Long Range Planning, Vol. 22, No. 1, pp. 118 to 123, 1989
Printed in Great Britain

0024-6301/89 $3.00 + .00
Pergamon Press plc

Toward the Library of the Future

E. A. Feigenbaum

Over the past decade libraries have become increasingly aware of the revolutionary impact of developments in information technology on their key function. The new developments challenge the library's traditional role as electronic information products and services open up a previously unimaginable array of options. The author contemplates a vision of the future in which the unpredictability of the human mind interacts with an expanding array of new technologies and libraries strive to develop an information infrastructure to serve teaching and research in universities.

In this article I would like to discuss knowledge processing. We are making the transition from the 'file servers' that we see in the computing area to 'knowledge servers', or perhaps, more broadly, 'The Library of the Future'. I would like to examine the future of computer science and technology. When people make forecasts, they tend to overestimate what can be done in the short term and underestimate what can be achieved in the long run. Looking back on computer science and technology over the past 40 years, our gains in this area have been substantial; perhaps many of us would think impressive. However, we have hardly begun, and it is important to us not to lose sight of the place to which we are heading, however distant that place may seem.

Today, we are witnessing a transition from 'data processing'—the computing of the past—to 'knowledge processing'. The key tool of our speciality is the digital computer: the most complex and yet the most general and universal machine that has ever been invented. Though the computer is a universal symbol-processing device, until now we have exploited only its mundane capabilities to do arithmetic, and to file and retrieve data. The researchers in the part of computer science called 'artificial intelligence' have been studying for several decades the techniques for representing human knowledge for computer use, and the methods by which that knowledge can be used to reason toward the solution of problems, the formation of hypotheses, and the discovery of new concepts and new knowledge.

Know-bots

These researchers have been inventing the knowledge servers of the future. I like to use the term 'know-bots' to distinguish these from the 'robots' that do factory automation. We have been inventing the 'know-bots' of the future. Scientists and technologists, like all people who create, must dream, must put forth a vision of their future, or else they relegate their work to a kind of near-pointless incrementalism. So please allow me to sketch my vision of the future of artificial intelligence research and development over the next several decades, and the knowledge systems that can be produced thereby to assist the modern knowledge worker. Let me discuss first the beginnings of this dream.

Fifty years ago, before the modern era of computation began, Alan Turing's theorems and abstract machines gave hint of the fundamental idea that the computer could be used to model the symbol-manipulating processes that make up that most human of all endeavours: thinking. Thirty years ago, this work began in earnest. 1986 is generally taken to be the thirtieth birthday of the beginning of real work in artificial intelligence, dating from the Dartmouth summer conference on artificial intelligence in 1956. The founding principle of the artificial intelligence research area is really an article of faith, that the digital computer has the necessary and sufficient means for intelligent action. This first principle has a name. It is called 'The Physical Symbol System Hypothesis'. It was so named by Carnegie–Mellon University scientists Newell and Herbert Simon, when they gave their Turing Award Lecture of the Association for Computing Machinery.

The early dreaming included intelligent behaviour at very high levels of competence. Turing himself speculated on wide-ranging conversations between people and machines, and also on chess-playing

This article is based on a lecture given at Aston University in November 1986 to mark the opening of the University's new Computing Suite. Professor E. A. Feigenbaum is Professor of Computer Science at Stanford University, California.

programs. Later on, Newell and Simon themselves worked on and wrote about champion-level chess-playing programs, and began their work towards that end. Samuel worked on draughts; Gelernter worked on geometry theorem-proving programs; many other people contributed early efforts towards the achievement of high levels of competence by computer programs in solving problems. At Stanford, Joshua Lederberg and I chose 'reasoning in science' as our task, and we began work with Buchanan and Djerassi, a chemist, on building a program that would solve chemical structure elucidation problems at high levels of competence. These became known as the Dendral collection of programs. What emerged from the many experiments with Dendral was an empirical hypothesis that the source of the program's power to solve chemical structure problems from spectral data was its knowledge of basic and spectral chemistry.

