Impacts of
Artificial
Intelligence

R. Trappl
editor

IMPACTS OF ARTIFICIAL INTELLIGENCE

*Scientific, Technological, Military, Economic,
Societal, Cultural, and Political*

Edited by

R. TRAPPL

*Austrian Research Institute
for Artificial Intelligence
Vienna
Austria*

With Contributions by

Michael A. Arbib, Margaret A. Boden, Stefano A. Cerri,
Ivan M. Havel, Werner Horn, Makoto Nagao, Nils J. Nilsson,
Roger C. Schank, Stephen Slade, Ingeborg Steinacker,
Robert Trappl, and Tibor Vámos

Acknowledgments

Contents

CHAPTER 1

Introduction

Robert Trappl *

1 Background to the Study

How is my job likely to change under the impacts of new technologies, and how will these affect the persons close to me? Will I soon lose my privacy as computer surveillance becomes more widespread? More and more people are asking these questions as the coverage in the media of artificial intelligence, and particularly the threats it may pose, increases steadily. At the same time, politicians, civil servants, and executives in funding and investment institutions are asking whether artificial intelligence will really yield a high return on the money and other resources invested.

Researchers who are themselves active in AI note that an increasing number of their colleagues are being employed by defense agencies of one sort or another. Even those working in universities or private companies are concerned about whether the results of their work will be of real benefit to mankind, or wonder if their positions are not uncomfortably close to that of J. Robert Oppenheimer during the Manhattan Project.

There are certainly grounds for all these questions and concerns, and they have been strengthened by the recent explosion of growth in spending on AI research. The Japanese Fifth Generation Computer Systems Project was established in 1981 with a three-year budget of US $45 million. This was swiftly countered by a group of US companies who

*Department of Medical Cybernetics and Artificial Intelligence, University of Vienna, Freyung 6, A-1010 Vienna, Austria.

founded the Microelectronics and Computer Corporation (MCC); the MCC project ALFAOMEGA will receive the bulk of the Corporation's funds, with US $15 million per year for ten years' work on AI studies, new computer designs, and models and structures for data bases. The US Department of Defense has since requested US $600 million for military research, mostly in the field of AI. And more than US $100 million has already been raised in the United States as venture capital for AI software companies.

In Europe, the UK Government has started its Alvey Project, one of whose four main research areas is AI, with funding equivalent to US $300 million. The Federal Republic of Germany followed suit with a US $50 million research program on pattern recognition and knowledge processing. In Spring 1984 the European Community decided to launch ESPRIT, a European Strategic Program in Information Technology; ESPRIT has an enormous US $1.5 billion budget to be spent in five areas, and the largest of these (accounting for even more than microelectronics) will be advanced information processing, defined in the program as "machine intelligence."

After several years of working in the area of AI, I also became more and more troubled about the possible impacts of our research, but felt that the topic needed more systematic attention than was possible in occasional gatherings at conferences. I therefore approached the Austrian Federal Ministry for Science and Research to support a study of the potential impacts of AI. It seemed absolutely essential that the study should not be limited to a literature search and perusal of the available material but that it should also provide the opportunity for leading AI researchers to discuss the topic at some length in a suitable atmosphere. Very fortunately, the Ministry agreed with these proposals.

2 Organizational Framework

The goal of our study, and therefore also of the Meeting that gave rise to this volume, was defined as follows:

> "... to determine and assess impacts of AI in order to avoid potential damage and to encourage socially helpful and economically useful AI research and applications."

The results of our work should be of interest to three types of audience:

- concerned AI researchers;
- decision makers in government, funding institutions, and private companies; and
- "educated laymen" who are aware of the potential changes that AI may bring about in their lives.

Together with the International Institute for Applied Systems Analysis, in Laxenburg, Austria, which agreed to co-organize and host a three-day Task Force Meeting, I tried to identify leading AI researchers who had already expressed interest in the impacts of AI, were working in a range of different areas of AI, and came from countries with a variety of social and economic systems.

After many long-distance calls, letters of invitation were finally sent to Professors Michael A. Arbib (United States), Margaret Boden (UK), Ronald Brachman (United States), Stefano Cerri (Italy), Ivan M. Havel (Czechoslovakia), Friedhart Klix (GDR), Makoto Nagao (Japan), Nils J. Nilsson (United States), D.A. Pospelov (Soviet Union), Peter Raulefs (FRG), Roger C. Schank (United States), and Tibor Vámos (Hungary) to attend this Meeting.

Each was invited to prepare a position paper reflecting his or her own views concerning the future of AI according to the following guidelines:

- Development within the next ten years:
    1. What are the most likely core research areas in AI and what results do you expect?
    2. What impacts of AI on other sciences or on technology do you expect?
    3. What social impacts do you expect?
    4. What economic impacts do you expect?
- If you were to speculate, what might be the long-range research results and impacts of AI?
- Which particular research areas of AI would you, as a member of the AI community, support mostly, and why?

The position papers we received (Roger Schank's was co-written with Stephen Slade) were distributed among the participants before the Meeting. Only AI researchers were invited to attend, thus running the risk of overspecialization and omissions, but ensuring that it was possible to conduct involved discussions from a broadly shared background.

The Meeting of AI researchers described above and its physical outcome, this book, comprised the first stage of the study. In the second stage there will be a conference based on the book, which will

actively involve not only AI researchers but also sociologists, political scientists, and decision makers.

3 Structure of the Book

This book opens with a short, "one-hour course" in AI, which is intended to provide a nontechnical but hopefully informative introduction to the material that follows. Next comes an overview chapter, which is based on an extensive literature search, the position papers, and the discussions during the Meeting. The overview is not a comprehensive summary of the Meeting but represents my own — naturally somewhat biased — personal views on the subject.

The next section of the book, Chapters 4–11, contains the position papers (revised and updated by the authors), whose richness and diversity illustrate the wealth of opinions and research directions that today fall under the umbrella term "AI research." The papers are followed by a select bibliography containing nearly 700 books, articles, and research memoranda on AI-related topics, together with a thesaurus and KWIC index to facilitate the retrieval of information. The book closes with an index and two appendices, one listing the names and addresses of all the contributing scientists and the other giving details of the AI curriculum at the University of Vienna.

This book is intended to serve as the first stage in an iterative process of detecting, predicting, and assessing the impacts of AI. If, after reading it, you would like to contact either myself or any of the other contributors with your opinions or questions about AI, you are welcome to do so. We look forward to receiving your reactions.

CHAPTER 2

Artificial Intelligence: A One-Hour Course

Robert Trappl *

1 Summary

After reading this chapter, which will take about one hour, you will know

- how artificial intelligence is defined,
- something about the three AI product groups that are already commercially successful,
- how some of the principles of AI really function,
- which software tools and which dedicated machines AI research requires,
- in which directions AI may develop in the future, and
- where to look for more information about AI.

2 Definitions of Artificial Intelligence

The definition of artificial intelligence most widely accepted is *making computers smart*. Another definition, also widely accepted, is *making computer models of human intelligence*. And a third one might be summarized as *building machines that simulate human intelligent behavior*.

---

*Department of Medical Cybernetics and Artificial Intelligence, University of Vienna, Freyung 6, A-1010 Vienna, Austria.

A survey conducted at the Eighth International Joint Conference on Artificial Intelligence in 1983 revealed 143 different definitions of AI, the majority of which could be related to one of the three presented here. Interestingly, there is also a correlation with world regions: the first definition is generally accepted in the United States, while in Europe the distribution is more uniform (Negrotti 1983).

These three definitions correspond to the three main ways of approaching artificial intelligence: the *behavior-oriented* approach, which attempts to program computers to behave in an intelligent, or "smart" way; the *cognitive* approach, which tries to model human thought processes in order to understand the human mind better; and the *robotic* approach, which is interested not only in programs, but also in building the machines.

The term "artificial intelligence" was coined by John McCarthy in 1956 for a conference called "The Dartmouth Summer Research Project on Artificial Intelligence." Many researchers at that time and ever since have felt that the name is not the best possible description for what they are doing; but in the absence of any general agreement on other candidates, such as "machine intelligence" or "intellectic," artificial intelligence has persisted.

AI researchers have different opinions about how to define the subject, and therefore stress different areas and efforts of AI (and ignore others), but most of them will agree that what they are doing is AI and they will also accept the work of their colleagues as AI. Therefore, AI curricula at different universities don't differ as much as might be expected. Since the curricula also give a good overview of *which* topics are generally agreed to belong to AI, I have included in Appendix 2 details of our own version here in Vienna, where the AI curriculum is an elective specialization of the computer science study program.

3 Three Commercially Successful AI Product Groups

3.1 Expert systems

An expert system is a large computer program that makes available to its user rare or important expertise, in other words the knowledge and reasoning of one or more experts. Since good experts are expensive, most AI companies concentrate on expert systems. Probably more than 1000 expert systems have already been developed, in areas like medicine, chemistry, geology, molecular genetics, electronics, law, computer configuration, and elsewhere.

To give some idea of what a dialogue with an expert system looks like, Figure 1 lists a dialogue with the famous medical expert system

- What is the patient's name?
  John Doe

- Male or female?
  Male.

- Age?
  He is 55.

- Have you obtained positive cultures indicating general type?
  Yes.

- What type of infection is it?
  Primary Bacteremia.

- When did symptoms first appear?
  May 5, 1975.

- Let's call the most recent positive culture CULTURE-1. From what site was CULTURE-1 taken?
  From the blood.

- When?
  May 9, 1975.

- Let's call the first significant organism from this blood culture ORGANISM-1. Do you know the identity of ORGANISM-1?
  No.

- Is ORGANISM-1 a rod or a coccus or something else?
  Rod.

- What is the gram stain of ORGANISM-1?
  Gramnegative.

- Have there been positive cultures from sites which are not associated with infections about which you are seeking advice?
  No.

- Did ORGANISM-1 grow aerobically?
  Yes.

- Did ORGANISM-1 grow anaerobically?
  Yes.

- What is the suspected portal of entry of ORGANISM-1?
  The gastrointestinal tract.

- My first recommendation is as follows: In order to cover for items 1, 2, 3, 5, and 6, give Gentamycin using a dose of 119 mg (1.7 mg/kg) q8h IV [or IM] for 10 days. Modify dose in renal failure. And in order to cover for item 4, give Clindamycin using a dose of 595 mg (8.5 mg/kg) q6h IV [or IM] for 14 days. If diarrhea or other gastro-intestinal symptoms develop, check for pseudomembranous colitis.

- Do you wish to see the second choice?

Figure 1 Sample dialogue with the medical expert system MYCIN (adapted from Winston 1977).

MYCIN (Buchanan and Shortliffe 1984). MYCIN acts as a medical expert for the diagnosis and therapy of infectious diseases; the dialogue reproduced here is slightly modified from Winston (1977).

In order to perform these tasks, an expert system consists of five subsystems, as shown in Figure 2. The *knowledge base* contains the knowledge of the expert(s) plus the facts provided by the user during the interaction with the system, while the *inference engine* infers or deduces knowledge that is not explicitly contained in the knowledge base. The *user interface* formulates recommendations and questions to the user and, if it is an intelligent interface, will also understand rather more than simple answers like "Yes" or "No" or numbers. In an ideal case it will be a natural language interface (but more about this in later sections). The *explanation part*, which explains to the user how the system arrived at a specific conclusion, has proved to be of special importance: it not only increases acceptance by the user but also helps identify and correct logical errors or simple bugs in the system. In most present-day expert systems the knowledge of the expert is first obtained through a dialogue with a *knowledge acquisition module*, then structured, and finally loaded into the computer.

Several examples of successful expert systems are mentioned later in Chapter 3.

Figure 2 The five main components of an expert system.

3.2 Natural language systems

Natural language systems understand, process, or produce natural language, such as English, as opposed to "artificial" or "technical" languages, like programming languages. Natural language systems should not be confused with speech understanding systems, which remain in their infancy.

Important areas of application of present-day natural language systems include the following:

- *Machine translation*. The first translation programs performed only word-to-word translation. This led to very curious examples like "water-goat" for "hydraulic ram" or "invisible idiot" for "out of sight, out of mind". Present-day translation programs consider not only the syntax, but also the meaning, the semantics, of a text. However, it is still not possible to make consistently acceptable translations, so a two-stage procedure, involving machine translation followed by editing, is used.
- *Understanding of texts*, e.g. executing orders, making summaries, or drawing conclusions. In order to perform such tasks, programs must know something about human beliefs, goals, and plans. A good example is CYRUS, developed by Schank, which searched UPI wires for information about the then US Secretary of State, Cyrus Vance. When asked if the wife of Mr. Vance ever met the wife of the then Israeli Prime Minister, Menachem Begin, it concluded – correctly – from the fact that the two men had on several occasions attended the same state banquets that their wives had very likely accompanied them and thus met (Alexander 1982).
- *Generation of texts*, e.g. for product descriptions or user manuals that can be adapted to the specific needs of a user group. It is also possible to generate (simple) fairy tales (Haase-Rapoport and Pospelov 1984). It may sound very easy to generate such texts, but it is often very difficult to make them acceptable. When do you end a sentence? How often should you repeat a person's name and when is it better to use just the personal pronoun? And so on.
- *Natural language interfaces*. These are used in data-base systems, question-answering systems, and expert systems, and they can greatly increase the degree of acceptance of all three. The best-known natural language interface, INTELLECT^{TM} from Artificial Intelligence Corporation, has been installed in more than 300 systems. Interfaces that develop a model of the user and therefore adapt to his or her personal needs, are under development.

3.3 Robotics

Present-day robots don't look much like the two on the left-hand side of Figure 3; they're usually much more prosaic, like the industrial robot shown on the right-hand side of the figure.

Figure 3 Robots: conception and reality. On the left are R2D2 and C3P0 from *Star Wars* while the picture on the right shows a real industrial robot from Zeppelin–Hitachi.

Robots used in industry are manipulators that cannot move as freely as men, and neither can they sense, reason, or plan. They consist in principle of two components: a moving part with arm, wrist, and hand, and a control system. Robots are programmed, either by keying in a program or by guiding the arm manually along a specific track and then instructing the robot to repeat this track. The control system may also be directly connected to a computer.

While robots are not yet human-like, they already can and do outperform men in many respects (Engelberger 1980): they can work in hostile environments, where noise, vibration, smells, or danger would have an adverse effect on the human system. The use of hydraulic or electric power gives a robot more potential "muscle" than a man. And they don't stop to eat, they never want to go home, and they don't go on strike.

4 AI Methods

If you are interested in how AI programs function and what the underlying principles are, this section will give you a short introduction to a few important and fascinating AI methods. If on the other hand this doesn't interest you, just skip this section and reduce the "one-hour course" to thirty minutes.

4.1 Knowledge representation

In order to act "intelligently," you have to know something. One common form of representation involves *semantic networks*. Let's suppose that we want to represent the fact that "all robins are birds" in a semantic network (modified from Barr and Feigenbaum 1981). We can do this by creating a simple graph in which the *nodes* (here ovals) represent the objects and the links the "is-a" relation between them:

If "Clyde" is a particular robin, we can easily express this fact by adding a node and a link to our graph:

If we want to add our knowledge that birds (or at least most of them) have wings, we add a node for "wings" and introduce another type of link, namely a "has-part" relation:

This representation enables us to deduce facts that are not explicitly stated in the network, e.g. that robins have wings since they are birds, and so does Clyde since he is a robin. This feature is called *property inheritance*.

If we now wanted to express the fact that "Clyde owns a nest" and tried to represent "own" by another link, we would run into trouble if we built in the very plausible condition that Clyde owns his nest only from the spring to the fall. Since we have at our disposal only binary relations, i.e. only links between pairs of nodes, we couldn't connect all these facts. We solve this problem by realising that "ownership" is also a fact that we can express by a node, and the fact that Clyde owns something is an *instance* of this general fact. Also, the specific nest belonging to Clyde is an instance of the general "nest." Our graph now looks as follows:

We also have expressed the fact that "ownership" is a "situation." Since "own" implies both an owner and an ownee, the next time the information that something is owned is input, the system can automatically search for someone (an animate object, person, or legal entity) who is the corresponding owner. This sort of procedure could also, for example, be of great help in sentence analysis.

Knowledge can also be represented as *rules*. The knowledge in the medical expert system we saw in the preceding section is stored in the form of rules. Such rules could look like this:

* *If* stain is grampositive *then* organism is streptococcus.
* *If* stain is gramnegative *then* organism is E. coli.
* *If* organism is streptococcus *or* bacteroid *then* penicillin is indicated.
* *If* drug is indicated *and* don't know whether allergic to drug *then* ask whether allergic to drug.

```
Generic RESTAURANT Frame

    Specialization-of:  Business-Establishment
    Types:
          range:        (Cafeteria, Seat-Yourself, Wait-to-be-Seated)
          default:      Wait-to-be-Seated
          if-needed:    IF plastic-orange-counter THEN Fast-Food,
                        IF stack-of-trays THEN Cafeteria,
                        IF wait-for-waitress-sign or reservations-made
                            THEN Wait-To-Be-Seated,
                        OTHERWISE Seat-Yourself.
    Location:
          range:        an ADDRESS
          if-needed:    (Look at the MENU)
    Name:
          if-needed:    (Look at the MENU)
    Food-Style:
          range:        (Burgers, Chinese, American, Seafood, French)
          default:      American
          if-added:     (Update Alternatives of Restaurant)
    Times-of-Operation:
          range:        a Time-of-Day
          default:      open evenings except Mondays
    Payment-Form:
          range:        (Cash, CreditCard, Check, Washing-Dishes Script)
    Event-Sequence:
          default:      Eat-at-Restaurant Script
    Alternatives:
          range:        all restaurants with same FoodStyle
          if-needed:    (Find all Restaurants with the same FoodStyle)
```

Figure 4 The Restaurant frame (Barr and Feigenbaum 1981).

Rules consist of first an *antecedent*, i.e. a condition, a situation, or a premise, and second a *consequent*, which can be an action or a conclusion. The antecedent may consist of several conditions logically connected by "and(s)" and/or "or(s)."

Frames (Minsky 1975) are a special type of network structure used to represent complex knowledge. They feature *slots*, which are places reserved for special information. A well-known example is the Restaurant frame (Barr and Feigenbaum 1981) shown in Figure 4. The slots in this particular frame are Specialization-of, Types, Location, etc. "Specialization-of" allows for property inheritance, as in a semantic network. The slot "Types" gives a *range* of types to which this particular restaurant may belong. If this information is not requested, a *default value*, essentially the most common value, is assumed; if it is requested, procedures are given that help to elicit this information (*procedural attachment*).

And what about complex actions, such as eating at a restaurant, which would be a fairly predictable sequel to a restaurant frame? The representation of such a sequence of events is called a *script*, like a movie script. Our self-explanatory "Eat-at-Restaurant" script (Barr and Feigenbaum 1981) might look like Figure 5:

EAT-AT-RESTAURANT Script

Props:	(Restaurant, Money, Food, Menu, Tables, Chairs)
Roles:	(Hungry-Persons, Wait-Persons, Chef-Persons)
Point-of-View:	Hungry-Persons
Time-of-Occurrence:	
	(Times-of-Operation of Restaurant)
Place-of-Occurrence:	
	(Location of Restaurant)
Event-Sequence:	
first:	Enter-Restaurant Script
then:	if (Wait-To-Be-Seated-Sign or Reservations) then Get-Maitre-d's-Attention Script
then:	Please-Be-Seated Script
then:	Order-Food Script
then:	Eat-Food Script unless (Long-Wait) when Exit-Restaurant-Angry Script
then:	if (Food-Quality was better than Palatable) then Compliments-To-The-Chef Script
then:	Pay-For-It Script
finally:	Leave-Restaurant Script

Figure 5 Expansion of the "Eat-at-Restaurant" script from Figure 4 (Barr and Feigenbaum 1981).

4.2 Search

Search is a very important tool for problem solving, inference, planning, and reasoning. As an example, consider ways in which we might go about solving the "8-puzzle." The problem is to transform a configuration like this

2	8	3
1	6	4
7	■	5

into the solution

1	2	3
8	■	4
7	6	5

in which the empty field — denoted by a black square — is in the middle of the puzzle. Three moves are possible from our given starting position: we could move 7 to the right, 6 down, or 5 to the left. Since none of the resulting situations is the solution, we must then examine every possible move from each of the three intermediate positions. We can represent the problem by a *search tree*, which is *expanded* from the given starting position. If we expand the tree to its full breadth on each level and check to see if the solution is contained in that level, this is known as a *breadth-first search*. The fully expanded tree is shown in Figure 6 (from Nilsson 1971).

It is quite clear from Figure 6 that we need to develop some means of reducing the seach efforts involved; otherwise it would, for example, be impossible to program even a theoretically simple end-game in chess.

Figure 6 Fully expanded breadth-first search tree for the 8-puzzle (Nilsson 1971).

This reduction can be performed by evaluating each intermediate position in turn. Since it is rarely possible to specify a precise metric, we simply use *rules of thumb* to estimate what looks promising in the search for the solution and what does not. For example, the more levels the tree already has, the "worse" is the situation. And the more elements that are misplaced in any given position, i.e. not where they should be in the solution, the worse the position is. We simply add up the number of levels already developed and the numbers of elements misplaced. Before developing our tree to the next level, we evaluate each position in this way and then expand only the one that has the lowest value of the *evaluation function* we have developed. This procedure is continued on each successive level until the solution is obtained. The tree produced by this type of search is depicted in Figure 7 (from Nilsson 1971).

Figure 7 Simplified search tree for the 8-puzzle, resulting from the application of a simple evaluation function (Nilsson 1971).

We see that this new tree is far simpler than the one shown in Figure 6. This reduction results from the *heuristics* of the evaluation function. Had we defined a more restrictive evaluation function, then the tree might have been even smaller, but we would have run the risk of expanding it in the wrong direction. Had we defined the function less restrictively, the tree would inevitably have become larger.

4.3 Reasoning and planning

We have now seen several ways of representing knowledge. These representations are the basis for reasoning and planning.

Suppose we want to answer the question "What does Clyde own?" using our knowledge representation in the form of the semantic network developed earlier. The first step will consist in restructuring the question into a *network fragment*:

In the second step we look for a similar structure (of nodes and links) in our semantic network (which in a "life-representation" will be considerably larger!). This process is performed by a *pattern matcher*. As soon as the pattern matcher has successfully performed its task, we find, at the location of the question mark, the answer "NEST-1."

If our program had to reply to a more complex question like "Is there a bird who owns a nest?" the first step, again, would be restructuring the question into the following network fragment:

In this case we will not find an identical pattern in our earlier semantic network since CLYDE is not directly connected to BIRD by an "is-a"-relation but only via the intermediate ROBIN node. The pattern matcher thus has to make *inferences* on the basis of the semantic network, which in this example means deducing from the two "is-a"-relations that Clyde is a bird, as we did in Section 4.1.

How can we reason using *rule-based knowledge representation*? Let's have another look at the rules introduced in Section 4.1:

- *If* stain is grampositive *then* organism is streptococcus.
- *If* stain is gramnegative *then* organism is E. coli.
- *If* organism is streptococcus *or* bacteroid *then* penicillin is indicated.
- *If* drug is indicated *and* don't know whether allergic to drug *then* ask whether allergic to drug.

In proceeding toward a therapy recommendation, we may start from the fact that the stain of an unidentified organism is grampositive (whatever that means). We then match this fact with all the antecedents until we find a rule in which this condition is fulfilled. Thus we can deduce "organism is streptococcus" as a new fact.

We then compare this new fact with all the antecedents until we find one in which this condition holds, and so on. Since the true antecedents—consequents—antecedents— follow each other like the links of a chain, this reasoning process is called *forward chaining*. This is, in principle, one way in which expert systems reason. But since expert systems often consist of several thousands of rules, additional methods are needed to derive a subset of all the rules that is still able to select the important facts.

Alternatively, we could start from a particular therapy, by asking e.g. "When would you prescribe penicillin?" Here the program has to search the rules "backward," from consequent to antecedent, and so on; this method is therefore called *backward reasoning*.

Chaining methods are also used for *planning*, in which we start by defining an initial state and a goal state. In rare cases it may be possible to go directly from the initial state to the goal state, but in such a simple situation we clearly wouldn't need planning. Usually *operators* are available that enable one state to be transformed into another (nearby) state, thus producing a series of *intermediate states*. Operators are selected with respect to reducing the distance to the initial state or to the goal state. Two modes of operation are possible: either from the initial state to the goal state by forward chaining, or *vice versa*. The decision on how to proceed will depend on the particular *state space* concerned (we have already encountered

one state space when considering the 8-puzzle). Obviously, we try to proceed in directions where the chance of encountering *dead-ends*, i.e. states that are not the goal and from which proceeding to the goal is impossible, is smaller. Dead-ends make it necessary to *backtrack* to former intermediate states and then to select another path. In the 8-puzzle example we used forward chaining. On the basis of the principles described here, several programs for planning have been developed.

4.4 Learning

Suppose that we want the computer to learn to "understand" the concept of an arch:

I am aware that this arch lacks something of the beauty of the arches of the Alhambra. It is, however, frequently necessary to reduce complex objects from the real world into simpler examples: an environment sometimes used in AI is the *blocks world* representation seen here.

One method of teaching the concept of an arch involves first presenting to the computer an accurate realization of the concept, followed by a number of *near misses* – realizations that each deviate in one particular aspect from an arch – (Winston 1975), in a "didactic" sequence. This sequence of presentations enables the program to improve its representation in a stepwise fashion (modified from Winston 1977).

We assume that the program has prior knowledge of the fact that an arch consists of three elements (X, Y, Z), all of which are blocks:

Furthermore, we assume the program knows various categories for which it must look, such as support, touch, and orientation, which it can then add to its representation by corresponding links. If we now present to the program this near-miss representation:

and inform the program that this is not an arch, it will realize from comparison between this and the original that in an arch two blocks are standing (vertical) and one is lying (horizontal), and improve its representation accordingly:

If we now present the near-miss:

and inform the program that this is also not an arch, it recognizes the difference related to "support" and adds two more links to its representation:

Finally, if we present the near-miss:

it will realize that, in an arch, the two standing blocks must not touch. We thus arrive at the final representation of an arch:

which would correspond to the arch originally depicted above.

Improved learning programs have succeeded, for example, in learning the descriptions of dozens of soybean diseases from several hundred previously identified samples (Michalski *et al.* 1983). These descriptions were in no case worse than those supplied by human "control" experts, and in some cases they were actually better (Winston 1984). Since soybeans are a primary nutritional factor in several countries, this represents important progress. The AM program of Lenat (Davis and Lenat 1982) has even discovered interesting new concepts in mathematics.

5 AI Software and Dedicated Machines

Most AI programs are written in the programming language *LISP* (from list processing). Actually, LISP was one of the first high-level programming languages like FORTRAN or COBOL; it was developed by John McCarthy and his colleagues in the late fifties. There is a very

simple way of recognizing if a program is written in LISP: many lines in a LISP program will start with a parenthesis "(".

This language is (almost) ideally suited to expressing and manipulating the concepts of AI, i.e. for symbolic processing. For example, examine how the definition of an arch in the world of blocks is expressed in LISP:

```
(ARCH (PARTS X Y Z)
        (X IS-A BLOCK)
        (Y IS-A BLOCK)
        (Z IS-A BLOCK)
        (X HAS-ORIENTATION STANDING)
        (Y HAS-ORIENTATION STANDING)
        (Z HAS-ORIENTATION LYING)
        (X SUPPORTS Z)
        (Y SUPPORTS Z)
        (X MUST-NOT-TOUCH Y)
        (Y MUST-NOT-TOUCH X))
```

Another example of the representation of knowledge in LISP is as follows:

```
(ARAI (A-KIND-OF RESEARCH-INSTITUTE)
        (FULLNAME AUSTRIAN_RESEARCH_INSTITUTE_
        FOR_ARTIFICIAL_INTELLIGENCE)
        (PART-OF AUSTRIAN_SOCIETY_FOR_
        CYBERNETIC_STUDIES)
        (LOCATION (SCHOTTENGASSE_3 VIENNA_1 AUSTRIA)
        (PHONE 6632810)
        (ACTIVITIES (TEACHING
                        CONSULTING
                        RESEARCH))
        (RESEARCH_AREAS (EXPERT_SYSTEMS
                            NATURAL_LANGUAGE_SYSTEMS
                            AI_TOOLS
                            AI_IMPACTS)))
```

The same structures hold for rules like the examples in our expert system:

```
(RULE12
        (IF (SAME STAIN CONTEXT GRAMPOS)
        (THEN (CONCLUDE ORGANISM CONTEXT STREP 1.0)))
```

The main purpose of LISP is not "number crunching" but evaluating symbolic expressions. Since there is no difference between program and data, a LISP program can regard another program or even itself as a symbolic expression and evaluate it accordingly. Unfortunately, there is not just one LISP language, but rather several "dialects" xLISP (where x = INTER, MAC, FRANZ, ZETA, etc.). This means that LISP programs are not very easily portable. However, there is a good chance that COMMON LISP may become widely accepted as a standard.

Recently, another language called *PROLOG* (from programming in logic), developed by Colmerauer and Kowalski, has found adherents in Europe and Japan. PROLOG received a special boost when it was decided to use it as the principal language of Japan's Fifth Generation Computer Systems Project. A good introduction is given by Clocksin and Mellish (1981). A PROLOG statement of the fact that John is the father of Mary would read

    Father (John, Mary)

In order to pose the question "Who is the father of Mary?" the input would have to be

    ? - father (X,Mary)

to which the system would reply

    X = John

It is also possible to define conditions, such as "X is the parent of Y if X is the father of Y," which in PROLOG would read

    parent (X,Y) :- father (X,Y)

Thus, PROLOG can clearly be used to answer more complicated questions. PROLOG is very useful if the problem can be cast into logical forms of the type described above: first, state facts and conditions in logical forms and then prove Z by showing that A and B and C and D, and so on, are fulfilled.

The pros and cons of LISP and PROLOG have been extensively discussed, e.g. in O'Keefe (1982) or Foster (1984). There have already been some attempts to combine the advantages of both languages, for example in LOGLisp by Robinson (Harris 1983) or POPLOG (Sloman *et al.* 1983). One definite advantage of LISP is that there are, at present, far more tools available that can simplify life considerably when

developing a large program. These cover not only software but also hardware tools: *LISP machines* have been developed by several companies (Symbolics, Lisp Machines Inc., Xerox, Texas Instruments, etc.), and these offer extremely fast execution of LISP programs in comparison with normal computers together with a wide variety of software features (*environments*). Some of these machines also support the execution of PROLOG programs. Contrary to earlier expectations (see e.g. Alexander 1982), the price of many of these machines has not dropped much from the initial bracket of US $100,000. Only recently have machines in the US $60,000 range been announced (Texas Instruments). Since all of these machines (with one exception) are single-user computers, between 60,000 and 100,000 dollars must be spent to provide computing facilities for each AI researcher. An excellent overview of the options and factors to be considered when planning AI facilities can be found in Fahlman and Steele (1982). But the costs of a working place in an AI research department are still considerably lower than in other high-tech development areas.

6 Future Trends in AI

Expert systems will become larger, and natural language systems will approximate natural language *and* speech, both for input and output. Automatic programming will make the need to actually learn a programming language obsolete for most people. Robots will sense their environments, and in particular hear, see, feel, and move around freely. They will not only be able to play chess; they will also be able to find the chess board!

Will these tasks be achieved just by using "more of the same" or will it be necessary to develop qualitatively different approaches? Take for example the medical expert system: it knows that "If organism is streptococcus or bacteroid then penicillin is indicated" but it does not know how penicillin interferes with the metabolism of a bacterium. Therefore, will its present *surface knowledge* suffice, for example to recognize where its knowledge ends and where it may be safer not to give any recommendation? Many scientists such as Hart (1982) argue that expert systems should also have knowledge of the underlying structure: it is not accidental that physicians are first trained in anatomy and physiology before they start diagnosing and treating disease. We will thus have to incorporate a vast additional amount of structural knowledge into *deep systems*, and as yet it is quite uncertain whether our present methods of representation and inference will prove to be adequate.

At present, AI systems have to be told what to do, but it is a considerable waste of time and energy that a knowledge engineer has to structure the knowledge of an expert and that a programmer has to expand the vocabulary of a natural language system. Therefore a core research area within AI (some scientists even say *the* area) will be *learning*. Section 4.5 presented a simple learning program based on comparing samples; however, we need to develop programs that *learn by discovery*, by interacting with the environment on their own initiative. A few steps have already been taken in this direction (e.g. Michalski *et al.* 1982, Michalski 1983), but the majority of the work still lies before us.

Compared with human beings, AI programs have a very short and limited existence. They cannot collect information as freely as we can and they operate for only a few hours a day. Richard Weyrauch has proposed a new class of AI programs with a continuous existence, so-called *computer individuals*. These programs should never be turned off, and should be able to develop a model of the world and of their users that could undergo modifications in the same way as human beings learn from experience. Their interaction with the world would, at first, consist of dialogues in natural language, but in later stages other sensory information could be added. Such programs would force us to combine and extend our present knowledge in AI tremendously; they would also be the forerunners of Neil Frude's "intimate machines" (1983).

Both the increase in the size of the systems and qualitative improvements made to them will considerably increase the number of computations required to perform a task. In order that the patience of the user should not be overstrained by excessively long response times, the processors in the computers will have to become faster or a new technology will have to be employed. Both alternatives have adherents, but the second, though as yet only operational in experimental settings (e.g. Stolfo *et al.* 1983), is of greater interest. The main contender here is *parallel processing*: instead of performing all the computations sequentially using a single processor, as practically all present computers (von Neumann machines) do, parallel processing would split the procedures in such a way that they could be performed in parallel by many, perhaps thousands of processors in so-called "non-von" computers.

If, for example, we want to answer the question "Is there an animal which has wings and owns a nest?" we have to make a great many comparisons in a large semantic network. It is, however, not strictly necessary to make one comparison after the other; most of the comparisons could be made at the same time, in parallel, and this would considerably decrease response time. Clearly, in many cases it will be extremely difficult to adapt our present procedures to this new

technology. We will therefore have to develop a theory of *cooperative processing*. It is interesting to note that it is considerably easier to adapt PROLOG to parallel processing than LISP; this is one of the reasons why Japanese scientists chose it as the language for their Fifth Generation Project.

Finally, we still tend to think of AI as something remote, resident in some distant computer. However, with the advent of the personal or even the home computer, those formerly distant machines have already started to enter our lives far more closely. We are usually unaware that we are using a tiny computer, a microprocessor, when we program our videotape recorder or step on a gas pedal in a car with fuel injection. Since LISP machines will very soon be squeezed onto a single silicon chip, AI programs will enter our everyday life in the near future.

7 Sources of Information about AI

An excellent introduction to AI is provided by the book by Winston (1984), while other very good introductory texts are those by Raphael (1976), Nilsson (1980), and Rich (1983). If you are more interested in a nontechnical overview from the commercial point of view with detailed descriptions of applications, Winston and Prendergast (1984) is the best bet. If your interest is more philosophical, Boden (1977) is excellent reading. A book about the development of AI containing lots of interviews with eminent AI researchers is McCorduck (1982). Hofstadter (1979) is the first (disguised) introduction to AI, and it won a Pulitzer Prize. The best in-depth presentation is given by the three-volume *Handbook of Artificial Intelligence* (Barr and Feigenbaum 1981, 1982, Cohen and Feigenbaum 1982). Furthermore, there are the "traditional" AI journals *Artificial Intelligence* and *Cognitive Science*.

An excellent overview of the latest research is provided by the biennial (in odd years) International Joint Conference on Artificial Intelligence (IJCAI) and the biennial (in even years) European Conference on Artificial Intelligence. The Conference of the American Association for Artificial Intelligence (AAAI) is held annually except in those years when the IJCAI takes place in the United States. The proceedings of these conferences present a wealth of information on the state of the art.

General-interest articles, conference reports, titles of recent papers, etc., can be found in newsletters of various national or multinational organizations, such as the *AI Magazine* of the AAAI, the *SIGART Newsletter* of the Association for Computing Machinery (ACM), the *AISB Quarterly* of the Society for the Study of Artificial Intelligence and

the Simulation of Behaviour, and the *ECCAI Newsletter* of the European Coordinating Committee for Artificial Intelligence. More newsletters are published by several other AI societies. These societies also frequently offer introductory seminars and short courses on AI. Very probably your nearest university already offers courses or even a full curriculum in artificial intelligence.

If you want to start right now with AI programming on your personal computer, there is a book containing simple AI programs in BASIC (Krutch 1981), but you can even build your own expert system in BASIC with the help of a very funny book by Naylor (1983). However, if you are more seriously interested in AI, you should learn LISP, which will also give you the flavor of the subject; the best introduction here is the book by Winston and Horn (1984). There are already several LISP interpreters available for personal computers, e.g. TLC-LISP^{TM} from the LISP Company for Z80 systems or IQLISP^{TM} from Integral Quality for the IBM PC^{TM}. A good introduction to LISP is Golden COMMON LISP^{TM} from Gold Hill Computers, since it presents the probable coming standard COMMON LISP on a micro and also offers a tutorial. LISP from ExperTelligence should also be available soon for Apple's Macintosh^{TM}.

In any event, if this chapter has whetted your appetite for AI and you want further information, just contact me and I will try to help.

References

The references cited in this chapter can be found in the Artificial Intelligence Bibliography, List of References, pp. 221–251 of this volume.

CHAPTER 3

Impacts of Artificial Intelligence: An Overview

Robert Trappl *

1 Summary

In this chapter I introduce and discuss the probable impacts of AI in the areas of the sciences, technology, military applications, economics, society, culture, and politics. In my opinion, the eight most important impacts are as follows:

1. Less and different work, through machine substitution for physical and mental labor.
2. New industries, parts-on-demand factories, trading of "knowledge" rather than just "information" (knowledge economics).
3. Widening or narrowing of the North–South gap.
4. More "efficient," but also more vulnerable weapons systems.
5. "Improved" computer surveillance.
6. Computational paradigm penetrating most of the biological and social sciences where information processing occurs.
7. Computers/robots with "individuality."
8. Drastic changes in human self-understanding.

The extent and directions of these and other impacts will depend on several factors, some of which can be influenced by you, the reader (see Conclusions).

*Department of Medical Cybernetics and Artificial Intelligence, University of Vienna, Freyung 6, A-1010 Vienna, Austria.

2 Impacts in Different Areas

2.1 The sciences

Within the next 50 years, AI may well come to be viewed as a type of metascience, like mathematics; that is, AI will be used as a tool directed toward problems in numerous scientific and technical domains (Schank and Slade, Chapter 10).*

AI has already introduced a "qualitative" trend into sciences: while the computer up to now has encouraged quantitative aspects, namely numerical calculations, AI promotes the procedural aspects. A good example is naive physics: we all know that when a glass of water is moved to the corner of a table it will eventually fall off, but we also know that a glass cannot fall through the table, unless the table has a hole in it. Conventional physics never took the trouble to formalize these relations, because they were felt to be "self evident"; but a program that has to handle solid objects, perhaps sometimes even objects filled with liquids, must have procedural knowledge about this "naive physics." However, this view is not always well received, as can be seen by remarks about "computational imperialism" (Turkle 1984).

The relation between *philosophy* and AI has been treated in several books (e.g. Ringle 1979). AI has been viewed as applied epistemology (Havel, Chapter 7), while the construction of a semantic net relies heavily on ontology, or even defines one (Trappl *et al.* 1982).

If AI can be regarded as a modern methodological tool being used in the ancient enterprise of the study of mind (Schank 1983), then it should have a significant impact on *psychology*. Schank and several of his colleagues have focused attention on many facets of cognition, including language ability, learning, memory, motivation, emotions, interpersonal relations, beliefs, and planning (Schank and Abelson 1977, Schank 1982). Minsky (1982) gives a whole list of new concepts that are familiar to AI programmers and cognitive psychologists, but that have no concrete technical counterparts in traditional psychology or philosophy.

Does all this have a humanizing effect on psychology? Boden's reply to this question would be "Yes!". "Mind" and "mental processes" have once again become respectable. AI has had a considerable influence on the refutation of behaviorism (Boden, Chapter 5). The "Self" has entered into the focus of scientific interest (Havel, Chapter 7). The issue of free will has taken on a new flavor, together with the concepts of self-determined action, authentic choice, and intentionality

---

*Editor's note: Wherever authors are quoted with a chapter number but without a date, I refer to the position papers that comprise Chapters 4–11 of this volume.

(Havel, Chapter 7). However, the participants in the Task Force Meeting, including myself, did not leave Havel's ideas undisputed, especially since it was unclear what exactly was meant by "free" will. It is interesting to note that Turkle (1984) also refers frequently to free will and people; she concludes, essentially, that many people don't really grasp the meaning of this concept, a fact that Turkle deeply regrets. Unfortunately, she never informs the reader exactly which interpretation *she* gives to "free will."

Since learning and memory organization will very probably be the core areas of AI in the next decade, the results of this research should have a marked impact on psychology. The scientific rather than the technological aspects are expected to be of the greatest benefit to mankind: How do people learn? Make decisions? Think? (Schank and Slade, Chapter 10).

However, one could also argue in the opposite direction: as long as man tried to build airplanes by imitating the solutions nature had found, i.e. wings, the attempts were unsuccessful. Only when a solution appropriate to the material used was developed did airplanes become viable. The same may well be true for computers: as long as we try to build programs that more or less simulate human mental processes, our AI programs may not be very successful. Only when we enable programs to learn by themselves and thus improve themselves will they begin to achieve their true potential.

There has been a similar controversy in *linguistics*, as to whether or not AI research has important implications for the central questions of that discipline. Dresher and Hornstein (1976) thought the answer was "No," while Winograd, in a reply published the following year, said "Yes." The dispute was mediated by Lakoff (1978), who felt that AI "provides a formal framework for investigating interesting problems," with special reference to scripts. Wilks (1981) again stressed the important models that AI presents in the field of natural language understanding.

The computational paradigm (Boden, Chapter 5) has influenced not only psychology but also *biology*: the functional or developmental significance of a given metabolite may vary widely across species or even within an individual at different stages of the life-cycle. So the appropriate question might not be "what is the molecular biology of this substance?" but rather "what computation is this substance performing for the organism?" (Boden 1981).

Turning to *medicine*, there is already a relatively long tradition of modeling medical expertise in expert systems. Among the first to appear in the early seventies were expert systems for infectious-disease diagnosis and therapy selection (MYCIN), for glaucoma (CASNET), and for internal medicine (INTERNIST). Since then,

medicine has been one of the main areas of application of AI; the American Association for Artificial Intelligence has even established a special group for "Artificial Intelligence in Medicine." An excellent summary of the main papers in this field was edited recently by Clancey and Shortliffe (1984).

But how well do such programs perform? Four systems, the three mentioned above plus PUFF (pulmonary function) have been evaluated in statistical studies (Duda and Shortliffe 1983). About 80% of their diagnostic behavior was accepted by the evaluators; and this percentage corresponded roughly to the extent of agreement among the evaluators themselves.

Are any of these programs already in regular use? The answer here is "Yes" for PUFF, EXPERT (applied to serum protein electrophoresis), and ONCOCIN (cancer chemotherapy). And it is very likely that the number of medical expert systems available will increase very rapidly: since these tools were first developed, the time and effort needed to create an expert system has been drastically reduced from about 45 man-years to less than ten (Davis 1984).

But will they be accepted by the medical establishment? Pople, who developed both INTERNIST and its successor, CADUCAEUS, is very positive: "Physicians are quick adopters of new technologies that can be demonstrated to be of value to them" (discussion, in Winston and Prendergast 1984, p. 97). It is interesting to note that EXPERT (electrophoresis) was originally developed on a large mainframe, but then transferred to a microcomputer. It is thus very likely that, with the continued increase in the power of micros and progress in the development of small expert systems (expert system "shells" have already been implemented on micros, see e.g. Widmer 1984, Widmer and Horn 1984), expert systems will also become available to the general practitioner. Furthermore, pattern analysis and vision will help in the analysis of images in medicine, e.g. ECG, EEG, X-ray, tomography, and ultrasound pictures (Trappl *et al.* 1983).

