in most forms of mental activity. In fact, one suspects that our present machines would be able to do this now were we but smart enough to write the right kinds of programs. The limitation is not in the machine but in man. Here then is a paradox. In order to make machines which appear smarter than man, man himself must be smarter than the machine". This, I think, is the proper answer to those people who decry the way that computers are permeating our society and who express the fear that the machine will eventually dominate man.

The second paper, published in 1964, and called "The Banishment of Paper Work", was an attempt to predict the state of affairs in 1984. I want to review these predictions both to assess our progress and as a prelude to making some new predictions for the year 2000.

A few excerpts will give the flavor of these predictions. In one place I say "Given computers that are perhaps 100 to 1000 times as fast as the present day computers [they now bracket this range], computers with large memories [and they have gotten large, in spite of von Neumann's insistence that 1000 words of fast storage was all that would ever be needed], computers which occupy perhaps one one-hundredth the volume that they now do [again, just about right], computers that are much cheaper [the current factor is about 1000] and finally computers that learn from their experience and which can converse freely with their masters—what can we predict?" All of this has become more or less true, except for the final prediction that computers would learn from their experience and that they would be able to converse freely with their masters.

I went on to predict a number of specific things about computers. I will mention only three:

1. That we would see a dichotomy in the development of very large computers acting as number crunchers and as data bases with the widespread use of small personal computers that were quite respectable computers in their own right but that would also serve as terminals for access to the central data bases.

2. That all the accounting work of the world would be done by computer, hence the title of *The Banishment of Paper-Work*.

3. That process control with the attending automation would have reached a very high degree of development.

As for my unrealized predictions regarding artificial intelligence, I went on to say that: "...we have good reason for predicting that two rather basic problems will have been solved. The first of these has to do with learning... The second difficulty resides in the nature of the instructions which must be given ... in short one cannot converse with a computer." I still believe that I was right in assessing these two problems as the two most basic problems that stood in the way of progress but I was completely wrong in my belief that they would be solved in twenty years.

One can quibble with the soundness of my reasoning when I predicted that these problems would be solved. In one sense, problems relating to the real world are never solved, since each new contribution to our understanding usually leads to the unfolding of facets of the problem that we did not previously know existed. This has happened in the field of AI. What I did mean was that our lack of understanding would no longer stand in the way of our effective use of machine learning and of man-computer communications in the application of AI to problems of the real world. We are still some distance from this goal.

The situation with respect to machine learning is particularly distressing. What does *The Handbook of Artificial Intelligence* have to say about the matter? I quote "But in general, learning is not noticeable in AI systems."

One might question the correctness of this stricture when one notes that the third volume of this same handbook devotes some 188 pages, out of the three-volume total of 1300 pages to the subject of "Learning and Inductive Inference". Much of this work is really quite good. An

embarrassingly large amount of space, however, is used to describe some of my own work, work that was done twenty and more years ago and work that should have long since been superseded.

The stricture in volume 1 is essentially correct, since it is referring to the use of AI learning methods in application programs and not to the research on learning as such, and not also to the transfer of man derived knowledge to computer programs as is usually done in most Expert Systems applications. In fact the stricture appears even more damning when one considers the large amount of good research work that has been done, and the extent to which Knowledge Engineering and Expert Systems have been adopted by practitioners.

A considerable amount of work has been done on the problem of man-machine communication, and in making use of this work in AI systems, if we include, under this umbrella, the work on menu-driven techniques, the use of questioning techniques in which the computer restricts the range of the human responses by asking questions, the work on character reading, and the more profound work on speech recognition, and on language understanding, In spite of all of this work, we are still a long way from having a complete answer to the problem of man-machine communication. I believe that the desired progress in man-machine communication will not be made until it is treated as a learning problem.

To illustrate my point, let me raise the question as to how a child learns to talk. You will note an implication in the very way that we phrase the question. The question is: How does a child learn to talk?, not, How do we store in the child's brain all the information that he or she will need in order to be able to talk? Children learn to talk willy-nilly, whether there is a conscious effort on the part of their parents to teach them or not. We explain this by saying that children are born with an instinctive desire to learn, and with an instinctive ability to learn. Computers do not have instinctive desires but they are obedient servants and if we tell them to learn and if we tell them how to learn, they will do it. So all we have to do is to tell them how, either by wiring this knowledge into their very structure or by supplying them with the necessary software. This we have failed to do.