For Dendral, knowledge was power. Now is that an obvious hypothesis? Well, in retrospect, perhaps. In retrospect, most scientific propositions turn out to be obvious. But the prevailing view in artificial intelligence at the time, and that was 1965, ascribed power to the reasoning process—in modern terms, to the inference engine—not to the knowledge base. Thus, in the late 1960s, the 'knowledge is power' hypothesis stood as a contra-hypothesis awaiting further test, and the accumulation of more evidence. Much of this evidence came in the 1970s. Medical problem-solving provided the springboard. We did a program at Stanford called Mycin, the work of Dr Shortliffe and others. It was the prototype of what is now called the 'expert-level advisory system', or the 'expert system'. The core of Mycin was its knowledge base of rules for infectious disease diagnosis and therapy. Its reasoning process was simple, namely backward chaining. It was even *ad hoc* in parts. But Mycin was built as an integrated package of intellectual abilities. It could interact with a professional, in the professional jargon of the speciality; it could explain its line of reasoning. It even had a subsystem that could aid in the acquisition of new knowledge by guiding an expert to find defects in the stored knowledge. Overall, Mycin provided strong confirmation of the 'knowledge is power' hypothesis.

At nearly the same time, other efforts in medical problem-solving were providing similar results. At the University of Pittsburgh, the focus of the Internist project of Pople, Myers and Miller was the construction of an enormous electronic textbook of the knowledge of internal medicine for use by problem-solving programs. With its current knowledge base of 572 diseases, 4500 manifestations of disease, and hundreds of thousands of links between these, Internist has provided the strongest confirmation yet of the 'knowledge is power' hypothesis. In the late 1970s, in fields other than medicine, an explosion of expert systems was taking place: in Engineering, in Manufacturing, in Geology, in Molecular Biology, in Financial Services, in Diagnostic Servicing of Machinery, in Military Signal Processing and in many other areas.

I take great pleasure in talking about the proliferation of expert systems in the industrial and commercial world today. It is the feeling that a proud parent has for the success of a child. Recently I read a story in the *Financial Times* entitled 'Expert Systems are now finding practical applications in industry and commerce', and it describes the use of expert systems in sales and marketing, in running department stores. There is one U.K. product that was shown in a San Diego show: an expert system which makes it simple for its customers to reconfigure their data communications networks. It is designed to solve the problem experienced by every operator of large data communications networks of how best to handle the regular and inevitable changes in the physical and logical structure of the network. Expert system work has been proliferating from extremely complex applications, such as assisting fighter pilots in piloting high-speed aircraft, all the way to assisting foreign exchange traders in the enormous financial transaction manipulations they do every evening, and to helping people manage the reservations and other transactions taking place with the public in a resort hotel in the Bahamas. In other words, expert systems are finding applications as wide as the entire world of work, which is exactly what you would expect.

Expert systems are aids to human professional endeavour, and wherever you find humans at work doing knowledge work, intellectual work, you will find an expert system application. There is little that ties these areas together other than this. In each of them high-quality problem-solving is guided by experiential, qualitative, what we call 'heuristic' knowledge. The explosion of applications has created a new type of professional, the knowledge engineer, now in very short supply, and a new industry, the expert systems industry, now expanding rapidly. The scientific generalization from this frenzy of activity is simply this: massive additional confirmation of the 'knowledge is power' hypothesis. The reasoning procedures associated with all of these systems are weak; in their knowledge basis lies their power. Not by design are these reasoning procedures weak, let me hasten to add, but because of our own limited scientific and engineering insight as to how to build more powerful reasoning processes. We would if we could, but right now knowledge *is* power.