I very much doubt that physicians in hospitals or general practitioners will be substituted by computers, certainly not in the foreseeable future and perhaps never. Nevertheless, I feel that their influence, and perhaps also their number, may decrease, as a result of patients (and healthy persons) using medical expert systems on their own personal computers. Medical self-care is on the upsurge, as a result of both the increasing costs of and a widespread distrust in established medicine; and self-care books are bestsellers (1.5 million copies of one of these books, by Vickery and Fries, have already been sold). These books are inherently stupid, as compared to expert systems, since they always have to start from one primary symptom, and then cannot develop their analysis very deeply. At present the only

medical systems available for personal computers are simple question-answering programs, but I predict that the situation will be totally different five years from now (Trappl 1985).

2.2 Technology

The improvement of *man–machine interactions* in order to increase the satisfaction or at any rate decrease the boredom or fatigue people experience when working with a computer is an important task for AI. Natural language interfaces (NLI) enable people to interact with computers using their everyday language (Trost and Trappl 1985). As soon as prestel or videotex(t) systems are available on a wider basis, people will neither want to use boring search-trees nor wish to learn a specific query language to extract information from a data base. For the English language, INTELLECT^{TM}, developed by Artificial Intelligence Corporation, is the first NLI and is already widely used. As far as I am aware, at present there exists no other NLI comparable in performance to INTELLECT, and I don't know of NLIs in languages other than English, though many institutions are working hard in this direction, for example, the Austrian Research Institute for Artificial Intelligence.

A potential drawback of verbal interaction may lie in the necessary limitations of such a system: while it might be acceptable to use only simple sentences when retrieving information from a data bank, more general application of question/answer systems would prevent people from using the full richness and subtlety of a living language. Straightforward price considerations will prohibit the development of programs that incorporate all the possibilities of words and sentences that a person reasonably fluent in English could easily grasp.

There are two other ways to make interaction more user-friendly. Keying in information could be replaced by speaking to a computer (Ballantine 1980). And, in the long run, interfaces could be written for each individual, personalized to fit the needs of an individual tailor, musician, secretary, etc., and stored on a pocket-sized "smart" card. The user would then only have to plug the card into any terminal and it could respond to the user's individual profile (Norman, in Joyce and Wingerson 1983).

CAD and *CAM* (computer aided design and computer aided manufacturing) are important applications of AI to the construction and production of goods (Teicholz 1984, Uttal 1984). An interesting aspect is VLSI* design, since this implies that the computer can help construct

---

*Very large scale integration.

an improvement to itself (Newell *et al.* 1983, Teresko 1983). However, CAD could also be used not just for specification but for suggestion, in the sense that the design program could originate novel ideas – ideas that are not merely quantitatively different from previous specifications (Boden, Chapter 5).

AI will continue to change *automation*: the incorporation of AI in the key points of production processes contributes to increased productivity, enhanced production quality, improved working environments, energy savings, and the suppression of labor-cost increases (Nagao, Chapter 8). Furthermore, AI could make possible the efficient production of a variety of objects in small quantities ("parts-on-demand facilities"), create varieties of new products, and enhance the function and reliability of existing products (Nagao, Chapter 8). The largest benefits will be for small and medium sized companies (Nagao, Chapter 8).

The most important contribution of AI to automation is in *robotics*. There is at present some disappointment with robots, largely due to inflated expectations, but the next generation of free-moving robots, equipped with tactile sensors and especially with vision (Arbib *et al.* 1983, Braggins 1983), should meet the needs (Engelberger 1980, Albus 1981).

2.3 Military applications

An increasing number of AI researchers are working on *military* projects. Very little is known about military applications in the Soviet Union, although it is thought that Soviet scientists are working to detect automatically the "scarring" or surface disturbance left by the underwater motion of submarines (Garrison 1983). However, that is very probably not the only research project under way.

More is known about research in the United States, which is focussed mainly on speech recognition, vision, and expert systems. One project aims at the identification of ships by microphones placed in harbors. The Strategic Computer Program recently developed by the US Defense Advanced Research Projects Agency (DARPA) plans the development of autonomous land vehicles equipped with advanced vision and expert systems capabilities, a pilot's associate system to help a fighter pilot manage his aircraft's flight and weapons systems under battle conditions, and a naval battle management system that would forecast likely events, suggest different courses of action, develop detailed action plans, resolve conflicts between competing goals, and react to changing battle developments (Schatz and Verity 1984).

It is interesting to note that the use of such systems would make it easier for an opponent to predict the behavior of his antagonist (and to take counter-measures) than is presently the case with human decision makers, provided he knew the rules of the expert systems used *and* had the faster computers. A second danger lies in the fact that it is almost impossible to fully debug such complicated programs, and under battle conditions there would be too little time to rectify obviously incorrect actions. Charles Hoare, when receiving the ACM Turing Award in 1980 for his contributions to the definition and design of programming languages, gave an urgent warning in his Award Lecture with regard to the computer language ADA, developed on behalf of the US Department of Defense: "In this last resort, I appeal to you, representatives of the programming professions of the United States, and citizens concerned with the welfare and safety of your own country and of mankind: Do not allow this language in its present state to be used in applications where reliability is critical, i.e., nuclear power stations, cruise missiles, early warning systems, anti-ballistic missile defense systems" (Hoare 1981).

Ironically, the back cover of the journal in which Hoare's lecture is printed displays an advertisement from a software company claiming "exceptionally fail-safe military applications software," referring particularly to their use of ADA!

In recent years about 400 computer scientists have joined CPSR (Computer Professionals for Social Responsibility), which wants, among other goals, to challenge the overdependence of the US and Soviet military systems on unreliable, computer-based systems (Verity 1984).

2.4 Economics

Within the foreseeable future, expert systems will be able to carry out *management* at middle or lower levels, which frequently involves routine tasks in which decisions are based on a limited number of rules. However, whether these management echelons really will be replaced will depend on the almost unpredictable "dynamic conservatism" of organizations (Schon, cited in Pitt and Booth 1983). Several systems are already operational, e.g. XCON, developed by McDermott (1982) to help Digital Equipment Company (DEC) in the configuration of VAX computer systems from the point when an order is received. Since 1980 this has been in daily use, has analyzed nearly 20,000 unique orders, and is running with 95-98% accuracy (Kraft 1984).

The effects of XCON serve as a good illustration of the potential benefits of future systems:

1.   An increased order throughput rate.
2.   A reduction in costs and delay attributable to errors (which were previously inevitable).
3.   Better use of materials, since the system rapidly accepts and acts upon constraints on component availability.
4.   Better use of personnel, through redeployment of highly skilled senior technicians to the most difficult tasks that are unsuitable for XCON.
5.   Rapid transformation of component requirement data into effective purchasing and manufacturing information.

DEC has not published any figures on exactly how much the program has saved, but it must be substantial, since the company otherwise would not have become such a firm believer in AI.

In the meantime, DEC has developed another program named XSEL (McDermott 1982) that helps a salesperson to prepare the order for Digital's computer systems. One might imagine that such a clever program would not recommend parts that are at present out of stock or that, for some reason, currently have a low profit margin — things a good salesperson would also do. The annual savings to DEC are estimated to be more than $10 million (*Business Week*, July 9, 1984). Further systems are TAXADVISOR, developed by Dungan, and AUDITOR, developed by Michaelsen (Michaelsen and Michie 1983). The names of these systems speak for themselves.

But what about high-level managers? There are quite a number of decision support systems available for managers (e.g. Davis 1984), but there is as yet no AI program with a technique that could substitute for a manager. Current expert systems are inadequate for this task because they are monotonic, i.e. they cannot withdraw earlier conclusions in the light of new information. Since approximately 20% of the knowledge base of a manager is altered every year, nonmonotonicity is an important prerequisite for successful managerial activity (Lee 1983). However, several attempts have already been made to formalize this aspect (McDermott and Doyle 1980, McDermott 1982, Moore 1983), so that expert systems with nonmonotonic reasoning can be expected soon. Anyway, serious scientists with long experience in the business world such as Hertz (1983), predict that "As AI moves to commercial and industrial applications, the scientific and medical applications of recent years will seem like child's play."

The most important impact of AI on economics will be in the area of *employment* and *employment structure*. However, even among AI researchers, the opinions expressed differ widely: Vámos (Chapter 11) states that unemployment is an incurable disease when it affects poor countries; he feels it is a problem of economic strength, and economic

strength is mainly based on the level of technology and automation. Nagao (Chapter 8) writes that many more jobs will be created by the development of intelligent systems than will be lost by the introduction of such systems. These opinions were shared by the then British Secretary of State for Industry (Jenkin 1983).

Schank and Slade (Chapter 10) are of the opinion that, over the next ten years, it is unlikely that AI-related unemployment will be much of a concern. If we view the workings of the economy as a resource allocation problem, AI programs will most probably be utilized for applications where it would be too expensive or too dangerous to use people. In the long term, AI programs will take over tasks that are better suited to computers, allowing humans to devote their time to personal services.

Boden (Chapter 5) states that traditional manufacturing and clerical–administrative jobs will be decimated, while some new jobs for computer engineers and programmers will be created. Whether there will be enough new jobs to compensate for the loss of the old ones (as has always happened in the past, at least eventually) is however unclear, for AI can potentially apply to all jobs where personal human contact is not essential. New methods of work sharing and income distribution will have to be developed (with income not necessarily being closely linked to jobs).

Nilsson (Chapter 9) also refers to the diminishing need for human labor. The majority of jobs created by automation will require only low-skilled labor. However, the world demand for "consumables" can (and ought to) increase dramatically. Nilsson also poses the provocative question, "What's so bad about unemployment?", noting that Adam and Eve apparently quite enjoyed it! Nilsson recognizes two fears, an economic and a socio-psychological fear. The solution to the economic fear lies not in inventing unnecessary jobs, but in approaches such as those proposed by Albus (1983). Albus suggests the formation of a National Mutual Fund that would use credit from the Federal Reserve System to finance private investment in automated industries. This extra investment would earn profits that would be distributed by the NMF to the general public as dividends to stockholders, so that everyone would receive a substantial income. Albus notes that separating income from employment explicitly acknowledges "...that the primary goal of an economic system is not to create work, but to create and distribute wealth, i.e., goods and services that people want and need."

At least two important questions remain open: first, who will be replaced earliest by computers, and second, is it possible to make quantitative forecasts of the likely impacts?

It is not so difficult to guess *who* will be replaced by a computer; but it is much more difficult to predict *when* this will take place.

Expert systems, natural language systems, and robotics taken together are already successful AI applications; they will be the first to have significant impacts on employment. While robotics will displace blue-collar workers, expert systems and natural language systems will replace the middle range of white-collar workers. But here too, distinctions must be made. It will be easier to replace by a robot the worker who puts the spare tire into the trunk than the gardener who trims the bushes in a park. Since it will be some time before robots with good vision, hearing, and motor coordination are available, the cut in white-collar workers could take place earlier than that affecting blue-collar staff. Or, as Herbert Simon once put it in a discussion (1982), "It is easier to replace a professor by a computer than a dredger operator". Luckily, for the professors, most of them already have tenure (I am one of them)!

Making quantitative forecasts is very difficult, as Rumberger and Levin (1984) remark in their project report on "Forecasting the Impact of New Technologies on the Future Job Market"; regarding robots, they feel that it is not at all easy to project future employment needs for robot technicians, because no one knows how fast or how widely robots will be introduced into the economy; furthermore, as robot technologies become more fully automated, labor requirements for building robots will probably decline.

It has often been argued that new technologies will create at least as many new jobs as will be destroyed by the introduction of these technologies. A short glance at Table 1 shows that this is very probably not true. Many of the jobs in the fastest relative growth area are in the computer industry. However, the percentage of the total number of jobs that they account for is very small. Furthermore, the jobs that are growing fastest are in the lower earnings range. This is a trend that has already been evident in the United States for some time. Many of the new service jobs pay far less and require fewer skills than the blue-collar occupations that have been dwindling (see Figure 1). As a result, the number of middle-class workers is steadily shrinking (Greenwald 1984). Less than a quarter of the jobs in high-tech industries require any substantial knowledge of technology. According to a recent report from Sweden, the introduction of numerical control equipment into industry has reduced markedly the level of skills required from experienced operators (cited in Anderson 1984).

Leontief and Duchin (1983) made employment forecasts based on three different scenarios: no technological diffusion, moderate technological change, and rapid technological change. Their projections suggest that rapid diffusion of technology could eliminate 20 million jobs by the year 2000, this representing 11% of all the jobs that would exist in the absence of further technological diffusion. And it is very likely

Table 1 Changing patterns of employment in the United States, 1982–95 (after Rumberger and Levin 1984).

Occupation	Employ-ment 1982 ($\times 10^3$)	Employment growth 1982–95		Required education[a] (years)	Relative earnings[b] (%)
		($\times 10^3$)	(%)		
Fastest relative growth[c]					
Computer service technicians	55	53	97	13–15	129
Legal assistants	45	43	94	13–15	86
Computer systems analysts[d]	254	217	85	16	86
Computer programmers[d]	266	205	77	16	122
Computer operators	211	159	76	12	82
Office machine repairers	56	39	72	12	106
Physical therapy assistants	33	22	68	12	57
Electrical engineers[d]	320	208	65	16	167
Civil engineering technicians[d]	35	23	64	13–15	121
Per. EDP equipment operators	49	31	63	13–15	73
Total	1,324	1,000	76	–	126
Fastest absolute growth[e]					
Building custodians	2,828	779	28	<12	69
Cashiers	1,570	744	47	12	69
Secretaries	2,441	719	30	12	67
General clerks, office staff	2,916	685	24	12	52
Nurses, professional	1,312	642	49	13–15	90
Waiters and waitresses	1,665	552	34	12	39
Teachers, kindergarten/ elementary	1,366	511	37	16	110
Truck drivers	1,604	425	27	12	117
Nurses, aides, and orderlies	1,218	423	35	12	58
Total	19,268	6,186	32	–	70
All occupations	101,510	25,600	25	–	–

[a]The number of years of education completed by the majority of the workers employed in each occupation during the spring of 1980.
[b]The average weekly earnings during 1979 of workers in each occupation relative to the average weekly earnings of all workers.
[c]Based on the greatest percentage increase.
[d]High-tech occupations as defined by the Bureau of Labor Statistics.
[e]Based on the greatest increase in the number of new jobs.

Sources: G.T. Silvestri, J.M. Lukasiewicz, and M.E. Einstein, "Occupational Employment Projections Through 1995," *Monthly Labor Review*, 106 (November 1983), Table 1; calculations based on the 1980 Public Use Sample, US Bureau of the Census.

Figure 1 The relationship between increasing levels of automation and job skill requirements.

that a considerable percentage of the remaining jobs will require less knowledge than those that existed in 1985.

How reliable are all these forecasts? This is very hard to estimate because most of the parameters used to evaluate the impacts of new technologies involve guesswork to a greater or lesser extent.

2.5 Society

Turning to our *working life*, the hope that "computers will enrich our jobs, free us from disagreeable labor, and increase our productivity" (Lucky 1983) is overshadowed by the fact that the introduction of high technology has already displaced a large number of workers or degraded the jobs of many of them. There seem to be basic contradictions between improving the quality of working life and organizational effectiveness and introducing new technology within existing bureaucratic organizational frameworks. Williams (1983) has listed the main contradictions, which are reproduced in Table 2. Williams proposes a participative sociotechnical redesign, which requires first, management commitment to involving all parties affected by the proposed innovation in every planning stage, and second, the development of a continuous learning process, through which employees who traditionally would be excluded from planning can acquire the understanding and confidence necessary for them to participate effectively in designing the new sociotechnical system.

Table 2 Conflicts between the improvement of working life and the introduction of new technologies (after Williams 1983).

Requirements for improving the quality of working life	Likely effects of new technology within existing organizational framework
Decentralized decision making and democratization of work	Centralization of decision making and less democracy
Optimal division of labor with employees working on whole tasks	Microscopic division of labor
Opportunities to learn on the job, develop skills, and acquire new skills	Skill requirements and learning opportunities in many occupations are almost wiped out
Increased employee interaction and joint self-management based on work teams	Reduced interaction and less scope for joint self-management
Challenging work and making a significant contribution to the job	Meaningless work and little or no sense of contribution

That this process can be carried out effectively has already been shown by several examples. Wilkinson (1982) describes the application of new technology in an optical firm, which led to a degrading of many jobs and diminished job satisfaction: the solution lay in extensive job rotation. Melman (Joyce and Wingerson 1983) cites the experience of an American car parts manufacturer where one plant elected to perpetuate Taylorism when it installed computers and the other agreed to train workers to program. In the first plant, workers are designated "operators"; when their machines malfunction, they flip the on–off switch and call in a repair man. In the second, they are "journeyman–machinists," paid higher rates than the operators in the first plant, and are trained to maintain and program their own machines. As Melman says, in the first plant machines are down for at least half of the time, while in the second, where worker turnover is almost nil, downtime amounts to only 3% of the total.

Modern technology will also enable white-collar employees to work at home and thus to intermingle working life with private life. It is with good reason that many unions strongly object to these developments: they are afraid that many social rights that employees have obtained through the efforts of the unions will be lost because white-collar workers at home will be unionized to a far smaller extent.

It is often supposed that computer literacy will be a necessary prerequisite for most of the jobs in the next decade. This is one of the reasons why currently so much attention is being paid by education

authorities to computer languages (especially BASIC), even in elementary school training. But it is in fact very likely that the underlying assumption will not hold true. The first people who bought cars, in the late 19th and early 20th centuries, had also to be mechanics in order to be able to repair their cars. Hardly anybody today really knows how a car works, but she or he can still drive it very easily. With the advent of intelligent interfaces we will need to know less and less about what the computer really does. An excellent example in this direction is provided by Apple's Macintosh^{TM} microcomputer.

Robots are not always used to take over unhealthy, hazardous, or boring work. Rosenbrock (1982) describes a plant making electric light bulbs. While the plant was almost completely automatic, it also employed a woman who picked out a short length of aluminum wire from a box with tweezers, held it by one end, and then inserted it delicately inside a coil that vaporized the wire to produce the reflector. This task was repeated every four and a half seconds. When the specialist who had automated this factory was asked why this specific task was not also performed by a machine, he replied that "to bring in a universal robot would mean using a machine with many abilities to do a single job which may require only one ability."

Clearly, it will only be through our own efforts that Lucky's hopes for enriched jobs, freedom from disagreeable toil, and increased productivity will come true.

How about *social life*?

> "Tahitians do not have to spend their days working in field or factory, since their island is so fertile that ripe mangos (still) lay uneaten under the wayside trees. They have developed a culture in which personal relations are supreme, in which families are extended, sexual life is adventurous, and social occasions are marked by the subtle rituals of the dance. These facts are connected, for Tahiti's mango trees provide another precious resource besides food. Dancing, offering the cup of Kava, playing with children, sharing a friend's joys and anxieties, or exploring one's own — all these social skills take time. AI could be the westerner's mango tree" (Boden 1983).

But in trying to reap the benefits of this new "mango tree" we encounter similar problems as with employment: how to make the transition from here to there? As Boden herself says, there will be enormous political and psychological problems in passing through the transition phase. Many people do not enjoy their work, but would they really prefer to be *without* work? Isn't their work already more than just a way of sustaining physical life through earning money? Jahoda (1982) has stressed that work is an indispensible part of human identity, which cannot easily be replaced even with the most meaningfully

used leisure time. Jenkins and Sherman (1981) have even entitled their book "The Leisure Shock."

Are we prepared for this increase in our leisure time, which will inevitably come if we distribute the remaining workload more equally? Artificial intelligence will present us with many toys to distract us from more human interactions. Some of these toys will be improved video games. Already, sales of game cartridges run to millions of copies, compared to about half a million for the best-selling non-game program (Lucky 1983). A current bestseller is Microsoft's flight simulator, which is nothing other than a sophisticated video game (I like to play with it myself!).

Another toy that AI could offer us is telepresence. Imagine a robot equipped with sensory organs, this robot being perhaps thousands of miles away, while you yourself are standing or sitting or lying in equipment that gives you the same sensations (visual, auditory, tactile, etc.) through videoscreens, loudspeakers, etc., as the robot has. You could have the illusion of being somewhere else or of having adventures that perhaps you would never otherwise dare.

But who or what will our companions be in the future? Frude (1983) argues that in the long run robots will be built that could be the ideal companion machines. They would not only look, feel, and sound friendly, but they would also be programmed to behave in a convivial manner. Each machine would appear to be charming, stimulating, and easy-going, and yet would remain slightly unpredictable, and therefore interesting. It would sometimes take the initiative and would have a personality of its own.

If we take this fantasy a little further, could not such machines also exhibit the character of persons who are already dead? It is already recognised as important for a natural language dialogue system to build up a model of its user (Kobsa *et al.* 1982, Kobsa 1984). Interacting with a human being, the computer could build up an image of her or his personality, which it could reproduce after the person was dead. Perhaps already in ten years from now we may not be looking at pictures or movies of our dearly departed, but instead carrying on dialogues with their personality representations in a computer. And we could even change their personalities. We could change the fascist father (mother, etc.) into a convincingly liberal one. We could even live out our Oedipus- (Jocasta-) complexes via simulated father- (mother-) murder through the program – but would this be necessary after our careful reprogramming of our father (mother)?

But, to be serious: are we really prepared for such an amount of leisure? And what do we have to do now, so that we, or at any rate our children, are prepared?

Education will be considerably changed by artificial intelligence. The first method to explicitly include ideas stemming from AI was LOGO, developed by Papert (1980) at the AI Laboratory of MIT. Papert showed how the fundamental concepts of mathematics could be understood and mastered by young children through the use of his LOGO program, by interaction with the computer. While computer-aided instruction is usually regarded as meaning the use of computers to program persons, Papert reversed this process, enabling the child actually to program the computer.

The original programmed instruction systems of the late sixties tried to force people to learn along very limited paths; these systems have fortunately since disappeared. *Intelligent computer assisted instruction* (ICAI) needs to have problem-solving expertise of its own, its own diagnostic for student modeling capabilities, and its own explanatory capabilities. In order to orchestrate these reasoning functions, it must also have explicit control or tutorial strategies that specify when to interrupt a student's problem-solving activity, what to say, and how best to say it, all in order to provide the student with instructionally effective advice (Sleeman and Brown 1982). With the advent of beautiful color graphics and even video disk presentations, it is very likely that such systems will find a place not only at schools or universities, but also in the training and retraining of employees.

ICAI has proved to be especially helpful with handicapped children. Sylvia Weir (cited by Boden), when working with severely handicapped, mainly autistic children, observed really significant intellectual and social–emotional improvements. The British Government has already decided to spend 2.5 million pounds on electronic aids for handicapped children at Britain's 700 special schools (*New Scientist*, 1983). Money will be spent, for example, on equipment like the turtle, a small dome-shaped device on wheels controlled through a microcomputer, which is an important part of the LOGO system.

Little can be said here about the potential impacts on *crime*. While the computer is certainly being used more and more for computer-based criminal fraud or blackmail (Gillard and Smith 1983, Lamb 1983), it is not yet generally used to improve the "effectiveness" of crime generally. AI, for example through the development of special expert systems, may unfortunately help.

2.6 Culture

Human self-understanding may be radically changed through AI. How will man's view of himself change when confronted with a superior intelligence?

Computers can be made to devote all their resources to a given task, and they can run the same program 100 or 1000 times. In the next century, an AI program may well be awarded a Nobel prize (Schank and Slade, Chapter 10).

But even short-term expectations are fairly radical: 200 or even 100 years ago, a person's most important attribute in getting a job was physical strength. With the introduction of the machine, this has drastically changed, and most people now get their jobs on the basis of their knowledge or their mental abilities. If those abilities are substituted by AI programs, will people start to perform "mental jogging" in order to remain mentally fit? And will such a situation help swing the pendulum from a somehow "over-intellectual" culture to another one where spontaneity, emotionality, and creativity are more highly valued?

2.7 Politics

Computer surveillance is something many people are afraid of. With the aid of machine vision, all pictures currently recorded by video cameras on the streets, in banks, subway stations, etc., could be automatically analyzed and the movements of individual people and their companions traced. Speech analysis makes it possible to analyze phone calls, and this would be especially easy in countries where telephone companies have been nationalized. Thus, AI could make the powerful people even more powerful. On the other hand, it has been argued that the use of computers to store information makes it easier to find out *which* information has been stored, compared to the era when documents were kept on obscure shelves (Arbib 1977).

AI could have a positive influence on politics. The revival of interest in human self may help to restore the respect for human individuality, privacy, and freedom of thought (Havel, Chapter 7). AI may also help to introduce participative democracy: people could not only find out which data are stored concerning themselves but they could also have direct access at any time to every case or administrative procedure concerning them, provided they had requested specific permission from the government (Vámos, Chapter 11).

In the long run, a political expert system could help a political candidate by devising campaign strategies, sending out press releases, drafting speeches, preparing position papers, conducting polls (via electronic mail), developing media commercials, and analyzing the opponent's record (Schank and Slade, Chapter 10). One more step: a computer program that displayed such encyclopedic knowledge of the campaign and politics generally should probably be a candidate itself!

While a proper distribution of AI products could help developing countries and thus decrease the *North–South gap* (Vámos, Chapter 11), it is not really clear *why* such a distribution should occur. On the contrary, it is likely that the gap will continue to widen. When comparing the cost of programmable industrial robots with US industrial wages, Pelton was able to show that the curves crossed between 1976 and 1977, which corresponded to the introduction of robots into industry on a larger scale (see Figure 2).

Figure 2 The costs of human labor *versus* those of industrial robots (after Pelton).

However, though the industrial wages in developing countries are far below those in the developed ones, even these wages will be higher than the costs of robots from about 1985 onwards. While at present many companies have production facilities in developing countries to take advantage of the lower local wages, it is very likely that in the near future the use of robots in the home countries of these companies will be the cheaper strategy. It is not difficult to imagine what this might mean for the employment situation in the developing countries.

AI has already led to an increase in *international competition*. While for many years the United States had a monopoly on AI (apart from a small group in the United Kingdom), a situation that was generally approved of or at least tolerated, the situation has changed radically since AI entered the economic market. Japan, with its Fifth Generation Computer Systems Project, provided the first challenge to the

established order. Western Europe then decided to start the ambitious ESPRIT project to counter these efforts. Both Japan and Western Europe are afraid that the United States may impose export restrictions — perhaps under the pretense of defense considerations — on AI software or dedicated hardware. Even if that does not turn out to be the case, the monopoly of one US company in dedicated AI hardware has led to an increase in the price of identical machines of 20% in just one year, while all other comparable machines have become considerably cheaper. Thus, AI will very likely *not* lead to closer cooperation between the United States, Japan, and Western Europe, and even tougher competition can be expected.

Those countries all belong — ironically! — to one bloc, that of the nonsocialist industrialized countries. Can AI be used at least to decrease the tension between the two major blocs, East and West? I would like to propose three projects for *cooperation* (for more details, see Trappl 1985).

The first project would be the development of a Russian–English/English–Russian translation program. Such programs are already under development both in the United States and in the Soviet Union; but I would propose specifically that the work be carried out by a mixed team, consisting of Soviet and American AI researchers, and that the program be made sufficiently portable so that it could be used on both US and Soviet computers. Such a program could, for example, help to avoid unwanted ambiguous phrasing by having an authorized translation available in advance or during emergency conversations over the "hot line," which is actually a telex connection.

Second, I would recommend a joint project between US and Soviet AI specialists on developing a real-world model as a semantic net. This might help clarify differences in viewpoints, but could also uncover similarities. Furthermore, it would also be very useful for the proposed Russian–English translation program. There is already a project underway, led by the US scientist Stuart Umpleby and Vadim Sadovsky from the USSR Academy of Sciences, which sets out to clarify the terminology of systems science (see Trappl 1985, Part II). Though I consider this a very ambitious project, its results will, at best, be verbal statements, while a semantic net would constitute a precise AI ontology.

Third, starting from the fact that many crises in history have led to war (although, luckily, not all of them), I propose a joint East–West project to study history from this aspect and to try to build a crisis-handling expert system to help prevent war, or, at least, to provide a common basis for discussions during a crisis. This may sound naive (and of course it is), but how wise is it to build an expert system for an autonomous cruise missile? And wouldn't it be worth investing the cost of a single missile in these projects?

3 Conclusions

Some of the many potential impacts of AI considered in this chapter look positive, a few neutral, and some negative. The extent and directions of these impacts will depend on several factors, some of which can be influenced by you, the reader:

- *The responsibility shown by AI researchers.* One may of course argue why the researchers should know any better than politicians, businessmen or military experts what is best for mankind. However, I personally see topics where I would refuse to cooperate and others where I have shown and will continue to show initiative. A concerned AI researcher should ask, before starting work and while working: What are the potential impacts of my work? And then decide whether she or he can justifiably continue.
- *The opinion of the "general public."* In many parts of Europe the "alternative" or "green" movement is very much opposed to high technology, including such seemingly innocent devices as microcomputers. The extent to which people are informed and involved in the decisions to be made will be of the utmost importance. This will, for example, influence considerably the ratio of money spent on AI research in the military *versus* the "civilian" domain. If the general public does not perceive the benefits of AI (which could mean that there are really not enough of them), even restrictive legislation may be favored.
- *Unions.* Unions have already conducted very hard negotiations over the introduction, for example, of new typesetting technologies. Unions – at least in those countries where they are fairly independent of the government – will play a major role in protecting employees from the drawbacks that the introduction of AI into the office and the work-place may bring; however, the unions could also hinder the introduction of AI technology for reasons of interest only to the union. For example, I am afraid that they will not be supporters of a work-independent income. In any case, the unions will have to persist with their efforts to reduce the number of working hours per person in order to prevent unemployment and to obtain a fair distribution of the benefits of AI.
- *Export restrictions.* It is not yet clear which policy the US Government will follow in future regarding the export of hard- and software. This will not only influence the relations between the United States and "the rest of the world," but will also have a considerable impact, positive or negative, on the North–South gap.

- *Technological progress.* It is very uncertain what sort of progress AI will make and at what pace; see for example the incorrect forecast of AI accomplishments made by Herbert Simon and Allan Newell in 1958 or the difficulties encountered when developing new hardware (good examples are the overoptimistic forecasts regarding technology based on the Josephson effect).

In his Presidential Address at the Annual Meeting of the American Association for Artificial Intelligence in 1983, reprinted in the *AI Magazine*, Winter 1983, Nils J. Nilsson concluded: "At best, Artificial Intelligence will both liberate us from unwelcome toil and provide us with the most detailed picture we have ever had of ourselves. Possibly no science has ever posed greater challenges than those."

References

The references cited in this chapter can be found in the Artificial Intelligence Bibliography, List of References, pp. 221–251 of this volume.

CHAPTER 4

On Being Human in the Computer Age*

*Michael A. Arbib***

1 The Computer Age

Goethe said, and it was something that my thesis advisor repeated to me when he wanted me to concentrate on my thesis, "If you would master the infinite, take the finite and master it from all sides." When I chose the topic "On Being Human in the Computer Age," it was with the aim of really covering two areas: one trying to assess how far progress in artificial intelligence, brain theory, and cognitive science could tell us to what extent we are machines, and the other trying to see the way in which computerization and automation were changing the workplace and the human condition. But I chose a rather grandiose title, and so I am afraid that there are times when I may get a little carried away and speak about things of which I know little. But if the curse of the modern academy is overspecialization, perhaps a little chutzpah will do no harm.

When I say "the computer age," I am very much aware that this is not *just* the computer age. When we try to describe this century of ours, there are many other terms besides 'computers' that come to mind. This is the century of horrors, of World Wars, genocide, and concentration camps. It is a century of staggering anachronisms, when Iranians who would return their state to the fundamental Shi'ite

religion of the 8th century use cassette recorders to spread the word of their Ayatollah, and when perhaps the group most adept at the art of TV networking and computer mailing in the United States comprises those who would have us adopt the most fundamentalist form of Christianity. It is also the century in which we have become more conscious than ever before that we live in this world and not just in a village or a town or a nation. I was most struck by this in 1969 when the *Apollo 11* astronauts returned from the moon, and splashed down in the Pacific Ocean 3,000 miles from the United States. The TV announcer said, "The astronauts are home!" For the first time in human history, just to be on the surface of this planet was to be home. Yet at the same time that we come to think of ourselves in this global perspective, we find that our world is threatened by overpopulation, by pollution, by the depletion of natural resources, and this very consciousness of one world is shattered by the resurgence of nationalism as each nation tries to grab for what it can get. So all this will be in the background as we discuss that finite topic, the impact of computers by the end of the 20th century.

When I talk about the impact of computers, I am going to be really thinking of the computer in two different ways: the computer as a tool, and the computer as a metaphor. In speaking of the computer as a tool, I will stress two particular aspects: the increase in automation, changing the nature of work in our society as more and more tasks become accessible to the computer; and the way in which the computer is changing our notion of literacy, as what it means to read and write and think gets changed by the computers that we can interact with and by the rise of the computer as a personal device that each of us, at least in the well-developed societies, can have in our own homes.

When I talk about the computer as metaphor, I am going to be looking at artificial intelligence, the attempt to program computers to do things that you would swear were intelligent until you knew that a computer had been successfully programmed to achieve them. The question I will ask is: to what extent, as we come to better understand the mind of the machine, can we come to see a continuity with our own minds? Is there an impermeable barrier between man and machine, or are we highly sophisticated mechanisms; and if so, what does that mean?

2 The Image of the Computer

But before I get into these questions, let me just step back a little and consider the image of the computer. For each of us the computer means something different, and I would like to ensure that we have certain agreements about what is being discussed. For those of us over 40,

the computer is probably the symbol of conformity — we remember when computers required information to be put on paper tapes or punched cards, and we remember the old jokes about "I am a human, do not fold, spindle or mutilate me." But if you are under 20, the computer is probably something quite different — what you play video games on — and for you a computer is fun, and a friend, and probably a lot better than your parents. We now see the computer going from the tool of the elite, as it were, to something that is quite universal, and I think that we will see a transition to computer literacy.

There was an article a few years ago in the *New York Times* about how much we take literacy for granted. It contained interviews with various people who managed to survive in New York without being literate: the desperate ploys when in a restaurant of saying "What looks good to you on the menu?" and then following suit; or of trying to get the instructions to a friend's house that did not require one to read the street signs. The ability to read and write has become part of our subconscious; we no longer think about it. Yet before the invention of the printing press, to be literate was to be an exceptional person, to be part of a small elite. I was struck, when reading the biography by John Gardner of Geoffrey Chaucer, that in the middle of the 14th century Chaucer was very fortunate to have a tutor who was exceptionally rich and possessed all of 50 books. Yet each of us takes for granted that we can accumulate hundreds of books of our own, just as we take for granted the use of reading to find our way around our world.

Just as we have become used to street signs and menus, so we will become used to the computer; and, in the words of my colleague Conrad Wogrin, we will come to find the computer as "invisible" as reading and writing. We have long been used to sharing our thoughts with distant friends by writing letters to them; it is not all that long ago that people found it frightening to communicate by telephone, and now we all take this to be a matter of course; and we are beginning to use electronic computer-mediated mail. The computer is going to change the way in which we think and solve problems. There will be no need to be *very* good at mental arithmetic when we can use the computer to keep track of the arithmetic details and let us concentrate on the problem-solving. The ability to write down our thoughts has long since relieved us of having to acquire the skills of rote memory that were common in preliterate societies. With computers we will develop yet different memory skills to learn about retrieval, and about how best to organize information within our computer network so that we can find what we want.

There is a professor in the School of Education at the University of Massachusetts who believes that "the three R's" are the tool of the elite and looks forward to the day when Ph.D.s will not be required to be literate any more. But I suspect that he is mistaken. There will still

be the need to articulate carefully what we have to say, to edit it, to accumulate it, so that we can best express our thoughts. To what extent those skills will require us to write in longhand as opposed to typing into our console and editing, I do not know. With the development of computers we will see not so much a decay of skills that we have, but rather their enhancement. But there is still the fear that the computers are warping our lives in ways that we do not really understand.

E.G. Schumacher, author of *Small is Beautiful*, distinguished machines that serve men from machines that require men to serve them. Unfortunately, the distinction is a very hard one to hold. What is the telephone? When you want to contact a distant friend, it is a machine that serves man. When you are happily ensconced in the shower and the phone rings, it is a machine that forces men to serve it. Similarly for the car, and similarly for the computer.

We cannot yet think through well enough what it is we really want for ourselves. We are seeing increasing automation taking over not only the blue-collar jobs but also the white-collar jobs. In some ways this is good, because suddenly each of us can afford the services that before only the elite could afford: we can have check accounts, we can have dishwashers, all because we have mechanical slaves. The only trouble is that there are a lot of people who made their living doing the drudgery that those machines are taking over. How, then, do we build an adaptive dynamic society in which people are not defined by some particular social role and then thrown aside when society discovers that that particular skill is no longer needed? The problems that we face, while they may be exacerbated or brought into focus by the computer, are much broader than that. They are human problems, but human problems taking a different shape at this stage in our history.

3 On Being Human

There has been much work in philosophy and religion and literature and poetry and drama and art to try and address the issue of what it is to be human, and I shall certainly not try to recapitulate all of it right now!

Those of us working in artificial intelligence have tended to stress things like problem solving, game playing, question answering, ability to use some aspects of a natural language, and elements of vision in our study of 'intelligence'. We are encouraged by our increasing ability to write computer programs that we think provide some insight into the human mind. But we must not pretend that those dimensions exhaust what it is to be human. True, Aristotle and many since him defined man

(and woman) to be the rational animal. But there is a lot more to being a human than being rational: there are the dimensions of emotion, love, compassion, of having a family, being part of a community, belonging to a society. Unfortunately, I will not say much about these aspects.

These are the positive sides of being human. A title like "On Being Human in the Computer Age" has a utopian ring to it − Let us all be better humans through love of our machines! But there are other things about being human: there is evil, hatred, jealousy, and war. In fact, if we were to bet on two things that will continue right through the computer age, they might well be war and hatred. But we will be a little more utopian, to see some way in which we can steer our way through the shoals of the remaining 6,470 days of the 20th century. (This is how machines change you: I happen to have a calculator which allows one to subtract one date from another − if it wasn't for that machine, I would never have known that there were 6,470 days left in the 20th century. This is perhaps a reason to despair about the over-prevalence of computers in today's environment!)

Although much of what I say will be focussed on one of those aspects of our humanity most affected by current developments in computers, I am still trying, in those words of Goethe, to get that handle on the infinite. I do not think that being human is some sort of unitary thing that we already know. I do not believe that there already exists an ethical or religious system that exhausts what it is to be human, and that our look at the computer age can fit neatly into that system. I think that there is no ultimate reality in being human, and that if you believe there is an ultimate reality, you are doomed to delusion or to despair. Rather, I see our reality as a contingent one: we do the best we can. And that best that we can do is conditioned both by our intuitive feeling as humans, as members of our society, and by our rational analysis of social, cognitive, and psychological "forces."

4 Man is a Machine

Humans are very resistant to new ideas, yet after we become used to them we forget that there was ever any problem in accepting them. Copernicus advocated the idea, one that had lain dormant for about 2,000 years, that the earth actually moved. This was obviously a stupid idea − the earth is stationary, it is the sun that rises, not the earth that turns, as you notice every morning if you get up early enough. When Copernicus first came out with this idea, although it created a certain stir, the Church was not very worried. But the religion of that time, Christianity as reflected in Dante's universe, had the tiers of hell

within the earth, and then — in spheres around the stationary earth — the planets, the stars, and beyond that God's heaven. But if the earth was just another orb moving through the heavens, and if the stars could extend to infinity, where were heaven and hell? So by the time Galileo came on the scene, his theories were ripe matter for the Inquisition.

We have assimilated all that. Although we still say the sun rises and probably mean it, we understand that this is just a way of speaking and that, yes, it is the earth that rotates as it moves in its elliptical orbit about the sun. It would be hard to imagine that a legislature in the United States would enact a law saying that equal time should be given to the view that Dante's universe is an acceptable scientific theory. It is strange, isn't it, that Darwin's theory has not fared so well, so that it is still a matter of heated debate in certain state legislatures as to whether or not we really did descend from (horror) apes. In fact, if one looks at history, it is probably bad enough that we descended from some of those humans that were around a little while ago...

What I want to argue is that our study of minds, brains, and computers is bringing us to a similar conceptual revolution. Freud showed us that much of our apparently free behavior has deep roots within our unconscious. Freud started as a neurologist, and he began to develop a mechanistic view for understanding what might be apparently irrational in our behavior. We are now coming to better and better understand the mechanisms of mind, the mechanisms of the brain. I would like to think (perhaps wrongly) that 100 years from now it will seem as absurd that there were people back in those benighted days of the 20th century who doubted that the mind was a mechanistic phenomenon as we now find it strange that people should doubt the concept of evolution or that the earth does indeed move.

I should offer a caveat here about what I take a scientist to be — I assure you I am not quite as dogmatic as I sound. There is one image of science as progressing in terms of "normal science," some theory that explains just about everything, until too many facts accumulate that cannot fit into this framework. Then for a while there is turmoil, until these new facts are assimilated in a new theory. This is the Kuhnian model of scientific revolutions, moving from one paradigm to another. In fact, my experience — reflected in a book like Feyerabend's *Against Method* — is that science is much more pluralistic, even perhaps anarchistic. Not that scientists go around with bombs, throwing them into each other's meetings, but anarchistic in the sense that science does not progress by the Academy saying "This is the true theory," but rather as a result of many people trying different theories. Now let us see the dichotomy in views of the relation between mind and brain. Some scientists, perhaps starting from a certain religious world view, are struck by how huge the gap is between those aspects of mind that

we can now understand in terms of computer programs or neural net-
works — certain aspects of vision, memory, and motor control — and the
full richness of being human to which we have already alluded. There
are others of us who note that it took billions of years for an amoeba to
evolve to a human, and that without the fossil record and a bit of good
theory, you would not really believe that such an evolution could hap-
pen and then argue that, similarly, there is no discontinuity that we
can yet see that would distinguish machines in their evolution from
humans.

In our Department of Computer and Information Science at my
University, and in our interactions with our colleagues in the Cognitive
Science Program in Psychology and Linguistics and Philosophy and else-
where, we have begun to explore a number of ways in which the work in
computer modeling and brain modeling can give us an insight into cogni-
tive mental activities. We have robots that can pick up eggs without
breaking them. We understand how it is that a frog can detour around
a barrier to get at a worm. We know how a computer, if asked questions
about a relatively simple domain of knowledge, but asked in English
rather than a programming language, could come up with the correct
answer. So we see this progress: we are well beyond the amoeba,
perhaps we have already reached the worm. Some would say the worm
is a long way from the human. Others of us say that we have not been
stumped yet, so let us keep going.

In this latter spirit, I would suggest that the notion of mechanism
can help us understand our humanity. I said before that our under-
standing was pluralistic: if we are going to talk about our emotions,
then for much of the discussion our normal everyday discourse will be
fine; but if we are looking for an effective drug therapy, then the
mechanistic underpinnings become crucial. Again, for many human
problems the language of sociology, of people in interaction in society,
becomes crucial. So when I say that man is a machine, I am not trying
to restrict us to some narrow form of mechanism, but rather to suggest
that for a rational analysis of the human condition, an evolving notion
of mechanism is what is required.

Such conceptual evolution has already occurred in physics. It is
almost 300 years since Newton published the *Principia*, and at that
stage we began to see the Universe as one vast machine. But in the
20th century, the definition of physical machine has undergone its
revolutions — the phenomenon of relativity and the quantum-mechanical
change from determinism to probability amplitudes have changed our
notion of mechanism. I do not see any reason why our increasing
attempts to build better computers and to better understand ourselves
won't in the same way synergize to grow yet new notions of machine.