In fact, I look on learning as the central problem of AI research, and I believe that our failure to address this problem with enough vigor is the chief reason why the entire field of AI has failed to come up to my expectations. Oh, of course, you will say that my expectations were highly inflated and based on a gross failure on my part to realize how difficult the problems really are. The fact remains, that my 1963 AI predictions are nearly as far from realization today as they were in 1963, and this in spite of the fact that there has been much more research, and good research, done on AI in general, than I envisioned possible in 1963.

As an example, I said in 1963 that; "Computers will perform yet another function—that of language translation. Not only will we be able to obtain information from

the central files in the language of one's choice, but automatic translation via the telephone will also have come into use—although perhaps not general use, because of the cost and because of the gradual drift toward a universal language. It will, nevertheless, be possible to dial anywhere in the world and converse with anyone speaking a different language with only a slight transmission delay to allow for the differences in sentence structure and word ordering between the languages." The drift toward a universal language has been painfully slow and does not lead to any simple solution because of linguistic loyalties although English has become the *lingua franca* in some situations as evidenced by this conference that is being conducted in English in a German speaking country. There is still need for a telephone translation system which would be technically feasible to implement today if we only knew how.

Perhaps it is time to pause and see where all of this is leading us.

In the first place, there is an apparent inconsistency in my remarks. I am advocating an increase in research effort in the field of machine learning and yet the deficiencies that I have been citing all seem to be in the application of learning principles to problems of the real world. I think that this can be explained by remarking that the users of AI systems and particularly the designers of these systems would be anxious to add more elaborate learning schemes to their systems if they could. This, they are not able to do because there seem to be missing links in our knowledge that it is your business to supply.

Let me see if I can point out some of these missing links. I will do this under the four headings used in *The Handbook of Artificial Intelligence*. These headings are: 1. Rote learning, 2. Learning by being told, 3. Learning from examples, and 4. Learning by analogy.

Rote Learning

Rote learning was one of the first forms to be studied. As the Handbook points out, rote learning is useful only if it takes less time to retrieve the learned information than it does to recompute it.

There is an interesting parallel, where, in the early days, the Mark I computer at Harvard was constructed and used primarily for the purpose of computing mathematical tables that were bound into books and offered for sale. Most of these books were never sold and those that were sold were never used because the computational speeds of the, soon to follow, true electronic computers made it easier to recompute mathematical functions as needed than to store and retrieve them. We did go through an interesting transition state in which key tabular values were frequently stored and used with interpolation routines, but even this use of stored values has largely disappeared.

So I look on rote learning as ancient history.

Learning by Being Told

Learning by being told also appeared on the AI scene at an early date with the publication of John McCarthy's paper on advice-taking systems. But here, future work followed two diverging paths, the one envisioned by McCarthy in which the system would be able to accept abstract, high-

level advice and turn this advice into rules that would govern its subsequent behavior, and the second in which the emphasis has been on the development of tools that make it easy for an expert to transfer his detailed and extensive knowledge to the system.

This division between the Advice Taker and the more recent work on Expert Systems may seem unnatural and may even be resented by devotees of the expert system approach. In support of my view, I can reference a recent rather long review paper on "Expert Computer Systems" that lists some fifty references and that neither references nor mentions McCarthy's work or any of the much more recent work typified by the work of Mostow, Hayes-Roth, Klahr, and Burge that can be considered as a follow up on McCarthy's original ideas.

By and large, there has been a hiatus in the work on McCarthy's Advice Taker that is, in my opinion, very unfortunate. There is some work being done, as previously noted, but not nearly enough. This is not the place to try to develop a procedure for rectifying this deficiency, and indeed I am afraid that I have nothing constructive to offer.

A great many tools have been and are being developed to simplify the task of transporting specific and sometimes abstractly formulated expertise from a human expert to a computer program. In fact the entire field of Knowledge Engineering and of Expert Systems, as commonly understood, is based on this work. I certainly do not want to belittle this work but much of it is application engineering and not research.

Learning from Examples

If we again judge the effort in a field by the number of pages devoted to the field, in the Handbook of Artificial Intelligence, then certainly learning from examples is the most extensively studied phase of learning, with 150 out of 188 pages devoted to it. So I believe we can conclude that this form of learning is being given enough attention.

Learning by Analogy

Referring to the Handbook, which is my bible on these matters, we note that it leaves this form of learning undiscussed "...since this area has not received much attention." So here is a neglected area, that is, if it is important. I happen to think that it is.

Perhaps there is a slight misconception here that confuses the issue. By and large, people do not often learn by analogy. People reason by analogy in order to try to make sense out of a strange situation but then, if they are smart, they try to confirm the correctness of the analogy by some independent method of evaluation. If computers are to ape the way that people learn, they must somehow be given this added ability.