Other areas of artificial intelligence research made shifts to this knowledge-based viewpoint. It is now commonplace to say that a program for understanding natural language must have extensive knowledge of its domain of discourse. A vision program for image understanding must have knowledge of the world that it is intended to see. To take a much more complex case, learning programs must have a substantial body of knowledge from which to expand. We now understand that even a task like learning takes place at the fringes and the interstices of what it is we already know. Thus, the dream of the computer that performs at a high level of competence, over a wide variety of tasks that people perform well, seems to rest upon knowledge in the task areas. It is more powerful to be highly knowledgeable than to be a clever and astute reasoner.

The 'knowledge is power' hypothesis has received so much confirmation that we can now assert it as 'The Knowledge Principle'. Let me state the Knowledge Principle:

> A system exhibits intelligent understanding and action at a high level of competence primarily because of the specific knowledge that it contains about its domain of endeavour.

The reasoning processes of an intelligent system, being general, and therefore weak, are not the important source of power that leads to high levels of competence in behaviour. The Knowledge Principle simply says that, if a program is to perform well, it must know a great deal about the world in which it operates. In the absence of knowledge, reasoning will not help. The Knowledge Principle is the emblem of what I call the first era of articial intelligence: Act 1 of this dream. The Knowledge Principle should inform and influence every decision about what it is feasible to do in the science and technology of artificial intelligence.

Act 2: Interaction Through Natural Language
Now let me project into the second era of this field's work: what I call Act 2, the middle of the dream.

Today, our intelligent artefacts perform well, but only on specialized tasks within narrowly defined domains. As I mentioned earlier, an industry has been formed to put this technological understanding to work, and widespread transfer of this technology has been achieved. Although the first era of the intelligent machine is ending, many problems remain to be solved.

One of these problems is naturalness. The intelligent agent should interact with its human user in a fluid and flexible manner that appears natural to the person, but the systems of the first era share with the majority of computer systems an intolerable rigidity of stylistic expression, vocabulary and concept. For example, programs rarely accept synonyms, and they cannot interpret and use metaphors. They always interact in a rigid, grammatical strait-jacket. Incidentally, notice the metaphor in the use of the term 'strait-jacket'.

In the second era of artificial intelligence, we will see the evolution of the natural interface. The processes controlling the interaction will make greater use of the domain of knowledge of the system, and of knowledge of how to conduct fluid discourse. Harbingers of naturalness already exist, and they are based to a large extent upon pictures. The Oncocin project team at Stanford, a team that has built an expert system for advising physicians on the administration of cancer chemotherapy, invested a great effort in an electronic flowsheet to provide a seamless transition for the oncologist from the paper forms for patient data entry to electronic versions of these forms. The commercially-available expert system development software packages sometimes contain elegant and powerful packages for creating pictures that elucidate what the knowledge system is doing, and what the emerging solution looks like. (For example, Intellicorp's KEE Pictures and Active Images.)

Naturalness, of course, need not only rely upon pictures. The advances in natural language understanding have been quite substantial, particularly in the use of knowledge to facilitate understanding. In the search for naturalness, the need for metaphor to induce in the user a feeling of naturalness seems critical. Metaphorical reference appears to be omnipresent and virtually continuous in our use of language, and if you believe that our use of language reflects our underlying cognitive processes, then metaphor is a basic ideational process. Thus, in the second era, it will become commonplace for knowledge systems to interact with humans in natural language within the scope of the systems' knowledge.

The interaction systems of the second era will increasingly rely on continuous natural speech. In person-to-person interactions, people generally say what they want, rather than type it. Typing is useful, but unnatural. Speech-understanding systems of wide applicability, based on the knowledge principle, are coming. For example, at Stanford we are beginning experiments with an experimental commercial system interfaced to the cancer chemotherapy advisory system Oncocin. To summarize these projections regarding naturalness of the interaction between person and machine in the next two decades, we will see highly fluid natural language dialogues based on knowledge, great facility with interactive picture systems, and knowledge-based continuous natural speech interaction.