To conclude this part of my paper, let me simply assert (the argument would take us too far afield here) that there are notions of free will and social responsibility that are in no way incompatible with what I take to be the mechanistic view. What does it mean to have free will? I think it means that we can act in accord with our own complexities; not that we do something that is totally unrelated to our entire preceding life. We are shocked when we cannot begin to understand the roots of our actions. Freud has helped bridge that gap by showing us how to take the occasionally inexplicable and try to dredge up forgotten (or repressed) experiences to better understand how we behave. Some people have thought that quantum mechanics takes us from determinism to free will, but if we always jump at the toss of a die rather than thinking through what it is that we really want to do, we do not have free will. And so one can continue in this way, trying to think through the complexities of our behavior, trying to see how this behavior grows out of our own self, seeing how that self grows out of both our biological background and our social experience, to understand human responsibility within a mechanistic view of the human mind. I think we will be able to match this with a sufficiently complex view of mechanism.

5 The Computer in the Workplace

I have tried to suggest that the computer metaphor is not a static one. I do not say that we think in binary code, or that we cannot understand something unless it is fed through one ear on a paper tape. Rather, we are beginning to build a vocabulary that goes far beyond the normal vocabulary of mentalism, just as Freud had already begun to enrich our vocabulary for thinking about the mind.

With that I want to turn from the computer as metaphor and to think about the computer as a tool, and I want in particular to examine the nature of work. Some people, in defining freedom, have suggested that freedom from toil, freedom from work, was the goal of human life. But I think that, for most of us, work is part of the definition of our life so long as that work is "meaningful." Some people work because they know they will starve if they do not work, and that gives work meaning enough. Others of us, fat cats of the professoriate, like to think of our work not as a means of subsistence, but rather as a means of our development, etc. The point is that I think it really is part of being human not only to be a member of a family, to have friends, to be part of a society, but also to have some work that is meaningful for us, and to have some reasonable standard of physical well-being.

Marx, or at least the young Marx of the Grundrisse, observed that in feudal society people had fixed roles. They did not have to question

who they were: society was static, and God had decreed each person's status — whether bondsman or lord, he had his duties. When the machine came, many of those roles that had hitherto been fixed for humans became jobs that could be done by the machine. And this had effects both good and bad: it was liberating — suddenly people were not locked for life in the role that the fall of the hereditary dice had cast for them, but they could begin to explore new roles. But it also brought alienation because suddenly they were no longer possessed of a worth that all society recognized as a farmer or a craftsman, but they were on the labor market and they were worth what they could get in competition against the machine. Suddenly people could no longer count on that sense of meaning.

In the 20th century the pace of change has picked up. It takes perhaps ten years for some forms of employment to become outmoded and outdated. What we have to ask ourselves as we try to be human in the computer age is perhaps not so specifically tied to computers or automation, but rather "How are we going to cope with massive change?" I claimed that we are (in some sense) machines, but are we machines to be thrown out on the dust heap once we no longer play our economically useful role, or are we rather special machines like a fine clock that is so beautifully engineered that you will keep it even when it stops keeping accurate time?

It is easy to look at the temper of our age in terms of the pre-valence of machines, just as there are some people who say "The weather has been so bad lately because of the nuclear bomb," or "We wouldn't have had that snow if it hadn't been for the Falkland Islands." People are often possessed of a very simple view of causality, trying to find just one cause for each thing that ails them. It seems to me that if we look at the history of modern society we can see people displaced from their jobs because of machines but we can also see people dis-placed from their jobs because of a certain economic system. Depend-ing on who you are, you can see technology or capitalism or communism as the culprit. One of the things that we learn in cybernetics is the notion of mutual causality, that you cannot just isolate one item and say "Here is where the chain of cause and effect starts, here is the cause and there is the effect, and that's it." So it is that we must come to understand a complex interacting system in which the presence of the computer, and the presence of automation, is changing the idea of what people can achieve as a whole. There are things that we can do in the way of predicting the weather, of mass-producing goods that were hith-erto available only for the wealthy, and of letting everybody have a hi-fi and a stereo, that would not have been possible without the machine. But at the same time, many people who had a job that gave some mean-ing to their life are now without it. Some people show little sympathy,

stating that if these people no longer have an economic role, they should go out and do something that "shows their initiative." However, I believe that we must redesign our social structures so that the benefits of technology "trickle down" to all these people. But, of course, it will not just "trickle down" — it requires some coherent vision of social justice.

I cannot offer such a vision here. But what I want to say very strongly is that we are in a situation where the nature of the workplace means that some of us, with good connections or good education or just good luck, can benefit tremendously from the new opportunities of the machine, to enjoy privileges that only a few could have enjoyed in the past. Remember that tutor of Geoffrey Chaucer's who was rich indeed to own 50 books back in 1354. Tragically, we have seen an increasing level of unemployment amongst minorities, and especially amongst black youths. This is not something that responsible citizens can allow to continue. We have to understand, then, how we can provide opportunities for these people to enable them to make their contribution to society. In no small part, that has something to do with education.

6 Reconstructing Society

We have been talking about these problems as if they were really quite new, but they are not. The introduction of cotton-spinning machines in Italy in the 1300s caused a great dislocation of the industrial base in medieval Europe. If one reads about the Industrial Revolution of 150 to 200 years ago, one is struck at how many of their societal problems are ones that, if only one changes the names of the machines, we recognize today. The 19th century came up with a great invention — the labor union. Men recognized that the new machinery was concentrating economic power in the hands of the few. That was not new, because economic power had been in the hands of a hereditary aristocracy, but now it came into the hands of a mercantile aristocracy. And so what workers did was to band together to say that "You think that you should not pay us more than an equivalent machine would have to be paid in terms of power and maintenance. But we are humans. We need more than that. We cannot live on these wages." And so this great adaptive social invention, a response to the first round of mechanization in modern times, was the labor union. However, I fear that at the end of the 20th century the union is becoming maladaptive, because unions have become identified with large bodies of workers who themselves are identified in terms of some particular technology or trade. As that technology or trade changes drastically, then too often the

work of the union is to turn back the clock to say "You must keep using this obsolete, outdated technology which gives our people employment."

I think we have to see a move toward a more adaptive society in which we still have special interest groups, advocates for the welfare of groups of workers who are skilled or unskilled, who are skilled for the future or only skilled for the past. But these groups must come to define their obligation to people not in terms of the demarcation dispute of holding to a static framework, but by helping people understand what is going on, so that they can adapt to change. Here, interestingly enough, is where I think the computer may be developing to a stage at which it can help us with these new social needs. Many of our concerns with the social or industrial machine are concerns with bigness, concerns that power has been concentrated in the hands of large companies or big government, with workers organizing in large groups to meet this. I would argue that we can now move into a pluralistic world where, because we have computer networks, because people can have access to large data bases no matter where they live, with sophisticated expert systems, knowledge networks, and visual processing devices, we can begin to have people defining their skills and their trades in terms of much greater diversification. The effective unit of organization can be relatively small and thus restore more of the human scale.

I do not see technology as determinate, in the sense that technology must make us all better people, or technology as such that it must inevitably ruin our social world. In the hands of a secret police force, a data bank can be an invaluable tool of oppression. In a democracy, access to home computers hooked into a computer network of data banks can be part of creating a new level of informed citizenry. But even though computers can be used for good and for bad, I still think we must face up to the fact that in this age of increasing automation (increasing, I would remind you — automation has been with us since the first irrigation ditch, the first water wheel), technology contributes to many social changes that are threatening to disenfranchize many of us, to deprive many people of their livelihood.

But again, I do not think we can see this in a unicausal way, of the machine *causing* the social malaise. We might say that greed, need, and power drive our economic structures, and it is the economic demands of big business or big government that make a cheap machine more desirable than a well-fed person. And the computer extends the reach of this automation, so when we talk about being human in the computer age, we are not simply saying "Let's program our computers to do what we want them to do," nor are we throwing up our hands in despair and saying that the computer will dehumanize us all. We have to think in a multicausal fashion about the ways in which, as we come to

better understand ourselves, as we come to better understand the tools that machines give us, as we better understand our interactions within society, we can use this knowledge to begin to address the problems of a changing society.

Using that knowledge means that we have to be far more self-conscious. Here's a current example: a few years ago there was an energy crisis. Petroleum was running out, and we were going to have to find alternate sources of energy. But then the price of gas went down 10 cents a gallon, and people forgot the energy crisis. Well, we cannot solve our pressing social problems with that sort of instant amnesia. We are going to have to learn to assimilate the complexities of our world. We are going to actually have to talk to people in other departments. We are going to have to not only talk about our little specializations, but try to understand whether they mean anything in the larger context of society's problems.

We tend to think in very narrow terms, whether we are humanists or technologists. We tend to define a problematique in terms of something that we can publish in a paper within a year; we have all too little time to step back and really wrestle with these larger issues and try to understand how what we know can make sense outside our speciality. I recently gave a series of lectures on "The Construction of Reality." I do not want to develop the theme of those lectures here, but I do want to suggest that the range of topics we have discussed reminds us that we have a reality to construct, a social reality. Even though I have talked about man as a machine, even though I have talked about deterministic models helping us understand ourselves, I have also said that within that framework we can understand free will -- we are responsible machines. And as responsible machines who care for our children we must try to bring all our rationality to bear to ensure that as this computer age continues, hopefully long after the threat of nuclear war is a memory of a strange and aberrant past, we can still be truly human.

CHAPTER 5

Impacts of Artificial Intelligence

Margaret A. Boden *

1 AI Development in the Next Ten Years

1.1 Core research areas and likely results

Several core research areas of artificial intelligence (AI) are likely to make solid progress within the next decade. Each of these is already being worked on in various countries, and progress does not depend solely upon the success of Japan's ambitious "Fifth Generation" project (though it might be accelerated by associated hardware and software developments).

One of the areas is low-level vision, based on techniques using parallel hardware and cooperative processing. Current "connectionist" research in this area differs in its approach from work on 2D pattern-recognition by "property-lists", and from the top-down "scene-analysis" of 3D scenes. Based on detailed studies of image formation, it is able to extract from the ambient light information about 3D features (such as shape, depth, texture, and surface orientation), which in previous approaches could have been computed only, if at all, by way of high-level knowledge of the expected scene. Some of this work is being done in the context of human psychology and neurophysiology, some in a more technological context. Dedicated (massively parallel) machines

---

are being designed for this research, and major advances depend upon such hardware.

A second area in which we can expect significant progress is robotics. This includes problems of movement control, trajectory planning, and visumotor coordination (where it will probably build on advances in low-level vision). As in the case of vision, some projects will rely on "artificial" means to ensure success (such as light-stripes for automatic welding machines, capable of recognizing different sorts of weld-joint and guiding the welder accordingly), while others will relate more closely to psychophysiological theories of motor control and visumotor coordination in living organisms.

Intelligent knowledge-based "expert" systems (also known as IKBS systems) will multiply enormously in the next decade, not least because there is considerable commercial interest in them. Different domains of human expertise may require different approaches to knowledge-engineering. In domains less fully covered by an explicit scientific theory, it may be easier to extract knowledge from human experts who are competent but who have not yet achieved the "intuitive" mastery of the domain that topflight experts enjoy. The latter give the right answer more often, but cannot easily introspect their reasoning processes, which happen very rapidly and are not consciously accessible. The former take time to come to a decision, often consciously weighing distinct considerations against each other and verbally identifying areas of unclarity. Domains (such as medical radiology) that depend on the comparison and interpretation of complex visual images are especially difficult to automate, since low-level visual processes are not open to voluntary inspection or control. Indeed, experts often give highly misleading advice about how they may be carrying out the relevant comparisons (eye-movement studies show, for instance, that expert radiologists do not scan X-ray photographs in the way they say that they do). In tandem with the increasing experience of AI-trained knowledge-engineers, further psychological studies of the organization of knowledge in different domains should be useful.

Research on expert systems will also focus on the computational architecture required to deal with large, complex, knowledge bases. Current systems are relatively simple and inflexible, and restricted to very narrow domains. They can be incrementally improved, but only up to a point. Eventually, the interactions between the increasing number of independently-added rules become too difficult to control, and the system's reliability and intelligibility are jeopardized. Current systems have no access to higher-level representations of the knowledge domain and their own problem-solving activity (see below). Special problems arise if a system has to work in real-time, where unexpected events can require quick switching from the current activity to some other. The

next ten years will see some general work on powerful IKBS architectures (as well as the production of more examples of specific commercially useful systems), including parallel-processing devices.

Progress can be expected also in natural language processing, both of individual sentences and of texts. Key issues include syntactic parsing, the integration of syntax with semantics, and the understanding of connected text. Machine translation could in principle benefit from advances both in single-sentence parsing and in text analysis.

Current work on parsing is motivated both by theoretical (linguistic) interests, and by the hope of improving the man—machine interface so as to make it possible for nonspecialist users to communicate with programs in (some reasonable subset of) natural language. Where a program is used for some specific purpose, semantic factors can more readily be used to help in the parsing and disambiguation of queries and instructions input by the user. Verbal interchanges about lunar geology, or about airline reservations, are already reasonably "natural" because of the exploitation of semantic constraints, and further domain-specific semantics will be developed over the next decade. More generally applicable (theoretical) research will continue into the best point at which to use semantics in parsing: from the beginning of the sentence, or spreading out from the middle, or only after an initial parse of the entire sentence?

Text-analysis programs can already give a précis of most short news-stories about specific topics (such as earthquakes, hijackings, and road accidents). But they rely on rigid, preprogrammed schemas, which provide the semantic skeleton of the types of stories concerned. Some recent research is aimed at enabling a text-analysis program to learn new schemas for itself, to integrate one schema with another so as to understand a story combining both, and to use a given schema to reason analogically in an unfamiliar context. A high degree of success cannot be expected within the next ten years, but our understanding of the relevant problems should be advanced.

A variety of educational applications are already receiving attention. Some are focussed on particular curricular subjects, and require both a model of the theory of that subject and a model of the student's knowledge of it (which varies in level and in organization, from person to person, and from time to time). Others are less specific, and aim to use AI-based techniques to improve the pupil's attitude to intelligence in general. There is some evidence that both normal and handicapped students can attain greater self-confidence and intellectual achievement by experience with these specially-designed programming environments. Controlled research into the classroom effects of AI-based systems has recently been initiated, and this can be expected to bear fruit within the next decade.

An extremely important area, which is increasingly being studied because of recent hardware developments, concerns the computational properties of large parallel systems. At present, we understand very little of the potential and limitations of such systems. Some of the connectionist work mentioned above suggests that cooperative processing may have some highly surprising properties. For example, the number of individual processors required to make the "human" range of visual shape-discriminations appears to be markedly less than one would naturally assume. Again, making a connectionist system stochastic rather than deterministic *improves* its chance of finding an optimal solution. The computational properties of parallel machines will not be well understood for a long time, but experience with these new systems in the near future will doubtless lead to some advance.

Five topics studied recently in AI, and which will be further developed over the next ten years, are nonmonotonic reasoning, naive physics, self-updating memory, creativity, and machine-learning. I shall refer to these difficult problems in the section on "Long-Range AI Research" below.

1.2 Impacts on other sciences and technologies

The impacts of AI on other technologies will include many different examples of applications to individual problems. For example, an olfactory chip is being designed using AI techniques of pattern recognition. Given advances in very large-scale integration (VLSI), instruments and products of many different kinds will come to include chips whose design makes use of AI methods. Any commercial or industrial task that could benefit from even a limited degree of intelligence could in principle be performed better with the help of AI, so that the technological applications of AI will be extremely diverse.

Turning from technology to science, AI will influence other sciences in their general philosophical approach as well as their specific theoretical content. Indeed, psychology and (to a lesser degree) biology have already been affected by computational ideas. And, contrary to what most people assume, AI has had a humanizing effect in psychology. The behaviorists had outlawed reference to "mind" and "mental processes" as unscientific and mystifying, but AI — based as it is on the concept of *representation* — has made these concepts theoretically respectable again.

AI's influence will be especially strong in the psychology of vision and language, and (as noted above) it is likely that robotics will engage with the psychophysiology of movement. Psychological research will feed back into AI; for example, insofar as psychologists arrive at a

better understanding of the organization of knowledge their work may be useful in designing computerized expert systems. Cooperative interdisciplinary research should be encouraged: the institutional separation of empirical psychology and AI or computer science has hindered fruitful collaboration between these groups.

1.3 Social implications

Social impacts will be of various types. First, there will be effects on individuals and institutions brought about by specific applications of AI, such as expert systems for medical diagnosis, legal and financial advice, or educational help.

These programs will not merely provide a service (whose adequacy should be very carefully monitored), but will very likely change the social relations of the profession or institution concerned. For example, if general practitioners, or nurses, can use an AI program to aid in various aspects of patient-care, the social image of the specialist physician may be profoundly affected. (And legal responsibilities for medical decisions may be assigned in a way very different from today.) Likewise, legal programs may undermine the status of lawyers, and alter the nature of their work. In both cases, while the mystique of the human experts may be lessened, their opportunity for exercising their specifically human powers may be increased.

The general public might come to be less dependent on human experts than they are today. Reducing the power of professionals such as doctors, lawyers, and teachers would certainly have advantages. But replacing human professional advice by computer programs will be dangerous to the extent that AI systems in public use are inadequate − and/or ill-understood. Systems that have taken several man-years to develop (and whose original programmers may be retired, or dead) are often very difficult to evaluate or alter, because even computer scientists do not fully understand how they work. (It follows that attention should be given to methods of perspicuous documentation, to help make clear what it is that a given program is actually doing, and how.)

A second type of social impact concerns general social trends brought about by applications of AI and information technology. These include changes in the proportion of the workforce in service and leisure industries, changes in the division of labor and sexual roles, and changes in general lifestyles and patterns of interaction.

For example, males will be increasingly freed to take up jobs in the "caring" professions (such as nursing, education, and social welfare). This could change the general evaluation of emotionality in the

masculine role, an effect that could also be encouraged by men's having increased leisure time to spend with family and friends. Such an effect could be liberating and humanizing, leading to a more convivial society than we have today.

But other potential consequences of AI point in the opposite direction. The widespread use of home terminals, for instance, threatens to have an isolating influence even more powerful than that of television. If people are encouraged to work, and to shop, from their sitting-rooms, there may be unfortunate psychological effects in terms of personal stress and loneliness. Community computer-centers could offset these effects to some extent, providing a social meeting-place outside the confines of the home and nuclear family. Some writers even predict that commercially available (and highly profitable) AI-systems will be heavily used, not only in task-oriented ways but as surrogates for human contact. On this view, the strong tendency to anthropomorphism that most of us share will result in patterns of interaction being skewed away from human beings, and towards quasihuman computer systems (with naturalistic "voices," and sometimes even "bodies"). Although such forecasts grossly underestimate the technological difficulties involved in building programmed "friends-off-the-shelf," they do suggest that human interactions could be impoverished to some degree in the future.

These contrasting examples show that widespread application of AI will have subtle, and varying, influences in society. Moreover, AI could foster a general view of humanity as either "mechanistic" or "nonmechanistic," depending on how it is interpreted by the public. The commonest interpretation is that AI presents us as "mere machines," with no free choice or moral responsibility. Since this image of man could have socially pernicious effects, people should be helped to understand that it is fundamentally mistaken. The education or computer literacy discussed below could help here. More generally, we should start thinking now about what the optimal social arrangements might be for a postindustrial society.

1.4 Economic aspects

The economic impacts will be far-reaching. Traditional manufacturing and clerical—administrative jobs will be decimated. But new jobs will be created: some directly connected with new technology (like computer engineers and programmers), others made possible because people are freed to devote their time to services (caring professions, education, leisure). Whether there will be enough new jobs to compensate for the loss of old ones (as has always happened in the past, at

least eventually) is however unclear, for AI can potentially apply to all jobs where personal human contact is not essential. New methods of work-sharing and income-distribution will have to be worked out (with income not necessarily being closely linked to jobs). Radical structural changes in society are likely, and the transition phase will not be easy.

2 Long-Range AI Research

There will be "more of the same," in that the areas mentioned above will provide perplexing problems for many years to come. Especially hard problems include learning, high-level vision, naive physics, and abstract work in computational logic – including the development of a taxonomy of representations and computational processes, showing the potential and limitations of distinct types.

I referred earlier to short-term research on IKBS architecture. But the deep problems involved in the organization and control of large knowledge bases will not be solved within a decade. And this is quite independent of the fact that parallel machines may support forms of inference radically different from those implemented today.

For instance, expert systems are at present unable to explain their reasoning except by "backwards-chaining": giving a resumé of the chain of inferences (rules) that led up to their conclusion. They cannot relate their conclusion to the domain in general, nor rely on an overview of the problem to assess the relative theoretical reliability of different hypotheses (probabilities are of course built into rules, but are assessed for each rule individually or in relation to a small number of other specific rules). Nor can they monitor and adjust the structure of their own problem-solving, for they have no high-level representation of it. They are also unable to integrate different knowledge-domains, or to use concepts and patterns of inference taken from one domain to reason (analogically) in another. Nor can current systems explain their conclusions differently to different users, taking account of the specific user's knowledge. The user can ask for a "deeper" explanation (a more detailed inference-resumé), but the program has no user-model in terms of which to adjust its explanations to the human's particular range and level of knowledge. For this reason also, the pattern of interaction between user and system is at present very limited. The user cannot offer his own conclusions for comment and criticism, for example, as students can do with human teachers.

All of these abilities that current programs lack will need a richer understanding of the structure and content of different knowledge-domains. Some of the projects mentioned elsewhere (such as research on naive physics) are highly relevant to many domains, and

psychological research into human reasoning processes could be useful also. This is just one illustration of the fact that theoretical and empirical research may be needed for radical improvements in techno-logical applications.

The need for a model of the user's knowledge also delays advances in educational programs such as those mentioned above. In principle, computer-assisted instruction based on AI techniques could be highly flexible, and subtly attuned to the student's particular strengths and weaknesses. But this requires that the program be equipped with a representation of the rich content and inferential organization of human knowledge in the relevant domains (which, in turn, requires psychological understanding of a high degree). To achieve this will be a difficult task, for the long term rather than tomorrow.

A special case of human knowledge is "naive physics," one's every-day knowledge of the properties and behavior of different sorts of physical substances, and the nature of the causal relations between them. This knowledge enters into vision and motor control, and also into natural language. For example, a language-using program would have to understand the differences in meaning between verbs such as *pour, flow, spill, drop*, and the like, if it were to give instructions or understand texts about activities dealing with liquids. Similarly, a robot capable of seeing that a container was just about to spill its con-tents onto the object below, and of adjusting its movements accord-ingly, would need some representation of the behavior of fluids. (It might of course be programmed to halt movement if it saw an unex-pected patch appearing at the rim of the container, but that is a dif-ferent matter.) Very little work has been done on these issues so far, and they are likely to provide a challenge for many years to come.

Another topic that is likely to receive much attention in the future is truth-maintenance when using nonmonotonic reasoning. Tradi-tional logical systems are monotonic, in the sense that propositions are proved once and for all: if a proposition has been inferred as true (or as false) on one occasion, its truth-value cannot change thereafter. But in commonsense reasoning, a proposition may be taken as true for very good reasons, but later found (or inferred) to be false. AI systems dealing with complex problems involving incomplete knowledge similarly require nonmonotonic reasoning, and new canons of inference are needed to control such knowledge-systems, and to prevent them from falling into absurdities.

The development of self-updating computer memories is closely related to the issue of nonmonotonic reasoning. But in addition to allowing changing truth-values, such a memory-system needs to be able to make inferences of many different sorts on being told "one" new fact. Human beings do this every day. For example, if one is told that

an acquaintance is a supporter of a particular political group, one's internal representation of that person may "unthinkingly" alter in many different ways (and one's attitudes and future behavior regarding the person will be influenced accordingly). An intelligent program presented with new information ought to be able to do the same kind of thing. Although some preliminary AI work has been done on this problem, it is not yet well understood.

"Computer-aided design" is typically thought of as involving the graphical display of precise three-dimensional specifications of various products (from machine-tools, through cars, to buildings), taking into account a wide range of values of many parameters. But a recent form of computer-aided design involves suggestion rather than specification, in the sense that the design program originates novel ideas – which are not merely quantitatively different from previous specifications. For example, heuristic programs are already being used to suggest novel experiments (described at the intramolecular level) in genetic engineering, or to help design new sorts of three-dimensional silicon chips. These programs were developed in tandem with a closely similar system that originates interesting mathematical ideas from a basis of elementary set-theory.

The potential of systems like these should be further explored. The computer modeling of creative thinking will require long-term research, especially with respect to domains whose crucial concepts cannot be so readily defined as the concepts of molecular biology, chip-circuitry, or set-theory. We need a better understanding of how conceptual structures (and the inference-patterns associated with them) can be explored, represented (on varying levels of abstraction), compared, and indexed. Highly interconnected processing networks may turn out to be useful for mediating "unexpected" conceptual associations. But association is not enough: associations need to be evaluated and controlled, and integrated into previously existing cognitive structures. As yet we have little idea how to do this so as to model creative thinking.

Machine-learning is a pressing problem for the future. If a program cannot learn for itself, its development is limited by the time and ability of the programmer to provide it with new information and ideas. The system should be able to induce regularities for itself. Some progress is being made in enabling computer systems to learn about specific aspects of particular domains. But "open-ended" learning, where what is to be learnt is not defined beforehand, is especially intractable.

Some connectionist workers have presented a general "learning algorithm," claiming that the input of large numbers of instances (of visual scenes, for example) could enable a connectionist system to learn

to recognize the structure of the input class, irrespective of what that structure is. However, to say that something can be done in principle is not to provide a practical, usable way of doing it. These claims cannot be further explored until suitable hardware is available (dedicated machinery is currently being designed).

In general, the properties of parallel computation will be a focus of research in the long term as well as in the next few years. It remains to be seen whether the Japanese hopes concerning VLSI and PROLOG (a logic-based, nonprocedural, programming language that is thought to be especially suited to a parallel architecture) will be achieved. But massively parallel hardware will increasingly become available, and will enable AI to progress in ways that are impossible using traditional types of machine.

3 Some Areas Worthy of Particular Support

Research areas can be supported for their intrinsic scientific interest and/or for their social usefulness. Low-level vision and robotics include work qualifying on both counts, and solid progress is likely within the next ten years. Naive physics is less well developed, but is likely to be important not only for advanced robotics but for language-understanding too.

Research in computational linguistics and speech-understanding merits support for its practical uses and theoretical interest. User-friendly programming environments and man—machine interfaces require natural-language "front ends." Although these do not need to handle every linguistic subtlety, and so can ignore many problems that are of more theoretical interest, there is still much room for improvement.

Support for IKBS should encourage basic research into general issues of system-architecture and nonmonotonic reasoning, rather than leading to the proliferation of the relatively simplistic systems available today. This is a long-term project, but essential if AI systems are to be widely used in decision-making contexts.

More research is needed on the educational applications of AI. A few groups have already started to study the effects of giving children (of various ages) access to the "LOGO" programming environment in the classroom. LOGO is a programming language specially developed for use even by young children; it utilizes a mechanical "turtle" (or an equivalent on a TV screen) to draw pictures of various sorts. Some experience is also being gained in using LOGO to help gravely handicapped children. As noted above, preliminary results suggest that this programming environment helps both normal and handicapped children

to express and develop their intelligence, emotional relations, and self-confidence. As with new educational methods in general, it may be the enthusiasm and commitment of the pioneers involved that is crucial. Carefully controlled studies in a range of schools, involving a range of teachers, are needed to evaluate the claims that have been made in this context.

Psychological research into the organization and use of knowledge in different domains could contribute usefully to applications of AI. As mentioned above, both educational and "expert" AI programs will need an internal model of the student/user to enable them to interact in flexibly appropriate ways.

The general problem of computation in parallel systems has been referred to several times already. It is clearly an important area. For a few years yet, we can expect exploration rather than exploitation. But this exploration of the potential and limitations of such systems is essential.

Funds should also be made available for combatting the ignorance and sensationalism that attends AI today. Research on friendly programming environments, and on interactive "programmer's apprentices," should be supported. This involves not only work on natural-language interfaces, but also psychological studies of how people learn to program and (what is not the same thing) how they carry out and interpret an interaction with a quasi-intelligent program. It may be that certain words or phrases, and certain ways of structuring the interaction, help users to appreciate the specific limitations of the program they are using, and remind them that they are interacting not with a person but with an artifact. Some universities have already begun to develop programming environments and exercises designed primarily to awaken naive users to the potential *and the limitations* of AI programs, and the general educational value of such experiences should be explored.

One might ask why widespread ignorance about AI *matters*. Part of the answer is obvious: in a society where most jobs involve access to computerized facilities making use of AI techniques, individuals without any understanding of AI will be at a disadvantage (and the more of them there are, the more social unrest is likely). But there is another important consideration, which can be illustrated by an advertisement recently shown widely on British television.

The advertisement showed six people sitting at six computers, each sold by a different manufacturer. The "voice-over" message said something to this effect: "We provided details of the performance and cost of six different computers to the six computers themselves, and asked them to choose the best. The X chose the X (I shall not advertize the firm further by giving its name here) – and so did all the

others. It makes you think that a person ought to choose the *X* too."

This type of persuasion is pernicious, for it deliberately obscures the fact that each machine was running the same choosing-program, which someone had to write in the first place (the "someone" in question being, of course, an employee of firm *X*). People who do not understand what a program is – who do not realize that not only its data, but also its inferential or evaluative processes, are in principle open to challenge – may indeed be gulled into believing that "If computers choose something, then we should choose it too." If the choice merely concerns the purchase of one commodity rather than another, this is perhaps not too worrying. But if it concerns more socially or politically relevant problems, such mystlfication could be most unfortunate.

Sensationalism feeds on ignorance, and many descriptions of artificial intelligence in the media, and in "popular" books about the subject, are sensationalist in nature. Whether proclaiming the "wonders" or the "dangers" of AI, they are not only uninformative but highly misleading – and socially dangerous to boot. They suggest that things can be done, or will be done tomorrow, which in fact will be feasible only (if ever) after decades of research (including the "long-range research" mentioned above). And they underplay the extent of human responsibility for these systems, in much the same way as the *X*-advertisement described earlier.

Unfortunately, these sensational reports are sometimes encouraged by ill-judged remarks from the AI community itself. A recent hour-long BBC-TV science program began and ended with a quote from a senior computer scientist at MIT, gleefully forecasting that the intelligent machines of the future would worry about all the really important problems for us (*for* us, not *with* us). As he put it (with apparent satisfaction): if we ever managed to teach chimps to speak, we wouldn't talk to them for long – for they would want to talk only about bananas; super-intelligent machines will be similarly bored by people, for we won't be capable of understanding the thoughts of the machines. His conclusion was that the super-intelligent AI-systems will justifiably ignore us, leaving us simply to play among ourselves.

Humanity has of course been advised before to neglect the difficult moral and philosophical questions, to live life on the principle that *"Il faut cultiver son jardin."* But that was said in a rather more ironic spirit. Enthusiasts evaluating AI's contribution to society would do well to emulate the common sense, if not the skepticism, of Voltaire.

References

Ballard, D.H. and C.M. Brown (Eds.) (1982) *Computer Vision.* Englewood Cliffs, NJ: Prentice-Hall.

Boden, M.A. (1977) *Artificial Intelligence and Natural Man.* Brighton, UK: Harvester.

Boden, M.A. (1983) Educational Implications of Artificial Intelligence. In W. Maxwell (Ed.), *Thinking: The Frontier Expands.* Pittsburgh, PA: Franklin Press, pp. 227–236.

Boden, M.A. (1984) Artificial Intelligence and Biological Reductionism. In M. Ho and P. Saunders (Eds.), *Beyond Neo-Darwinism.* London: Academic Press, in press.

Bond, A.H. (Ed.) (1981) *Machine Intelligence* (Infotech State of the Art Report, Series 9, No. 3). Maidenhead, UK: Pergamon.

Brady, J.M. (Ed.) (1981) *Computer Vision.* Amsterdam: North-Holland.

Davis, R. and D.B. Lenat (Eds.) (1982) *Knowledge-Based Systems in Artificial Intelligence.* New York: McGraw-Hill.

Erman, L.D., P.E. London, and S.F. Fickas (1981) The Design and an Example Use of Hearsay III. *Proceedings of the International Joint Conference on AI, Seventh,* Vancouver, BC.

Feigenbaum, E.A. and A. Barr (1981) *The Handbook of Artificial Intelligence* (3 vols.). London: Pitman.

Feigenbaum, E.A. and P. McCorduck (1983) *The Fifth Generation.* New York: Addison-Wesley.

Hayes, P.J. (1979) The Naive Physics Manifesto. In D. Michie (Ed.), *Expert Systems in the Micro-Electronic Age.* Edinburgh: Edinburgh University Press, pp. 242–270.

Hayes-Roth, F., D.A. Waterman, and D.B. Lenat (Eds.) (1983) *Building Expert Systems.* Reading, MA: Addison-Wesley.

Hinton, G.E. (1981) Shape Representation in Parallel Systems. *Proceedings of the International Joint Conference on AI, Seventh,* Vancouver, BC.

Hinton, G.E. (1983) Optimal Perceptual Inference. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.* Washington, DC, June.

Hinton, G.E. and J.A. Anderson (Eds.) (1981) *Parallel Models of Associative Memory.* Hillsdale, NJ: Erlbaum.

Jenkins, C. and B. Sherman (1981) *The Leisure Shock.* London: Eyre Methuen.

Lehnert, W.G. and M.H. Ringle (Eds.) (1982) *Strategies for Natural Language Processing.* Hillsdale, NJ: Erlbaum.

Lenat, D.B. (1982, 1983) The Nature of Heuristics (three papers). *Artificial Intelligence,* 19 and 20.

McCarthy, J. (1980) Circumscription – A Form of Non-Monotonic Reasoning. *Artificial Intelligence,* 13, Nos. 1 and 2.

McDermott, D. and J. Doyle (1980) Non-Monotonic Logic. *Artificial Intelligence,* 13, Nos. 1 and 2.

Marr, D. (1982) *Vision.* San Francisco, CA: Freeman.

Nilsson, N.J. (1980) *Principles of Artificial Intelligence.* Palo Alto, CA: Tioga.

Papert, S. (1980) *Mindstorms: Children, Computers, and Powerful Ideas.* Brighton, UK: Harvester.

Paul, R.P. (1981) *Robot Manipulators: Mathematics, Programming and Control.* Cambridge, MA: MIT Press.

Rich, E. (1983) *Artificial Intelligence.* New York: McGraw-Hill.

Schank, R.C. and C.K. Riesbeck (Eds.) (1981) *Inside Computer Understanding: Five Programs Plus Miniatures.* Hillsdale, NJ: Erlbaum.

Selfridge, O., M. Arbib, and E. Rissland (Eds.) (1984) *Adaptive Control in Ill-Defined Systems.*

Sleeman, D.H. and J.S. Brown (Eds.) (1982) *Intelligent Tutoring Systems.* London: Academic Press.

Stefik, M., J. Aikins, R. Balzer, J. Benoit, L. Birnbaum, F. Hayes-Roth, and E.D. Sacerdoti (1982) The Organization of Expert Systems. *Artificial Intelligence,* 18:135–173.

Waterman, D.A. and F. Hayes-Roth (Eds.) (1978) *Pattern-Directed Inference Systems.* New York: Academic Press.

Winograd, T. (1983) *Language as a Cognitive Process: Syntax.* Reading, MA: Addison-Wesley, (chapter on computer models).

Winston, P.H. and R.H. Brown (Eds.) (1979) *Artificial Intelligence: An MIT Perspective* (2 vols.). Cambridge, MA: MIT Press.

Proceedings of the International Joint Conference on Artificial Intelligence (biennial, odd-numbered years).

Artificial Intelligence (quarterly journal).

Proceedings of the American Association for Artificial Intelligence (annual).

Proceedings of the European Society for the Study of Artificial Intelligence and Simulation of Behaviour: AISB (biennial, even-numbered years).

International Journal of Robotics Research (quarterly journal).

CHAPTER 6

The Intersection of AI and Education: Investing in Thought-Intensive Endeavors

Stefano A. Cerri *

1 Introduction

In this paper we examine how people, groups, and countries that do not have access to large human and financial resources can, nevertheless, contribute to and benefit from high-level scientific and technological developments. In our opinion, artificial intelligence (AI) promises to generate rewarding, if unpredictable, results provided that some (relatively limited) resources are invested in medium- to long-term R&D projects with thought- or expertise-intensive goals.

The optimism of this statement goes hand in hand with a fundamental pessimism arising from other considerations – for example, the conviction that only the United States and Japan are at present able to envisage any real expansion of their production, because they alone seem to have access to the high technologies needed to cope with the keen competition in the modern international market. This implies that even countries in Western Europe (and obviously all those countries that are less industrialized) will play an ever more subordinate role over the years to come. Such a picture should not appeal even to the US and Japanese authorities and, if needed, we are ready to present many reasons to show why this sort of situation is, in fact, convenient for no one.

*Department of Informatics, University of Pisa, Corso Italia 40, 56100 Pisa, Italy.

So, let us start from the hypothesis that a more even distribution of technological – and hence economic and social – development in the world is beneficial for everyone. "Even distribution" and the associated "fair growth" will be referred to generally as "our purpose" from now on.

One of the most important recent developments in this context is in information technologies. AI and its related methods, techniques, and products are among the most promising branches of these technologies (it is stated as a basic premise for launching the ESPRIT project in the EEC, the Alvey project in the UK, the Japanese 5th Generation project, and the recent US investments in these fields).

Assuming that one agrees that AI is of great strategic importance, the question arises as to whether it is realistic to start any R&D endeavor in a technological field, with limited financial resources and expertise, and hope to compete in a reasonable time with countries, groups, and people that have access to greater funds and a better established system of scientific education. We do not propose to attempt any comprehensive answer to this question here. Instead, we would like to formulate the problem and indicate possible solutions by examining the three major intersections between AI and education: namely, AI education, learning and the automatic acquisition of knowledge by machines, and intelligent tutoring systems.

The line of inquiry will be as follows:

- *Economic Impacts*. Assuming that the development of AI R&D projects is strategically important for the development of information technologies, and therefore has a great potential economic impact, are there aspects of AI R&D that distinguish it from other R&D projects, in such a way that – tentatively – we can conclude that AI R&D is potentially suited to "our purpose" (even distribution and fair growth)?
- *Social Impacts*. Assuming that the growth of information technologies implies a corresponding growth in their use, that both are somehow unavoidable, and that they will certainly have a non-trivial social impact, are there aspects of AI R&D that single it out in such a way that – tentatively – we can conclude that this social impact may be positive?

2 Information Technologies: Computer Science *versus* Artificial Intelligence

Many criteria can be considered to discriminate between traditional computer science and the AI approach to informatics (see for

example Newell 1982). In this paper we cannot explicitly discuss Newell's considerations, but we consider his work to be particularly illuminating in helping us to understand the history and the nature of AI as compared to traditional computer science.

In brief, whereas traditional computer science is mainly concerned with numerical and well-defined problems to be solved algorithmically by constructing efficient computing systems, AI is interested in symbolic, ill-defined problems that are hard to formalize, but that can be solved by designing adequate representations of the knowledge available, possibly in terms of a hierarchy of levels or viewpoints, and nondeterministic search strategies and reasoning systems.

The fundamental question is not whether such AI systems are consistent, complete, or efficient, but whether they provide an adequate solution to the original problem. At a later stage it is necessary to verify the consistency and completeness of the system in order to ensure the correctness of the solutions. Finally, efficiency considerations allow one to assess whether or not the system is (economically) viable.

The shift from the efficiency to the adequacy of the solution as a primary goal has some immediate consequences:

(a)   it reverses the traditional approach to information technology: it is now necessary to design special-purpose software and hardware starting from the specific needs of the problem and ending up with a complete system;

(b)   it requires considerable efforts and expertise on the part of both the expert in the application domain and the expert in the formalization of the knowledge needed to solve a problem in that domain;

(c)   it can have a considerable impact, not only on applications in the domain concerned, but also on computer technology more generally, provided the solution found can be shown to be valid in other domains.

In passing I would like to note that I believe the knowledge framework and the knowledge-engineering approach are promising candidates for solving many of these problems and that they will remain relevant for some years to come. However, since these issues have been considered by many other contributors, I will not discuss them further here. Instead, I shall concentrate on the need for high-level competence of a qualitatively new *type* in the construction of knowledge-based systems.

3 Needs for Specific Know-How in AI

Whereas one can expect a traditional computer scientist to be capable of designing an efficient algorithm for the solution of a well-formulated problem, it is hard to imagine that the same skills would be sufficient to solve problems for which no formalization is available, because the solution can only be induced from the behavior (possibly unconscious) of an expert in the object domain. In other words, the traditional computer scientist does not *in principle* have an adequate education (mastery of concepts and skills, and ability to communicate) for the tasks fundamental to AI.

Thus, *education* is central to any attempt to take up this challenge. We emphasize this point because our experience to date in European countries has been somewhat negative and we believe that a different attitude, from both public and private institutions and companies, could significantly modify the situation in a relatively short time.

There are at least five main reasons why AI education has not yet gained much attention in Europe:

(1) AI has not been considered a "formal enough" field to be a candidate for academic interest;

(2) for all relatively new fields, Europe has tended to respond more slowly than the United States. In the best cases, AI education is offered in Europe as a "side" option in computer science courses in the last years of the undergraduate curriculum, while the conceptual content could and should be assimilated at a much earlier stage, possibly even in high school;*

(3) there has been little or no interest from the consumers, i.e., industries;**

(4) even when there has been interest from industries in some fields of science or technology, the reaction of the academic sector has normally been very slow: industrial–academic cooperation generally has been poor;

(5) since AI is interdisciplinary (involving computer science, psychology, linguistics), no traditional department has been able to catalyze AI (educational) activities, but each specific department has been strong enough to inhibit the emergence of the new, competitive discipline.

---

*An exception – proving the rule – is given by the Department of Artificial Intelligence in Edinburgh.

**Only recently I read for the first time an advertisement in an Italian newspaper calling for AI experts.

Let us now consider AI education as a prerequisite for successful AI R&D projects. We have claimed that AI requires a new *type* of know-how. We would now like to suggest that Europe on the one hand, and the United States and Japan on the other, possibly for different reasons, should both sponsor activities on AI education and fundamental AI research, i.e., projects for building up fundamental, specific know-how in artificial intelligence. For the Europeans, one very good reason might be represented by the statement that, while it is not easy to build up know-how, it is still cheaper than importing it! In fact, we believe that the technological gap between the United States and Japan and, say, Western and Eastern European countries cannot be bridged merely by studying existing prototypes embodying technologies that are, or will soon be, available in products on the commercial market. Europeans should enter a new area, and preferably the one where they seem to have the maximum chance of investing fruitfully, namely, thought-intensive R&D. Obviously, they must build *on*, and not ignore, the state-of-the-art technology.

Further investment in AI education and fundamental AI research would also pay dividends in the United States and Japan, because accelerating the process of "know-how formation" in a time of crisis can be vital, over the medium term, for economic expansion.

New and economically promising fields can be opened up while solving important current problems, such as the growing need for selective information exchange as an adequate basis for decision making (for example in offices, or in managing the explosion of scientific information contained in technical reports), access to natural resources (for example oil), the need for significant growth in agricultural production, and the need for democratic scrutiny of weapons control, which are all fields that could potentially benefit from basic AI advances.

In the United States and Japan, research and development in AI must address relevant issues in a systematic way; these are thought-intensive endeavors and not simple exploitations of existing tools (see Schank 1983). It is regrettable that most, if not all, of the US PhDs in artificial intelligence tend to move from research to the industrial world.

All of these considerations rely on economic evaluations that are far beyond the scope of this paper. However, we can estimate some costs for building and maintaining an AI R&D laboratory. Roughly, we need about two–four years' education for the average trainee, we would pay him/her at a rate that might be about 30–50% higher than the going rate for another scientist/technician – because the field has not yet stabilized – and we would envisage spending an average of $50 K per man in special hardware/software resources. We believe these costs

are affordable for any group that is seriously interested in high technology and sensitive to the arguments presented here.