And speaking of the way people learn, this brings me back to a point that I have been on the verge of making several times during this talk. When we speak of research in machine learning methods we are really concerned with a problem in human learning and in how we can best devise experiments that will make it possible for us to learn how to tell a computer how to learn. So we must, after all, look into the way people learn, taking advantage, when we can,

of the work in the emerging field of Cognitive Science.

Also, there is still a need for continuing work on toy problems, since it is often quite impossible to control, or even to know about, all of the variables that are involved when one deals with a real-life problem.

Perhaps I have said enough. It was my original intention to cite chapter and verse as to why research in artificial intelligence was in danger of being neglected, but you already know these reasons and they have even been discussed in the public press.

So it is time for me to make my predictions for the year 2000 and to sit down.

As far as computers are concerned, I predict:

1. The present trend toward personal computers will continue and two distinct types of small computers will be available, a home computer and a much smaller pocket computer.

2. The pocket computer will be a rather modest sixteen bit computer with, perhaps, 128 K bytes of a non-volatile memory. Its chief function will be to communicate by radio with one's home computer or one's office computer either directly or via an exchange system when far away. It will, in effect, be a very smart terminal. It will also double as a telephone so that one will be able to speak with anyone any place in the world who has a phone connection. One will also be able to send voice messages to those computer systems that can receive such messages although the pocket computer itself will not be able to receive voice communications. Video-phone equipment will still be rather bulky and so most people will not include this facility in their pocket computers.

3. The home computer will be a thirty-two-bit machine and it will have at least ten million bytes of memory. This may use a non-volatile storage medium although such a memory will have to compete with volatile storage schemes that consume so little power that they will normally be left on at all times.

This computer will be capable of communicating with the many large data banks that will be available and with other personal computers, all, of course, under suitable protocols which attend to such matters as file integrity and secrecy. It will also be capable of communicating via radio with one's pocket computer, as mentioned earlier.

The home computer will be capable of receiving and transmitting pictorial material as well as textual material, and it will function as a video-telephone when this is desired. Since there will still be an occasional need for printed output this computer will produce book quality output and adequate paper handling facilities will be available, probably at an extra cost, to enable one to receive printed and pictorial material, the morning paper perhaps, without having to be present at the computer when it is arriving.

The tendency will be to leave one's home computer operating at all times so that it can receive and store messages that may be sent to it and even acknowledge the receipt of such messages. Most home computers will be battery powered with a trickle charger so as to be independent of power failures.

4. Many large centrally located computers with associated large data files will be available for public use, rather later than I predicted in 1963, but differing from those predicted for 1984 by the routine use of the knowledge engineering principals that are currently being used in present day Expert Systems, even augmented by principles yet to be discovered through future research. These systems will be, in effect, expert systems for the non-expert.

As I said in 1963, "One will be able to browse through the fiction section of the central library, enjoy an evening's light entertainment viewing any movie that has ever been produced (for a suitable fee, of course, since Hollywood will still be commercial) or inquire as to the previous day's production figures for tin in Bolivia—all for the asking via one's remote terminal".

But even more than that, these systems will be able to acquire knowledge and to learn not only from experts but also from their experience in dealing with the general public. Information acquired from the public, and generalizations based on this information, will probably still have to be approved by a panel of experts before being incorporated into the data base available to the public.

5. Large centrally located computers will still be used for large calculations but they will have reached a practical limit in size as set by velocity of light limitations. The general trend will be either to partition large problems so that they can be solved in pieces or to tie many separate machines together on a temporary basis for truly large problems. Newer methods of storage will have made it possible to construct very large memories so that this will not be a bottleneck.

6. I am less sanguine now than I was in 1963, about the prediction that: "Libraries for books will have ceased to exist in the more advanced countries except for a few which will be preserved as museums" although, I do believe now as I then said, "and most of the world's knowledge will be in machine readable form".

Having lived with computers for many years, with a terminal in my office and one in my study at home, I still find printed books to be very comforting. I suspect that many people feel this way and I now wonder if this feeling will somehow keep libraries in existence in spite of the fact that, and I again quote, "... the storage problem will make it imperative that a more condensed form of recording be used,..." Almost all of the new knowledge that is being generated, even at the present time, is now in machine readable form because of the wide spread use of word processing systems but we may still have a way to go before the machine-readable form will be the only form in which this knowledge is preserved.

7. I am much less sure as to where AI research will be in the year 2000. Perhaps I should leave this, as they say in some texts, as an exercise for the reader. If I let you make the predictions then, perhaps, you will feel a compulsion to do the necessary research to guarantee that your predictions come true. It is up to you to decide what the future will be.