Model-based Reasoning

A limitation of first-era systems is what I call their 'brittleness'. To switch metaphors, these systems operate on a high plateau of knowledge and competence until they reach the extremity of their knowledge. Then they fall off precipitously to levels of utter incompetence. People suffer from the same difficulty. They too cannot escape the Knowledge Principle, but their fall is more graceful. The cushion for the soft fall is the knowledge and use of weaker, but more general, models that underlie the highly specific and specialized knowledge of the plateau. For example, if an engineer is diagnosing the failure of an electronic circuit for which he has no specific knowledge, he can fall back upon his knowledge of electronics, on circuit analysis methods, and on handbook data for the components.

The capability for such model-based reasoning by machine is just now under study in many laboratories. For example, by Ramesh Patil at MIT, whose influential work on model-based reasoning in medicine has demonstrated the important role of underlying mechanistic and causal models in enhancing the performance of medical advice systems. His Abel system deals with the multilevel model of electrolyte and acid–base abnormalities. Simple patient-specific inferences could be reached with simple inference procedures, but more complex reasoning and explanations of observations require deeper and deeper analyses of underlying pathophysiologic mechanisms in the kidney and gåstrointestinal tract.

Common Sense Knowledge

Consider another expert system—one done at Stanford, called the Onyx system—by the same group that did Oncocin. Onyx is the model-based reasoning adjunct to Oncocin. Oncocin reasons with rules codified from the cancer treatment protocol. But what do you do when the rules do not fit, or when situations not covered by the protocol are encountered? Onyx can derive the necessary treatment strategy from more basic models, for example, models of drug distribution, models of the mechanism of action of these drugs, including the point of action in the cell cycle. In short, Onyx can reason creatively through difficulties encountered in employing the standard protocol. These capabilities will emerge as an important feature of second-era systems.

The capability does not come free; knowledge engineers must explicate and codify general models in a wide variety of task areas. I say 'task areas'. But what if there is no task? Can we envision the intelligent program that behaves with common sense at the interstices between tasks when task knowledge is completely lacking? Common sense is itself knowledge, an enormous body of knowledge distinguished by its ubiquity, and by the circumstance that it is rarely codified and passed on to others as is more formal knowledge.

To take an example, pregnancy is associated with females, not males. We know that. We do not need to write it down. We do not ordinarily communicate it. It is a piece of common sense knowledge. The extremely weak, but extremely general, forms of cognitive behaviour implied by the term 'common sense reasoning' constitute for many the ultimate goal in the quest for machine intelligence. Researchers are now beginning the arduous task of understanding the details of the logic and representation of common sense knowledge, and the codification of large bodies of common sense knowledge. The first fruits of this will appear in the later systems of the second era. Common sense reasoning will probably appear as an unexpected naturalness in the interaction with one's intelligent agent. As an example of this, to go back to my previous example of common sense knowledge, in consultation advisory systems: if pregnancy is mentioned early in the interaction, or can be readily inferred, the interaction shifts seamlessly to understanding that a female is involved. Magnify this example by 100,000 or 1,000,000 unspoken assumptions, and you will understand what I mean by a large knowledge base of common sense knowledge.

Reasoning by Analogy

As knowledge and systems expand, so does the scope for modes of reasoning that have so far eluded the designers of the knowledge systems. In my view, foremost among these are methods for reasoning by analogy and the sibling of reasoning by analogy—metaphorical reasoning. The essence of analogy has been evident for some time, but the details of analogizing have not been. An analogy is a partial match of the description of some current situation with one's stored knowledge. The extent of that match is critical. If the match is too partial, then the analogy is seen to be vacuous or far-fetched. If the match is too complete, then the analogy is seen as hardly an analogy at all. Analogizing broadens the relevance of the entire knowledge base. It can be used to construct interesting and novel interpretations of situations and data. It can be used to retrieve knowledge that has been stored, but not stored in the expected way. Analogizing can supply default values for attributes not evident in the description of the current situation. Analogizing can provide access to powerful methods that would otherwise not be evoked as relevant. To take one famous example from early 20th-century physics, Dirac made the analogy between quantum theory and mathematical group theory that allowed him to use the powerful methods of group theory to solve important problems in quantum physics.