4 Is AI Special in Any Way?

Artificial intelligence is still in its infancy, like a framework for future, dimly-perceived developments in high-level technology. Therefore a certain amount of instinctive judgment underlies any position concerning the real nature of AI and what its impact will be. One easy option would be to state merely that, since AI is new and appears promising, everyone should hurry so as not to miss the expected rewarding results. However, we feel it is more valuable here to make bolder and more specific statements in order to stimulate reactions. We would claim that artificial intelligence is quite distinct from other disciplines, and that its characteristics justify an optimistic view about the expected economic and social impact of AI with respect to what we previously called "our purpose."

Many scientists have complained that AI results have not yet been cumulative, in the sense that an existing "intelligent" program or achievement is of almost no use for a new and different project. This aspect of artificial intelligence is not a temporary phenomenon: we believe that it is actually inherent in the process. In fact, the formalization of a piece of knowledge or reasoning is a (relatively) new enterprise every time. If it were possible to build a completely additive theory of artificial intelligence, we would, in effect, have constructed a theory of thought, which remains very improbable over the next few centuries.

So, each new problem requires us to think hard, to attempt to formalize the knowledge at hand and to structure our own formalization process. Notions such as heuristics, and imprecise, approximate, or contradictory knowledge are very far removed from any general, logical—deductive theory that embodies every solution to every problem.

The fact that "languages" such as mathematics or logic are helpful in explaining how the knowledge manipulation in AI programs actually works does not contradict the statement that the *selection* of knowledge and knowledge-processing mechanisms is *not* a purely logical process. Just as in mathematics, the intuition behind a theorem is not as logical as the proof. Concerning this controversy, I favor Newell's point of view on Nilsson's statements about logic in AI (see Newell 1981, Nilsson 1981).

The inductive, experimental character of the construction of expert systems — the embodiment of AI methods and tools — by

formalizing the expertise of humans and transforming it into running systems, makes AI difficult for a beginner but accessible to anyone who has been trained in its methods and tools. Though AI seems conceptually difficult to a novice, we do not expect – perhaps surprisingly – that we will have to expend much effort in this training, because the structure of reasoning in AI is much more similar to everyday human thought than is the case, for instance, with traditional mathematics or even disciplines such as physics or chemistry. Moreover, intelligence is not a question of "all or nothing," so we can grade our ambitions accordingly. At the University of Pisa we gave two one-year (90-hour) courses to about 100 computer science students, and by the end of the courses we considered the participants to be at least potential AI workers. Some of them were so enthusiastic and worked so hard on their own that we are sure they are capable of producing high-level research products within a relatively short time and becoming tutors of other novices.

We would therefore contend that, though results in AI are not cumulative, AI education and training are very much so. As a general reason for this phenomenon, we can point to the fact that AI education cannot be purely formal, and that it therefore relies much more on commonsense intuition and specific engineering practice.

Nilsson (1981) states that "AI research should be more concerned with the general form and properties of representational languages and methods than it is with the content being described by these languages," but also that "notable exceptions involve 'commonsense' knowledge about the everyday world and metaknowledge."

Although we agree that the purpose of AI is to design general, formal methods and tools for representing and processing knowledge in machines, we find this distinction between form and content somewhat inconsistent. In fact, it seems to us that the content of research into metaknowledge is precisely the general form and properties of representational languages and methods when these represent knowledge about the world, and that the study of commonsense knowledge cannot be abstracted from its content, i.e. specific everyday problems about which one reasons using commonsense (heuristic) methods.

On this basis, we would argue that the form and content of representational languages and methods are inherently linked and that a more correct definition of the concerns of artificial intelligence would be the study of formal and commonsense reasoning about the world by embodying this reasoning in machines using knowledge-representation languages and methods and the study of metaknowledge.

Our point of view is that AI is basically an experimental discipline – the empirical study of knowledge; its theoretical framework is

constructed by applying the experimental methods and tools to AI itself. This experimental aspect of AI may also be beneficial in promoting better cooperation between industry and the academic world. The rigorously formal, logical deductivism of traditional science and engineering education (particularly in Europe) has damaged the creative, inductive processes that might otherwise be fruitfully employed in the experimental sciences and the design of technological artifacts.

There is another aspect that is rather specific to AI, which can be usefully applied for "our purpose." The cooperation between two types of source of knowledge in the construction of expert systems requires not only knowledge in the subject domain and in AI, but also creative skill in combining the two in a unique system that is only satisfactory if it proves adequate for the purpose (cf. the discussion of adequacy at the beginning of Section 2). This skill is missing in many industrial and academic contexts, but, once gained, it can be very widely applied. It requires only those resources that are likely to be available even in low-budget projects, i.e., intelligence and hard intellectual work.

We have presented above some of our convictions about the need for AI education. Many qualified persons have argued convincingly that the study of learning and automatic knowledge acquisition in expert systems should be a central *research* issue in all further AI developments (see for example Schank 1983).

We believe that the major components of intelligent tutoring systems (ITS) could represent key endeavors in AI *development*, i.e., building complex AI systems based on the state-of-the-art know-how. We explicitly refer here to such issues such as the analysis of correct and incorrect problem-solving behavior from human protocols, natural knowledge (linguistic, graphic, and possibly pictorial) understanding, representation and production, cooperative mixed-initiative behavior in a "mutual (incomplete) knowledge" framework, integrated (hardware—software) knowledge compilation, etc. When building an "intelligent" tutor one can begin with a traditional CAL*-like system and progressively add "intelligent" modules to make the interaction more "user-friendly" and effective. ITS are good testbeds for AI development projects, because they can address a wide variety of subject fields — including AI itself — and they can easily embody a spectrum of levels of "intelligence," without necessarily losing their effectiveness. The development of ITS during the seventies was one of the major sources of ideas, methods, and tools for AI in general.

---

*Computer assisted learning.

Finally, we would like to integrate the three fields, AI education, learning, and ITS, with the remark that the systematic organization of prescriptions for expert system design, needed to render the results cumulative, is based on the sort of creative skill that is normally attributed to a good teacher; this can be embodied in an intelligent system teaching AI notions, and this system might be able, in the future, to learn from its own experience.

5 Conclusions

We have tried to promote an optimistic view of AI as an activity with positive economic and social impacts even for low-budget countries or groups by examining whether AI can be considered as distinctly different from other disciplines. In a discussion of the three intersections between AI and education, namely AI education, research on learning, and the development of intelligent tutoring systems, we have identified a number of characteristics of AI. These are the *"lack of cumulativeness"* of the results, the need for *cooperation* between experts in the problem domain and knowledge engineers, the *experimental* nature of AI, and the likely *costs* of an AI project.

Thus far, it may appear that the social impact of AI is of relatively little interest. This is in fact by no means the case, and we will now try to show that the arguments presented above to describe the particular advantages of AI as a vehicle for economic development − if accepted − can also form convincing starting points for a discussion of the social impacts of artificial intelligence.*

As we said in the Introduction, the social importance of AI can be considered to be a direct consequence of the strategic importance of AI in the development of advanced technologies. However, this does not mean that any type or degree of technological growth whatsoever is socially acceptable.

Secondly, one of the main concerns of AI is the construction of "user-friendly" systems. Actually, the shift from the "efficiency" to the "adequacy" of the problem solution has always been paralleled, in the history of AI, by the shift of concern from the "computer" to the "man" in interactive systems. Since we now accept that the computer will become an everyday tool for everyone's activities, we can easily infer that AI and its methods will be important in the "Information Society" of the future. But while this may seem to be the major aspect

---

*In a recent paper, Steels (1984) attributes the major potential impact of AI on society to the possibility of limiting what he calls "knowledge erosion" rather than to simple economic developments.

of the potential social impact of AI, we believe that the next two considerations deserve even more careful attention.

Thirdly, as AI is concerned with the formalization of commonsense reasoning in fields that are not yet formalized, the development of AI can play a central role in the scientific organization of knowledge (or theory formation) in traditional "humanistic" disciplines (e.g., linguistics and psychology) but also in some aspects of disciplines — such as law or medicine — where problem solving and decision taking has traditionally been left to the "clinical eye" of the expert. Right down to apparently "lower" aspects of the organization of work in the office, the access to information allowed by the development of "user-friendly" systems and the power of AI systems in problem solving and decision making can have a radical impact on society, not only because these systems propose solutions previously unexplored, but also because the solutions proposed represent a formalization of the problems themselves. However, these applications of AI to real problems should not be considered as a direct expansion of the application of traditional information technology. The characteristic difference, we believe, is reflected in the following area.

Because AI is also concerned with *itself* as a field of application (the meta-level), the formalization, generalization, and abstraction power reflected in AI systems allows one to imagine a technological framework within which, possibly for the first time, the accumulation of results, theoretically hard to justify, is obtained simply as a consequence of practice. From this viewpoint, work on meta-knowledge and meta-reasoning have an impact on the parameters defining the social impact of AI systems applied to the solution of "external" problems (e.g., the speed of introduction of new technologies in different, yet "analogous" domains).

One can argue that this characteristic is also present in traditional computer science. The difference is that the AI paradigm seems to attack in a unique general framework — that of knowledge — all those levels of the representation and processing of information that were previously left to a large set of specific formalisms and languages. We do not know whether the AI metaphor can also be considered useful for modeling human behavior. Certainly, it is the most advanced paradigm now available for *describing* it.

Therefore, we conclude by presenting one last provocative statement, namely that networks of "intelligent" computers will be able, in the future, to simulate (autonomously?) that part of human social behavior — down through the centuries — that has usually been called intellectual development.

References

Newell, A. (1981) The Knowledge Level. *AI Magazine*, 2/2 (Summer):1–20.

Newell, A. (1982) Intellectual Issues in the History of Artificial Intelligence. In F. Machlup and U. Mansfield (Eds.), *The Study of Information: Interdisciplinary Messages*. New York: Wiley.

Nilsson, N.J. (1981/82) Artificial Intelligence: Engineering, Science or Slogan? *AI Magazine*, 3/1 (Winter):2.

Schank, R.C. (1983) The Current State of AI: One Man's Opinion. *AI Magazine*, 4/1 (Winter/Spring):3.

Steels, L. (1984) *Do We Need Artificial Intelligence*? AI Laboratory News, No. 1. Free University of Brussels.

CHAPTER 7

Artificial Intelligence: A Lesson in Human Self-Understanding

Ivan M. Havel *

> "The substance of history consists in the experiences by which man gains understanding of his humanity and together with it understanding of its limits." (Voegelin, New Science of Politics.)

1 Beginnings and Potential Impacts of AI

The middle of the twentieth century, with its particular technological, cultural, and intellectual climate, gave birth to a new scientific discipline: artificial intelligence (AI). The newborn is growing relatively rapidly compared to other scientific disciplines, and is already approaching adolescence. We are therefore, quite naturally, becoming seriously concerned about its future, and we ask ourselves whether we should try to take it in hand and if so, how.

AI has its own characteristic inner development (Newell 1982) but it also has an outer life, consisting of actual or potential (and apparent or hidden) impacts on society, technology, and life-style, and last but not least on our way of thinking.

For the purposes of this study I found it useful to discriminate between three kinds of impacts that AI, like many other human activities, may have.

---

*META, Engelsovo nb.78, 12000 Prague 2, Czechoslovakia.

Impacts of the first kind are the *intended* outcomes of AI, i.e., those that actually motivate specific research projects (I have in mind particularly the engineering trends in AI). These outcomes are or will be realized in various products like automatic consultants, advisors, tutors, computer psychiatrists, marriage counsellors, sophisticated knowledge bases, picture processors, identification and recognition systems, intelligent industrial and domestic robots, etc. Impacts of this kind were the subject of some earlier assessments (Firschein *et al.* 1973) and they are also a favorite topic for popular newspaper articles and pocket books.

As *impacts of the second kind* I classify those that are *unintended*. They may in principle be foreseeable by people with strongly developed imaginations, but they may also arrive quite unexpectedly. These kinds of impact are more likely to be valued negatively and as something beyond our control, but this is by no means a rule (similarly it is not a rule that impacts of the first kind are always positively valued).

The more obvious impacts of the second kind are possible side effects and/or misuses of the intended products of AI (for example speech-understanding systems used for Orwellian purposes like the surveillance of private conversations). There are, however, some deeper and less apparent impacts that were also not among the preconceived aims of AI. These latter take place in the "ideosphere" (Hofstadter's term) and involve shifts in viewpoints. For example, one of the main maxims of AI, that "machines may be smart," tends to reawaken the old myth of *l'homme machine* — with all its negative implications for human self-respect. Or, as Margaret Boden (1977) puts it:

> "If the public believes — rightly or wrongly — that science regards people as 'nothing but clockwork', then clockwork-people we may tend to become."

Impacts of both kinds have been widely disputed, within and outside of the AI community. Joseph Weizenbaum, for instance, is the main representative of insiders who have warned against irreversible dehumanizing or unethical effects of what he calls "instrumental reason" (Weizenbaum 1976).

Here I wish to draw attention to a further category of impacts, characteristic of AI, which I consider yet more fundamental than the previous ones. I call these *impacts of the third kind*.

2 AI Impacts of the Third Kind

Impacts of the third kind manifest themselves through *reflection*. That is, they can be discerned through full, conscious awareness of the new situation in which man finds himself, as a consequence of AI and the lessons it offers him.

The main intellectual contribution of AI is that it elaborates a computer metaphor for the mind. Thus, as a challenging theme, the mind, and together with it the mysterious "self," enter into the focus of scientific interest. This gives some grounds for hope that AI will over-come those much-talked-about dehumanizing effects we may be afraid of.

Let us illustrate this point by a passage from Boden's book:

> "Contrary to common opinion... the prime metaphysical signifi-
> cance of artificial intelligence is that it can *counteract* the
> subtly dehumanizing influence of natural science, of which so
> many cultural critics have complained. It does this by showing, in
> a scientifically acceptable manner, how it is possible for psycho-
> logical beings to be grounded in a material world and yet be prop-
> erly distinguished from 'mere matter'. Far from showing that
> human beings are 'nothing but machines', it confirms our
> insistence that we are essentially subjective creates living
> through our own mental constructions of reality (among which
> science itself is one). In addition, for those of us who are
> interested, it offers an illuminating theoretical metaphor for the
> mind that allows psychological questions to be posed with greater
> clarity than before. The more widely these points are recog-
> nized, both within and outside the profession, the less of a threat
> will artificial intelligence present to humane conceptions of self
> and society." (Boden 1977, p. 473.)

By virtue of the orientation toward the human self and due to the inherent value of any reflective activity, impacts of the third kind are fundamentally positive. Moreover, abstract as they may seem, they are more significant than "lower" kinds of impacts because they directly influence the social and intellectual climate – the very climate that gave birth to AI.

3 Metaphors, the Self, and the Workings of AI

Let us amplify some of the points touched upon above, as well as identifying some further aspects of impacts of the third kind.

3.1 AI elaborates a computer metaphor for the mind

There is a subtle but important difference between the metaphor and the model. Viewing the computer as a *model* of something involves the reductionist standpoint: the model provides the criteria for explanation, it forces us into its specific language. The modeled entity is virtually identified with the model; any difference is disregarded.

The metaphor, on the other hand, always reminds us of the difference: in fact, it is the tension between the identity and the difference that gives the metaphor its creative power. Metaphors are catalytic tools that, using the principle of analogy, lead to new ideas and to a new understanding. A metaphor initiates exploration, whereas a model terminates it. What is characteristic of metaphors is that they are not despotic; they never exclude other metaphors, even mutually incompatible ones.

We should bear in mind this distinction between modeling and metaphorical transfer, particularly when talking about computers. Take, for instance, the following statement:

> "...the computer is a powerful new metaphor for helping us to understand many aspects of the world, but ... it enslaves the mind that has no other metaphors and few other resources to call on." (Weizenbaum 1976, p. 277.)

Clearly, the second part of the sentence is talking about the computer as a model rather than as a metaphor.

3.2 AI and the self

The revival of interest in self, in human subjectivity, is already documented in the highly suggestive titles of books like *The Mind's I* (Hofstadter and Dennet 1981) or *The Brain and Its Self* (Popper and Eccles 1978).

AI is primarily a constructive discipline. It tries to construct something that would, at least by its performance, resemble, replace, or transcend human thinking, or in the case of robotics, resemble, replace, or transcend humans as such. Anyone who plans to launch such a project is inevitably confronted with questions like "What is thinking?" or "What is man?" As a matter of fact, since the questioner is himself human, he may pose the questions in another, self-referential way: "What is my thinking?" or "Who am I?" Now there is an important difference between these two ways of formulating the questions: the former is external, using the more common *outer view*, while the latter

is internal, relying on one's *inner view*. This is how *the self* enters into the picture.*

Is subjectivity realizable on the principles of AI? As I see it, the answer should be no. While constructing any artifact requires a prior *external* description, we experience our self only by means of the mind's *inner* eye. Our subjective self is more a view than a thing. Incidentally, this difference points to certain intrinsic limitations on writing computer programs that attempt to realize genuine intellectual activities. Even a very specific and rule-obeying activity like playing chess, if considered in depth, involves the subjective intentional self as an inseparable constituent.

But whether successful in its aims or not, AI research awakens our desire for greater self-understanding. This very desire is, I believe, the most positive impact of AI. The passage from Boden's book cited above illustrates the point. Its essence is that AI can counteract dehumanization and that this counteracting ability may be based on illustrating the difference between psychological beings and "mere matter." On the other hand, the claim that AI actually shows *how* the psychological being transcends matter is somewhat overstated. This "how" most probably refers to Boden's claim that

> "the crucial notion in understanding how subjectivity can be grounded in objective causal mechanism is the concept of an internal model or representation" (Boden 1977, p. 428.)

Maintaining an internal model of the world, perhaps even of hypothetical worlds and of the system's own capabilities, is one thing, but possessing inner *perspective* is quite another. What AI has shown up to now is just that computers, in a sense, transcend clockwork.**

3.3 Free will revisited

Self-determined action (Boden 1977), authentic choice (Weizenbaum 1976), or intentionality (Dennet 1978) – however we choose to define it – appears to be something other than just an opposite of determinism. Surely computers can be programmed to make deterministic decisions as well as to take random steps? But free choice is

---

*An exciting intellectual path from AI through the phenomenon of self-reference to the riddle of the self is one of the strands in Hofstadter's braid (Hofstadter 1979).

**It is worth mentioning that AI is not the only, nor even the first, scientifically based discipline that directs attention to the subjective self. One example is the new physics: "If the new physics has led us anywhere, it is back to ourselves, which, of course, is the only place we could go" (Zukav 1979, p. 114). Another example is the biological–cybernetic approach of Ruyer (1965).

something quite different from randomized decision making. It is a creative act guided by reason grounded in authentic individual experience. Yet this act is essentially unpredictable. A world picture based just on chance and necessity (Monod 1971) is no longer sufficient, for it leaves no room for intentionality.

3.4 Computational linguistics

Attempts to teach a machine to understand and talk in natural language have revealed that talking is something more than just exchanging coded messages. Language acts are deeds that affect things and events in the world in a way analogous to physical actions. Through language man enters into a complicated arena where the ideosphere pervades the biosphere.

One common opinion in AI circles is that a large knowledge base might solve the problem of natural language understanding by computer. This argument may perhaps hold for an *ad hoc* restricted language with rigid semantics, but actually talking in a human language is hardly feasible without actual human experience in the real world.

3.5 AI and new methodological approaches

Current research in AI is an unusual combination of constructive, empirical, and theoretical methods. Another new and powerful source of ideas in AI is introspection (which serves, for example, as the intuitive basis for Minsky's frame systems).

AI has given a concrete form to heuristic reasoning and reasoning by analogy; the creative role of metaphors was mentioned above. Related to metaphors is Pribram's (1981) concept of "abduction": this involves the creative transfer of a whole complex of concepts from one scientific field to another (for example, the holographic brain). Finally, we should not forget paradoxes, rather unconventional tools that, nevertheless, assist us in gaining deeper insight where other tools fail.

3.6 AI working through literature

Besides widely distributed science fiction and standard popularizing literature (which often gives a rather simplified and/or exaggerated picture of AI), I would like in particular to point to a new kind of literature best exemplified by the well-known book *Gödel, Escher, Bach: An Eternal Golden Braid* (Hofstadter 1979) or the anthology *The Mind's I* (Hofstadter and Dennet 1981). This is highly influential and stimulating material which, in addition to its entertainment value (or perhaps by

virtue of it), has profound philosophical and scientific significance. Quite possibly the impacts of AI induced by this type of literature will be more recognizable than impacts associated with concrete applications of AI.

3.7 The importance of reflection

Every human activity can be the subject of reflection. Both the activity itself (its motivations, purposes, means, etc.) and the active subject may be reflected upon. What is specific to AI is that both these types of reflections are, in a sense, already its constituents: man is, as we have seen, the prototype for AI, and among the goals of AI is the testing of its own limitations.

Any proper reflection must also consider the reflecting subject as well as the act of reflection itself – a self-referential loop that leaves little room for fallacies or ill effects: another argument for not being so afraid of the consequences of AI.

4 Developments Over the Next Ten Years

4.1 Core research areas and likely results

I do *not* expect that large knowledge systems (the present core research area in AI) will grow ever larger. The well-known "pig's principle" (if something is good, then more is better) certainly does not hold here. Perhaps more attention will be given to methods of converting data into knowledge. *Semantic search* (as opposed to combinatorial search) or *knowledge retrieval* may be new areas of interest. "Knowledge space" may be endowed with a topography, in which events like changing one's viewpoint may have a programmable analogy.

New conceptions of knowledge organization in the computer may be expected. For instance, using the visual field as a metaphor suggests that the system should have at every instant a "concentration point," which, nevertheless, cannot be isolated from its surrounding environment, i.e., the context which gives it proper meaning. A shift of topic is then represented by a transfer of the concentration point.

4.2 Impacts on other sciences and technologies

In view of what has been said above about the third kind of impact of AI, we can expect in the near future a substantial effect on psychology and related fields (in fact, the effect is already visible now). I do not mean dependence of psychology on computer modeling (if so, then

rather on the computer metaphor) but a more fundamental change: psychology will no longer be a science concerned with "objects" (habitually called "subjects") but a well-founded effort for the deeper self-understanding of man. Neuropsychology (a lower-level approach) and behavioral psychology (a higher-level "outer" approach) need a third companion (a higher-level "inner" approach), which uses introspection as its main source of knowledge.

Through its impacts of the third kind AI may also influence other social sciences and the humanities, perhaps even more than the natural and technical sciences.

A few words about mathematics. It seems that the current growth of interest among mathematicians and logicians in computer science will continue and that AI will provide them with a number of challenging topics for fundamental research. For instance, certain intuitively defined qualities relevant to AI can usefully be set up in opposition to analogous qualities that already have a mathematical explication: heuristic vs. algorithmic, vague vs. accurate, small (large) vs. finite (infinite), analogous vs. homomorphic, insightful vs. deductive, self-referential vs. recursive, paradoxical vs. contradictory, etc.

4.3 Social and economic impacts

In the near future impacts of the first and second kinds will more probably be associated with computer technology than specifically with AI. As for the impacts of the third kind, we should be able to feel them soon as slow changes in the intellectual climate. However, the way they may affect our daily life is difficult to assess.

Regarding the use of computers in management, I think that what we can learn from AI about the difference between programmable decision making and free authentic choice might spare us quite a lot of disillusion.

5 Long-Range AI Research

I have already expressed my skepticism concerning the possibility of fabricating an autonomous mind. A more promising long-range trend in AI seems to be using the computer as an on-line extension of the human brain (the occasionally used term "symbiotic AI" is inappropriate here because symbiosis involves *two* organisms in mutual interdependence). After all, in order to help somebody it is not necessary to replace him.

To speculate about long-range social impacts – and choosing the optimistic end of the scale – there is a chance that impacts of the

third kind, particularly the revival of interest in human self, may help to restore respect for human individuality, privacy, and freedom of thought.

Turning now to those AI research outcomes that I feel are most worthy of support, the impacts of the third kind seem to me to be the most valuable and least dangerous. Therefore I would favor those areas of AI that promise to yield new approaches to human self-understanding. It seems that theoretical investigations are, in this respect, more promising than *ad hoc* programming.

6 In Conclusion

A few final remarks on precautionary measures. Discussions concerning scientists' ethics often circle around the issue of whether certain research areas with potentially harmful applications should be abandoned and by whom (individuals, the scientific community, or governments). I doubt that the problem is properly posed. Far earlier than the stage where such issues arise in concrete form, the social climate has to have been already pathological, namely in producing the very *idea* of the harmful activity in question. Neither the inventor nor the manufacturer of a microphone is fully responsible for its possible misuse as a bugging device: the idea of intruding into one's privacy appeared much earlier (certainly well before Shakespeare wrote *Hamlet*, for example!).

Eliminating symptoms is not the most efficient therapy. In the case of possible harm caused by science and technology, it is the social climate itself that should be diagnosed and cured first. The growing concern within the AI community and the general public about the potential influence of AI (the *existence* of this concern, rather than specific measures it may suggest) indicates to me that the contemporary climate is resilient enough to look after itself.

References

Boden, M.A. (1977) *Artificial Intelligence and Natural Man.* New York: Basic Books.

Dennet, D.C. (1978) Intentional Systems. In D.C. Dennet, *Brainstorms: Philosophical Essays on Mind and Psychology*. Montgomery, VT: Bradford Books.

Firschein, O. *et al.* (1973) *Forecasting and Assessing the Impact of Artificial Intelligence on Society*. International Joint Conference on Artificial Intelligence (IJCAI), Palo Alto, CA.

Hofstadter, D.R. (1979) *Gödel, Escher, Bach: An Eternal Golden Braid*. New York: Basic Books.

Hofstadter, D.R. and D.C. Dennet (Eds.) (1981) *The Mind's I. Fantasies and Reflections on Self and Soul*. New York: Basic Books.

Monod, J. (1971) *Chance and Necessity*. New York: Random House.

Newell, A. (1982) Intellectual Issues in the History of Artificial Intelligence. In F. Machlup and U. Mansfield (Eds.), *The Study of Information: Interdisciplinary Messages*. New York: Wiley.

Popper, K. and J. Eccles (1978) *The Self and Its Brain*. New York: Springer-Verlag.

Pribram, K.H. (1981) The Brain, the Telephone, the Thermostat, the Computer, and the Hologram. *Cognition and Brain Theory*, 4(2):105–122.

Ruyer, R. (1965) *Parodoxes de la Conscience et les Limites de l'Automatisme*. Paris: Albin Michel.

Weizenbaum, J. (1976) *Computer Power and Human Reason*. San Francisco, CA: W.H. Freeman.

Zukav, G. (1979) *The Dancing Wu Li Masters: An Overview of the New Physics*. New York: Bantam Books.

CHAPTER 8

Social and Economic Impacts of Artificial Intelligence – A Japanese Perspective

Makoto Nagao *

1 Introduction

Artificial intelligence (AI) has already begun to influence various aspects of social activities, especially with the spread of intelligent information systems and industrial robots. Its true impact will come in the future, but there are already both great expectations for AI and serious fears surrounding its possible effects. We must endeavor to assess the potential impact of AI upon society from all conceivable points of view, whilst at the same time not forgetting that the utilization of modern science and technology has made a great contribution to living standards and human culture.

2 AI in Japanese Industry

Japanese industries have achieved very high standards in efficient and cost-effective manufacturing systems and in product quality. They also are relatively open to the concept of artificial intelligence, prepared to utilize AI technology in their manufacturing activities, and ready to embark on the creation of a new era for industrial society. But from a different standpoint it could also be said that they have no

*Department of Electrical Engineering, Kyoto University, Sakyo-ku, Kyoto, Japan.

alternatives for improving their industrial activities and innovating *other than* the use of artificial intelligence.

The incorporation of AI into key stages of production processes contributes to increased productivity, enhanced quality, improved working environments, energy savings, and the holding down of increases in labor costs. Furthermore, AI should make it possible to efficiently produce a variety of objects in small quantities, create ranges of new products, and enhance the function and reliability of existing products. Other advantages should derive from the higher added value of the products as well as lower production costs.

The utilization of AI, and especially robotics techniques, is expected to raise the standards of small and medium-sized enterprises in particular. Many of these enterprises presently suffer from a lack of skilled labor. The use of high-performance, low-cost, and easy-to-use robotic machines that incorporate AI can help solve the problem of labor shortages while increasing productivity. Over the longer term, this should lead to the modernization of these smaller enterprises.

3 AI in a Future Information Society

Japan is entering a postindustrial information society. Everybody wants to reap the benefits of easily accessible and high-quality information. The information systems that will be extended throughout Japan in the near future must be able to accommodate the demands of all of the people. Therefore the systems must be "people-friendly," and they will need AI approaches to achieve a quality that satisfies all the likely demands. The systems should be mature enough to accept natural language (speech, handwritten/typewritten characters, word processors, etc.), to provide the range of information that people want, and to both protect privacy and prevent unethical utilization of the information systems' potential.

4 Some Possible Configurations

Many of the ordinary, day-to-day tasks will be performed by some sort of intelligent machine. For example, *intellectual systems* will provide computer-aided instruction (CAI) in elementary and middle schools, technical education at university level, and technical training within industries; other applications include consultation systems for medical diagnosis, shopping, cooking, travel, and so on, as well as information retrieval systems of all kinds providing daily news, encyclopedic data, and library and archival information, etc.

On the other hand, *physical systems* will be available in the form of robots for housekeeping, hospital care, guides for the blind, fire fighting, deep-sea exploration, and so on; there will clearly be a wide range of industrial robots in use, and other physical systems will progressively take over more and more inspection and testing functions of every kind.

5 Unemployment Problems

The recent advance of microcomputers and microelectronics in Japan that has created many new industrial fields and new jobs, as well as improving efficiency, has caused relatively little unemployment (JIP-DEC 1980). There have been no particular labor problems caused by the introduction of factory and office automation systems in Japan, and it is believed that no serious problems will arise in the near future from the introduction of AI technology into industries and society in general.

Many more jobs will be created by the development of intelligent systems, and robotics systems of all kinds, than those destroyed by the introduction of such systems. AI applications will create new products, which will in turn also create new jobs.

6 The Need for Continuous Education

The introduction of many highly intelligent systems into every part of society may divide the quality of labor into two parts: those people who are intelligent and adaptable enough to either fabricate robots or use them for their own purposes, and those who are unable to adapt to the changed circumstances and are therefore driven into lower-level jobs where robots cannot be introduced. The advanced information society will also have the option of dividing itself into two groups: one that makes different kinds of profit by utilizing *all* the information available from the systems, and the other, which has no ability to enjoy or profit from such intelligent systems and therefore falls into some lower level.

This phenomenon, if it ever becomes a reality, would be very serious for the society of the future. We need to think about some means to raise the educational levels of the people so that they can adapt the intelligent systems for their individual purposes. A gradual and smooth transition of employment structures to forms appropriate for a future information society is essential, and this can only be achieved by the

continuous education of the general public (Versailles Declaration 1982).

7 Conclusion

The kind of intelligent information society discussed above will not be realized within a short time. Robots cannot become so intelligent as to displace many human jobs in the near future. Robots will be introduced first in simple, low-level jobs in certain controlled conditions. We still have time enough to think about and discuss the future potential of AI and how it can best be dedicated to human welfare and peace. Artificial intelligence still has a long way to go, but the main lines of development can already be discerned. It will progressively display its ability to function in five main areas: as a strong inference mechanism; as a powerful tool for the representation of knowledge; as a means of increasing flexibility in decision making by providing organic, human-like, decision-support systems; as a means of unifying language, images, speech, and other sensory information through knowledge support; and as a means of developing strong software support for increasingly complex tasks (Nagao 1983).

References

JIPDEC (1980) *Report of the Investigation into the Effects of Microelectronics on Employment*. Tokyo: Japan Information Processing Development Center (JIPDEC), January. (In Japanese.)

Nagao, M. (1983) *Exploration of Knowledge and the Algorithm in Human Brain*. Paper presented to the Fourth Convocation of Engineering Academies, Stockholm, May 30–June 1.

Versailles Declaration (1982) *Report of the Working Group on Technology, Growth, and Employment*. Declaration of the Seven Heads of State and Government and Representatives of the European Communities. Chateau of Versailles, June 4–6.

CHAPTER 9

Artificial Intelligence, Employment, and Income*

*Nils J. Nilsson***

1 Introduction

Artificial intelligence (AI) and other developments in computer science are giving birth to a dramatically different class of machines — machines that can perform tasks requiring reasoning, judgment, and perception that previously could be done only by humans. Will these machines reduce the need for human toil and thus cause unemployment? There are two opposing views in response to this question. Some claim that AI is not really very different from other technologies that have supported automation and increased productivity — technologies such as mechanical engineering, electronics, control engineering, and operations research. Like them, AI may also lead ultimately to an expanding economy with a concomitant expansion of employment opportunities. At worst, according to this view, there will be some, perhaps even substantial shifts in the *types* of jobs, but certainly no overall reduction in the total *number* of jobs. In my opinion, however, such an outcome is based on an overly conservative appraisal of the real potential of artificial intelligence.

Others accept a rather strong hypothesis with regard to AI — one that sets AI far apart from previous labor-saving technologies. Quite simply, this hypothesis affirms that anything people can do, AI can do

as well. Certainly AI has not yet achieved human-level performance in many important functions, but many AI scientists believe that artificial intelligence inevitably will equal and surpass human mental abilities – if not in twenty years, then surely in fifty. The main conclusion of this view of AI is that, even if AI does create more work, this work can *also* be performed by AI devices without necessarily implying more jobs for humans.

Of course, the mere fact that some work *can* be performed automatically does not make it inevitable that it *will* be. Automation depends on many factors – economic, political, and social. The major economic parameter would seem to be the relative cost of having either people or machines execute a given task (at a specified rate and level of quality). In this respect too, AI differs from many previous labor-saving technologies in that it is relatively very inexpensive and will undoubtedly become even more so in the future. Yet, even granting an economic rationale for replacing human labor with machines, we as a society may choose not to do so. That is, we may decide to continue to employ humans in jobs "next to the window" (as the Japanese say), simply as a way to distribute income and to give people something tangible to do.

In this paper I examine the potential economic effects of artificial intelligence. I conclude that AI does indeed offer the potential for achieving massive reductions in the amount of human labor needed to produce the world's goods and services. While acknowledging that there are understandable reasons people might feel threatened by this outcome, it seems to me that we should view it as a blessing rather than a curse. As John Maynard Keynes said over fifty years ago: "All this means in the long run [is] *that mankind is solving its economic problem... The economic problem is not – if we look into the future – the permanent problem of the human race*" (Keynes 1933). From this standpoint, I then review some suggestions for dissociating income from employment so that people will be able to benefit from the elimination of unnecessary toil.

2 The Diminishing Need for Human Labor

2.1 Commonsense arguments

Before beginning a more technical discussion of the economic effects of AI, it is worth considering a few general statements that are being made about the consequences of automation.

First, let's look at some of the arguments supporting the view that automation (including AI) will *not* result in unemployment. In a recent

interview, James Albus, a leading researcher in robotics, made several important points. He stated, for example: "There is no historical evidence that rapid productivity growth leads to loss of jobs. In fact, quite the contrary. In general, industries that use the most efficient production techniques grow and prosper, and hire more workers. Markets for their products expand and they diversify into new product lines" (Albus 1983).

A related argument is based on the observation that unemployment is worse in the developing countries than in the industrialized ones. Since automation is much less pervasive in the Third World and unemployment is still so acute there, automation obviously cannot be the principal cause of unemployment.

Even if automation makes it possible to perform every task with fewer workers, there are a great many needs to satisfy. Albus expands on this point by observing that "...there is not a fixed amount of work. More work can always be created....Work is easy to create....there is always more work to be done than people to do it...The problem is not in finding plenty of work for both humans and robots. The problem is in finding mechanisms by which the wealth created by robot technology can be distributed as income to the people who need it. If this were done, markets would explode, demand would increase, and there would be plenty of work for all able-bodied humans, plus as many robots as we could build" (Albus 1983).

There are several industries that have pursued automation aggressively without reducing overall employment. In US banking, for example, because the increased productivity resulting from automation has been accompanied by a relatively even higher demand for bank services, employment grew by 50 percent between 1970 and 1980 (Ernest 1982). (On the other hand, however, we note that most of the jobs in the banking industry involve "knowledge work" of one sort or another – i.e., the very category that is succumbing most rapidly to automation by AI techniques. In fact, the Bank of America recently announced that it is now seeking to reduce its employment levels significantly (Gartner 1984).)

Even if automation proceeds rapidly, the task of converting to automatic factories and offices will itself require considerable labor. According to Albus, "...building the automatic factories...is a Herculean task that will provide employment to millions of workers for several generations" (Albus 1983).

Critics of the hypothesis that artificial intelligence will be able to do *anything* argue that there is a large number of tasks that simply can *never* be completely automated. For instance, some people believe that it will prove impossible or undesirable to automate such services as marriage counseling, child care, and primary-school teaching. They

might also claim that machines will never be able to generate truly excellent music, literature, and other art forms. Some, such as Professor Thorne McCarty of Rutgers, have suggested that the economy of the future might be based on these specialized kinds of "human-oriented" and creative services, just as much of our present-day economy is based on a more general array of services.

On the other hand, those who argue that the more advanced forms of automation (like robotics and AI) will cause *increasing* unemployment have several reasonable arguments on their side. For example, they point to the fact that over the past decade or so unemployment in the technically advanced societies *does* appear to have grown. With each successive business cycle, the "troughs" in the unemployment graph move upward. Although many people lose their jobs at times of recession, there are many others among the jobless who can blame their plight on robots and other automatic devices. Some economists think that we are already in the initial stages of a critical period in which large-scale unemployment due to automation is inevitable. For example, Dr. Gail Garfield Schwartz, an economic consultant in Washington, DC, was recently quoted as saying "perhaps as much as 20% or more of the work force will be out of work in a generation" (Neikirk 1982).

Nobel prize-winning economist Wassily Leontief, director of the Institute for Economic Analysis at New York University, adds weight to this prediction. He says that "we are beginning a gradual process whereby over the next 30–40 years many people will be displaced, creating massive problems of unemployment and dislocation....In the last century, there was an analogous problem with horses. They became unnecessary with the advent of tractors, automobiles, and trucks....So what happened to horses will happen to people, unless the government can redistribute the fruits of the new technology" (Leontief 1983).

We should also realize that employment data, as collected and published by the Bureau of Labor Statistics of the US Department of Labor, include all the people who are ordinarily considered to be working. We must be honest enough with ourselves to admit that probably not *all* of these people are really presently *required* to produce the goods and services that we need. Some might not actually be needed – but are being paid anyway because of labor contracts that have set excessive standards as to the number of persons it takes to perform certain jobs. Most of us are being paid, quite legitimately perhaps, for vacations – which is one way of spreading the available work around. Some are being paid because various governmental bodies have been persuaded that certain goods and services are "needed," despite the fact that they are quite controversial and might not even be desired by a majority of those of us who are "buying" them. Some are being paid

for crops *not* produced. Some are being paid because of eliminable inefficiencies that we prefer to continue tolerating. Because statistical employment does not necessarily mean real employment, the magnitude of the *unemployment* problem may already be greater than we realize.

Another factor pointing toward future reductions (or at least shifts) in the labor force is the rapid progress in automating much "white-collar" work. It has been estimated that more than half of all US workers are engaged in "information-processing activities." Included in this category are many managerial functions, such as decision-making, reporting, communicating and coordinating, fact-gathering, and the supervision of similar activities by subordinates. Also included are many paper-handling clerical functions. The "expert systems" and automatic planning programs currently being developed in AI research laboratories will be able to perform many of these tasks, with a consequent drastic reduction in the need for human involvement or intervention.

Others have argued that the majority of new jobs created by auto-mation will require only low-skilled labor. In a recent Stanford study Levin and Rumberger (1983) conclude the following:

> "Most new jobs will not be in high-technology occupations, nor will the application of high technology in existing jobs require a vast upgrading of the skills of the American labor force. To the contrary, the expansion of the lowest-skilled jobs in the American economy will vastly outstrip the growth of high-technology ones. And the proliferation of high-technology industries and their products is far more likely to reduce the skill requirements for jobs in the US economy than to upgrade them....About 150,000 new jobs for computer programmers are expected to emerge during [the next] 12-year period, a level of growth vastly outpaced by the 800,000 new jobs expected for fast-food workers and kitchen helpers....Past applications of technology in the workplace as well as present evidence suggest that future technologies will further simplify and routinize work tasks and reduce opportunities for worker individuality and judgment. Moreover, the displacement in jobs and the downgrading of skill requirements for most of the new positions will undermine employment generally, and especially the employment of skilled workers."

So we see that there are many more or less reasonable arguments on both sides of this issue. It is likely that we will continue, almost daily, to hear conflicting opinions about the prospective impact of AI on employment. For a more prescient analysis, however, we must turn to more technical economic arguments.

2.2 Technical arguments

The economists whom we consult also have differing opinions on this matter. In fact, two Nobel laureates in economics, namely Herbert Simon and Wassily Leontief, seem to be on opposite sides. On the one hand, Herbert Simon (who, of course, is also a distinguished AI scientist) invokes the law of comparative advantage, which is, as he states, the standard rebuttal of economic theory to the concern that mechanization causes technological unemployment. Simon claims that this law "...shows, essentially, that both people and machines can be fully employed regardless of their relative productivity. By adjustment in the relative price of labor and capital, respectively, it will come about that labor will be employed in those processes in which it is relatively more productive, capital in the processes in which *it* is relatively more productive." Simon does admit, however, that the law of comparative advantage "...does not settle all the essential issues. In particular, although it shows that at some wage all labor would be employed in equilibrium, no matter how efficient machines become, it does not predict what that wage would be. It does not guarantee that real wages will not drop as the economy's productivity improves through mechanization. It does not even guarantee that real wages will remain above the subsistence level" (Simon 1977, p. 145).

To pursue this matter further, Simon analyzes in more detail the effects of technological progress in a model economy (Simon 1977, pp. 147–160). Equilibrium conditions in his model are given by the following equation:

$$w_L \bar{t}_L + (1 + r_I)\bar{c} = 1$$

where w_L is the labor wage rate, $\bar{t}_L$ the average labor time required to produce one unit of output, r_I the interest rate, and $\bar{c}$ the average capital required to produce one unit of output. A labor-saving technological change in Simon's model economy would lower the average labor time required to produce one unit of output. He concludes that "...so long as the rate of interest remains constant, an advance in technology can only produce a rising level of real wages." He then gives several reasons, some historical, why he believes it likely that (real) interest rates will remain constant.

Simon summarizes his position by saying that "...nothing about the current advances in automation indicates that these advances will have any different economic effects than earlier industrialization and mechanization. The main long-run effect of increasing productivity is to increase real wages — a conclusion that is historically true and analytically demonstrable."

In Simon's model, the entire cost of capital (interest on borrowed capital plus return on invested capital) is combined as "interest." One wonders whether the predictions of his model would survive an analysis in which technology makes possible a shift from salaries paid to humans toward "salaries" paid to machines (in the form of higher profits to the robot owners). Would he include these "salaries" as "wages" or as "interest"?

It is true, as Simon notes, that the cost of capital (the real interest rate) has remained fairly level over recent decades. Consequently, as our economy has become more productive, both the labor wage rate and total labor wages have risen. But, as Kelso and Adler (1958, pp. 40—43) argue, it is machines, not labor, that have become more productive; therefore, it is capital, rather than labor, that might more appropriately have received the extra reward. Our society has simply "chosen" (through various social and political means) to distribute income to consumers through higher wage rates, rather than through higher return on investment. However, there is no compelling reason to believe that, as capital becomes even more productive (through the use of AI devices), equilibrium won't ultimately be attained by rewarding capital more handsomely than labor. (For example, could not the recent upward pressure on real interest rates be partially explained by the fact that it is becoming ever more difficult for society to deal with the increased productivity of capital by repeatedly granting further rises in wage rates?)

Taking a somewhat different approach, Wassily Leontief and his colleagues have used arguments based on input/output analysis to show that automation will lead to a *reduced* need for human labor. (Input/output analysis is a computer-based technique for quantitatively analyzing how the various sectors of an economy supply and depend on one another for goods and services (Leontief 1966).) A series of input/output simulations of the Austrian economy, for example, forecast that intensive automation would increase Austria's customarily low (2%) unemployment rate to 10% by 1990 unless the reduced need for work were offset by a shorter work week (Leontief 1982, pp. 188—204).

Together with Faye Duchin and other colleagues, Leontief has recently developed a dynamic input/output model of the American economy that has been used to make predictions about employment levels on the basis of three different assumptions about technological change. They conclude that a rather "intensive use of automation over the next twenty years will make it possible to conserve about 10% of the labor that would have been required to produce the same bill of goods in the absence of automation. The impacts are specific to different types of work and will involve a significant increase in professionals as a proportion of the labor force and a steep decline in the relative number of

clerical workers" (Leontief and Duchin 1983). After discussing these results with Faye Duchin, I think it is likely that AI technology will allow even more automation than that assumed in their "intensive" scenario (Duchin 1983, private communication).

One could cite other economists who take opposing positions on the effects of automation. In my opinion, however, those who expect AI and automation to expand the number of jobs available fail to take into account some important distinctions between AI and previous technologies. I believe that a relatively straightforward chain of economic reasoning, based on the special nature of artificial intelligence, leads us directly to the conclusion that the total amount of human labor used to produce our goods and services will decline markedly.

Our terminology will be simplified if we use the word "consumables" in place of the phrase "goods and services." By "consumable," I mean to include specific physical objects, such as automobiles, cameras, shoes, apple juice, missiles, and all the other tangible items that might be needed (purchased) from time to time by or for persons, animals, processes, corporations, governments, or other consumers. I also include all kinds of services, such as rides on jet planes, hairdressing, the production and issuance of environmental impact reports, deciding about factory locations, financial management, poetry composition and readings, and such other services as might be performed in a society. Although there may be intermediate nonhuman consumers, such as corporations, processes, or robots, the most important class of consumers consists of people – i.e., of living, breathing individuals. If consumption by a nonhuman does not, in some manner, relate ultimately to the consumption of *something* by a human, we might well ask (anthropocentrically): "What's the point?"

People are both "consumers" and "producers" of consumables. Economic systems and economic theory seek to establish relationships between these two activities. Let us focus first on consumption.

It is particularly difficult to analyze our "needs" for consumables. Consumables can often be replaced by substitutes without objection or detriment. Cheap plastic parts can be used instead of more expensive metal parts, resulting in a different consumable. If the substitution is successful, our "need" for the original consumable disappears. Sometimes people can be convinced that they really do not need consumables formerly considered essential. (What has happened to "house calls" by medical doctors?) Previously contented, satisfied people can also come to believe that they can no longer be satisfied without acquiring additional or different consumables. Suddenly we "must have" a new kind of toothpaste, a backyard swimming pool, a particular weight-reducing program, and so on. Times, styles, and values change – and, as they do, our needs expand and contract correspondingly.

Another important point is that, regardless of changing needs in the industrialized countries, most of the people in the world presently consume relatively little; consequently, the global demand for consumables can (and ought to) increase dramatically. James Albus has noted the effect that this accelerating demand will have on employment: "The world is filled with need. It is premature to worry about robots eliminating work as long as there exist such overwhelming problems as providing food, clothing, shelter, education, and medical care for millions of people living in desperate poverty" (Albus 1983).

Nevertheless, one should not assume that the potential for human beings to absorb consumables is *infinite*. It may in fact be very large, but it *is* finite. There is a "law of diminishing returns" for consumption as well. After a certain level of consumption has been reached, people just don't have the extra time, attention, or desire to absorb even more goods and services, regardless of their cost. Those consumables that are purchased, but which are then literally forgotten and abandoned before they are used at all, might just as well not have been produced – except, possibly, for the role their creation might have played in providing jobs. But such jobs, we would all agree, would certainly be of the "make-work" category. It should also be clear that the finite nature of material and energy resources might effectively limit the production of some consumables even before people became sated with them.

The reason for pressing this point is that some might argue that there will always be a requirement for large amounts of human labor if demand is potentially infinite. Automation will simply never catch up with continually increasing demand. Now, we might agree that the ultimate ceiling on demand could be very high indeed – far above current levels. But AI applications will lead to *increases in the exponent* of a productivity that is already rising exponentially, and such "superexponential" increases overtake finite (even if large) limits very quickly. (A similar argument supporting continuing requirements for human labor might be based on a constantly shifting but finite set of needs. Here again automation might have difficulty keeping pace with constantly shifting needs – but so might human workers.)

We should also question whether we, as a society, really *want* to persuade ourselves to increase consumption *mainly* for the purpose of providing employment. If drastically increased consumption does not result in more fulfilling and rewarding lives, the extra production seems rather pointless. I am reminded of the naturalist John Muir's statement, "Why, I am richer than [railroad magnate Edward] Harriman. I have all the money I want, and he doesn't."

Now let us consider production. Consumables are produced by processes that change from time to time. Each process involves an amalgamation of human labor (of various types), capital (machinery and

plant), and materials (including energy). In a free-market, competitive economy, suppliers of consumables seek to reduce the individual costs of these components. We might call this tendency the "law of reduc- tion." Reduction is often achieved through "technology." (Of course, suppliers might also seek to substitute a different, lower-cost consum- able by attempting to convince purchasers that they need the substi- tute instead.) Most attention is paid to that component of a consumable's production which offers the best expectation of achieving maximum overall cost reduction.

Given the law of reduction (and a strong belief in the power and economy of artificial intelligence), we can conclude that the cost of the human-labor component of *any* consumable will fall until it is no longer a significant percentage of the cost of the consumable. The law of reduction, in this case, works as follows. So long as the cost of the human-labor component of a consumable remains a significant factor in the latter's total cost, producers will want to substitute lower-cost machines for human labor. This desire will motivate attempts to increase the power and decrease the cost of computer hardware and AI software. It is technologically inevitable that such attempts will succeed and will continue to do so until the cost of the human-labor component of consumables becomes insignificant. Essentially what we shall be witnessing will be the replacement of expensive human labor by ever-cheaper AI "labor" until a consumable's ultimate cost approaches the sum of the costs of the non-AI capital, materials, and taxes. (Since AI is so inexpensive, we are assuming that the cost of the AI capital is also insignificant.) Thus, the total cost of human labor needed to produce all consumables will become small in comparison with the sum of the other costs.

An economic tautology that links input and output tells us that the total income earned by people through labor is the same as the total human-labor cost in producing all consumables, and that the total income earned by people through means other than labor (such as transfer payments and return on investment) is the same as the total nonlabor costs of producing all consumables. Therefore, the total income people earn through their labor can be expected to fall until it is a small percentage of the total income earned by people through other means.

These trends are already evident in our economy for many indivi- dual industries, such as agriculture — even if they have not yet been prominent in the aggregate. The primary income source of a typical family farmer, for example, has shifted from his own labor to a return on his investment in laborsaving farm machinery. In general, however, labor has been able to move to areas in which its technological replace- ment by lower-cost machines has not yet occurred, such as service

industries and white-collar work. This shift has been possible because demand has expanded to include consumables produced by such labor. But the law of reduction is still operative, chipping away at the labor component of these consumables as well; and, through AI, their labor costs too will ultimately be drastically reduced. With inexpensive AI in relentless pursuit, it seems likely that there will be few sanctuaries left where income is derived from toil. The only way in which the labor fraction of overall income will not be lowered is if many new "automation-proof" consumables are added to those we need. For example, some people think that several human-service occupations (nursing, teaching, psychiatry, marriage counseling, etc.) will remain forever beyond the reach of automation, and that our future economy might be based entirely on that sort of work. It seems to me, though, that there are too few intrinsically nonautomatable services on which to base an entire economy.

If the small percentage represented by overall labor income is concentrated among a small percentage of the total population, then the rest of the population will be unemployed. Alternatively, if that small amount of overall labor income is evenly spread among the population, everyone will obtain only a small percentage of his income through work, there will be a disincentive to work long hours, and everyone will be largely unemployed. In either case we have unemployment. By "unemployed" I do not mean unoccupied. Nor do I mean to imply that people will regard their unemployment as in any way undesirable. I merely mean that people's time will not be spent predominantly in working for an income. Income will come from other sources. I shall discuss this issue further in the next section.

3 What's so Bad about Unemployment?

3.1 Two fears

Instead of welcoming the arrival of mechanical slaves to perform much of the world's toil, most people view the prospect of increasing unemployment with great alarm. Leontief puts this paradox in sharp relief: "Adam and Eve enjoyed, before they were expelled from Paradise, a high standard of living without working. After their expulsion, they and their successors were condemned to eke out a miserable existence, working from dawn to dusk. The history of technological progress over the past 200 years is essentially the story of the human species working its way slowly and steadily back into Paradise. What would happen, however, if we suddenly found ourselves in it? With all goods and services provided without work, no one would be gainfully

employed. Being unemployed means receiving no wages. As a result, until appropriate new income policies were formulated to fit the changed technological conditions, everyone would starve in Paradise" (Leontief 1982, p. 192).

Leontief's story highlights one of the fears that people have about unemployment, namely, that they will lose their incomes. Presumably this economic fear would evaporate if people could obtain an income in some other manner so that they could purchase goods and services produced by the machines. Many economists, as well as others, have proposed various schemes that separate income from employment; I shall examine some of them in this section.

Another cause for apprehension has to do with social and psychological needs of human beings rather than with their economic requirements. What will people do with their "free time"? What activities will be as fulfilling and rewarding as jobs? Some people are pessimistic about the ability of their fellows (but not of themselves) to adjust to "becoming rich." Others, like John McCarthy, an AI pioneer, opine facetiously that this adjustment "...could take all of ten minutes" (McCarthy 1983). I shall also have some comments about this problem.

3.2 Allaying the economic fear

There are several ways of dealing with the economically motivated fear of unemployment. They range from rather crude approaches, like attempting to slow down or halt technological change so as to delay or prevent unemployment, to more sophisticated and possibly impractical reorientations of our economic system.

Placing obstacles in the path of either using or abetting technology might be called a "Luddite approach" to the economic problem of unemployment. This approach is unfair to humanity because it condemns us to continue toiling when toil is technologically unnecessary. To use Leontief's metaphor, it is equivalent to disrupting our attempts to re-enter Paradise. In any case, the approach would inevitably fail because, fortunately, no government or other group has sufficient repressive power to prevent technical progress. Even if technology were temporarily slowed in one country, so much the worse for that country; its foreign competitors would soon outrace it and it would have unemployment anyway – unemployment *and* poverty.

Another way to solve the economic problem posed by technological unemployment is to invent jobs that are really either unnecessary (that is, they do not contribute to absorbable consumables) or could be performed by machines. This approach may be one way of distributing income, but it is unfair because it condemns some people to

unnecessary toil. There is reason to fear that the conventional goal of *full employment* (espoused by both political parties in the United States) can be achieved only through such "make-work" schemes.

Separating income from employment would seem a better way to solve the economic problem of unemployment. This solution actually suggests itself as a corollary of our earlier economic analysis; if income is not derived principally from labor, it must instead come from either capital investments, sale of materials, or transfer payments.

As regards transfer payments, the industrialized nations already have a great deal of experience with government techniques for distributing income independently of work. Social security, "welfare" payments of various kinds, farm subsidies, and the "negative income tax" have all been used in the United States. Expansion of these programs is one possible way of decoupling income from work. I shall not attempt here to give arguments for or against transfer payments, except to note that many people fear the pernicious effect of some types of transfer payments – i.e., that they might subvert the American ideal of a free and independent citizenry.

There have been many intriguing proposals for more of us to obtain more of our income from a return on capital investment. Louis Kelso and Mortimer Adler have written a book that proposes an imaginative, capitalistic "...society in which machines do all or most of the mechanical work that must be done to provide the wealth necessary both for subsistence and for civilization." They recommend a diffuse, private ownership of the means of production so that

> "...every man, or every family, has a sufficient share in the private ownership of machines to derive sufficient subsistence from their productivity. In this automated industrial society, each man, as an owner of machines, would be in the same position as an owner of slaves in a slave society. As a capitalist, he would be an economically free man, free from exploitation by other men, free from destitution or want, free from the drudgery of mechanical work – and so free to live well if he has the virtue to do so" (Kelso and Adler 1958, p. 28).

Kelso and Adler envision that people would receive most of their income from dividends on common stock. To achieve this situation, they make proposals that would

(1) Broaden the ownership of existing enterprises.
(2) Encourage the formation of new capital and the organization of new enterprises owned by new capitalists.

(3)   Discourage concentration of the ownership of capital by house-
      holds in which such concentration has passed beyond the point
      determined to be the maximum consistent with the just organiza-
      tion of a completely capitalistic economy (Kelso and Adler 1958, p.
      169).

      They also list specific recommendations, including the use of tax
and credit devices, whereby families may begin to accumulate stock
ownership in corporations.
      James Albus has suggested the formation of a National Mutual Fund
(NMF), which would use credit from the Federal Reserve System to
finance private investment in automated industries. Ultimately this
fund would invest about $300 billion a year, which would double the
then current (1980) rate of investment in plant and equipment. This
extra investment in private companies would earn profits that would be
distributed by the NMF to the general public as dividends to stockhold-
ers. "Everyone would receive a substantial income from invested capi-
tal. Everyone would be a capitalist, not just the wealthy" (Albus 1981).
      To offset the short-term inflationary effect caused by the invest-
ment of this newly created money, Albus suggests that short-term
demand be restrained through a mandatory savings bond program.
These bonds would bear interest and be redeemable after five years.
"The key idea in this plan, which might be called an Industrial Develop-
ment Bond program, is to index the mandatory savings rate to the lead-
ing indicators for inflation on a monthly basis. If inflation is predicted,
mandatory savings go up for the next month and reduce consumer
demand. As soon as prices stabilize or decline, mandatory savings are
reduced. This policy would effectively divert short-term demand from
consumption into savings and compensate for increased investment. At
the same time, it would assure that the purchasing power to distribute
the fruits of investment in highly productive technology would be avail-
able once the new plants and modernized machinery began to produce
increased output."
      As Albus notes, separating income from employment explicitly ac-
knowledges "...that the primary goal of an economic system is not to
create work, but to create and distribute wealth, i.e., goods and ser-
vices that people want and need." He goes on to say, "I believe we have
it within our power to create an everyperson's aristocracy based on
robot labor" (Albus 1983).
      The process of converting to an economic system that separates
income from employment will face major, perhaps insurmountable politi-
cal, psychological, and social obstacles. "Earning a living" is a very
deeply ingrained notion in our culture. Different levels of skill, luck,
and hard work in earning a living allow a spread of incomes, from low to

high, that many people regard as equitable and desirable. Even if most consumables were being produced automatically, material and energy limits might not allow everyone to consume at the rate he would like. Since some of us (perhaps many) will still need to work, the lure of higher incomes might provide the necessary incentive – even as it so often does now.

3.3 Allaying the sociopsychological fear

There are many people who are fortunate enough to gain many psychologically valuable benefits from their jobs in addition to those of a strictly economic nature. Job satisfaction, the joy of achievement, an enhanced personal identity, opportunities for growth and learning, and social interaction are among the things that many of us derive from our work. Clearly, humans need such beneficial activities, but must they be tied to the production of income? There is already a large number of people who gain fulfillment and psychological rewards from activities they pursue in retirement (at which time their income is derived from pensions, social security, investments, etc.), or from volunteer or public-service activities (with income perhaps provided by a spouse's job or from inherited wealth). Many people also forego a chance at higher incomes so that they can fulfill themselves in artistic and creative pursuits in which the potential for income might be very low or even nonexistent.

Although many of us fear the prospect of losing a job, do we *really* fear more the loss of psychological rewards than economic ones? One simple test is to ask, "Suppose you inherited one million dollars. Would you go back to your old job, or would you do something else with your time?" Probably not many people are fortunate enough to have a job they would want to continue if they were suddenly to become wealthy.

Margaret Boden argues quite convincingly that the new age of automation could be "rehumanizing" rather than "dehumanizing." She foresees a "Polynesian-type" culture based on artificial intelligence. In Polynesia (at least in precolonial times) no one worried very much about the fact that freely available mangos caused unemployment. Professor Boden states: "AI could be the Westerner's mango tree. Its contribution to our food, shelter, and manufactured goods, and to the running of our administrative bureaucracies can free us not only *from* drudgery but *for* humanity. It will lead to an increased number of "service" jobs – in the caring professions, education, craft, sport, and entertainment. Such jobs are human rather than inhuman, giving satisfaction not only to those for whom the service is provided, but also to

those who provide it. And because even these jobs will very likely not be full-time, people both in and out of work will have time to devote to each other which today they do not enjoy. Friendship could become a living art again" (Boden 1983). It should also be noted that, besides providing people with time for human-oriented activities, automatic devices can be utilized in support of these activities to make them richer and more enjoyable.

Thus, it seems that there is no real reason to believe that a paying job is essential for a rewarding life. There is abundant evidence that people can receive important life-fulfilling benefits from a wide variety of activities that do not generate income. Some, like Willis Harman, envision a new conception of work made possible by our growing ability to produce goods and services automatically (Harman 1981).

Before leaving this topic, however, we might mention another possible function of employment. In addition to the positive benefits that accrue to a job holder, some observers, citing the correlation between crime and unemployment rates, see compulsory employment as a way to keep people out of trouble. Such a view not only seems inordinately pessimistic with respect to human nature and the human potential, but is probably wrong about the underlying causes of criminal activity. In any case, there are probably more humane ways to maintain civil tranquility than chaining people to work they dislike. Also, as Herb Simon has pointed out "...most people who are alarmed at [the prospect of too much leisure time] do not find that they themselves are endowed with too much leisure. But there are 'many people,' it is argued, who would not know what to do with leisure time, and who, presumably, would lend their hands to the Devil" (Simon 1977, p. 142).

4 The Transition

For those who are willing to grant that artificial intelligence and related technologies will eventually reduce the total need for human labor and that there are stable and desirable socioeconomic systems that separate employment from income, there still remains one very difficult question: how do we get there from here? Now some might say that we have plenty of time to worry about that problem and that now is too early to think about a transition. In the first place, it might be a long time before we develop the ultimate systems that will be able to perform the new jobs created by currently emerging AI systems. Secondly, a huge amount of human labor will be required to convert present-day industrial societies to fully automated ones (not to mention the labor needed to lift the living standards of the Third World).

Nevertheless, I think there are good reasons people should now start concerning themselves with this problem. First, the pace of technical change is accelerating. While it is true that the technical problems involved in creating artificially intelligent systems are still immense, we may solve most of them within the next generation. Second, if we begin to welcome rather than fear the "unemployment" consequences of AI, we can avoid the technological lethargy that unwarranted anxiety might otherwise induce. Third, socioeconomic changes are extremely slow (compared with technical progress). We must allow time for the several stages needed for the transition to new systems of distributing income. There will be at least five to ten years of discussion and argument among intellectuals and other social thinkers. Next, the voting public must have sufficient confidence in some of these ideas to approve any necessary legislation. At the same time, we must anticipate an inevitable reaction against these changes, stimulated by a general yearning to return to the "good old days" in which everyone did an honest day's work for an honest day's pay. People may blame these economic experiments for one or more of the expected future slumps in the business cycle. Taking all of these processes into account, it may well require one or two generations before the necessary changes can be made in our economic system, even if concerned people begin thinking earnestly about the problem right now.

The fourth reason for starting to think about the problem and instituting some transitional measures now is that this will minimize the discomfort of workers who are already being affected by automation. There are grounds for believing that the current high unemployment rates of the industrialized countries are not completely explainable by business cycles and will be cured by neither "supply-side" nor "demand-side" economic policies. This unemployment is rather a symptom of the "new automation," and it will continue to worsen even as business conditions improve. If the root causes of high unemployment are in actuality related to automation, policies that recognize this fact will have a better chance of alleviating the misery and poverty of the unemployed.

There are several things that we can begin to do now to prepare for these effects of AI. First, we must convince our leaders that they should give up the notion of "full employment" as a goal for the postindustrial economic system. It is unachievable, unnecessary, and undesirable — and it keeps us from focussing on the real problem.

Retraining is critically important, but we must not assume that everyone who now holds a blue-collar or middle-management job can become a computer scientist or programmer when his present job disappears. We must begin training for such automation-resistant

"human service" jobs as teaching, family counseling, daycare, and health care. We must also educate people in arts, crafts, literature, writing, and sports so that they will benefit more from their increasing leisure time. People cannot become "Polynesians" without training. Many community colleges already give adult education courses with this orientation; these programs should be expanded.

We should also begin to work much more earnestly on the many transition projects required to move us into the computerized, postindustrial age. Probably our most important task is to improve the living standards of people in Third World countries. I agree with James Albus that "Without rapid economic growth, a world of growing shortages will become an increasingly dangerous place. Nations competing over a shrinking stock of wealth and resources will inevitably come to military confrontation. The world's best hope is a great surge of industrial productivity that can outstrip the present population explosion and give us one more period of affluence in which we will have another chance at bringing the human population into stable equilibrium with the finite living space aboard the planet Earth" (Albus 1981). Another transition task is to design and build new automated equipment and factories. This work should be preceded by national projects, like those sponsored by the Japanese, that plan and develop the necessary technology. Additional projects could be initiated to improve education and health care in all parts of the world. Communities throughout the United States have been concerned about the problem of aging highways, bridges, and other transportation and communication facilities. Upgrading this "infrastructure" would absorb surplus labor during the transition stage. The postindustrial, information age will need another infrastructure – one consisting of computer systems, data bases, and networks. Putting all of this in place and maintaining it will require human labor for several more years.

Much of the work I have just mentioned can be funded only by governments. Because such work accomplishes goals that need to be satisfied, it should not be thought of as "made-work." But it does have the desirable side-effect of giving people employment during the transition from an economy in which most income is derived from employment to one in which most income is derived from other sources. In order to undertake these large public projects, we need to think differently about the matter of spending public funds. Instead of asking the rather outmoded question, "Can we *afford* such expenditures?" we need to learn to ask instead, "Should otherwise idle human labor be employed to accomplish socially desirable tasks?" The use of terms like "spending public money" and "affordability" focuses on arbitrary *accounting conventions* rather than on *reality*. What *really* counts is not an *abstraction* like money but whether or not people who could

be working on these projects would otherwise be idle and whether or not the rest of society can produce enough goods and services to satisfy the demand of people working on public projects.

In seeking to analyze the financial aspect of these public works, one could begin by observing that laborers are idle because the type of work they would have been doing in producing consumables in the private sector is now being done by machines. During the transition – that is, before these idle laborers receive fully compensating income from sources other than employment – part of the automation-derived benefits realized by business and consumers should be used to help pay the salaries of the workers on public projects. Simply put, the public projects can be financed by taxes levied on automation and consumption. Salaries paid to workers on public projects will increase demand (beyond what it would have been if those workers had remained unemployed and unpaid), but this demand can be met by the increased productivity of the automated industries. Of course, the taxes levied on automation must not be so high as to destroy the incentive to automate. Furthermore, they should decline as the cost of labor for the public projects goes down because of automation.

As automation takes over more and more of the work heretofore performable only by humans, we need to take steps to ensure that people become unemployed in a gradual and nondisruptive fashion. New approaches to work, such as job-sharing, should be encouraged. Shrinking of the workweek and a compensating increase in income derived from nonemployment sources, such as stock ownership and transfer payments, should proceed in step.

5 Conclusions

To recapitulate, I have argued that artificial intelligence is quite different from previous automation technologies in that it will lead to machines capable of performing quite inexpensively most of the tasks that now require or are best done by human labor. Business people in a free economy will then use these machines in preference to humans because this substitution will lower the cost of production – and simultaneously raise the quality – of goods and services. The overall consequence of using machines instead of people to produce goods and services will, of course, be unemployment.

Even though we have historically thought of unemployment as a serious problem to be corrected, the new unemployment might better be thought of as a liberating development. It will permit people to spend time on activities that will be more gratifying and humane than most "jobs." Nor need this unemployment be accompanied by reduced

levels of production and consumption. Various changes in our economic system can be suggested that will allow an ever-expanding flow of goods and services to be distributed equitably – not only in the industrialized nations, but in the "underdeveloped" ones as well.

I acknowledge that these developments will take time. The technology of artificial intelligence is still young; perhaps a generation or two must pass before its full impact on employment will be felt. Progress can of course be accelerated somewhat, if we so desire. Similarly, it will take at least a generation or two to make appropriate changes in our economic system – provided we start thinking seriously about these problems now. In the meantime, since AI and related technologies are undoubtedly already beginning to have some effect on employment, our concern is hardly premature.

It would be both foolish and tragic for us to slow our progress toward automation because of concern about unemployment. The world needs all the productive capacity it can create if its inhabitants are ever to live as human beings should. In fact, we may now be entering a very narrow and perhaps last-chance time window in which to make the transition from poverty, despair, and constant revolution and warfare to a more stable, just, and prosperous world society. Besides the will to accomplish this transition, we shall need all the help automation can give us.

6 Acknowledgments

References

Albus, J. (1981) *Brains, Behavior and Robotics*. Peterborough, NH: BYTE
    Books (a subsidiary of McGraw-Hill), p. 12.
Albus, J. (1983) The Robot Revolution: An Interview with James Albus. *Com-
    munications of the ACM*, March:179–180.
Boden, M.A. (1983) Artificial Intelligence as a Humanizing Force. In:
    *Proceedings of the International Joint Conference on Artificial
    Intelligence*. Los Altos, CA: William Kaufmann, Inc.
Ernest, M.L. (1982) The Mechanization of Commerce. *Scientific American*,
    September:133–145.
Gartner, T. (1984) B of A is Offering to Pay Workers Who Resign Early. *San
    Francisco Chronicle*, 17 January.
Harman, W.W. (1981) Work. In: A. Villoldo and K. Dychtwald (Eds.), *Mil-
    lenium*. Los Angeles, CA: Jeremy Tarchure.
Kelso, L.O. and M.J. Adler (1958) *The Capitalist Manifesto*. New York: Ran-
    dom House.
Keynes, J.M. (1933) Economic Possibilities for our Grandchildren (1930). In:
    Essays in Persuasion. London: Macmillan.
Leontief, W.W. (1966) *Input–Output Economics*. New York: Oxford Univer-
    sity Press.
Leontief, W.W. (1982) The Distribution of Work and Income. *Scientific Amer-
    ican*, September:188–204.
Leontief, W.W. (1983) The New New Age That's Coming is Already Here. *Bot-
    tom Line/Personal*, 4(8) (April):1–2.
Leontief, W.W. and F. Duchin (1983) *The Impacts of Automation on Employ-
    ment, 1963–2000*. Technical Report (draft). New York: Institute for
    Economic Analysis, New York University.
Levin, H.M. and R.W. Rumberger (1983) *The Educational Implications of
    High Technology*. Technical Report. Stanford, CA: Institute for
    Research on Educational Finance and Government, School of Education,
    Stanford University.
McCarthy, J. (1983) Interview with John McCarthy. *Omni*, April:100–122.
Neikirk, B. (1982) Recovery Could Be a Jobless One. *Chicago Tribune*, 15
    August.
Simon, H.A. (1977) *The New Science of Management Decision*. New York:
    Prentice-Hall.

CHAPTER 10

Social and Economic Impacts of Artificial Intelligence

Roger C. Schank * and *Stephen Slade* *

1 AI in the 1990s

There are many branches of artificial intelligence (AI) and each has its own research agenda. However, there are several general problems that affect efforts in the main areas of AI: natural language, vision, robotics, speech recognition, and expert systems. The first is learning. It is widely accepted that AI programs require very large amounts of knowledge in order to be effective. Once we know how that knowledge is to be represented, in general, we must then fill the computer with specific facts. These facts about the world are innumerable and ever changing. A smart program must be a learning program. It must be able to respond as a person does to new information and situations. It must adapt to changes in its environment. It must be able to grow on its own.

Another aspect of the same problem of the large amount of knowledge in these programs is both theoretical and technological. The theoretical problem is *How can that knowledge be organized?* The main current computer systems that contain and handle tremendous amounts of information are very large data bases. These systems tend to be fairly homogeneous and inflexible. By contrast, AI knowledge bases tend to be heterogeneous and variable. What is an appropriate

*Computer Science Department, Yale University, PO Box 2158, Yale Station, New Haven, CT 06520, USA.

model, then, for making these AI programs orders of magnitude larger? This leads to the technological question of *How can we actually build such large, complex programs*? Our machines are getting more powerful and this allows us to build larger programs. But it is still not certain that the machines of tomorrow will be able to keep up with a geometric growth in AI programming requirements. There will always be a need for more powerful machines for AI research. The problems of knowledge representation, acquisition, and organization will continue to be at the core of AI for decades to come.

We have taken the view in our own research that large, complex computer programs should be psychologically motivated. That is, the computer programs should be models of human cognitive behavior. There are several reasons for this point of view. First, people provide us with an existence proof for intelligence. We know that people can understand language, solve complex problems, and learn through experience. Therefore, we know that such behavior is *possible*. There may be other ways of developing intelligent computer programs, but we wish to try to model human cognition in as direct a way as possible.

Second, since our programs are psychological models, we can take advantage of psychological experiments and methodology in formulating our theories. We can appeal to psychological evidence in developing our programs, and, in turn, our programs and theories can prompt psychologists to try new experiments.

Third, and most important, we want our programs to be psychologically correct so that we can learn more about the human mind. As computer scientists, we are well aware of the intrinsic importance of intelligent machines, but we must aver that an understanding of human intelligence is a far greater goal than the creation of smart computers. Our work in modeling human cognitive processes has focussed attention on many facets of cognition including language ability, learning, memory, motivation, emotions, interpersonal relations, beliefs, and planning (Schank and Abelson 1977, Schank 1982).

In the next decade we expect that the two central problems of learning and memory organization will be explored in greater depth. It is very difficult to predict exactly what the results might be, since the answers are obviously unknown. However, these two areas should have the greatest impact on the field as a whole, since they touch on practically all areas of AI.

2 AI and Future Technology

The initial effect of AI on other scientific enterprises has been the use of AI's tools and methods. One can already witness the application of LISP software-engineering techniques and hardware to other

areas, such as CAD/CAM, text editors, and compilers. At Yale, we have ongoing research efforts in all three, which are quite separate from our AI research program.

Just as AI can no longer claim sole ownership of LISP, neither can it be proprietary regarding its basic paradigm of process models of human cognition. Numerous other fields, including psychology, philosophy, linguistics, and anthropology, are investigating AI theories and models of cognitive processes. This new interdisciplinary focus on cognitive behavior is called cognitive science. We see this as a very exciting research area, ancillary and complementary to AI.

Expert systems offer another way of relating to other fields. The AI methodology of modeling expert behavior can be applied to other scientific endeavors. Until now, the achievements in other fields by AI programs have been rather modest, when compared to human performance. We would expect though that future programs — especially those that can learn from experience — will be able to solve problems that might have baffled the human expert.

At Yale, we are developing several such expert systems that are meant to adapt to new situations based on prior experience. A fundamental assumption of this research is that the basis of expertise is experience itself. Our computer programs build a large memory of previous cases in a specific domain on which they can draw when a new problem arises. The previous cases are the source of general rules that can handle the normal situations. This is comparable to the common rule-based model of expert systems. However, when a novel situation occurs, the program should be able to analyze the new event in terms of previous events that shared similar features. Thus, even though no particular rule was invoked, the program would still be able to reason about the problem at another level.

In the next 50 years, AI may well come to be viewed as a type of meta-science, like mathematics. That is, AI will be used as a tool directed toward problems in numerous scientific and technical domains. A program that can reason and reformulate its knowledge according to experience will be a powerful tool in any application.

There is a belief among many people in AI that computers will be able to become far more intelligent than people. This is a debatable proposition, but there is one aspect of the argument that may be particularly relevant here. Computers can be made to devote all their resources to a given task. They are not subject to the usual array of human distractions such as hunger, thirst, boredom, or exhaustion. Furthermore, if you have developed a program to explore a certain problem, it is very straightforward to run that same program on 100 or 1000 computers at once — until one of them comes up with the answer. It will most likely happen that a computer program will make significant

discoveries in medicine or economics or mathematics. In the next century, an AI program might very well be awarded a Nobel prize. A *person* will have written the program, but the *program* will have made the award-winning discovery.

3 AI and Society

In the next ten years, there will probably be no significant social effects from AI. Certainly, the proliferation of home computers will continue and the public will become more aware of computers. However, the home computers will not be able to support the large, complex AI programs. People will begin to realize the great difficulty in building a computer that can reason as a person can.

People will have much greater exposure to computers, both at home and at work, and there will be a greater need to make computers more accessible. The design of the man—machine interface will be a vehicle for AI technology in general and natural language processing in particular. The computer should communicate effectively with the user, understand the intentions of the user, and learn about particular users from repeated exposure.

Thus, the relationship between man and computers will continue to evolve over the next ten years, and AI will be able to offer ways to normalize that relationship by making it easier for people to work with computers. There will be a greater public understanding of the problems of AI and also its possible applications.

One possible social impact of AI is unemployment. Over the next ten years it is not likely that this will be much of a concern. Viewed as a resource-allocation problem, the use of AI programs will most likely be for applications that are either too expensive or too dangerous to use people. AI programs will be fairly expensive (compared to word processing and other personal computer applications) due to the hardware requirements and the fact that there are not many people trained in AI.

In the long term, AI programs will take care of tasks that are better suited for computers, allowing humans to devote their time to personal services. For example, computers will play a major role in education in the coming years and the role of the teacher will change considerably. We would expect that the teacher in the next 10 or 20 years will be able to concentrate less on preparing lessons and correcting homework, and spend more time dealing with the interpersonal and social aspects of school. AI will play a growing role in education. Current educational software is at a very primitive stage. These programs often consist of little more than an electronic book in which the child types RETURN to see the next page. One active area of AI

research is ICAI (intelligent computer assisted instruction). An ICAI program builds a model of the student and tries to understand the misconceptions the student may have about the problem. Natural language abilities will clearly be of great importance to the educational programs of the future.

4 AI and Economics

AI programs require quite large amounts of computer power. As hardware becomes less expensive, more powerful, and more widely available, AI programs can finally become economically feasible in the marketplace.

As mentioned above, AI plays a growing role in the design of software systems – particularly the user interface. Given the increasing power of hardware, these commercial systems will become ever more common and widespread. The demand for these intelligent systems will be tremendous. In ten years, an AI interface will be the *sine qua non* of any major software system.

There is, however, a problem with this scenario. There will not be enough trained people available to build the systems. Even today, there is a critical shortage of manpower in AI. Only a limited number of people across the world are currently actively engaged in AI endeavors. Expansion of the field will certainly occur, but it is likely that the average quality of the work will diminish. It may seem odd to some people to state this problem. Aren't there thousands and thousands of computer programmers in the world? We see job training centers popping up to turn unemployed steel workers into programmers. We recognize the demand for programmers and apparently there are many efforts underway to meet this demand.

However, the situation is quite different in AI. An AI programmer requires considerably more training and ability than an applications or systems programmer. The problems are not well formulated. The answers are not absolute. The usual training for an AI programmer (who should already be a proficient programmer) involves a period of apprenticeship at a research lab. This is vaguely comparable to an internship or residency in medical education. The problem, though, is that most medical schools in the United States graduate more doctors each year than the cumulative total number of AI researchers trained in the entire world.

Thus, the amount of progress in the field and the application of results is severely limited by the number of researchers. This shortage of trained researchers and practitioners will have a noticeable damping effect on the economic impacts of AI.

5 AI in the 21st Century

Imagine the world 100 years from now. There will be robots that can understand speech, read the newspaper (if there are newspapers), recite Shakespeare, and assemble a bicycle. These robots will be closer to R2D2 or C3PO than to any current machine. They will be our servants.

This is the science-fiction vision. It is probably true, but it is surely a small part of the picture. Saying that AI research will culminate primarily in robot slaves is like saying that the main result of the discovery of electromagnetism is color television.

Clearly there is a technological aspect of AI that drives much of the research. However, there is a very important scientific aspect of AI that involves the study of the mind. What is the organizational structure of human memory? How do people learn? How do people make decisions? How do people think?

It is the answers to these admittedly grand questions that hold the greatest potential benefits to mankind. AI as a science may be able to point the way to discoveries in other areas. It may be possible to apply such advanced AI programs to a variety of pressing human problems – medical, educational, political, economic, and social.

As a case in point, imagine what a political candidate might expect from AI in 100 years. There could be a program that could advise the leader on the course of action to take in a given situation. One might think of this program as a political expert system. The program could devise campaign strategy, draft speeches, send out press releases, prepare position papers, conduct polls (via electronic mail), develop media commercials, and analyze the opponent's record.

Now let us take it one more step. A computer program that could display such an encyclopedic knowledge of the campaign and politics generally should be a candidate *itself*. While we do not wish to suggest that people will end up voting for computers, it is reasonable to predict that computers, with the help of AI programs, will play a growing role in the way in which our governments are run. For example, an AI expert system on preparing the budget will most likely evolve over the next century. Such a program would not be a mere spread-sheet calculation program, but would analyze the line items of the budget from a political perspective. A country's budget is the national agenda, an allocation of the country's resources, and a reflection of the political values of the country. An AI program designed to prepare the budget should know what the goals of the leadership are and how those goals can best be enacted through the budget. The program should know about the political feasibility of passing various provisions and what impact they will have on society.

This political advisor is but one of many types of application of AI that could have a much greater impact on people's lives than a robot that brings you your coffee every morning. The horizons for profound applications of AI are wide.

6 AI Today

As indicated above, the areas of learning and knowledge organization are of central importance. These touch on many of the subfields within AI, including natural language processing, vision, speech recognition, expert systems, and robotics. We have also mentioned the reasons for developing a psychological methodology for modeling cognitive phenomena.

One main requirement for productive and useful AI research is the capability to work on large, real-world problems. There is a broad consensus that no major breakthrough in AI will emerge from a small, concise program or from toy domains. AI programs should be explorations in complexity and should tackle the world as it exists.

One extension to this argument is that AI researchers should look at problems that span more than one subfield. Thus, there should be efforts in integrating expert systems and natural language, or speech recognition and vision. Researchers tend to shy away from these large problems, saying that there are still many small problems that have yet to be solved. That may be, but by looking at larger problems, we may gain important insights into the nature of intelligence and cognition.

References

Schank, R. (1982) *Dynamic Memory: A Theory of Learning in Computers and People*. Cambridge University Press.

Schank, R. and R. Abelson (1977) *Scripts, Plans, Goals and Understanding*. Hillsdale, NJ: Lawrence Erlbaum Associates.

CHAPTER 11

AI — Subjective Views, Future, and Impacts

Tibor Vámos *

"If your intellect lets you down, approach experience, which is a weaker device and not so lofty..." (Montaigne: On Experience, Essais III, Ch.15.)

1 What's in a Name?

At the outset I would be happy if we could agree on a less gorgeous but more solid name for our subject. I fully realize that artificial intelligence (AI) is now an accepted and impressive-sounding title — as cybernetics was, 20–30 years ago — but it is giving rise to sci-fi expectations, especially among those people to whom we would like to sell our achievements and ideas, and it has an unfortunate tendency to introduce charlatans to a field that is, in fact, a real R&D area for hard work, the typical mixture of 1% inspiration and 99% perspiration, with discipline and controlled self-assessment. And misleading our audience is even more dangerous in a period when for the first time moderate successes, at the limits of current possibilities, have been achieved by people devoted to this area of problem solving. "Machine intelligence," for example, would be more attractive from the viewpoint of not attracting cheap propaganda or of raising false expectations.

*President, International Federation for Automatic Control (IFAC), Computer and Automation Institute, Hungarian Academy of Sciences, Kende utca 13-17, Budapest, Hungary.

This short note on my philosophical approach also indicates the direction of my further remarks. I would like to select for examination those areas where relevant results are being obtained, i.e., real problems that can be solved or at least have reached the stage of realistic formulation. In this context I will emphasize not only the primacy of practical applications (*not* demonstrations) but also the degree of problem complexity. This is where the real importance of AI is to be found: as an approach for solving or helping to solve problems of real-life complexity, which are either almost or completely insoluble by conventional, straightforward computational methods.

Like any scientific discipline, our field is growing in two directions: methods (methodologies) and applications. If these are not related (not necessarily simultaneously or directly), both lose their relevance.

In my opinion, two major fields have emerged in the recent past as successful, relevant, and promising directions: vision and knowledge engineering (the latter including the term "expert systems" and, somehow, the problem of understanding). I will discuss each of these areas in turn, before examining possible social and economic implications of AI.

2 Perception: Vision and Voice Recognition

2.1 Vision

The use of artificial vision devices is already a practical proposition, with hardware and software on the market for several different purposes. The most important are as follows:

Picture processing

- special applications for remote sensing, used in meteorology, agriculture, natural resource surveys, archeology, etc. (devices intended to facilitate automated massacres are beyond my scope);
- biomedical applications (especially recognition of morphological characteristics, such as shape and texture);
- inspection functions, especially in technological processes (roughly the same tasks).

Robot vision

- object recognition (2D, 3D, color, partially invisible, shaded, and overlapping objects, etc.);

–   situation recognition (range finding, objects in motion, relations between several objects).

Here a few remarks are in order. It is quite typical yet perhaps still surprising that research projects start out with very ambitious (and highly intelligent) general goals, but as they approach realization, attention is progressively restricted to rather pragmatic and economic solutions. Most vision applications would not be considered as real AI results by the more extravagant AI philosophers. These applications use several simple edge-detection, region-discrimination, and selection techniques, combined with the calculation of easily computable geometric characteristics (area/perimeter, momentum, connectivity, etc.). Easy teaching methods, standard object modeling, and the calculation of texture features (intersections, correlations, etc.) are the major tools here.

Even the search procedures are mostly very simple for decision in case of ambiguities. The alternatives are well-known in advance, a learning period evaluates the most relevant features of particular cases, some statistical or fuzzy parameters can help in branching and accelerating the search, and if the machine decision is not sufficiently convincing (beyond some preset level of certainty), the human operator must intervene. Most of these applications have evolved toward feasibility through the availability of microprocessor-based work stations, i.e., by solutions whose costs lie in the US $10,000–20,000 range, when camera, other peripherals, and software are included.

A rapid further advance in this field seems very likely. As the applications proliferate, cheaper, strictly dedicated, and limited recognition and inspection systems are expected to replace many monotonous, and not really human workplaces in the areas of inspection, selection, simple evaluation, transfer, assembly, and other industrial operations (painting, welding, etc.). Vision is replacing such expensive tools and fixtures, and it promotes the flexibility of manufacturing processes. Most probably in the very near future (5–15 years) machine vision will be as natural an ingredient of work stations as microprocessors are today; big industries with a wide variety of products built of standard modules will be the norm. Applications will be extended to commerce and the services. Some people dream of a household robot by the 1990s; being personally engaged in doing such work regularly at home, I am unfortunately not so optimistic!

What kind of research will be needed in order to achieve these goals? The computer industry is performing its role in providing hardware – the present-day microprocessors are more than sufficient for most applications, while very fast parallel processors for picture processing (especially for tracking moving targets and such like, which

are not easily recognizable) are on the horizon. The optical industry mated with semiconductors promises new input devices. Areas where I would intensify efforts are the following:

- parallel recognition algorithms;
- character recognition, cheaper and more powerful optical character readers;
- a more systematic attack on morphology. This would mean a real AI-like approach: to transcribe the classical morphological descriptions of biologists, geologists, etc., into a formal meta-language and transfer this procedure into recognition and discrimination;
- more 3D research, knowledge-based help for the guidance of moving objects (transfer machinery) in the industrial environment and for robot assembly and other kinds of complex manipulation that need flexible adaptation;
- a computational link between computer-aided design on displays and automated operations (machinery, robot operations). This task leads toward knowledge-engineering but it also has a relevant visual interface, the understanding of images composed by various methods (computer graphics, 2D–3D transformations, photos, etc.).