Analogizing is seen also as an important process in knowledge acquisition, which is another name that we give to learning. In first-era systems, adding knowledge to knowledge bases has almost always been a manual process. People codify knowledge and place it in knowledge structures. Recent experiments by Douglas Lenat at the MCC in Austin, Texas, have shown that this laborious process can be partially automated, can be facilitated by an analogizing program. The program suggests the relevant analogy to a new situation, and the knowledge engineer manually fills in the details. In the second era, we will see programs that acquire the details with less, or no, human help. Many other techniques for automatic learning will find their way into second-era systems. For example, we are currently seeing the early experiments on what are called 'Learning Apprentices', that carefully observe people performing complex tasks and infer thereby the knowledge needed for competent performance. The second era will also see, I predict, the first successful systems that couple language understanding with learning, so that knowledge bases can be augmented by the reading of text. Quite likely these will be specialized texts in narrow areas at the outset.

To summarize briefly, because of the increasing power of our concepts and our tools, and the advent of automatic learning methods, we can expect that, during the second era, the knowledge bases of intelligence systems will become very large, representing therein hundreds of thousands, perhaps millions of facts, heuristics, concepts, relationships and models. Automatic learning will be facilitated thereby, since, by virtue of the Knowledge Principle, the task of adding knowledge is itself performed more competently the more knowledge is available. That is, the more we know, the easier it is to know more.

A Broader View of Expert Systems
Finally, in the second era, we will achieve a broad reconceptualization of what we mean by a knowledge system. In this broader concept, what we mean by 'the system' will be conceived as the relationship between an intelligent computer agent, and an intelligent person, or persons. Each will perform tasks that he, she or it does best, and the intelligence of the system will emerge from that collaboration. If the interaction is indeed seamless and natural, then it may hardly matter whether the relevant knowledge or the reasoning skills needed are in the head of the person or in the knowledge structures of the computer.

The Library of the Future

I have been describing the first era and the coming second era of knowledge systems. Now let me move on to what I call the far side of the dream. I like to use slogans and, needing a slogan, I created one for this purpose: 'The Library of the Future'. In the U.S.A. we have now a small national working group discussing large knowledge bases and the library of the future, and planning the structure, the personnel, the intellectual problems and the funding for such an enormous project. This small working group includes Professor Marvin Minsky, one of the founders of the artificial intelligence field, and several other renowned individuals in the field. One day, we were having a wide-ranging discussion of the intellectual concepts behind the library of the future. We were putting ourselves into that library 50 years from now, looking back on today. Professor Minsky said something that I thought in jest, but which was quite profound. In his life out in the future, he said 'Can you imagine, they used to have libraries where the books didn't talk to each other?' The libraries of today are warehouses of passive objects. The books and journals sit on shelves waiting for us to use our intelligence to find them, to interpret them, and cause them finally to divulge their stored knowledge. The so-called 'electronic' libraries of today are no better. Their pages are pages of data files, but the electronic page images are equally passive.

Books that Interact
Now imagine the library as an active intelligent knowledge server. It stores knowledge of the disciplines in complex knowledge structures, perhaps in a knowledge representation formalism yet to be discovered or invented. It can reason with this knowledge to satisfy the needs of its users. These needs are expressed naturally with fluid discourse. The system can, of course, retrieve and exhibit. That is, it can act as an electronic textbook, but it can also collect relevant information, it can summarize, it can pursue relationships. It acts as a consultant on specific problems, offering advice on particular solutions, justifying those solutions with citations, or with a fabric of general reasoning. If the user can suggest a solution or an hypothesis, it can check it. It can even suggest extensions, or it can criticise the user's viewpoint with a detailed rationale of its agreement or disagreement. It pursues relational paths of associations, to suggest to the user previously-unseen connections. Collaborating with the user, it uses its processes of association and analogizing to brainstorm for remote or novel concepts. With more autonomy, but with some guidance from the user, it uses criteria of 'interestingness' to discover new concepts, new methods, new theories, new measurements.