Many of these tasks will imply improved methods of handling pictorial data. Some people think that data-base management is the same for any kind of data, but I do not believe in the efficiency of generalized methods in this respect because the requirements are especially stringent and specific to the problem.

2.2 Voice recognition

Vision has been singled out here from other kinds of perception. This is natural in that it is the most relevant aid to human activity and the most advanced, too. The other important field is voice recognition. We can rather clearly perceive distinct boundaries as the order of complexity increases. A recognizer for a few tens of words is commercially possible and an easy task for any expert engineer. The next order of magnitude (a few hundred words) is rather expensive but still feasible. One order of magnitude higher and the problem already appears to be a Himalaya. The reasons are now clear and have been discussed by many authors – including those who have attained altitude records by dint of terrible efforts and also those who simply examined the basic, underlying problems of complexity. My personal impression is that in the next

decade we shall see a lot of very practical applications of strict, uncontexted (or simply contexted) military commands like spoken languages, a new tool that could help people in communication with computer-based devices and also leave their hands free for action, but a feedback, checking, and confirmation operation will be needed in every sensitive case. The replacement of the good stenographer–typist is , still far distant. This means that the main emphasis of activity should be on a low-context target of, at maximum, several hundred words. At many work stations and in other personal use situations this will trigger the next revolution: the mass replacement of keyboards, but *not* of human writing.

3 Knowledge Engineering

The second major area I feel to be very promising is knowledge engineering (expert, knowledge-based systems and related topics). This new technology of including human expertise in computer-aided systems will not be a replacement of human activity but a higher-level man–machine interface. To cite the conclusion of another recent paper of mine:

> "We cannot predict what the future perspective of the expert system is. They are not going to replace genuine human intellectual activity but are promising helpers in coming man–machine systems. Most probably a merging of recent trends in information systems and expert systems provides a new stage of information technology. Relational and distributed database, user-friendly man–machine communication, information service-networks will provide a practical application spectrum without the gorgeous nomenclatures which are nowadays popular in selling science and products." (Vámos 1983.)

Recognition will continue to be an important step in any kind of automation. In process control this role was played by the process instrumentation (thermometers, pressure, composition, and other sensors, transducers), but for most activities this link failed, since it really needed a part of human intelligence (perception) but was never itself creative, or really human. The mass processing of inputs can also revolutionize the potential of numerous sciences through the storage, processing, and retrieval of immense quantities of samples (medical, historical, archeological documents, etc.). It can help in mass screening of populations, in exploration for mineral resources, etc.

Concerning expert systems I can only repeat my last sentence. These systems will offer a new and substantial tool in the rethinking of existing knowledge in various fields, in the distribution of this

knowledge all over the world, and in extending the global human knowledge base to a much broader community with real-time access. They will be a great help (if properly used) in education, training, in any kind of maintenance (including biomedical), and in improving services; they will constitute a logical further development of the computer and communication revolution, if mankind will only realize that modern technology can also be used *for* the people rather than predominantly against it!

4 Social and Economic Implications

Many of the foregoing remarks relate also to the social impacts of AI. I have had several opportunities to express my opinion on the problems of unemployment related to automation. Unemployment is an incurable disease when it afflicts poor countries that cannot provide working places, nor afford appropriate training or retraining for the people. But in countries that are more well off, the political/economic decision as to whether or not to rearrange human activities as a consequence of progress in technology rests with the rich. The new needs for more social care, infrastructure, and services are far from saturation. It is really all a problem of economic strength, and economic strength is mainly based on the level of technology — in other words, automation.

The other relevant social question is whether these new trends in automation characteristic of AI can be helpful in strengthening democracy and increasing individual freedom, or whether, on the contrary, they will be new weapons for Big Brothers against the individual and new means for the "uniformization" and manipulation of the masses. This is an open question and cannot be answered by our science. We can only say that both directions are feasible. Because many papers and books have already been written on the disastrous alternative, I will discuss the optimistic one. I would like to quote here from a paper on the "State and Citizen," presented at a General Assembly Meeting of the Hungarian Academy of Sciences:

> "Our entire administrative structure is affected by what we developed somewhere along the line with the appearance of the written word, in an era when the clerks were literate and the people not. The development of information science, telecommunication networks, data banks and information systems makes both possible and imperative an administration that is entirely different from the existing one. On the one hand this will be a prerequisite if we are to be able to work competitively, and on the other hand, it offers new possibilities for improving the

quality of life, developing the socialist democracy, and increasing the attractiveness of our society.

In connection with computerized information, dangers have been stressed for too long concerning its possible use by a centralized, bureaucratic state to control its citizens. Much less emphasis has been laid on the opposite outlook, i.e., the extension of the rights of citizens, the transformation of public administration into a real service; after the initially centralizing tendency of computerization a much more forceful decentralization. With the development of nation- and world-wide information systems, centralization and decentralization become modified to an extent where it is easy to imagine, using current terms, countries without capitals and international systems functioning without centers, which are not hierarchic but truly cooperative as hoped for by the pioneers of socialism.

All these are not the distant dreams of futurologists or sci-fi novelists but real possibilities for the next one, two, or at most three decades; preparations made today are already late rather than timely. We are very close to a situation in which the actual statutes, regulations, and those decisions of government that are of interest are permanently and openly accessible to every citizen via the screen of his home television set. *The imaginary curtain erected over several thousand years by the alienated state and its clerks and offices can finally be drawn back for ever.*

In this way decisions and regulations can be made accessible and understandable to everybody; questions can be specified exactly and automated on the basis of judgmental criteria accessible and understandable to anyone, so that "arrangements," trickery, "influence," and the granting of privileges will become obsolete; it will be possible to separate out those decisions that require human judgment so that the person or body taking responsibility can be clearly identified. Just as socialist production and distribution can be realized in practice only at a high level of material production, computerized information opens up new possibilities for socialist administration.

This is how production—social structure—organization—administration can arrive at a new synthesis. For us to be able to survive world competition over the coming decades, this link in the chain will be decisive.

We are facing the possibility of a society of a qualitatively different composition and of a qualitatively differently mode of administration. The change has already commenced, and we cannot remain passive concerning it."

We need to examine a problem that has been rather neglected until now. In our efforts to humanize work, we focused our attention on avoiding tedious, monotonous jobs, and in that way increased the creative content of the work. We expected a fairly unlimited extension of human intellectual activity for people who had not before

experienced it. This expectation is mostly realistic: we see how man-kind has stepped out of illiteracy, how quickly children can master computers, how many useful skills can be learned if they are taught in an appropriate way, at an appropriate age, in an appropriate environ-ment. On the other hand there is a significant proportion of the popu-lation whose intellectual ability or development is retarded for prena-tal or perinatal reasons, and who at present can only be integrated into society in terms of low-level useful activities. Maybe this will be a major new task for both medical people and sociologists — the problems of the specially handicapped in a new society. I am involved now in a project that is attacking one part of this problem and is using expert systems for help.

Reference has already been made to the economy. I have men-tioned the evolution of future industries (e.g., vision) and new profes-sions are emerging (e.g., the knowledge engineer). Most important is the fact that wealth and welfare stem from higher technology; we have observed the fragility of fortunes based on the lucky availability of natural resources but lacking the "culture" of technology, as compared with the relative stability of the highly advanced countries in spite of all the unpleasant symptoms of the recent economic recession. Our field is one of the most promising; it influences and promotes other dis-ciplines, and it advances technologies in such a way that its impact on the economy ought to be mostly positive. One real *caveat* concerns its one-sided introduction into only the advanced countries: this runs the risk of making the gap between developed and developing countries even wider, which is undesirable for both. But proper distribution of these results could help the developing countries, especially in train-ing and in the service and maintenance of more advanced equipment and systems. The open, commercial availability of computer-accessible expertise can promote international trade, standardization, and the introduction of new equipment, and thus still higher levels of trade. As software becomes a major commercial item, a more sophisticated version than a network-service available expertise will be even more — *if* (and I can't emphasize this enough) properly used!

References

Boden, M.A. (1977) *Artificial Intelligence and Natural Man*. Brighton, UK: Harvester.

Duda, R.O. and J.G. Gasching (1981) Knowledge-Based Expert Systems Come of Age. *BYTE*, 6/9 (September).

Lenat, D.B. (1982) The Nature of Heuristics. *Artificial Intelligence*, 19:189–249.

Lenat, D.B. (1983a) Theory Formation by Heuristic Search. The Nature of Heuristics, II: Background and Examples. *Artificial Intelligence*, 21:39–59.

Lenat, D.B. (1983b) EURISKO: A Program that Learns New Heuristics and Domain Concepts. The Nature of Heuristics, III: Program Design and Results. *Artificial Intelligence*, 21:61–98.

McCarthy, J. and P.J. Hayes (1979) Some Philosophical Problems from the Standpoint of Artificial Intelligence. In B. Meltzer and D. Michie (Eds.), *Machine Intelligence*, 9. New York: Wiley.

Michie, D. (Ed.) (1979) *Expert Systems in the Microelectronic Age*. Edinburgh: Edinburgh University Press.

Minsky, M. (1977) Frame-System Theory. In P.N. Johnson-Laird and P.C. Wason (Eds.), *Thinking*. Cambridge: Cambridge University Press.

Nilsson, N.J. (1980) *Principles of Artificial Intelligence*. Palo Alto, CA: Tioga.

Vámos, T. (1981) Hungary and Technical Progress. *The New Hungarian Quarterly*, 22.1981.84:19–32.

Vámos, T. (1983) *Cybernetics: The Border of Intelligence – A Faustian Adventure*. Presented at the Seminar on Science and Technology by the End of the XXth Century, Sitges, Spain.

Vámos, T. (1984) *Neuere Erfahrungen mit Expertsystemen*. INFO '84, Dresden.

Vámos, T. and M. Báthor (1982) A Robot Vision Lab Concept. In D. Michie (Ed.), *Machine Intelligence*, 10. New York: Wiley, pp. 211–226.

Vámos, T. and L. Mérő (1981) Medium Level Vision. In A. Rosenfeld and L. Lanal (Eds.), *Progress in Pattern Recognition*, 1. New York: North-Holland, pp. 93–122.

140

CHAPTER 12

An Artificial Intelligence Bibliography

Ingeborg Steinacker, * *Robert Trappl,* * *and Werner Horn* *

This bibliography contains 672 books, articles, and research memoranda on the future and potential impacts of artificial intelligence. To make the information more easily accessible, the bibliography is divided into three parts:

- a thesaurus of the more important terms,
- a KWIC Index,
- a complete alphabetical (by author) list of the references.

Entries may be easily cross-referred using the numbers in square brackets [·].

The structure of the thesaurus has been influenced by a framework developed by Sugiyama (1982) for technological innovation in computation, communication, and control, and by comments made on an earlier version by Helga Nowotny.

Although we have endeavored to make this bibliography as comprehensive as possible, there may still be many references missing. We would be most grateful for your help in improving future versions.

*Department of Medical Cybernetics and Artificial Intelligence, University of Vienna, Freyung 6, A-1010 Vienna, Austria.

Thesaurus

Future of AI

General [22, 106, 129, 213, 244, 255, 262, 285, 308, 389, 404, 414, 439, 440, 455, 463, 475, 507, 536, 603, 603, 606, 612, 615, 616]
Hardware [7, 153, 210, 221, 244, 646, 664, 666]
  LISP machines [378, 379, 574, 587]
  Parallel processors [58, 152, 204, 224, 251, 504, 554, 563]
Software [107, 274, 606]
  AI languages [32, 154]
    Micro-LISP [176, 185, 239, 274, 497, 557]
    PROLOG [95, 103, 133, 159, 287, 357, 479, 494]
    LOGO [58, 268, 416]
    KR languages [49, 70, 154, 195, 196, 306, 511]
  AI environment [129, 186, 207, 378, 476, 523, 561, 573]
  Cooperative processes [58, 115, 177, 552]
AI theory [28, 30, 86, 96, 97, 403, 500]
  Nonmonotonic reasoning [58, 331, 349, 358, 359, 381]
Knowledge representation [30, 58, 63, 218, 395, 470, 588]
Learning [58, 96, 132, 257, 366, 463, 464]
Expert systems [8, 9, 68, 75, 84, 120, 140, 146, 160, 179, 216, 263, 310, 414, 431, 458, 464, 500, 554, 571, 588, 640]
  Transformation to other domains [58, 243, 457, 561, 595]
  Knowledge engineering [58, 134, 149, 155, 500, 554]
  Explanation [58, 90, 509]
  Knowledge acquisition [27, 58, 119, 277, 400, 464, 561]
  Deep systems [58, 86, 213, 233, 368]
  Structural, strategic, meta-knowledge [58, 91, 121]

Completeness and consistency [508]
Natural language [30, 54, 61, 198, 216, 245, 277, 307, 307, 426, 464, 513, 579, 581, 584, 588]
  Text analysis [58, 141, 201]
  Machine translation [269, 270, 329, 332, 445, 593, 607, 625]
  Dialog [25, 72, 279]
  DB interfaces [59, 200, 259, 539]
Vision [26, 65, 66, 67, 96, 278, 554, 587]
  Low-level [58, 115]
  3-D [71, 258, 554]
Speech understanding [25, 30, 58, 110, 191, 216, 293, 325, 338, 554]
Signal detection [401]
Robotics [4, 21, 58, 67, 107, 129, 131, 144, 164, 170, 262,2 78, 380, 387, 464, 501, 522, 554, 601]
Automatic programming [1, 32, 33, 107]
Education in AI [84, 404, 464, 564]
AI critique [129, 139, 215, 444, 575]

Impacts of AI

Sciences [32, 51, 98, 129, 171, 175, 207, 309, 310, 337, 396, 404, 414, 419, 490, 595]
  Philosophy [51, 52, 126, 128, 230, 231, 348, 382, 388, 441, 471, 486, 531, 538]
  Psychology [4, 18, 31, 34, 51, 52, 53, 54, 57, 58, 61, 81, 89, 135, 184, 187, 188, 216, 231, 249, 326, 361, 370, 372, 409, 414, 443, 462, 464, 471, 530, 535, 541, 615]
  Linguistics [137, 138, 188, 292, 426, 460, 489, 581, 581, 583]
  Biology [12, 16, 52, 419, 440, 488, 499, 666]

KWIC Index

Author	Title	Ref.
Boden M.A.	Artificial Intelligence and Biological Reductionions	[55]
Boden M.A.	The Educational Implications of Artificial Intelligence	[56]
Boden M.A.	Artificial Intelligence and Animal Psychology	[57]
Boden M.A.	Impacts of Artificial Intelligence - An Introduction	[58]
Charniak E.	Artificial Intelligence	[87]
Clancey W.J.	Readings in Medical Artificial Intelligence, Vol.3	[92]
Cohen P.R.	The Handbook of Artificial Intelligence,	[98]
Coles L.S.	The Application of Artificial Intelligence to Medicine	[101]
Cook S.	The Applications of Artificial Intelligence to Law: A Survey of Six Cu	[105]
Crumley D.V.	Concepts for Army Use of Robotic-Knowledge-Based Systems in Artificial Intelligence in the 21st Century	[112]
Davis R.	Artificial Intelligence	[121]
Dresher B.E.	On Some Supposed Contributions of Artificial Intelligence to the Scientific Study of Reason	[137]
Dreyfus H.L.	What Computers Can't Do. A Critique of Artificial Intelligence	[139]
Enyart J.M.	Artificial Intelligence Applied to the Command, Co	[145]
Feigenbaum E.A.	The Art of Artificial Intelligence: Thames and Case Studies o	[155]
Feigenbaum E.A.	The Fifth Generation. Artificial Intelligence and Japan's Computer Chall	[156]
Firschein O.	Forecasting and Assessing the Impact of Artificial Intelligence on Society	[161]
Foster E.	Artificial Intelligence Faces a Crossroad	[165]
Freiherr G.	The Seeds of Artificial Intelligence, SUMEX-AIM	[168]
Gause C.D.	Cybernetics and Artificial Intelligence	[175]
Gini G.	Artificial Intelligence: Impact on Computer Assist	[183]
Gleick J.	Artificial Intelligence: Exploring Labyrinth of th	[184]
Goldstein I.	Artificial Intelligence, Language and the Study of	[187]
Goldstein I.	Artificial Intelligence, Language, and the Study o	[188]
Groveston P.K.	Overview of Artificial Intelligence-Based Teaching Aids	[199]
Havel I.M.	Artificial Intelligence: A Lesson in Human Self-Un	[216]
Hertz D.R.	Artificial Intelligence and the Business Manager	[223]
Jackson P.C.	Introduction to Artificial Intelligence	[236]
James M.	Artificial Intelligence in BASIC	[238]
Kahn K.M.	valuation, Programming Methodology, and Artificial Intelligence	[257]
Krutch J.	Experiments in Artificial Intelligence for Small Computers	[289]
Kulikowski C.A.	Artificial Intelligence Methods and Systems for Me	[291]
Lehman-Wilzig S.	Unbound, Towards a Legal Definition of Artificial Intelligence	[305]
Lindsey R.K.	Applications of Artificial Intelligence for Organic Chemistry: The	[321]
Longuet-Higgins	Artificial Intelligence - A New Theoretical Psycho	[326]
McCarthy J.	ophical Problems from the Standpoint of Artificial Intelligence	[348]
Michaelis P.R.	Artificial Intelligence and Human Factors Engineer	[363]
Michalski R.S.	Machine Learning: An Artificial Intelligence Approach	[366]
Nageo M.	Social and Economic Impacts of Artificial Intelligence	[387]
Newell A.	Intellectual Issues in the History of Artificial Intelligence	[396]

Author	Title (keyword index)	Ref.
Kato I.	Coming Robot New Era (A Viewpoint)	[262]
Leontief W.W.	The New New Age That's Coming Is Already Here	[315]
Pitt D.	Paradigms lost? Reflections on the Coming Organizational 'Revolution'	[424]
Rossman P.	The Coming Great Electronic Encyclopedia	[450]
Bock P.	English as a Command Language for Robot Control (HIROB)	[50]
Davis P.K.	as a Technique for Exploring Strategic Command and Control Issues	[117]
Enyart J.N.	Artificial Intelligence Applied to the Command , Control, Communications, and Intellige	[145]
Enyart J.N.	ns, and Intelligence of the U.S.Central Command	[145]
Ernest M.L.	The Mechanization of Commerce	[148]
Dray J.	hose Time has Come or a Pipe Dream with Commercial Overtones?	[136]
Kaplan R.	Commercial Sector: AI in the Pharmaceutical Indust	[260]
Minsky M.	Look Ahead, AI Profs go Commercial	[374]
Davey P.G.	Robots with Common Sense? The Research Scene To-Day	[113]
Bruce B.C.	Natural Communication Between Person and Computer	[72]
Mariani J.	A Man-Machine Speech Communication System Including Word-Based Recognition	[338]
NN	, Scholarly Communication : Online Publishing and Education	[654]
Enyart J.N.	Igence Applied to the Command, Control, Communications , and Intelligence of the U.S.Central Co	[145]
Nemes L.	and Evaluation Methods for Man-Machine Communications on the Shop Floor	[394]
NN	ment and Representative of the European Communities	[605]
Hammond J.H.	Change: Information Technology and the Community	[209]
NN	High-Tech Companies Team Up in the R&D Race	[637]
NN	High-Tech Companies Team up in the R&D Race	[638]
Jenny C.J.	Expert Systems as Seen by an Insurance Company	[243]
Dieterich T.G.	l Descriptions: Evaluation Criteria and Comparative Review of Selected Methods	[132]
O'Keefe R.A.	PROLOG Compared with LISP	[412]
Lucas R.W.	ients with Alcohol-Related Illnesses: A Comparison	[328]
Suwa M.	An Approach to Verify Completeness and Consistency in a Rule-Based Expert	[506]
Newman N.	Europe's Catch-Up Complex	[399]
Trappl R.(ed.)	utonomy, Utopia: New Approaches Towards Complex Systems	[537]
Gould J.D.	Composing Letters with a Simulated Listening Type	[191]
Dyer M.G.	of Integrated Processing for Narrative Comprehension	[141]
Kotelnikov V.A.	Comprehensive Programme of Scientific and Technologic	[285]
Arbib M.A.	Cooperative Computation and the Cybernetic Society	[20]
Barr A.	Computation	[31]
Turkle S.	Artificial Intelligence: Cognition as Computation	[541]
Batali J.	: A Study in the Psychology of Personal Computation Introspection	[34]
Boden M.A.	echanisms: Philosophical Psychology and Computational Models	[52]
Manthey M.J.	The Computational Metaphor and Quantum Physics	[337]
Meltzer B.	The Unconscious in Psychology and Computational Modelling	[361]
Smith B.C.	The Computational Metaphor	[488]

Lucas R.W.	:	Psychiatrists and a Computer as Interrogators of Patients with Alcoh	[328]
Lucky R.	:	The Social Impact of the Computer	[330]
Macaulay L.A.	:	Computer Education for Everyman	[333]
Malik R.	:	Japan's Fifth Generation Computer Project	[335]
Marsh P.	:	A Brighter Future for the Computer Kids	[340]
Marsh P.	:	Computer Kids Head for the Big Time	[341]
Mason	:	Alvey Lines up Computer Research Projects	[344]
Mason R.	:	The Computer Home: Will Tomorrow's Housing Come Aliv	[345]
Modesitt K.L.	:	An Academic Meets Industry - Rethinking Computer-Based Education and Personalised System	[375]
Moto-Oka T.(ed.)	:	national Conference on Fifth Generation Computer Systems	[384]
Pagels H.R.(ed.)	:	Computer Culture: The Scientific, Intellectual a	[414]
Pagels H.R.(ed.)	:	, Intellectual and Social Impact of the Computer	[414]
Powers D.M.	:	cs and Psycholinguistics as a Basis for Computer Acquisition of Natural Language	[426]
Pribram K.H.	:	ain, the Telephone, the Thermostat, the Computer , and the Hologram.	[427]
Prigmore C.A.	:	A Personal View of Computer Assisted Learning	[428]
Raphael B.	:	The Thinking Computer - Mind Inside Matter	[429]
Reggia J.A.	:	Computer-Assisted Medical Decision Making: A Cri	[433]
Rifkin J.	:	The Other Half of the Computer Revolution	[440]
Rushby N.J.(ed.)	:	Computer Based Learning	[453]
Schank R.C.	:	Inside Computer Understanding	[461]
Shneiderman B.	:	Software Psychology. Human Factors in Computer and Information Systems	[477]
Shortliffe E.H.	:	or Medical Decision Making: A Review of Computer-Based Clinical Decision Aids	[478]
Simon H.A.	:	Veraendert der Computer unser Leben? Zehn Thesen des Nobelpreis	[482]
Stolfo S.J.	:	cations of DADO: A Large-Scale Parallel Computer for Artificial Intelligence	[504]
Sugiyama K.	:	Q-Analysis (I): A Case Study of Future Computer Systems Developments in Japan	[506]
Szabo M.	:	n Assessment of the State of the Art of Computer based Learning Systems Design	[510]
Teicholz R.	:	Computer Integrated Manufacturing	[519]
Toda M.	:	ysis (II): Strategy Analysis for Future Computer Systems Development in Japan	[528]
Trappl R.	:	Computer Psychotherapy: Is it Acceptable, Feasib	[530]
Trappl R.	:	Ontology and Semantics in the Computer : A Study in the Psychology of Personal	[531]
Turkle S.	:	The Subjective Computer	[541]
Uttal B.	:	Here comes Computer Inc.	[547]
Uttal B.	:	Free-For-All in Computer-Aided Engineering	[548]
Verity J.W.	:	Nuclear War & The Computer	[560]
Weizenbaum J.	:	Computer Power and Human Reason - From Judgement	[575]
Westin A.F.	:	New Issues of Computer Privacy in the Eighties	[576]
Winston P.H.(ed.)	:	The Psychology of Computer Vision	[585]
Yazdani M.	:	Novel Ways of using a Computer in an Educational Institution	[600]
NN	:	Artificial Intelligence, The Second Computer Age Begins	[606]
NN	:	he Fifth Generation: Dawn of the Second Computer Age	[614]

Author		Keyword		Ref.
Haugeland J. (ed	Mind	Design		[215]
Hogan D.I.	Survivability	Design	Guidelines for Fly-by-Wire Flight Contr	[232]
Jones C.B.	Specification and	Design	of (Parallel) Programs	[251]
Jones H.A.	Cybernetic	Design	Features for Computerised Business Info	[252]
Leal A.	Theoretical Foundations, Expert System	Design	and Experimental Plan	[302]
Miller D.A.		Design	for an On-Line Computerized System for	[369]
Nemes L.		Design	Criteria and Evaluation Methods for Man	[394]
Newell S.B.		Design	Automation for Integrated Circuits	[397]
Steeb R.		Design	for an Advanced Red Agent for the Rand	[496]
Steele G.L.		Design	of a LISP-Based Microprocessor	[497]
Szabo M.	Art of Computer based Learning Systems	Design		[510]
NN	Getting Computers to	develop	their own Chips	[634]
NN	Look Ahead, AI Markets	Developing		[641]
Aiken R.M.	Microcomputers as Educational Tools in	Developing	Countries	[2]
Gupta G.K.	Computing in the	Developing	Countries of ASIA	[203]
Sugiyama K.	Toward	Development	a Framework for Information Technology-	[505]
Davis P.K.	d Interim Conclusions about Utility and	Development	Options	[118]
Hogan D.I.	for Fly-by-Wire Flight Control Systems	Development		[232]
Sloman A.	POPLOG: A Multilanguage Program	Development	Environment	[487]
Toda M.	gy Analysis for Future Computer Systems	Development	In Japan	[528]
NN	an Strategic Programme for Research and	Developments	In Information Technologies (ESPRIT)	[649]
Austing R.H.	Recent	Developments	In Computers and Society Research and E	[23]
Sugiyama K.	A Case Study of Future Computer Systems	Developments	In Japan	[506]
Gini G.	ce: Impact on Computer Assisted Medical	Diagnosis		[183]
Kobsa A.	erangepasstes Dialogverhalten such ohne	Dialogpartnermodell	moeglich?	[279]
Guida G.		Dialogue	with Data Bases: An Effective Natural L	[200]
Lamson R.W.	Mapping Policy	Dialogues	: A Role for Computer Graphics	[300]
Kobsa A.	ist benutzerangepasstes	Dialogverhalten	such ohne Dialogpartnermodell moeglich?	[279]
NN	"Computer People": Yes, They really Are	Different		[672]
Kildall G.		Digital	Research's DR Logo	[268]
Kraft A.	XCON: An Expert Configuration System at	Digital	Equipment Corporation	[288]
Kanade T.	Recovery of the Three-	Dimensional	Shape of an Object from a Single View	[258]
Waterman D.A.	Pattern-	Directed	Inference Systems	[571]
Hart P.E.		Direction	for AI in the Eighties	[213]
Lenat D.B.	The Ubiquity of	Discovery		[309]
Robertson P.	Non-Temporal Prediction - A	Distributed	System for Concept Acquisition	[443]
Leontief W.W.	The	Distribution	of Work and Income	[314]
Arbib M.A.	elligence and Brain Theory: Unties and	Diversities		[18]
Davis R.	Expert Systems: Where Are We? And Where	Do	We Go From Here?	[120]
Dreyfus H.L.	What Computers Can't	Do	. A Critique of Artificial Reason	[139]

Nilsson N.J.	:	Principles of Artificial Intelligence	[403]
Nilsson N.J.	:	Artificial Intelligence: Engineering, Science, or Slogan?	[404]
Nilsson N.J.	:	Artificial Intelligence, Employment, and Income	[405]
Nilsson N.J.	:	Artificial Intelligence, Employment, and Income	[406]
Nitzan D.	:	Machine Intelligence Research Applied to Industrial Automati	[408]
Rich E.	:	Artificial Intelligence	[437]
Ringle M.	:	hilosophical Perspectives in Artificial Intelligence	[441]
Rose J.(ed.)	:	and Research in Robotics and Artificial Intelligence	[447]
Schank R.C.	:	cial and Economic Impacts of Artificial Intelligence	[464]
Shapiro S.C.	:	Artificial Intelligence Applications	[475]
Sloman A.	:	Epistemology and Artificial Intelligence	[486]
Smith G.W.	:	Conturen: Mass Producing Artificial Intelligence for a Rational World	[490]
Smith L.C.	:	Implications of Artificial Intelligence for End User Use of Online Systems	[491]
Stolfo S.J.	:	-Scale Parallel Computer for Artificial Intelligence	[504]
Szolovits P.(ed.)	:	Artificial Intelligence in Medicine	[512]
Trappl R.	:	ine with Special Emphasis on Artificial Intelligence Methods. Report to the Austrian Federal	[532]
Trappl R.	:	nternational Tension through Artificial Intelligence: A Proposal for 3 Projects	[534]
Trappl R.	:	Impacts of Artificial Intelligence: An Overview	[535]
Trappl R.	:	Artificial Intelligence: A One-Hour Course	[536]
Trost H.	:	ine Interaction: the Role of Artificial Intelligence, Industrial Liaison Session	[539]
Wahlster W.	:	Tutorial on Artificial Intelligence in Parallel	[562]
Waldrop M.M.	:	Artificial Intelligence Applications Called Dead End	[563]
Wallich P.	:	Artificial Intelligence	[564]
Waltz D.L.	:	Artificial Intelligence	[566]
Wilks Y.	:	l Language Understanding and Artificial Intelligence	[581]
Winston P.H.	:	Artificial Intelligence: An MIT Perspective, Vol.2	[586]
Winston P.H.	:	Artificial Intelligence: An MIT Perspective, Vol.1	[587]
Winston P.H.	:	Artificial Intelligence. 2nd Ed.	[588]
Winston P.H.	:	Artificial Intelligence	[589]
Wong C.M.	:	Application of Artificial Intelligence to Triple Quadrupole Mass Spectrometry	[595]
NN	:	Artificial Intelligence. The Second Computer Age Begins	[606]
NN	:	spects for New Technologies: Artificial Intelligence	[612]
NN	:	Program of Conference on Artificial Intelligence	[648]
NN	:	Researchers Cash in on Artificial Intelligence	[651]
NN	:	The Artificial Intelligence Report	[659]
Clancey W.J.	:	Methodology for Building an Intelligent Tutoring System	[90]
Fox M.S.	:	ling, Simulation, and Scheduling in the Intelligent Management System	[166]
Fox M.S.	:	Intelligent Management System: An Overview	[167]
Jones K.P.	:	Searching Techniques: Time to Replace Intelligent Operators by Intelligent Systems	[253]
Jones K.P.	:	ime to Replace Intelligent Operators by Intelligent Systems	[253]

Author	Title	Ref.
Bramer M.A.	A Survey and Critical Review of Expert Systems Research	[68]
Buchanan B.G.	Research on Expert Systems	[75]
Buchanan B.G.	Rule-Based Expert Systems - The MYCIN Experiments of the Stanford	[76]
Davis R.	Knowledge Acquisition in Rule-Based Systems - Knowledge about Representations as a	[119]
Davis R.	Knowledge-Based Systems : Where Are We? And Where Do We Go From	[120]
Davis R.	Amplifying Expertise with Expert Systems in Artificial Intelligence	[121]
Dennet D.C.	Intentional Systems	[122]
Duda R.O.	Expert Systems Research	[125]
Epstein S.	Challenges - Expert Systems	[140]
Fergusen G.R.	Aircraft Maintenance Expert Systems	[146]
Finin T.W.	viding Help and Advice in Task Oriented Systems	[156]
Gerring P.E.	System Responsiveness in Interactive AI-Systems	[160]
Cevarter W.B.	An Overview of Expert Systems	[177]
Hatvany J.	Advanced Manufacturing Systems	[179]
Hayes-Roth F.	Building Expert Systems In Modern Society	[214]
Hilts S.R.	Viewing Computing Systems	[219]
Hinton G.E.	Shape Representation in Parallel Systems within a Social Context	[225]
Hogan D.I.	idelines for Fly-by-Wire Flight Control Systems	[226]
Janda A.(ed.)	Human Factors in Computing Systems Development	[232]
Jenny C.J.	Requirements on Expert Systems	[239]
Jones H.A.	s for Computerised Business Information Systems as Seen by an Insurance Company	[249]
Jones K.P.	ce Intelligent Operators by Intelligent Systems	[252]
Kosy D.W.	Knowledge-Based Support Systems	[253]
Kulikowski C.A.	Artificial Intelligence Methods and Systems for Long Range Planning	[284]
Leal A.	ectiveness of Military Decision Support Systems for Medical Consultation	[291]
Lee R.M.	mological Aspects of Management Support Systems . Theoretical Foundations, Expert System	[302]
Lee R.M.	mantics: Expert versus Decision Support Systems	[303]
Michaelsen R.	Expert Systems	[304]
Modesitt K.L.	mputer-Based Education and Personalised Systems In Business	[364]
Moto-Oka T.(ed.)	Conference on Fifth Generation Computer Systems of Instruction	[375]
O'Shea T.	Intelligent Systems	[384]
Raulefs P.	On the Requirements of Future Expert Systems In Education	[413]
Sauers R.	Adaptive Control in Ill-Defined Systems	[458]
Selfridge O.	man Factors in Computer and Information Systems : State of the Art and Future Prospects	[472]
Shneiderman B.	Intelligent Tutoring Systems	[477]
Sleeman D.	Intelligence for End User Use of Online Systems	[485]
Smith L.C.	The Production Systems	[491]
Spacek L.A.	The Organisation of Expert Systems and PROLOG	[494]
Stefik M.	The Organisation of Expert Systems . A Tutorial	[500]

List of References

[1]   Abbot R.J. (1983): Program Design by Informal English Descriptions, Communications of the ACM, 26(11)882-894.

[2]   Aiken R.M., Qualid A. (1983): Microcomputers as Educational Tools in Developing Countries, Information Processing, 471-475.

[3]   Albus J. (1983): The Robot Revolution: An Interview with James Albus, Comm.ACM, March, 179-180.

[4]   Albus J.S. (1981): Brains, Behavior, and Robotics, BYTE Books, McGraw-Hill, Petersborough, N.H..

[5]   Alexander C.P. (1983): The New Economy - Technology Has Set Off a Scramble for Jobs, Profits and Global Markets, Time, May 30.

[6]   Alexander T. (1982): Teaching Computers the Art of Reason, Fortune, May 17,83-160.

[7]   Alexander T. (1984): Reinventing the Computer, Fortune, pp.62-70, March 5.

[8]   Amarel S., Brown J.S., Buchanan B., Hart P., Kulikowski C., Martin W., Pople H. (1977): Panel on Applications of Artificial Intelligence, in Proceedings of the 5th International Joint Conference on Artificial Intelligence, MIT, Cambridge, USA.

[9]   Amarel S. (1981): Review of Characteristics of Current Expert Systems, CBM-TM-89, Rutgers Univ., New Brunswick, NJ.

[10] Anderson A. (1983): Industrial Pets for the Workers, New Scientist, pp.46-47, July 7.

[11] Anderson I. (1983): Californian Schools Reap Bumper Harvest of Apples, New Scientist, p.571, March 3.

[12] Anderson I. (1983): Dawn of the Biological Computer is Delayed, New Scientist, p.329, November 3.

[13] Anderson I., Lamb J. (1983): Silicon Valley's Black Friday, New Scientist, October 13.

[14] Anderson I. (1984): New Technology will not Provide Jobs, New Scientist, p.24, May 10.

[15] Anderson J., Rushby N.(eds.) (1981): Selected Readings in Computer-Based Learning, Kogan Page, London.

[16] Anderson J. (1983): Dawn of the Biological Computer is Delayed, New Scientist, p.329, November 3.

[17] Aquesbi A., Bocquet J.C., Fouet J.M., Tichkiewitch S., Reynier M., Trau P. (1983): An Expert System for Computer Aided Mechanical Design, in Mason R.(ed.): Information Processing 83, North-Holland, Amsterdam.

[18] Arbib M.A. (1975): Artificial Intelligence and Brain Theory: Unities and Diversities, Annals of Biomed. Engin., 3,238-274.

[19] Arbib M.A. (1977): Computers and the Cybernetic Society, Academic Press, New York.

[20] Arbib M.A. (1982): Cooperative Computation and the Cybernetic Society, in Trappl R., et al.(eds.), Progress in Cybernetics and Systems Research, Vol.VIII-XI, Hemisphere, Washington, DC.

[21] Arbib M.A., Lawton D.T., Overton K.J. (1983): Perceptual Systems for Robots, University of Massachusetts, Amherst, COINS Technical Report 83-24.

[22]   Arbib M.A. (1985): On Being Human in The Computer Age, in Trappl
       R.(ed.), The Impacts of Artificial Intelligence, North-Holland, Amster-
       dam.

[23]   Austing R.H., Engel G.L. (1979): Recent Developments in Computers
       and Society Research and Education, AFIPS National Computer
       Conference.

[24]   Ayres R.U., Miller S.M. (1983): Robotics: Applications and Social
       Implications, Ballinger, Cambridge, Massachusetts.

[25]   Ballantine M. (1980): Conversing with Computers - The Dream and the
       Controversy, Ergonomics, 23 (9) 935-945.

[26]   Ballard D.H., Brown C.M.(eds.) (1982): Computer Vision, Prentice-Hall,
       Englewood Cliffs, NJ.

[27]   Balzer R., Erman L., London P., Williams Ch. (1980): HEARSAY-III: A
       Domain-Independent Framework for Expert System, in Proceedings of
       the First Annual National Conference on Artificial Intelligence, AAAI,
       Stanford University, pp.108-110.

[28]   Banerji R.B. (1980): Artificial Intelligence: a Theoretical Approach,
       Elsevier, New York.

[29]   Barber G. (1983): Semantic Support for Work in Organizations, Infor-
       mation Processing, pp.561-566.

[30]   Barr A., Feigenbaum E.A.(eds.) (1981): The Handbook of Artificial
       Intelligence, Vol.1, William Kaufmann, Los Altos, Calif..

[31]   Barr A. (1982): Artificial Intelligence: Cognition as Computation,
       Report STAN-CS-82-956, Stanford Univ., Dept. of Computer Science.

[32]   Barr A., Feigenbaum E.A.(eds.) (1982): The Handbook of Artificial
       Intelligence, Vol.2, William Kaufmann, Los Altos, Calif..

[33]   Barstow D.R. (1979): Knowledge-Based Program Construction, North-
       Holland, New York.

[34]   Batali J. (1983): Computational Introspection, MIT, AI-Lab, Rep. AI-M-
       701.

[35]   Belkin N., Hennings -R.-D., Seeger T. (1983): Mass-Informatics and
       their Implication for Everyday Life, Information Processing, pp.583-
       587.

[36]   Bell J. (1983): We'll have Robots like Ladies have Hats, New Scientist,
       February 24,528-529.

[37]   Bell J. (1983): Robot Prices Slashed, New Scientist, pp.366-367,
       November 3.

[38]   Bell T.E. (1983): My Computer, My Teacher, Personal Computing,
       pp.120-127, June.

[39]   Bell T.E.(ed.) (1983): My Computer, My Teacher, Personal Computing,
       pp.120-127, June.

[40]   Benoit J.W., Lemmon A.V., Selander J.M. (1982): KBS: An Expert Plan-
       ning System for Crisis Response, The MITRE Corporation, Report
       RADC-TR-82-217, McLean, Virginia.

[41]   Benson I.A. (1979): Planning for the Computerized Future - A Trade
       Union View, in Samet P.A. (ed.): Proceedings of the EURO IFIP,
       North-Holland Publishing Company, Amsterdam.

[42]   Ben-Bassat M., Freedy A. (1983): Knowledge Requirements and Manage-
       ment in Expert Decision Support Systems for (Military) Situation
       Assessment, Israel Inst. of Business, Tel-Aviv, Report 46/81.

[43]   Bernus P., Hatvany J., Nemes L. (1981): The Computer Aided Design of
       a Manufactoring System - Some Experimental Results, Proceedings of
       the 8th Triennial World Congress: Control Science and Technology for
       the Progress of Society, 14, Kyoto.
[44]   Best F. (1984): Technology and the Changing World of Work, The Futu-
       rist, pp.61-66, April.
[45]   Bidmead C. (1983): Teaching the IT Generation, New Scientist, p.445,
       November 10.
[46]   Bikson T.K., Gutek B.A., Mankin D.A. (1981): Implementation of Infor-
       mation Technology in Office Settings: Review of Relevant Literature,
       P-6697, Rand Corp..
[47]   Bisseret A. (1983): Psychology for Man Computer Cooperation in
       Knowledge Processing, in Mason R. (ed.): Information Processing 83,
       North-Holland, Amsterdam.
[48]   Blanning R.W. (1984): Conversing with Management Information Sys-
       tems in Natural Language, Comm. ACM 27(3)201-207.
[49]   Bobrow D.G., Winograd T. (1979): KRL: Another Perspective, Cognitive
       Science, 3(1)29-42.
[50]   Bock P. (1983): English as a Command Language for Robot Control
       (HIROB), Robotica, 1(1)19-24.
[51]   Boden M.A. (1977): Artificial Intelligence and Natural Man, The Har-
       vester Press, Hassocks.
[52]   Boden M.A. (1981): Minds and Mechanisms: Philosophical Psychology
       and Computational Models, Cornell University Press, Ithaka, New
       York.
[53]   Boden M.A. (1983): Artificial Intelligence as a Humanizing Force, in
       Proceedings of the 8th International Joint Conference on Artificial
       Intelligence, Karlsruhe, West Germany, pp.1197-1198.
[54]   Boden M.A. (1983): Computers Wait for Us to Understand, New Scien-
       tist, p.123, July 14.
[55]   Boden M.A. (1983): Artificial Intelligence and Biological Reductionism,
       CSRP 021,University of Sussex,UK.
[56]   Boden M.A. (1983): The Educational Implications of Artificial Intelli-
       gence, CSRP 017,University of Sussex,UK.
[57]   Boden M.A. (1983): Artificial Intelligence and Animal Psychology,
       CSRP 016,University of Sussex, UK.
[58]   Boden M.A. (1985): Impacts of Artificial Intelligence, in Trappl R.(ed.),
       The Impacts of Artificial Intelligence, North-Holland, Amsterdam.
[59]   Boguraev B.K., Sparck Jones K. (1982): How to Drive a Database Front
       End Using General Semantic Information, University of Cambridge,
       Computer Laboratory, TR-32.
[60]   Bond A.H. (1981): Machine Intelligence, Infotech State of the Art
       Report,9(3),Pergamon,Maidenhead,UK.
[61]   Boot M., Koppelar H. (1982): The Natural Computer or Paradigm Lost,
       in Goetschalckx J., Rolling L. (eds.): Lexicography in the Electronic
       Age, North-Holland, Amsterdam.
[62]   Boraiko A. (1982): The Chip, Electronic Mini-Marvel that is Changing
       Your Life, National Geographic, 162(4)421-456.
[63]   Brachman R.J., Fikes R.E., Levesque H.E. (1983): KRYPTON: A Func-
       tional Approach to Knowledge Representation, FLAIR Technical
       Report 16.

[64]   Brady M. (1981): Special Issue on Computer Vision: Preface - The Changing Shape of Computer Vision, Artificial Intelligence, 17,1.

[65]   Brady M.(ed.) (1981): Special Issue on Computer Vision, Artificial Intelligence, 17.

[66]   Brady M. (1982): Computer Vision (Correspondent's Report), Artificial Intelligence, 19(1)7-16.

[67]   Braggins D. (1983): Robots Shapen up their Vision, New Scientist, 811, December 15.