The New Publishing
The user of the library of the future need not be a person. The user may be another knowledge system, that is, any intelligent agent with a need for knowledge. Thus, *the library of the future will be a network of knowledge systems in which people and machines collaborate*. Publishing is an activity transformed. Authors may bypass text, adding their increment to human knowledge directly to the knowledge structures. Since the thread of responsibility must be maintained, and since there may be disagreement as knowledge grows, the contributions are authored knowledge based maintenance, that is, the updating of knowledge, itself becomes a vigorous part of the new publishing industry. This dream can take many forms and many dimensions. I have briefly sketched only a few in the limited space that I have, but I invite all of you to exercise your own imagination to sketch your own dream. The question here is not 'if' or 'whether' these things will come about, but only 'when'. Since projects of this magnitude do not happen spontaneously, the

further question is: under whose leadership? How will we get to the Library of the Future?

Library of Medicine 2006
Answers to these questions are now beginning to emerge, and in a surprising place: under the leadership of the National Library of Medicine in the United States; in particular under the leadership of its imaginative Director, Dr Donald Lindberg, formerly a Professor at the University of Missouri, and Chairman of the National Advisory Group for Artificial Intelligence in Medicine for a decade before he became Director. He has organized and guided panels of distinguished scientists, technologists and librarians to produce a vision of the National Library of Medicine of the year 2006 and beyond, and a plan of action to achieve it. Of special interest here is the plan for basic research in medical aspects of computer science. That remarkable report offers scenarios that are sketched in some enticing detail of the services that would be offered by the Medical Library of the Future.

Let me give one example. An industrial gas that had been stored for half a century is undergoing detoxification when an accident occurs. While the victims are being rushed to a hospital, mass spectral analysis is automatically performed on samples of the gas. Measurements monitoring the states of the victims are made and interpreted, including blood sample information. The toxicology information bank is automatically accessed, and personal medical histories are integrated immediately into the analysis. Since the toxic gas has been encountered only rarely in several decades, the system reasons from what little case history already exists, and from basic medical models, to offer treatment recommendations for both the emergency care *en route* and the emergency care when the victims enter the hospital.

Power Tools for the Mind
The encompassing vision is what I call 'power tools for the mind'. The first industrial revolution understood the concept of power tools. It understood the concept of amplification of human abilities using artefacts. The hard work of the farmer was revolutionized by agricultural machinery. The labour of the industrial worker was revolutionized by engines and heavy machinery. As we move into the post-industrial period of the 21st century, as work becomes increasingly the work of professionals and knowledge workers, the power tools are digital computers. The economic and social well-being of our advanced societies is increasingly the result of working smarter, rather than working harder. Computers are the agents of that change. Knowledge is power in human affairs, and knowledge systems are amplifiers of human thought and action.

I was struck when I visited Beijing, 3 years ago, by something I read in an English language newspaper, the *China Daily*. You would expect a Third World nation, an under-developed nation, vastly over-populated, not to appreciate the 'knowledge is power' idea, but this is what I read in an editorial on the advancement of technology, in particular, information technology: 'The saying goes, "The wise can support one thousand people, while the strong may support only one". Therefore, the key to success lies in knowledge and techniques, and not merely the number of workers' This was said in another way by the Chief Executive Officer of Olivetti, Carlo De Benedetti: 'For the first time in the history of mankind, innovation is the fundamental raw material. Real strategic resources are no longer represented by coal, steel or oil, but by the cleverness and cognitive capability of man.'