[68]   Bramer M.A. (1982): A Survey and Critical Review of Expert Systems Research, in Michie D.(ed.), Introductory Readings in Expert Systems, Gordon & Breach, New York.

[69]   Brandin D.H. (1982): The Challenge of the Fifth Generation, Communications of the ACM, 25(8)509-510.

[70]   Brodie M.L., Zilles S.N.(eds.) (1981): Proceedings of Workshop on Data Abstraction, Databases and Conceptual Modelling, Sigart Newsletter 74.

[71]   Brooks R.A. (1981): Symbolic Reasoning Among 3-D Models and 2-D Images, Artificial Intelligence, 17,285-348.

[72]   Bruce B.C. (1982): Natural Communication Between Person and Computer, in Lehnert W.G., Ringle M.H.(eds.), Strategies for Natural Language Processing, Lawrence Erlbaum Ass., Hillsdale.

[73]   Bruegge P. (1982): Sagen wir lieber nicht Humanitaet, Der Spiegel,36,76-112.

[74]   Bruno C. (1984): Labor Relations in the Age of Robotics, Datamation, March.

[75]   Buchanan B.G. (1981): Research on Expert Systems, Stanford University, Report STAN-CS-81-837.

[76]   Buchanan B.G., Shortliffe E.H.(eds.) (1984): Rule-Based Expert Systems - The MYCIN Experiments of the Stanford Programming Project, Addison-Wesley, Reading, Mass..

[77]   Burbridge M. (1983): Roundup in Rio, Datamation, 29(3)189-198.

[78]   Burstein D. (1983): Counting to a Billion, Datamation, 183-187, March.

[79]   Bylinsky G. (1981): A New Industrial Revolution is on the Way, Fortune, October 5.

[80]   Campbell J.A. (1982): The Expert Computer and Professional Negligence: Who is Liable?, Univ.Exeter, Dept.Computer Science, Report R.109.

[81]   Card S.K., Newell A. (1981): Are we Ready for Cognitive Engineering?, Proceedings of the Third Annual Conference of the Cognitive Science Society, Berkeley.

[82]   Cashman M. (1983): Computing at Tokyo U., Datamation, 203-205, March.

[83]   Caulkin S. (1982): The Apple Millionaires, Management Today, pp.42ff., October.

[84]   Cerri S.A. (1985): The Intersection of AI and Education: Investing in Thought-Intensive Endeavours, in Trappl R.(ed.), The Impacts of Artificial Intelligence, North-Holland, Amsterdam.

[85]   Cetron M., O'Toole T. (1982): Careers with a Future: Where the Jobs will be in the 1990s, The Futurist, 16(3)11-19.

[86] Chandrasekaran B. (1983): Towards a Taxonomy of Problem Solving Types, AI Magazine, 4(1)9-17.

[87] Charniak E. (1983): Artificial Intelligence - An Introduction, Tutorial No.1, AAAI-83, Washington, D.C..

[88] Christie B. (1982): Psychology at the User-System Interface, International Online Information Meeting 6, Institut f.Zukunftsforschung, Berlin, pp.39-47.

[89] Chytil M.K. (1980): Towards Cognitive Science and Cognitive Engineering, Teorie rozvoje vedy, 4(4)1-17.

[90] Clancey W.J. (1981): Methodology for Building an Intelligent Tutoring System, Stanford University, Report STAN-CS-81-894.

[91] Clancey W.J. (1983): The Epistemology of a Rule-Based Expert System - A Framework for Explanation, Artificial Intelligence, 20(3)215-251.

[92] Clancey W.J., Shortliffe E.H. (1984): Readings in Medical Artificial Intelligence, Addison-Wesley, Reading, MA.

[93] Cleveland H. (1982): Information as a Resource, The Futurist, 16(6)34-39.

[94] Clocksin W.F., Mellish C.S. (1981): Programming in Prolog, Springer, Berlin. ·

[95] Coelho H. (1983): PROLOG for Databases, Computers and Artificial Intelligence 2(1)35-46.

[96] Cohen P.R., Feigenbaum E.A.(eds.) (1982): The Handbook of Artificial Intelligence, Vol.3, William Kaufmann, Los Altos, Calif..

[97] Cohen P.R., Grinberg M.R. (1983): A Theory of Heuristic Reasoning about Uncertainty, AI-Magazine, 4(2)17-24.

[98] Cohen R.S. (1982): Science and Technology in Global Perspective, Int. Soc. Sci.J., 34(1)61-70.

[99] Colby K.M., Gilbert J.P. (1964): Programming a Computer Model of Neurosis, J.Math.Psych., 1,405-417.

[100] Colby K.M. (1981): Modeling a Paranoid Mind, BBS 4(4)1981, pp515-560.

[101] Coles L.S. (1977): The Application of Artificial Intelligence to Medicine, Futures.

[102] Collingridge D. (1980): The Social Control of Technology, The Open University Press, Milton Keynes.

[103] Colmerauer A. (1983): Prolog in Ten Figures, in Proceedings of the 8th International Joint Conference on Artificial Intelligence, Karlsruhe, West Germany.

[104] Connor S. (1984): The Invisible Border Guard, New Scientist, January 5, 9-14.

[105] Cook S., Hafner C.D., McCarty L.T., Meldman J.A., Peterson M., Sprowl J.R., Sridharan N.S., Waterman D.A. (1981): The Applications of Artificial Intelligence to Law: A Survey of Six Current Projects, AFIPS National Computer Conference.

[106] Corenson B. (1984): The Outlook for Integrated Software, Datamation, pp.225/9-14, March.

[107] Cornish B.M. (1981): The Smart Machines: Implications for Society of Tomorrow, The Futurist, 15(4)5-13.

[108] Cornish E. (1981): The Coming of an Information Society, The Futurist, 15(2)14-21.

[109] Corrigan R. (1983): To Trust or Antitrust, Technology Review, pp.49-68, July.

[110] Crepy H., Denoix B., Destombes F., Rouquie G., Tubach J-P. (1983): Speech Processing on a Personal Computer to Help Deaf Children, Information Processing, 669-671.

[111] Cross M. (1983): Human Cost of Cheap Electronics, New Scientist, 250, October 27.

[112] Crumley D.V. (1982): Concepts for Army Use of Robotic-Artificial Intelligence in the 21st Century, Army War Coll. Strategic Studies Inst., Carlisle Barracks, PA, USA.

[113] Davey P.G. (1983): Robots with Common Sense? The Research Scene To-Day, Journal of the Royal Society of Arts, 5327, 671-685.

[114] Davies D.J.M. (1982): Emotional Robots will be Almost Human, in Proceedings of the Fourth Biennial Conference of the Candian Society for Computational Studies of Intelligence (CSCSI/SCEIO), Saskatoon, Sakatchewan.

[115] Davis L.S., Rosenfeld A. (1981): Cooperating Processes for Low-Level Vision: A Survey, Artif. Intell., 17, 245-263.

[116] Davis M.W. (1984): Anatomy of Decision Support, Datamation, pp.201-208, June 15.

[117] Davis P.K., Stan P.J.E., Bennett B.W. (1983): Automated War Gaming as a Technique for Exploring Strategic Command and Control Issues, RAND Corp., Santa Monica, CA.

[118] Davis P.K., Winnefeld J.A. (1983): Rand Strategy Assessment Center: An Overview and Interim Conclusions about Utility and Development Options, RAND Corp.,Santa Monica, CA.

[119] Davis R. (1978): Knowledge Acquisition in Rule-Based Systems - Knowledge about Representations as a Basis for System Construction and Maintenance, in Waterman D.A., Hayes-Roth F.(eds.), Pattern-Directed Inference Systems, Academic Press, New York.

[120] Davis R. (1982): Expert Systems: Where Are We? And Where Do We Go From Here?, AI Magazine, pp.3-22, Spring.

[121] Davis R., Lenat D.B. (1982): Knowledge-Based Systems in Artificial Intelligence, McGraw-Hill, New York.

[122] Davis R. (1984): Amplifying Expertise with Expert Systems, in Winston P.H., Prendergast K.A.(eds.), The AI Business, MIT Press, Cambridge, MA.

[123] Dell'Mour R., Fleissner P., Sint P.P. (1982): Microelectronics and Unemployment in Austria, IIASA, Laxenburg, CP-82-41.

[124] Dell'mour R., Fleissner P., Sint P.P.: An Input-Output Approach for the Assessment of Technological Change, Oesterr.Akademie der Wissenschaften, Inst.f.Sozio-Oekonomische Entwicklungsforschung, Wien.

[125] Dennet D.C. (1978): Intentional Systems, in Brainstorms, Bradford Books, Montgomery.

[126] Dennett D.C. (1982): Recent Work in Philosophy of Interest to AI, Artificial Intelligence, 19(1)3-5.

[127] Dennett D.C. (1983): Formulating Human Purposes, Information Processing,253-258.

[128] Dennett D.C. (1984): The Role of the Computer Metaphor in Under-
standing the Mind, in Computer Culture (forthcoming), New York
Academy of Sciences, N.Y..

[129] Dertouzos M.L., Moses J. (eds.) (1979): The Computer Age: A Twenty-
Year View, MIT-Press, Cambridge, Ma.

[130] DeSanctis G. (1983): A Telecommuting Primer, Datamation, pp.214-220,
October.

[131] Descotte Y., Latombe J.-C. (1981): GARI: A Problem Solver That Plans
How to Machine Mechanical Parts, in Proceedings of the 7th Interna-
tional Joint Conference on Artificial Intelligence, Univ.British Colum-
bia, Vancouver, Canada.

[132] Dietterich T.G., Michalski R.S. (1981): Inductive Learning of Struc-
tural Descriptions: Evaluation Criteria and Comparative Review of
Selected Methods, Artif. Intell., 16(3)257-294.

[133] Doemoelki B., Szeredi P. (1983): Prolog in Practice, Information Pro-
cessing, 627-636.

[134] Doyle J. (1983): Methodological Simplicity in Expert System Construc-
tion: The Case of Judgements and Reasoned Assumptions, AI-Magazine,
4(2)39-43.

[135] Doyle J. (1983): A Society of Mind - Multiple Perspectives, Reasoned
Assumptions, and Virtual Copies, in Proceedings of the 8th Interna-
tional Joint Conference on Artificial Intelligence, Karlsruhe, West
Germany, pp.309-314.

[136] Dray J., Menosky J.A. (1983): Is Adapting Computers to the Various
Cultures of the World an Idea whose Time has Come or a Pipe Dream
with Commercial Overtones?, Technology Review, pp.12-16, May/June.

[137] Dresher B.E., Hornstein N. (1976): On Some Supposed Contributions of
Artificial Intelligence to the Scientific Study of Language, Cognition,
4,321-398.

[138] Dresher B.E. (1977): Reply to Schank and Wilensky, Cognition 5, 147-
149.

[139] Dreyfus H.L. (1972): What Computers Can't Do. A Critique of Artificial
Reason, Harper & Row, New York.

[140] Duda R.O., Shortliffe E.H. (1983): Expert Systems Research, Science,
220(4594)261-268.

[141] Dyer M.G. (1982): In-Depth Understanding of Integrated Processing
for Narrative Comprehension, RR #219, Yale Univ., Dept. of
Comp.Sci..

[142] Elmer-DeWitt P. (1983): Finishing First with the Fifth, Time, p.53,
August 1.

[143] Elmer-DeWitt P. (1984): Cracking Down - Hackers Face Tough New
Laws, Time, May 14.

[144] Engelberger J.F. (1980): Robotics in Practice, Kogan Page, London.

[145] Enyart J.N. (1983): Artificial Intelligence Applied to the Command,
Control, Communications, and Intelligence of the U.S.Central Command,
Army War Coll., Carlisle Barracks, PA.

[146] Epstein S. (1983): Challenges - Expert Systems, Sigart Newsletter,
83,16-18.

[147] Erman L., London P., Fickas St. (1981): The Design and an Example Use of Hearsay-III, in Proceedings of the 7th International Joint Conference on Artificial Intelligence, Univ.British Columbia, Vancouver, Canada.

[148] Ernest M.L. (1982): The Mechanization of Commerce, Scientific American,September,133-145.

[149] Ernst G.W., Banerji R.B. (1983): On the Relationship between Strong and Weak Problem Solvers, AI-Magazine, 4(2)25-29.

[150] Eustace P. (1983): The Great Brain Robbery, The Engineer, pp.16-17, July 21.

[151] Faflick P. (1982): Brutal Game of Survival, Time, August 16.

[152] Fahlman S.E. (1980): Design Sketch for a Million-Element NETL Machine, in Proceedings of the First Annual National Conference on Artificial Intelligence, AAAI, Stanford University.

[153] Fahlman S.E. (1980): Computing Facilities for AI: A Survey of Present and Near-Future Options, AI Magazine, 2(1)16-23.

[154] Fahlman S.E., Steele G.L. (1982): Tutorial on AI Programming Technology: Languages and Machines, AAAI Conference Tutorial Program.

[155] Feigenbaum E.A. (1977): The Art of Artificial Intelligence: Themes and Case Studies of Knowledge Engineering, in Michie D.(ed.), Expert Systems in the Micro Electronic Age, Edinburgh Univ.Press, Edinburgh (also Memo HPP-77-25, Stanford University).

[156] Feigenbaum E.A., McCorduck P. (1983): The Fifth Generation. Artificial Intelligence and Japan's Computer Challenge to the World, Addison-Wesley, Reading, MA.

[157] Feingold S.N. (1984): Emerging Careers, The Futurist, pp.9-16, Feb..

[158] Fergusen G.R. (1983): Aircraft Maintenance Expert Systems, Air Force Inst.of Techn., Wright-Patterson APB, OH, School of Engineering, USA.

[159] Ferguson R. (1981): PROLOG: A Step Toward the Ultimate Computer Language, Byte, 6(11)384-399.

[160] Finin T.W. (1983): Providing Help and Advice in Task Oriented Systems, in Proceedings of the 8th International Joint Conference on Artificial Intelligence, Karlsruhe, West Germany.

[161] Firschein O., Fischler M.A., Coles L.S., Tenenbaum J.M. (1973): Forecasting and Assessing the Impact of Artificial Intelligence on Society, Third International Joint Conference on Artificial Intelligence.

[162] Fischbach F. (1983): Audio-Terminals in der Rehabilitation, Die Computer Zeitung, pp.16-17, August 3.

[163] Fleissner P. (1982): Makrooekonomische Aspekte der Mikroelektronik in Oesterreich – ein Simulationsmodell, Kybernetik und Systemtheorie, Erich Schmidt Verlag, pp.157-172.

[164] Foith J. (1981): Robotics Research: From Toy Worlds to Industrial Applications, in Siekmann J.H.(ed.): GWAI-81, German Workshop on Artificial Intelligence, Springer, Berlin.

[165] Foster E. (1984): Artificial Intelligence Faces a Crossroad, Mini-Micro Systems, pp.119-125, May.

[166] Fox M.S. (1981): Factory Modelling, Simulation, and Scheduling in the Intelligent Management System, in Proceedings of the 7th International Joint Conference on Artificial Intelligence, Univ.British Columbia, Vancouver, Canada.

[167] Fox M.S. (1982): Intelligent Management System: An Overview, Report CMU-RI-TR-81-4, Carnegie Mellon Univ., Pittsburgh.

[168] Freiherr G. (1980): The Seeds of Artificial Intelligence, SUMEX-AIM, National Institutes of Health, Bethseda, Maryland, Report No.80-2071.

[169] Froehlich L. (1981): Give Tchaikovsky the News, Datamation, 27(11)130-140.

[170] Froehlich L. (1983): Mighty Oaks Take Time, Datamation, pp.113-118, January.

[171] Frude N. (1983): The Intimate Machine, Century Publishing, London.

[172] Gallaire H. (1983): Panel on the Fifth Generation Project, in Proceedings of the 8th International Joint Conference on Artificial Intelligence, Karlsruhe, West Germany, pp.1193-1194.

[173] Garrison L. (1983): A Fool's Game?, Time, p.8, July 11.

[174] Gartner T. (1984): B of A is Offering to Pay Workers Who Resign Early, San Francisco Chronicle, 17 January.

[175] Gause C.D., Rogers G. (1983): Cybernetics and Artificial Intelligence, in Trappl R.(ed.), Cybernetics: Theory and Applications, Hemisphere, Washington, DC and Springer, Berlin.

[176] Gehrke M., Primio F.di (1983): IQLISP - Ein LISP fuer den IBM-PC, KI-Rundbrief 30, Juni.

[177] Gerring P.E., Shortliffe E.H., van Melle W. (1982): The Interviewer/Reasoner Model: An Approach to Improving System Responsiveness in Interactive AI-Systems, AI-Magazine, 3(4)24-27.

[178] Geschke M.J., Bullock R.A., Widmaier L.E. (1983): TACII An Expert Knowledge Based System for Tactical Decision Making, Naval Postgraduate School, Monterey, CA..

[179] Gevarter W.B. (1982): An Overview of Expert Systems, Nat. Bureau of Standards, U.S.Department of Commerce, Rep. NBSIR 82-2505.

[180] Giboney T.B. (1983): Myth of Omniscient Cybernetics, Army Command and General Staff Coll., Fort Leavenworth, KS.

[181] Gilchrist B., Dagli A., Shenkin A. (1983): The DP Population Boom, Datamation, pp.100-110, September.

[182] Gillard C., Smith J. (1983): Computer Crime: A Growing Threat, Byte Publications Inc., pp.398-424, October.

[183] Gini G., Gini M. (1973): Artificial Intelligence: Impact on Computer Assisted Medical Diagnosis, Istituto di Elettrotecnica del Politecnico di Milano.

[184] Gleick J. (1983): Artificial Intelligence: Exploring Labyrinth of the Mind, International Herald Tribune, p.6, September 1.

[185] Goerz G. (1983): A Programming Environment for TLC-LISP, KI-Rundbrief 30, Juni.

[186] Goldberg A. (1981): Introducing the Smalltalk-80 System, Byte Publications Inc., August.

[187] Goldstein I., Papert S. (1976): Artificial Intelligence, Language and the Study of Knowledge, MIT, AI-MEMO #337.

[188] Goldstein I., Papert S. (1977): Artificial Intelligence, Language, and the Study of Knowledge, Cognitive Science, 1(1)84-124.

[189] Goodman F.L. (1981): Computers and the Future of Literacy, AFIPS National Computer Conference.

[190] Gosnell K. (1983): Happy Days for Software Pirates, New Scientist, pp.376-377, May 12.
[191] Gould J.D., Conti J., Hovanyecz T. (1983): Composing Letters with a Simulated Listening Typewriter, Communications of the ACM,26(4)295-308.
[192] Graham R.L. (1984): The Legal Protection of Computer Software, Comm. ACM, 27(5)422-426.
[193] Grand S.L. (1982): Reply to Trappl R.: Computer Psychotherapy: Is it Feasible, Advisable?, Computing Reviews, 23(11)535-536.
[194] Greenwald J. (1984): A Remarkable Job Machine, Time, pp.36-38, June 25.
[195] Greiner R. (1980): RLL-1: A Representation Language Language, Report HPP-80-9, Dept. of Computer Science, Stanford Univ..
[196] Greiner R., Lenat D.B. (1980): A Representation Language Language, in Proceedings of the First Annual National Conference on Artificial Intelligence, AAAI, Stanford University.
[197] Grieves R.T. (1983): Short Circuiting Reference Books, Time, p.45, June 13.
[198] Grosz B.J. (1982): Team: A Transportable Natural-Language System, SRI, TN-263, Stanford, CA.
[199] Groveston P.K. (1982): Overview of Artificial Intelligence-Based Teaching Aids, WP, MITRE Corp., McLean, VA.
[200] Guida G., Tasso C. (1982): Dialogue with Data Bases: An Effective Natural Language Interface, Informatika-82, Ljubljana.
[201] Guida G., Tasso C. (1982): Forward and Backward Reasoning in Automatic Abstracting, COLING-82, Prague.
[202] Gunsch G.H., Hebert B.V. (1983): Proposed Military Planning Task Simulator Using ROSS Language, Report No.AFIT/GE/EE/83D-24, Air Force Inst. of Tech., Wright-Patterson AFB,USA.
[203] Gupta G.K. (1983): Computing in the Developing Countries of ASIA, Information Processing, 439-446.
[204] Gurd J., Watson J. (1983): Preliminary Evaluation of a Prototype Dataflow Computer, Information Processing, 545-551.
[205] Haase-Rapoport M.G., Pospelov D.A., Semenova Y.T. (1984): Generation of Fairy Tales, Computers and Artif. Intell., 3(4)281-304.
[206] Hackman S. (1983): Automatic Experts for Sale, Technology Review, p.84, May/June.
[207] Hald A.P. (1981): Towards the Information-Rich Society, The Futurist, 15(4)20-24.
[208] Hall D. (1983): My Life with Robots, New Scientist, pp.468-470, February 17.
[209] Hammond J.H. (1983): The Challenge of Change: Information Technology and the Community, Information Processing, pp.781-785.
[210] Hamrin R.D. (1981): The Information Economy; Exploiting an Infinite Resource, The Futurist, 15(4)25-30.
[211] Harman W.W. (1981): Work, in A.Villoldo, K.Dychtwald (eds.), Millenium, Jeremy Tarchure.
[212] Harris L.R. (1983): Fifth Generation Foundations, Datamation, pp.148-156, July.

[213] Hart P.E. (1982): Direction for AI in the Eighties, Sigart Newsletter, 79,11-16.

[214] Hatvany J., Bjorke O., Merchant M.E., Semenkov O.I., Yoshikawa H. (1983): Advanced Manufacturing Systems in Modern Society, in Ellis T., Semenkov O.I.(eds.): Advances in CAD/CAM, North-Holland, Amsterdam.

[215] Haugeland J. (ed.) (1981): Mind Design, The MIT Press, Cambridge, Ma..

[216] Havel I.M. (1985): Artificial Intelligence: A Lesson in Human Self-Understanding, in Trappl R.(ed.), The Impacts of Artificial Intelligence, North-Holland, Amsterdam.

[217] Hawkride D., Helm C. (1983): A Major Challenge to Schools, New Scientist, p.597, March 3.

[218] Hayes P.J. (1979): The Naive Physics Manifesto, in Michie D.(ed.), Expert Systems in the Micro Electronic Age, Edinburgh Univ.Press, Edinburgh.

[219] Hayes-Roth F., Waterman D., Lenat D.(eds.) (1982): Building Expert Systems, Addison-Wesley, Reading, Mass..

[220] Hazony Y. (1983): Toward the Factory of the Future, Perspectives in Computing, 3(4)4-11.

[221] Healey M. (1983): Junking the Mainframe, Datamation, pp.120-136, August.

[222] Herman R. (1983): How Research Failings Cripple Industry, New Scientist, p.4, July 7.

[223] Hertz D.R. (1983): Artificial Intelligence and the Business Manager, Computerworld, 17(43)22-26.

[224] Higbie L. (1983): A Vector Processing Tutorial, Datamation, pp.180-200, August.

[225] Hiltz S.R. (1983): Viewing Computing Systems within a Social Context, Information Processing,647-655.

[226] Hinton G.E. (1981): Shape Representation in Parallel Systems, in Proceedings of the 7th International Joint Conference on Artificial Intelligence, Univ.British Columbia, Vancouver, Canada.

[227] Hinton G.E., Anderson A.(eds.) (1981): Parallel Models of Associative Memory, Erlbaum, Hillsdale, NJ.

[228] Hinton G.E. (1983): Optimal Perceptual Inference, Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Washington DC,June..

[229] Hoare C.A.R. (1981): The Emperor's Old Clothes, Comm.ACM, 24(2)75-83.

[230] Hofstadter D.R. (1979): Goedel, Escher, Bach: an Eternal Golden Braid, Vintage Books, New York.

[231] Hofstadter D.R., Dennett D.C. (1981): The Mind's I - Fantasies and Reflections on Self and Soul, Penguin Books, Harmondsworth, England.

[232] Hogan D.I. (1983): Survivability Design Guidelines for Fly-by-Wire Flight Control Systems Development, Naval Postgraduate School, Monterey, CA.

[233] Horn W., Ammer K. (1984): Causal and Anatomical Reasononing in a Two-Level Medical Expert System, in O'Shea T.(ed.), ECAI-84: Advances in Artificial Intelligence, Elsevier, Amsterdam.

[234] Hudson C.A. (1982): Computers in Manufacturing, Science, 215,818-825.

[235] Hunt H.A., Hunt T.L. (1983): Human Resource Implications of Robotics, Upjohn Institute for Employment Research, Kalamazoo, Michigan.

[236] Jackson P.C. (1975): Introduction to Artificial Intelligence, Petrocelli, Princeton, NJ.

[237] Jahoda M. (1982): Employment and Unemployment, Cambridge University Press.

[238] James M. (1984): Artificial Intelligence in BASIC, Butterworth, Borough Green, Sevenoaks, Kent, England.

[239] Janda A.(ed.) (1983): Human Factors in Computing Systems, SIGCHI Bulletin, Special Issue.

[240] Jansen T.B. (1983): Information Technology, The Need for Social Experiments, Technological Forecasting and Social Change, 23,325-352.

[241] Jenkin P. (1983): The Unemployed Cannot Blame Automation, New Scientist, pp.526-527, February 24.

[242] Jenkins C., Sherman B. (1981): The Leisure Shock, Eyre Methuen, London.

[243] Jenny C.J.: Requirements on Expert Systems as Seen by an Insurance Company, "Zurich" Insurance Company, Zurich, Switzerland.

[244] Johnson J. (1982): Four Expert Opinions, Datamation, pp.143-150, December.

[245] Johnson J. (1983): Reading Users' Lips, Datamation, pp.92-97, August.

[246] Johnson J. (1983): Can They Do It?, Datamation, pp.161-170, July.

[247] Johnson J. (1984): America Answers Back, Datamation, pp.40–57, May.

[248] Johnson J. (1984): Easy Does it, Datamation, pp.48-60, June 15.

[249] Johnson-Laird P.N. (1981): Cognition, Computers and Mental Models, Cognition 10, 139-143.

[250] Johnston B. (1983): Japan Builds a Pillar of the Fifth Generation, New Scientist, p.193, July 21.

[251] Jones C.B. (1983): Specification and Design of (Parallel) Programs, Information Processing, pp.321-332.

[252] Jones H.A. (1982): Cybernetic Design Features for Computerised Business Information Systems, in Trappl R.(ed.): Cybernetics and Systems Research, pp.815-821, North-Holland, Amsterdam.

[253] Jones K.P. (1982): Searching Techniques: Time to Replace Intelligent Operators by Intelligent Systems, International Online Information Meeting, Institut f.Zukunftsforschung, Berlin, pp.49-59.

[254] Jones R. (1982): MAPE - Micros and Primary Education, Computer Education 42.

[255] Joyce C., Wingerson L. (1983): Can we Adjust to Computer Culture?, New Scientist, pp.72-73, April 14.

[256] Joyce J.D., Warn D.R. (1983): Human Factors Aspects of a Modern Data Base System, Inf. Management, 6,27-36.

[257] Kahn K.M. (1984): Partial Evaluation, Programming Methodology, and Artificial Intelligence, The AI Magazine, pp.53-57, Spring.

[258] Kanade T. (1981): Recovery of the Three-Dimensional Shape of an Object from a Single View, Artificial Intelligence, 17,409-460.

[259] Kaplan J. (1982): Natural Language in the DP World, Datamation, pp.114-120, August.

[260] Kaplan R. (1981): Commercial Sector: AI in the Pharmaceutical Industry, SIGBIO Newsletter, 5(4)4-9.

[261] Karlgren H., Walker D.E. (1981): The Economic Impact of Research on Natural Language, Int.Forum of Inf.Doc., 6(81)4.

[262] Kato I. (1983): Coming Robot New Era (A Viewpoint), Robotica, 1(1)9-14.

[263] Kehler T.P., Friedland P., Pople H., Reboh R., Rosenberg S. (1983): Industrial Strength Knowledge Bases: Issues and Experience, in Proceedings of the 8th International Joint Conference on Artificial Intelligence, Karlsruhe, West Germany.

[264] Keller L.S. (1982): The Fifth Generation: A Bleak Outlook, Datamation, pp.24-26, April.

[265] Kelso L.O., Adler M.J. (1958): The Capitalist Manifesto, Random House, New York.

[266] Keynes J.M. (1933): Economic Possibilities for our Grandchildren, in Essays in Persuasion, Macmillan, London.

[267] Kiechel W. (1984): Beat the Clock, Fortune, June 25, 94-95.

[268] Kildall G. (1983): Digital Research's DR Logo, Byte, pp.208-218, June.

[269] King M. (1982): EUROTRA: An Attempt to Achieve Multilingual MT, in Lawson V.(ed.): Practical Experience of Machine Translation, North Holland, Amsterdam.

[270] King M., Perschke S. (1982): Eurotra and its Objectives, Multilingua, 1-1,27-32.

[271] Kingsburg J.A., Pool D.A., Ghosh S.K., Malthouse N.S., Rouillier G.J. (1983): Automated En Route Air Traffic Control Algorithmic Specifications. Vol.1: Trajectory Estimation, Federal Aviation Administration, Report DOT/FAA/ES-83/4, Washington, DC.

[272] Kirchner E. (1982): At the Mercy of Machines, Datamation, pp.252-261, September.

[273] Kirkley J.L. (1982): The Leap Forward, Datamation, p.31, October.

[274] Kiss G. (1982): AI Programming Facilities on Microcomputers, European Conference on Artificial Intelligence, Discussion Papers, Orsay.

[275] Kling R. (1980): Social Analyses of Computing: Theoretical Perspectives in Recent Empirical Research, Computing Surveys, 12(1)61-110.

[276] Klingler D.E., Miller D.A., Johnson J.H., Williams T.A. (1977): Process Evaluation of an On-Line Computer-Assisted Unit for Intake Assessment of Mental Health Patients, Behavior Res.Meth.& Instr., 9(2)110-116.

[277] Klix F. (1983): Remarks Regarding Some Aspects of AI Research, in Trappl R., Vasko T.(eds.), Papers for Task Force Meeting on Future and Impacts of Artificial Intelligence, IIASA, Laxenburg, Austria, CP-83-36.

[278] Knuth D.E. (1981): The Art of Computer Programming, Vol.2: Seminumerical Algorithms, Addison-Wesley, Reading, MA.

[279] Kobsa A., Trost H., Trappl R. (1982): Ist benutzerangepasstes Dialogverhalten auch ohne Dialogpartnermodell moeglich?, Bericht 82-06, Inst.f.Med.Kybernetik, Univ.Wien.

[280] Kobsa A. (1984): Three Steps in Constructing Mutual Belief Models from User Assertions, Proc. of the 6th European Conference on Artificial Intelligence, Pisa.

[281] Kolata G. (1984): Order Out of Chaos in Computers, Science, 223,917-919, March 2.

[282] Kopec D., Michie D. (1983): Mismatch between Machine Representations and Human Concepts: Dangers and Remedies, Report EUR-8426-EN, Commission of the European Communities, Luxembourg.

[283] Kornbluh M. (1982): The Electronic Office - How it Will Change the Way You Work, The Futurist, 16(3)37-42.

[284] Kosy D.W., Dhar V. (1983): Knowledge-Based Support Systems for Long Range Planning, Report CMU-RI-TR-83-21, Carnegie-Mellon Univ., Pittsburgh, USA.

[285] Kotelnikov V.A. (1981): Comprehensive Programme of Scientific and Technological Progress of the USSR for 20 Years, Seminar on Scientific Forecasting and Human Needs, Tbilisi, December 6-10.

[286] Kowalski R. (1983): Panel on the Fifth Generation Project, in Proceedings of the 8th International Joint Conference on Artificial Intelligence, Karlsruhe, West Germany, pp.1192-1193.

[287] Kowalski R. (1983): Logic Programming, Information Procesjing, pp.133-145.

[288] Kraft A. (1984): XCON: An Expert Configeration System at Digital Equipment Corporation, in Winston P.H., Prendergast K.A.(eds.), The AI Business, MIT Press, Cambridge, MA.

[289] Krutch J. (1981): Experiments in Artificial Intelligence for Small Computers, Howard W. Sams, Indianapolis, USA.

[290] Kulikowski C.A. (1979): Computer-Based Medical Decision Making, Comp. in Ophthalmology, pp.86-92, April.

[291] Kulikowski C.A. (1980): Artificial Intelligence Methods and Systems for Medical Consultation, IEEE Trans. Pattern Anal. Machine Intell., PAMI 2(5)464-476.

[292] Lakoff G. (1978): Some Remarks on AI and Linguistics, Cognitive Science, 2(3)267-278.

[293] Lamb J. (1983): Computer Firms Struck Dump by Voice Research, New Scientist, p.8, July 7.

[294] Lamb J. (1983): Crime by Computer, New Scientist, pp.11-19, July 7.

[295] Lamb J. (1983): Banker's Padlock Fights Program Pirates, New Scientist, p.19, October 6.

[296] Lamb J. (1983): Britain Gambles on Super Chips, New Scientist, p.418, November 10.

[297] Lamb J. (1984): West Counts the Cost of Computer Ban, New Scientist, pp.12-13, February 23.

[298] Lamond F. (1982): The Founding of France Inc., Datamation, 25-26,164, December.

[299] Lamond F.E. (1983): Microprocessors as Agents of De-Industrialisation, Information Processing,775-780.

[300] Lamson R.W. (1982): Mapping Policy Dialogues: A Role for Computer Graphics, Proc.of the Third Annual Conference of the National Computer Graphics Association, Fairfax, Virginia, Vol.1.

[301] Laver M. (1980): Computers and Social Change, Cambridge University Press, New York.

[302] Leal A. (1982): Evaluating the Effectiveness of Military Decision Support Systems. Theoretical Foundations, Expert System Design, and Experimental Plan, Integrated Sciences Corp., Santa Monica, CA.

[303] Lee R.M. (1983): Data and Language in Organizations: Epistemological Aspects of Management Support Systems, IIASA, Laxenburg, WP-83-46.

[304] Lee R.M. (1983): On Information System Semantics: Expert versus Decision Support Systems, IIASA, Laxenburg, Austria.

[305] Lehman-Wilzig S.N. (1981): Frankenstein Unbound, Towards a Legal Definition of Artificial Intelligence, Futures, 6(81)442-457.

[306] Lehnert W.G., Wilks Y. (1979): A Critical Perspective on KRL, Cognitive Science, 3(1)1-28.

[307] Lehnert W.G., Ringle M.H.(eds.) (1982): Strategies for Natural Language Processing, Erlbaum, Hillsdale, NJ.

[308] Leinweber D. (1981): Future Computing, Report P-6680, Rand Corp..

[309] Lenat D.B. (1978): The Ubiquity of Discovery, Artificial Intelligence 9,257-285.

[310] Lenat D.B. (1982): The Nature of Heuristics, Artificial Intelligence, 19(2)189-249.

[311] Lenat D.B., Clarkson A., Kiremidjian G. (1983): An Expert System for Indications and Warning Analysis, in Proceedings of the 8th International Joint Conference on Artificial Intelligence, Karlsruhe, West Germany.

[312] Leonard M. (1983): Now Automation Moves in on Custom Woodworking, Business Week, November 7.

[313] Leontief W.W. (1966): Input-Output Economics, Oxford University Press, New York.

[314] Leontief W.W. (1982): The Distribution of Work and Income, Scientific American, 192, September.

[315] Leontief W.W. (1983): The New New Age That's Coming is Already Here, Bottom Line/Personal,4(8)1-2.

[316] Leontief W.W., Duchin F. (1983): The Impacts of Automation on Employment, 1963-2000, Technical Report,Institute for Economic Analysis, New York University.

[317] Levin H.M., Rumberger R.W. (1983): The Educational Implications of High Technology, IFG Institute for Research on Educational Finance and Governance,Report 83-A4.

[318] Licker P. (1983): The Japanese Approach: A Better Way to Manage Programmers?, Comm.ACM, 26(9)631-636.

[319] Lindamood G.E. (1983): Japanese Computer Project, Science, 221,1008.

[320] Lindamood G.E. (1984): Fifth Generation Episto-Encabulators, Datamation, pp.193-196, June.

[321] Lindsay R.K., Buchanan B.G., Feigenbaum E.A., Lederberg J. (1980): Applications of Artificial Intelligence for Organic Chemistry: The DENDRAL Project, McGraw Hill, New York.

[322] Ljungberg S. (1982): Intelligence Service - A Tool for Decision Makers, FID-Congress 41: The Challenge of Information Technology, Den Haag.

[323] Lloyd A. (1983): Europe Examines Electronic Democracy, New Scientist, pp.634-645, June 2.

[324] Lloyd A. (1983): Computer Spies Give Thanks to Personal Computers, New Scientist, p.572, March 3.

[325] Lloyd A. (1983): Eloquent Machines are 10 Years Away, New Scientist, p.499, November 17.

[326] Longuet-Higgins H.C. (1981): Artificial Intelligence - A New Theoretical Psychology?, Cognition, 10,197-200.

[327] Longworth N. (1981): We're Moving into the Information Society. What Shall We Teach the Children?, Computer Education, pp.17-19, June.

[328] Lucas R.W., Mullin P.J., Luna C.B.X., McInroy D.C. (1977): Psychiatrists and a Computer as Interrogators of Patients with Alcohol-Related Illnesses: A Comparison, Brit.J.Psychiatry, 131,160-167.

[329] Luckhardt H.-D. (1982): SUSY: Capabilities and Range of Application, Multilingua, 1-4,213-220.

[330] Lucky R. (1983): The Social Impact of the Computer, in H.R.Pagels (ed.): Computer Culture: The Scientific, Intellectual and Social Impact of the Computer, New York Academy of Sciences.

[331] Lukaszewicz W.: General Approach to Nonmonotonic Logics, Institute of Informatics, Warszaw.

[332] Lytinen S., Schank R. (1982): Representation and Translation, Yale Univ., Research Report No.234.

[333] Macaulay L.A. (1981): Computer Education for Everyman, Computer Education, pp.23-25, June.

[334] Mahn T.G. (1983): Sparking Urban Growth, Datamation, pp.223-240, October.

[335] Malik R. (1983): Japan's Fifth Generation Computer Project, Futures, pp.205-210, June.

[336] Mankin D., Bikson T.K., Gutek B. (1982): The Office of the Future - Prison or Paradise?, The Futurist, 16(3)33-36.

[337] Manthey M.J., Moret B.M. (1983): The Computational Metaphor and Quantum Physics, Comm. ACM, 26(2)137-145.

[338] Mariani J., Prouts B., Gauvain J.L., Gangolf J.J. (1983): A Man-Machine Speech Communication System including Word-Based Recognition and Text-to-Speech Synthesis, in Mason R. (ed.): Information Processing 83, North Holland, Amsterdam.

[339] Marr D. (1982): Vision, Freeman, San Francisco, CA.

[340] Marsh P. (1983): A Brighter Future for the Computer Kids, New Scientist, pp.462-466, August 18.

[341] Marsh P. (1983): Computer Kids Head for the Big Time, New Scientist, 179, July 21.

[342] Marsh P. (1983): Intelligence to Sell Micros, New Scientist, 267, October 27.

[343] Marsh P., Lloyd A. (1983): Europe Pools its Resources on R&D, New Scientist, pp.850-851, June 23.

[344] Mason (1983): Alvey Lines up Computer Research Projects, New Scientist, p.401, November 10.

[345] Mason R., Jennings L. (1982): The Computer Home: Will Tomorrow's Housing Come Alive?, The Futurist, 16(1)35-47.

[346] Mattill J. (1983): A Face-Off on Computers, Technology Review, pp.78-79, May/June.

[347] Maurer H. (1983): On Some Beneficial Implications of Videotex on Society, Institut fuer Informationsverarbeitung, Universitaet Graz, Bericht F-116.

[348] McCarthy J., Hayes P.J. (1969): Some Philosophical Problems from the Standpoint of Artificial Intelligence, Machine Intelligence 4, Edinburgh Univ.Press.

[349] McCarthy J. (1980): Circumscription - A Form of Non-Monotonic Reasoning, Artif. Intell., 13(1-2)27.

[350] McCarthy J. (1983): Interview with John McCarthy, Omni, April,100-122.

[351] McCartney L. (1983): Academia Inc., Datamation, 29(3)116-126, March.

[352] McCorduck P. (1979): Machines Who Think, Freeman, San Francisco, CA.

[353] McCorduck P. (1983): Humans That Think (A Future Trialogue), The AI Magazine, p.35, Fall.

[354] McCorduck P. (1983): An Introduction to the Fifth Generation, Commun. ACM, 26(9)629-630.

[355] McCusker T. (1980): Examining the Office of the Future, Datamation, pp.120-121, February.

[356] McCusker T. (1980): Office is Where it's Happening, Datamation, 26(2)65-68.

[357] McDermott D. (1980): The Prolog Phenomenon, Sigart Newsletter, 72,16-20, July.

[358] McDermott D., Doyle J. (1980): Non-Monotonic Logic I, Artificial Intelligence, 13(1-2)41-72.

[359] McDermott D. (1982): Non-Monotonic Logic II: Non-Monotonic Modal Theories, J.Ass.Comp.Mach., 29(1)33-57.

[360] Meleka J. (1983): Will the Robot Change the World or will the World Change the Robot? A User Point of View, Journal of the Royal Society of Arts, 5327,657-670.

[361] Meltzer B. (1982): The Unconscious in Psychology and Computational Modelling, AISB Quarterly, Autumn.

[362] Merchant M.E. (1983): Data Driven Automation I - Overview: Production: A Dynamic Challenge, IEEE Spectrum, pp.36-39, May.

[363] Michaelis P.R., Miller M.L., Hendler J.A. (1982): Artificial Intelligence and Human Factors Engineering: A Necessary Synergism in the Interface of the Future, in Badre A., Shneiderman B.(eds.): Directions in Human-Computer Interaction, Ablex, Norwood, NJ.

[364] Michaelsen R., Michie D. (1983): Expert Systems in Business, Datamation, pp.240-246, November.

[365] Michalski R.S., Chilausky R.L. (1980): Learning by Being Told and Learning from Examples, Int.J.Policy Analysis and Information Systems, 4,2.

[366] Michalski R.S., Carbonell J.G., Mitchell T.M.(eds.) (1982): Machine Learning: An Artificial Intelligence Approach, Tioga, Palo Alto, Calif..

[367] Michalski R.S.(ed.) (1983): Proceedings of the International Machine Learning Workshop, Dept. of Computer Science, Univ. of Illinois at Urbana-Champaign.

[368] Michie D. (1981): High-Road and Low-Road Programs, AI-Magazine, 3(1)21-22.

[369] Miller D.A., Johnson J.H., Klingler D.E., Williams T.A. (1977): Design for an On-Line Computerized System for MMPI Interpretation, Behavior Res.Meth.& Instr., 9(2)117-122.

[370] Miller G.A. (1981): Trends and Debates in Cognitive Psychology, Cognition, 10,215-225.

[371] Minsky M. (1975): A Framework for Representing Knowledge, in Winston P.H.(ed.), The Psychology of Computer Vision, McGraw-Hill, New York.

[372] Minsky M., Papert S. (1977): The Society of Mind, Unpublished manuscript.

[373] Minsky M. (1982): Why People Think Computers Can't, The AI Magazine, 3(4)3-15.

[374] Minsky M., Feigenbaum E.A. (1983): Look Ahead, AI Profs go Commercial, Datamation, p.13, July.

[375] Modesitt K.L. (1979): An Academic Meets Industry - Rethinking Computer-Based Education and Personalized Systems of Instruction, AFIPS National Computer Conference.

[376] Molitor G.T.T. (1981): The Information Society: The Path to Post-Industrial Growth, The Futurist, 15(2)23-30.

[377] Monod J. (1971): Chance and Necessity, Random House, New York.

[378] Moon D.A. (1981): Lisp Machine Choice Facilities, MIT, WP 208A.

[379] Moon D.A., Wechsler A.C. (1981): Operating the Lisp Machine, MIT, WP 209.

[380] Moore D.T. (1981): Will Robots Save Democracy?, The Futurist, 15(4)14-19.

[381] Moore R.C. (1983): Semantic Considerations on Nonmonotonic Logic, in Proceedings of the Eighth International Conference on Artificial Intelligence, Karlsruhe, FRG.

[382] Mori M. (1981): The Buddha in the Robot. A Robot Engineer's Thoughts on Science and Religion, Kosei Publishing Co., Tokyo.

[383] Morrison P.R. (1983): A Survey of Attitudes toward Computers, Comm. ACM, 12(26)1051-1057.

[384] Moto-Oka T.(ed.) (1981): Proceedings of International Conference on Fifth Generation Computer Systems, Japan Information Processing Development Center, Tokyo, October.

[385] Moto-Oka T., Fuchi K. (1983): The Architectures in the Fifth Generation Computers, Information Processing 83,589-602.

[386] Nagao M. (1983): Exploration of Knowledge and the Algorithm in Human Brain, Paper presented to the Fourth Convocation of Engineering Academies, Stockholm,May30-June 1..

[387] Nagao M. (1985): Social and Economic Impacts of Artificial Intelligence, in Trappl R.(ed.), The Impacts of Artificial Intelligence, North-Holland, Amsterdam.

[388] Narayanan A. (1981): Ascribing Mental Predicates to Computers, Research Report-102, Computer Science Department, University of Exeter.

[389] National Research Council (1982): Outlook for Science and Technology, Freeman, San Francisco.

[390] Naylor C. (1983): Build Your Own Expert System, John Wiley and Sons Ltd., Chichester.

[391] Negrotti M. (1983): How AI People Think, Report III/2203/En 83, Chair of Sociology of Knowledge, Univ.Genova, Italy.

[392] Neikirk B. (1982): Recovery Could Be a Jobless One, Chicago Tribune, 15 August.

[393] Neill D.S. (1981): Libraries in the Year 2010, The Futurist, 15(5)47-51.

[394] Nemes L., Hatvany J.: Design Criteria and Evaluation Methods for Man-Machine Communications on the Shop Floor, Computer and Automation Institute, Hungarian Academy of Science, Budapest.

[395] Newell A. (1981): The Knowledge Level, AI-Magazine, 2(2)1-20.

[396] Newell A. (1982): Intellectual Issues in the History of Artificial Intelligence, Carnegie-Mellon Univ., Comp.Sc.Dept., Report CMU-CS-82-142.

[397] Newell S.B., deGeus A.J., Rohrer R.A. (1983): Design Automation for Integrated Circuits, Science, 220(4596)465-471.

[398] Newman N. (1982): Britains Silicon Gulch, Management Today, December.

[399] Newman N. (1982): Europe's Catch-Up Complex, Management Today, pp.82ff., October.

[400] Nii H.P., Aiello N. (1979): AGE (Attempt to Generalize): A Knowledge-Based Program for Building Knowledge-Based Programs, in Proceedings of the 6th International Joint Conference on Artificial Intelligence, Tokyo, Japan, pp.645-655.

[401] Nii H.P., Feigenbaum E.A., Anton J.J., Rockmore A.J. (1982): Signal-to-Symbol Transformation: HASP/SIAP Case Study, AI-Magazine, 3(2)23-35.

[402] Nilsson N.J. (1971): Problem-Solving Methods in Artificial Intelligence, McGraw-Hill, New York.

[403] Nilsson N.J. (1980): Principles of Artificial Intelligence, Tioga, Palo Alto, CA.

[404] Nilsson N.J. (1981): Artificial Intelligence: Engineering, Science, or Slogan?, AI Magazine, 3(1)2-9.

[405] Nilsson N.J. (1984): Artificial Intelligence, Employment, and Income, SRI International, Menlo Park, CA.

[406] Nilsson N.J. (1985): Artificial Intelligence, Employment, and Income, in Trappl R.(ed.), The Impacts of Artificial Intelligence, North-Holland, Amsterdam.

[407] Nisenholtz M., Morphos E.M. (1983): Information Technology and New Forms of Participation in the Narrative, Information Processing, pp.727-731.

[408] Nitzan D., Barnard S., Bolles R., Cain R., Hannah M. (1983): Machine Intelligence Research Applied to Industrial Automation, SRI International, Menlo Park, CA.

[409] Norman D.A., Rumelhart D.E. (1981): The LNR Approach to Human Information Processing, Cognition, 10,235-240.

[410] Olson M.H., Lucas H.C. (1982): The Impact of Office Automation on the Organization: Some Implications for Research and Practice, Comm.ACM, 25(11)838-847.

[411] Olson M.H. (1983): Remote Office Work: Changing Work Patterns in Space and Time, Communications of the ACM, 26(3)182-187.

[412] O'Keefe R.A. (1982): PROLOG Compared with LISP, Research Paper No. 180, Dept. of AI, Univ. Edinburgh.

[413] O'Shea T. (1982): Intelligent Systems in Education, in Michie D.(ed.), Introductory Readings in Expert Systems, Gordon & Breach, New York.

[414] Pagels H.R.(ed.): Computer Culture: The Scientific, Intellectual and Social Impact of the Computer, New York Academy of Sciences, (in press).

[415] Pape A., Fjalestad J. (1979): Research on Social Aspects of Computerization and Democratization of Working Life, in P.A.Samet (ed.): Proceedings of the EURO IFIP, North-Holland Publishing Company, Amsterdam.

[416] Papert S. (1980): Mindstorms - Children, Computers, and Powerful Ideas, Harvester Press, Brighton.

[417] Pask G., Curran S. (1982): Micro Man: Computers and the Evolution of Consciousness, Macmillan, New York.

[418] Paul R.P. (1981): Robot Manipulators: Mathematics, Programming and Control, MIT Press, Cambridge, MA.

[419] Peltola H., Soederlund H., Ukkonen E. (1983): Algorithms for some String Matching Problems Arising in Molecular Genetics, Mason R. (ed.): Information Processing 83, North-Holland, Amsterdam.

[420] Pelton J.N.. (1982): Global Talk and the World of Telecomputerenergetics, in Didsbury H.F., Communications and the Future, World Future Society, Maryland.

[421] Peltu M. (1983): U.K. Eyes 5th Generation, Datamation, pp.67-72, July.

[422] Peltu M. (1984): The Human Factor, Datamation, 30(2)198/5-6.

[423] Petroski H. (1982): The Electronic Newspaper: An Easy Route to 1984?, The Futurist, 16(4)59-60.

[424] Pitt D., Booth S. (1983): Paradigms lost? Reflections on the Coming Organizational 'Revolution', Futures, pp.193-204, June.

[425] Popper K.R., Eccles J.C. (1977): The Self and its Brain, Springer, New York.

[426] Powers D.M. (1983): Neurolinguistics and Psycholinguistics as a Basis for Computer Acquisition of Natural Language, SIGART Newsletter, 84, April.

[427] Pribram K.H. (1981): The Brain, the Telephone, the Thermostat, the Computer, and the Hologram., Cognition and Brain Theory,4(2),105-122.

[428] Prigmore C.A. (1981): A Personal View of Computer Assisted Learning, Computer Education, pp.12-13, June.

[429] Raphael B. (1976): The Thinking Computer - Mind inside Matter, W.H.Freeman, San Francisco.

[430] Rathmill K. (1983): The Great Robotics Explosion, Management Today, July.

[431] Raulefs P. (1981): Expert Systems: State of the Art and Future Prospects, in Siekmann J.H.(ed.), GWAI-81, German Workshop on Artificial Intelligence, Springer, Berlin.

[432] Rauzino V. (1982): Conversation with an Intelligent Chaos, Datamation, pp.122-136, May.

[433] Reggia J.A. (1981): Computer-Assisted Medical Decision Making: A Critical Review, Ann. Biomed. Engn., 9(5-6)605-619.

[434] Reilly T.M. (1983): The Extended Classroom, Datamation, 29(4)199-200.

[435] Remer D. (1982): Legal Care for Your Software, Nolo Press, Berkeley, CA.

[436] Renfro W. (1982): Second Thoughts on Moving the Office Home, The Futurist, 16(3)43-38.

[437] Rich E. (1983): Artificial Intelligence, McGraw-Hill, New York.

[438] Richardson J.J. (1983): AI: An Analysis of Potential Applications to Training, Performance Measurement, and Job Performance Training, US Air Force Human Resource Laboratory, AFHRL-TP-83-28.

[439] Richmond P.A. (1983): Futuristic Aspects of Subject Access, in Library Resources and Technical Services, 27(1)88-93.

[440] Rifkin J. (1983): The Other Half of the Computer Revolution, Datamation, pp.260-278, May.

[441] Ringle M. (1979): Philosophical Perspectives in Artificial Intelligence, Harvester Press, Brighton.

[442] Robertson A. (1981): Introduction: Technological Innovations and their Social Impacts, Int.Soc.Sci.J., 33(3)431-446.

[443] Robertson P. (1982): Non-Temporal Prediction - A Distributed System for Concept Acquisition, in Proceedings of the Fourth Biennial Conference of the Candian Society for Computational Studies of Intelligence (CSCSI/SCEI0), Saskatoon, Sakatchewan.

[444] Rogers P. (1982): AI as a Dehumanizing Force, Univ.Exeter, Dept.Computer Science, Report W-104.

[445] Rolling L. (1983): Machine Translation's Big Ten Years: 1980-1990, in P.Taylor, B.Cronin (eds.): Proceedings of the 5th Conference on Information Management in Europe, Versailles.

[446] Ros H. (1983): PEST of the New Technology, New Scientist, p.524, November 17.

[447] Rose J.(ed.) (1983): Robotica - International Journal of Information, Education and Research in Robotics and Artificial Intelligence, Cambridge University Press, Cambridge.

[448] Rosenbrock H.H. (1981): Robots and People, Fourth Hartley Lecture.

[449] Rosenbrock H.H. (1982): Engineers and the Work that People Do, Invited Paper for Publication in IEEE Control Systems Magazine.

[450] Rossman P. (1982): The Coming Great Electronic Encyclopedia, The Futurist, 16(4)53-57.

[451] Rumberger R.W., Levin H.M. (1984): Forecasting the Impact of New Technologies on the Future Job Market, Report 84-A4, Stanford University.

[452] Runyan L. (1983): Saying Sayonara to the 19th Century, Datamation, pp.211-212, May.

[453] Rushby N.J.(ed.) (1983): Computer Based Learning, Pergamon Press, Oxford.

[454] Ruyer R. (1965): Parodoxes de la Conscience et les Limites de L'Automatisme, Albin Michel, Paris.

[455] Samuel A.L. (1983): AI, Where It has Been and Where It is Going, in Proceedings of the 8th International Joint Conference on Artificial Intelligence, Karlsruhe, West Germany, pp.1152-1157.

[456] Samuelson R.J. (1984): Job Change: Perception and Reality, Economic Impact, 45,40, January.

[457] Sanders D.T. (1981): Problem Solving - An Adaptable Art, Perspectives in Comp., 1(2)46-48.

[458] Sauers R., Walsh R. (1983): On the Requirements of Future Expert Systems, in Proceedings of the 8th International Joint Conference on Artificial Intelligence, Karlsruhe, West Germany.

[459] Schank R.C., Abelson R. (1977): Scripts, Plans, Goals and Understanding, Lawrence Erlbaum Associates, Hillsdale.

[460] Schank R.C., Wilensky R. (1977): Response to Dresher and Hornstein, Cognition, 5,133-145.

[461] Schank R.C., Riesbeck C.K. (1981): Inside Computer Understanding, Lawrence Erlbaum Associates, Hillsdale, NJ.

[462] Schank R.C. (1982): Dynamic Memory - A Theory of Reminding and Learning in Computers and People, Cambridge Univ. Press, Cambridge, Ma..

[463] Schank R.C. (1983): The Current State of AI: One Man's Opinion, AI Magazine, 4(1)3-8.

[464] Schank R.C., Slade S. (1985): Social and Economic Impacts of Artificial Intelligence, in Trappl R.(ed.), The Impacts of Artificial Intelligence, North-Holland, Amsterdam.

[465] Schatz W. (1983): IEEE vs. DARPA, Datamation, pp.61-67, July.

[466] Schatz W., Tate P. (1984): Export Laws on the Line, Datamation, pp.44-48, March.

[467] Schatz W., Verity J.W. (1984): DARPA's Big Push in AI, Datamation, pp.48-50, February.

[468] Schiller D. (1983): The Storming of the PTTS, Datamation, pp.155-158, May.

[469] Schlefer J. (1983): Office Automation and Bureaucracy, Technological Review, 86(5)32-40.

[470] Schmolze J.G., Brachman R.J. (1982): Proceedings of the 1981 KL-ONE Workshop, Fairchild FLAIR TR-4.

[471] Searle J.R. (1980): Minds, Brains and Programs, Behavioral and Brain Sciences, 3,417-457.

[472] Selfridge O., Arbib M., Rissland E. (1984): Adaptive Control in Ill-Defined Systems, in press.

[473] Shackel B. (1983): The BLEND System: Programme for the Study of Some Electronic Journals, J. Am. Soc. Inf. Sci., 34(1)22-30.

[474] Shapiro E.Y. (1983): The Fifth Generation Project - A Trip Report, Commun. ACM, 26(9)637-641.

[475] Shapiro S.C. (1981): Artificial Intelligence Applications, Computing Reviews, March.

[476] Sheil B. (1983): Power Tools for Programmers, Datamation, pp.131-144, February.

[477] Shneiderman B. (1980): Software Psychology. Human Factors in Computer and Information Systems, Winthrop Publishers Inc., Cambridge, Ma.

[478] Shortliffe E.H., Buchanan B.G., Feigenbaum E.A. (1979): Knowledge Engineering for Medical Decision Making: A Review of Computer-Based Clinical Decision Aids, Proc.IEEE, 67(9)1207-1224.

[479] Silva G., Dwiggins D. (1980): Toward a Prolog Text Grammar, Sigart Newsletter, 73,20-23, October.

[480] Simon H.A., Newell A. (1958): Heuristic Problem Solving: The Next Advance in Operations Research, Oper. Res. 6, 6.

[481] Simon H.A. (1977): The New Science of Management Decision, Prentice-Hall, New York.

[482] Simon H.A. (1982): Veraendert der Computer unser Leben? Zehn Thesen des Nobelpreistraegers, Bild der Wissenschaft, 6,62-72.

[483] Skarnulis L. (1982): Is There an Electronic Byline in Your Future?, The Futurist, 16(6)40-46.

[484] Skidmore M.P. (1983): Manufactoring Change - A Robot Revolution, J. Royal Soc. Arts, 5327,641-657.

[485] Sleeman D., Brown J.S.(eds.) (1982): Intelligent Tutoring Systems, Academic Press, London.

[486] Sloman A. (1979): Epistemology and Artificial Intelligence, in Michie D.(ed.), Expert Systems in the Micro Electronic Age, Edinburgh Univ.Press, Edinburgh.

[487] Sloman A., Hardy S., Gibson J. (1983): POPLOG: A Multilanguage Program Development Environment, Information Technology: Research and Development, 2, 109-122.

[488] Smith B.C. (1982): The Computational Metaphor, XEROX PARC, Unpublished Manuscript.

[489] Smith B.C. (1982): Linguistic and Computational Semantics, Xerox, Palo Alto, Ca.

[490] Smith G.W., Debenham J.D. (1983): Conturen: Mass Producing Intelligence for a Rational World, Futures, 1(83)33-46.

[491] Smith L.C. (1980): Implications of Artificial Intelligence for End User Use of Online Systems, Online Rev., 4(4)383-391.

[492] Sofer K. (1981): Art? Or not Art?, Datamation, 27(11)118-127.

[493] Soloway E., Rubin E., Woolf B., Bonar J., Johnson W.L. (1983): MENO II: An AI-Based Programming Tutor, Report YALEU/CSD/RR-258, Yale Univ., New Haven, CT.

[494] Spacek L.A. (1981): The Production Systems and PROLOG, University of Essex, Report CSM-46.

[495] Sprong J.F. (1979): Road to Slavery, in P.A.Samet (ed.): Proceedings of the EURO IFIP, North Holland Publishing Company, Amsterdam.

[496] Steeb R., Gillogly J. (1983): Design for an Advanced Red Agent for the Rand Strategy Assessment Center, RAND Corp., Santa Monica, CA.

[497] Steele G.L., Sussman G.J. (1980): Design of a LISP-Based Microprocessor, Comm.ACM, 23(11)628-645.

[498] Steeves A.E. (1982): Is there U.S. Paranoia about the Fifth Generation Project?, Comm.ACM, 25(11)854.

[499] Stefik M. (1981): Planning with Constraints (MOLGEN: Part 1), Artificial Intelligence, 16(2)111.

[500] Stefik M., Aikins J., Balzer R., Benoit J., Birnbaum L., Hayes-Roth F., Sacerdoti E. (1982): The Organization of Expert Systems. A Tutorial, Artificial Intelligence, 18(2)135-173.

[501] Steier R. (1983): The Robot Revolution: An Interview with James Albus, Commun. of the ACM, 26(3)179-180.

[502] Steier R.(ed.) (1983): Cooperation is the Key: An Interview with B.R. Inman, Commun. ACM, 26(9)642-645.

[503] Stengel R. (1982): Calling "Dr.SUMEX", Time, p.53, May 17.

[504] Stolfo S.J., Miranker D., Shaw D.E. (1983): Architecture and Applications of DADO: A Large-Scale Parallel Computer for Artificial Intelligence, in Proceedings of the 8th International Joint Conference on Artificial Intelligence, Karlsruhe, West Germany, pp.850-854.

[505] Sugiyama K. (1982): Toward Developing a Framework for Information Technology-Based Innovation: 3C Technological Innovation, International Institute for Applied Systems Analysis, Laxenburg, WP-82-33.

[506] Sugiyama K., Toda M. (1984): Visual Q-Analysis (I): A Case Study of Future Computer Systems Developments in Japan, Cybernetics & Systems, 14(2-4)185-228.

[507] Sumner F.H. (1982): The Future of Information Technology: A Personal View, Proceedings of the 55th ASLIB Annual Conference, Oxford.

[508] Suwa M., Scott A.C., Shortliffe E.H. (1982): An Approach to Verify Completeness and Consistency in a Rule-Based Expert System, Stanford Univ., Report STAN-CS-82-922.

[509] Swartout W.R. (1981): Producing Explanations and Justifications of Expert Consulting Programs, Report MIT/LCS/TR-251, MIT, Cambridge, Ma.

[510] Szabo M., Bent D.H., Bent Ph.D. (1983): An Assessment of the State of the Art of Computer based Learning Systems Design, Information Processing, pp.839-845.

[511] Szolovits P., Martin W.A. (1981): Brand X: Lisp Support for Semantic Networks, in Proceedings of the 7th International Joint Conference on Artificial Intelligence, Univ.British Columbia, Vancouver, Canada.

[512] Szolovits P.(ed.) (1982): Artificial Intelligence in Medicine, Westview Press, Boulder, Colorado.

[513] Tancig P. (1983): SOVA - General Purpose Software Environment for Natural Language Understanding Systems, in Trappl R.(ed.), Cybernetics and Systems Research, North-Holland, Amsterdam.

[514] Tate P. (1984): An Eye on AI, Datamation, pp.54-59, January.

[515] Taylor A.L., Grieves R.T., Levinson M. (1982): Dealing with Terminal Phobia, Time, p.52, July 19.

[516] Taylor A.L. (1983): Economy & Business - The Software Hard Sell, Time, pp.46-47, June 20.

[517] Taylor J.E. (1983): R&D in Human Aspects of Robotic Technology, Robotica, 1(1)15-17.

[518] Taylor R. (1982): The Future for Jobs, Management Today, pp.76-81, December.

[519] Teicholz E. (1984): Computer Integrated Manufacturing, Datamation, pp.169-174, March.

[520] Tennant H. (1981): Natural Language Processing, Petrocelli Books, New York.

[521] Teresko J. (1983): CAD/CAM Goes to Work, Industry Week, 2,40-47.

[522] Teresko J. (1983): Robots come of Age, Industry Week 1,35-42.

[523] Tesler L. (1981): The Smalltalk Environment, Byte 8,90-147.

[524] Thackray J. (1983): America's New Technology Trap, Management Today, pp.86ff., June.

[525] Thackray J. (1983): America's Robotic Rising, Management Today, pp.54ff., February.

[526] Thackray J. (1983): The Microchip Mayhem, Management Today, pp.78-84, May.

[527] Thoresen K., Maus A. (1979): Technological Unemployment and Working Conditions in the Computerized Office - Two Problem Areas for Trade Unions, in P.A.Samet (ed.): Proceedings of the EURO IFIP, North-Holland Publishing Company, Amsterdam.

[528] Toda M., Sugiyama K. (1984): Visual Q-Analysis (II): Strategy Analysis for Future Computer Systems Development in Japan, Cybernetics & Systems, 14(2-4)229-251.

[529] Trappl R. (1981): The Soft Side of Innovation: Medicine - A Hard Case, Report 81-01, Dept.Med.Cybernetics, Univ.Vienna, Austria.

[530] Trappl R. (1981): Computer Psychotherapy: Is it Acceptable, Feasible, Advisable?, Cybernetics and Systems, 12(4)385-394.

[531] Trappl R., Leinfellner E., Steinacker I., Trost H. (1982): Ontology and Semantics in the Computer, in Language and Ontology, Proc. 6th International Wittgenstein Symposium, Hoelder-Pichler-Tempsky, Vienna.

[532] Trappl R., Szer A., Huebner H. (1983): Software-Innovation in Medicine with Special Emphasis on Artificial Intelligence Methods. Report to the Austrian Federal Ministry of Science and Research, Austrian Society for Cybernetic Studies, Vienna.

[533] Trappl R. (1985): Can Computers Cure Health Care?, Cybernetics and Systems, 16, (in press).

[534] Trappl R. (1985): Reducing International Tension through Artificial Intelligence: A Proposal for 3 Projects, in Trappl R.(ed.), Power, Autonomy, Utopia: New Approaches Towards Complex Systems, Plenum, New York.

[535] Trappl R. (1985): Impacts of Artificial Intelligence: An Overview, in Trappl R.(ed.), The Impacts of Artificial Intelligence, North-Holland, Amsterdam.

[536] Trappl R. (1985): Artificial Intelligence: A One-Hour Course, in Trappl R.(ed.), The Impacts of Artificial Intelligence, North-Holland, Amsterdam.

[537] Trappl R.(ed.) (1985): Power, Autonomy, Utopia: New Approaches Towards Complex Systems, Plenum, New York.

[538] Trautteur G. (1983): A Novel Interaction between AI and Philosophy, SIGART Newsletter, 83.

[539] Trost H., Trappl R. (1985): Man-Machine Interaction: the Role of Artificial Intelligence, in K.Fedra, D.P.Loucks, J.S.Maini (eds.), Interactive Systems Analysis, Int. Inst. for Appl. Systems Analysis, Laxenburg, Austria (in prep.).

[540] Troy G. (1984): Thwarting the Hackers, Datamation, pp.117-127, July.

[541] Turkle S. (1982): The Subjective Computer: A Study in the Psychology of Personal Computation, Social Studies of Science, 12,173-205.

[542] Turkle S. (1984): The Second Self - Computers and the Human Spirit, Simon and Schuster, New York.

[543] Turoff M., Hiltz R. (1983): Working at Home or Living in the Office, Information Processing, pp.719-725.

[544] Tydeman J. (1982): Videotex: Ushering in the Electronic Household, The Futurist, 16(1)54-61.

[545] Tyler M. (1983): Software and Services, Datamation, pp.227-232, August.

[546] Tyler M. (1984): Only one per Customer, Datamation, pp.49-55, April 15.

[547] Uttal B. (1982): Here comes Computer Inc., Fortune, pp.82-90, October 4.

[548] Uttal B. (1984): Free-For-All in Computer-Aided Engineering, Fortune, pp.53-66, June 11.

[549] Vagin V.N., Klimov E.E., Lebedeva T.F. (1984): On a Formal Approach to Modelling Mechanical Assembly Arrangement in CAD/CAM, Computers and Artificial Intelligence, 3(3)273-278.

[550] Vamos T. (1982): Cooperative Systems, Invited Lecture at the 275th Anniversary Conference of the Prague University, June.

[551] Vamos T. (1983): AI as a Tool for Flexible Manufacturing, Working Paper, E20, MTA SZTAKI, Budapest.

[552] Vamos T. (1983): Cooperative Systems - An Evolution Perspective, IIASA, Laxenburg, CP-83-37.

[553] Vamos T. (1984): Neuere Erfahrungen mit Expertensystemen, INFO'84, Dresden.

[554] Vamos T. (1985): AI - Subjective Views, Future, and Impacts, in Trappl R.(ed.), The Impacts of Artificial Intelligence, North-Holland, Amsterdam.

[555] VanSlype G. (1982): Economic Aspects of Machine Translation, in V.Lawson (ed.): Practical Experience of Machine Translation, North-Holland, Amsterdam, pp.79-93.

[556] Vasko T. (1983): Economic and Social Impacts of Microelectronics, Project Draft, IIASA.

[557] Verity J.W. (1982): Learning to LISP, Datamation, pp.70-74, October.

[558] Verity J.W. (1982): Charity Begins at Home, Datamation, pp.50-51, November.

[559] Verity J.W. (1983): IEEE vs. DARPA, Datamation, pp.61-65, July.

[560] Verity J.W. (1984): Nuclear War & The Computer, Datamation, pp.50-61, February.

[561] van Melle W.J. (1981): System Aids in Constructing Consultation Programs, UMI Research Press, Ann Arbor, Michigan.

[562] Wahlster W. (1983): Tutorial on Artificial Intelligence, Industrial Liaison Session, AISB/GI: Tutorial on Artificial Intelligence, IJCAI-83.

[563] Waldrop M.M. (1984): Artificial Intelligence in Parallel, Science, 225, 608-610, August 10,.

[564] Wallich P. (1982): Artificial Intelligence Applications Called Dead End, News Supplement to IEEE Spectrum, The Institute, 6,10.

[565] Walton P., Tate P. (1984): Soviets Aim for 5th Gen, Datamation, pp.53-64, July.

[566] Waltz D.L. (1982): Artificial Intelligence, Scientific American, 247(4),101-122.

[567] Ware W.H. (1981): Security, Privacy, and New Technology, P-6606, Rand Corp..

[568] Warren D.H.D. (1982): A View of the Fifth Generation and its Impact, AI Magazine, 3(4)34-39.

[569] Warrick P.S. (1980): The Cybernetic Imagination in Science Fiction, The MIT Press, Cambridge, 11-12.

[570] Wasniowski R.: Learning for Global Survival. The Role of Science and Technology, Futures Research Center, Technical University, Poland.

[571] Waterman D.A., Hayes-Roth F.(eds.) (1978): Pattern-Directed Inference Systems, Academic Press, New York.

[572] Weil U. (1983): Evaluating the Japanese Challenge, Datamation, 29(1)124-133.

[573] Weinreb D., Moon D.A. (1981): Introduction to Using the Window System, WP 210, AI-Lab, MIT, Cambridge, Ma.

[574] Weinreb D., Moon D. (1982): LISP Machine Manual, 4th Edition, Symbolics Inc., Cambridge, Mass..

[575] Weizenbaum J. (1976): Computer Power and Human Reason – From Judgement to Calculation, Freeman, San Francisco.

[576] Westin A.F. (1983): New Issues of Computer Privacy in the Eighties, Information Processing, pp.733-739.

[577] Widmer G. (1984): MINIMYC: Ein Expertensystem fuer einen Mikrocomputer, Diplomarbeit, Inst.f.Med.Kybernetik, Univ.Wien.

[578] Widmer G., Horn W. (1984): VIE-PCX – The Vienna PC Expert System, Report 84-04, Dept.Med.Cybernetics, Univ.Vienna, Austria.

[579] Wilensky R. (1982): Talking to UNIX in English: An Overview of UC, University of California, Berkeley.

[580] Wilkinson B. (1982): Managing with New Technology, Management Today, pp.33-40, October.

[581] Wilks Y. (1981): A Position Note on Natural Language Understanding and Artificial Intelligence, Cognition, 10,337-340.

[582] Williams T.A. (1983): Technological Innovation and Futures of Work Organization, Technological Forecasting and Social Change, 24, 79-90.

[583] Winograd T. (1977): On Some Contested Suppositions of Generative Linguistics about the Scientific Study of Language, Cognition, 5,151- 179.

[584] Winograd T. (1983): Language as a Cognitive Process, 1: Syntax, Addison-Wesley, Reading.

[585] Winston P.H.(ed.) (1975): The Psychology of Computer Vision, McGraw-Hill, New York.

[586] Winston P.H. (1977): Artificial Intelligence, Addison-Wesley, Reading, Ma.

[587] Winston P.H., Brown R.H.(eds.) (1979): Artificial Intelligence: An MIT Perspective, Vol.2, MIT Press, Cambridge, Ma.

[588] Winston P.H., Brown R.H.(eds.) (1979): Artificial Intelligence: An MIT Perspective, Vol.1, MIT Press, Cambridge, Ma.

[589] Winston P.H. (1984): Artificial Intelligence. 2nd Ed., Addison-Wesley, Reading, MA.

[590] Winston P.H., Horn B.K.P. (1984): LISP (2nd ed.), Addison-Wesley, Reading, MA.

[591] Winston P.H., Prendergast K.A.(eds.) (1984): The AI Business, MIT Press, Cambridge, MA.

[592] Withington F.G. (1983): Winners and Losers in the Fifth Generation, Datamation, pp.193-209, December.

[593] Witkam A. (1982): An Alternative Strategy for Steady Growth towards High Quality Translation Networks, in P.Taylor, B.Cronin (eds.): Proceedings of the 5th Conference for Information Management in Europe, Versailles.

[594] Wohl A. (1980): A Review of Office Automation, Datamation 26(2)117-119.

[595] Wong C.M., Crawford R.W., Kunz J.C., Kehler T.P. (1983): Application of Artificial Intelligence to Triple Quadrupole Mass Spectrometry (TQMS), Lawrence Livermore Ntational Lab, CA, Rep. UCRL-89935.

[596] Wong K.Y., Casey R.G., Wahl F.M. (1982): Document Analysis System, IBM J.Res.Develop., 26(6)647-656.

[597] Yasaki E.K. (1982): Tokyo Looks to the '90s, Datamation, pp.110-115, January.

[598] Yasaki E.K. (1982): Japans's Research Realm, Datamation, pp.164/19-20, December.

[599] Yasaki E.K. (1983): R&D in Japan, Datamation, pp.91-95, July.

[600] Yazdani M. (1982): Novel Ways of using a Computer in an Educational Institution, Working Paper-108, Computer Science Department, University of Exeter.

[601] Yuta S., Iijima J., Kanayama Y. (1981): A Locomotion Control System for Mobile Robots, in Proceedings of the 7th International Joint Conference on Artificial Intelligence, Univ.British Columbia, Vancouver, Canada.

[602] Zukav G. (1979): The Dancing Wu Li Masters, William Morrow and Company,New York.

[603] Zvegintzov N. (1983): Nanotrends, Datamation, pp.108-116, August.

[604] − (1979): HUSAT: An interface between the Human Sciences and Advanced Technology, The HUSAT Research Group, Loughborough Univ.of Technology, Leicester.

[605] − (1980): Report of the Investigation of Effects of Microelectronics on Employments, Japan Information Processing Development Center (JIPDEC), January, (in Japanese).

[606] − (1982): Artificial Intelligence, The Second Computer Age Begins, Business Week, pp.46-51, March 8.

[607] − (1982): Automatic Translation: Peace on Earth?, Infotecture Europe 10.

[608] − (1982): Declaration of the Seven Heads of State and Government and Representative of the European Communities, Report of the Working Group on Technology, Growth, and Employment, Chateau of Versailles, June 4-6.

[609] − (1982): Japan fuer Internationale Computerentwicklung, Japan Heute, 249,7-9.

[610] − (1982): M. Francois Mitterand Propose un "Immense Effort" concerte d'Investissements et de Formation, Le Monde, p.2, June 6,7.

[611] − (1982): Office Automation in Europe: Issues and Trends, Special Advertising Section.

[612] − (1982): Prospects for New Technologies: Artificial Intelligence, in National Research Council, Outlook for Science and Technology, Freeman, San Francisco.

[613] — (1982): Robot Sabotage, Time, p.52, September 20.

[614] — (1982): The Fifth Generation: Dawn of the Second Computer Age, SPL International, Abingdon, Oxfordshire.

[615] — (1982): The Science of Cognition, in National Research Council, Outlook for Science and Technology, Freeman, San Francisco.

[616] — (1982): The Shape of Things to Come, Management Today, December.

[617] — (1982): World Computer Center Studied by Congress, Communications of the ACM, 25(8)596.

[618] — (1983): A European Robot That's Taking on the World, Business Week, pp.30-31, June 27.

[619] — (1983): A One-Stop Shopping Center in your Home, U.S. News & World Report 94(18)A27.

[620] — (1983): Anybody Want to Buy a Robot?, New Scientist, October 13.

[621] — (1983): British Firms move against Software Pirates, New Scientist, p.7, October 6.

[622] — (1983): Chip Wars: The Japanese Threat, Special Report, Business Week, pp.50-64, May 23.

[623] — (1983): Computer Security: What Can be Done^?, Information Processing, Business Week, pp.80-81, September 26.

[624] — (1983): Computer Team Charts Course for 1990, New Scientist, p.76, July 21.

[625] — (1983): Computer translators say goodbye to Water Goats, New Scientist,15 Dec,811.

[626] — (1983): Dawn of the Microchip University, New Scientist, p.836, September 22.

[627] — (1983): Defence Electronics for the Doves Give Peace a Chance, New Scientist, p.134, April 21.

[628] — (1983): Education? By Computer, Naturally, U.S. News & World Report, 94(18)A5.

[629] — (1983): Enter the Robot Librarian, New Scientist, p.586, March 3.

[630] — (1983): ESPRIT Workplan, EEG, Issue 1, Version 2, April 7.

[631] — (1983): European Data Firms Huddle Together, New Scientist, p.670, September 8.

[632] — (1983): France Mobilises Robotics Research, New Scientist, October 13.

[633] — (1983): General Computer Takes Charge, New Scientist, p.153, April 21.

[634] — (1983): Getting Computers to Design their own Chips, The Economist, 286(7281)99.

[635] — (1983): Government Gives a Boost to Computing, New Scientist, p.272, May 5.

[636] — (1983): Handicapped Children win Offer of Computers, New Scientist, p.527, August 25.

[637] — (1983): High-Tech Companies Team Up in the R&D Race, Business Week, pp.62-63, August 15.

[638] — (1983): High-Tech Companies Team up in the R&D Race, Business Week, pp.62-63, August 15.

[639] — (1983): How Computers Remake the Manager's Job, Business Week, April 25.

[640] — (1983): Introducing Expert-Ease, ESI - Export Software International Ltd., Edinburgh.
[641] — (1983): Look Ahead, AI Markets develop, Datamation, 14.
[642] — (1983): Looking for Oil? Ask the Expert System, New Scientist, p.270, July 28.
[643] — (1983): New Clothes for IT in Schools, New Scientist, pp.133-135, July 14.
[644] — (1983): Novel Ruling, Time, p.53, September 12.
[645] — (1983): Old University Ties give a Boost to Business, New Scientist, p.431, November 10.
[646] — (1983): Optical Memories that will give Computer Storage a big Boost, Business Week, p.98, May 2.
[647] — (1983): Pop Programs Need to Fall into Line, New Scientist, p.27, July 7.
[648] — (1983): Program of Conference on Artificial Intelligence, Oakland University Center for Robotics and Advanced Automation.
[649] — (1983): Proposal for a Council Decision adopting the first European Strategic Programme for Research and Development in Information Technologies (ESPRIT), Commission of the European Communities, Bruessel.
[650] — (1983): Relief for Chip-Shop Eyes, New Scientist, p.457, May 19.
[651] — (1983): Researchers Cash in on Artificial Intelligence, New Scientist, p.420, November 10.
[652] — (1983): Robot Turns into a Work-Mate, New Scientist, p.223, April 28.
[653] — (1983): Robots Bump into a Glutted Market, Business Week, p.26, April 4.
[654] — (1983): Scholarly Communication: Online Publishing and Education, Queens College, Flushing, New York, 1,4.
[655] — (1983): Software Rentals: Piracy is the Hot New Issue, Business Week, pp.62-63, August 1.
[656] — (1983): Special Issue on AI Applications in Medicine, Computers & Medicine, 12,7, July.
[657] — (1983): Students' Pockets Bulge with Silicon Profits, New Scientist, p.134, April 21.
[658] — (1983): Telecommunications versus Poverty, New Scientist, p.661, December 1.
[659] — (1983): The Artificial Intelligence Report, The Artificial Intelligence Publications, Los Altos, August.
[660] — (1983): The Battle of the Supercomputers: Japan's All-Out Challenge to the U.S., Information Processing, Business Week, pp.67-71, October '7.
[661] — (1983): The Computer as Teacher, Consumer Reports, pp.614-617, November.
[662] — (1983): The Computer Moves in, Time, January 3.
[663] — (1983): The Proper Use of Human Ability: A Challenge to Engineers, IFAC-Newsletter, October.
[664] — (1983): Videodiscs and Computers: A Dynamic Duo, Business Week, pp.63-64, February 7.

[665] — (1983): Women cannot Resist Employer's Demands, New Scientist, p.554, August 25.

[666] — (1983): Business Week, November 28.

[667] — (1984): Betting Big on a Computer System for Managers, Business Week, pp.123-124, March 26.

[668] — (1984): Can Software Makers Win the War against Piracy?, Business Week, pp.92-93, April 30.

[669] — (1984): Computer Professional's Job Attitudes: Read it and Weep, Management Review, 73,5.

[670] — (1984): Myths and Chips, New Scientist, p.2, March 8.

[671] — (1984): Robot Census is Bad News for Manufacturers, New Scientist, p.27, February 23.

[672] — (1984): "Computer People": Yes, They really Are Different, Business Week, pp.44-48, February 20.

APPENDIX 1

List of Contributing Scientists

The following list contains the names, affiliations, and addresses of the scientists who contributed a paper (P), participated in the discussions of the Task Force Meeting (T), or helped in the preparation of the AI Bibliography (B):

Professor Michael A. Arbib (P)
Computer and Information Science
   Department
University of Massachusetts
Amherst, MA 01003
USA

Professor Margaret A. Boden (P+T)
School of Social Sciences
University of Sussex
Falmer, Brighton
Sussex BN1 9QN
UK

Dr. Ronald J. Brachman (T)
Fairchild Laboratory for AI
   Research
4001 Miranda Avenue
Palo Alto, CA 94304
USA

Professor Stefano A. Cerri (P+T)
Department of Informatics
University of Pisa
Corso Italia 40
I-56100 Pisa
Italy

Dr. Ivan M. Havel (P)
Engelsovo nb. 78
120 00 Prague 2
Czechoslovakia

Dr. Werner Horn (B)
Department of Medical Cybernetics
   and Artificial Intelligence
University of Vienna
Freyung 6
A-1010 Vienna
Austria

Professor Friedhart Klix (T)
President, International Union of
   Psychological Associations
Department of Psychology
Humboldt University
Oranienburger Strasse 18
102 Berlin
GDR

Dr. O. Kovrigin (T)
Institute of Systems Studies
9, Prospekt 60 let Oktyabria
117312 Moscow
USSR

Professor Makoto Nagao (P)
Department of Electrical
   Engineering
Kyoto University
Sakyo-ku
Kyoto 606
Japan

Dr. Nils J. Nilsson (P+T)
President, American Association
   for Artificial Intelligence (AAAI)
Artificial Intelligence Center
SRI International
333 Ravenswood Avenue
Menlo Park, CA 94025
USA

Professor Peter Raulefs (T)
Department of Computer Science
University of Kaiserslautern
Erwin-Schroedinger-Strasse
D-6750 Kaiserslautern
FRG

Professor Roger C. Schank (P)
Department of Computer Science
Artificial Intelligence Project
Yale University
10 Hillhouse Avenue
New Haven, CT 06520
USA

Dr. Stephen Slade (P)
Department of Computer Science
Artificial Intelligence Project
Yale University
10 Hillhouse Avenue
New Haven, CT 06520
USA

Dr. Ingeborg Steinacker (B)
Department of Medical Cybernetics
   and Artificial Intelligence
University of Vienna
Freyung 6
A-1010 Vienna
Austria

Professor Robert Trappl (P+T+B)
President, International Federation
   for Systems Research (IFSR)
Department of Medical Cybernetics
   and Artificial Intelligence
University of Vienna
Freyung 6
A-1010 Vienna
Austria

Academician Tibor Vámos (P+T)
President, International Federation
   for Automatic Control (IFAC)
Computer and Automation Institute
Hungarian Academy of Sciences
Kende utca 13–17
H-1502 Budapest, XI
Hungary

Professor Tibor Vasko (T)
Clearinghouse Activity
Science and Technology Area
International Institute for
   Applied Systems Analysis (IIASA)
A-2361 Laxenburg
Austria

APPENDIX 2

AI Curriculum in Vienna

1 Introduction

This appendix describes the curriculum for the elective specialization in AI in the joint computer science study program of the University of Vienna and the Technical University of Vienna. A total of twenty-eight course units are offered, of which six are required options. The units are made up of lectures (LE), exercises (EX), seminars (SE), laboratory exercises (LU), and practical exercises (PR). They cover a wide range of topics, including an introduction to AI, methodology and applications of AI, and potential impacts of AI. The methodological units deal with the areas of knowledge representation, problem solving, searching, and planning in AI, logic and deduction, learning and knowledge acquisition, and AI programming. AI applications examined include expert systems, natural language understanding, vision, robotics, automatic programming, intelligent computer-aided learning systems, and cognitive science.

The structure of the AI curriculum is shown in more detail in Table A1 (overleaf).

2 Purpose and Scope

After successful completion of the course, the student is able to design and implement AI systems to be used in economics, research, and administration, apply AI concepts for the development of computer systems, and develop and improve AI methods.

Table A1 Layout of the artificial intelligence curriculum at the University of Vienna and the Technical University of Vienna.

Introduction to AI (1)

METHODOLOGY

Knowledge Representation
  Knowledge Representation
  (2+1)

Problem Solving, Search,
Planning
  Problem Solving, Search and
  Planning in AI (2)

Logic/Deduction
  Logic in AI (2)
  Formal Methods in AI (2 SE)
  Theorem Proving (2)
  Mathematical Logic 2 (2)
  Fuzzy Inferencing (1)

Learning and Knowledge
Acquisition
  Learning and Knowledge
  Acquisition (1)

AI Programming
  Functional Programming and
  its Application to AI (2+1)
  PROLOG and Logic Programming (2+1)

APPLICATIONS

Expert Systems
  Expert Systems ((2)+1)
  AI in Medicine (2 SE)

Natural Language Understanding
  Natural Language
  Understanding (2+1)
  Computers and Natural
  Language (4 SE)

Vision
  Scene Analysis (2)
  Pattern Recognition (2+2 SE)

Robotics
  Robotics (2)
  Laboratory Exercises on
  Robotics (4)

Automatic Programming
  Automatic Programming (1)
  Symbolic Mathematics (2)

Intelligent CAI
  ICAI and Tutorial Systems (1)

Cognitive Science
  Cognitive Science (1)
  Neuroinformatics (2 SE)

Seminar in AI (2)
Practical Exercises in AI (4)
Impacts of AI (1)

3 Description of the Course Units

The required units are as follows:

- Introduction to AI (1 LE),
- Knowledge Representation (2 LE),
- Problem Solving, Search, and Planning in AI (2 LE),
- *Either* Logic in AI (2 LE) *or* Theorem Proving (2 LE),
- Functional Programming and its Application to AI (2 LE + 1 EX).

Brief descriptions of each course unit now follow:

Introduction to AI (1 LE)

What is meant by "artificial intelligence (AI)"? What are the important subfields of AI? What are the basic concepts? Which methodologies are specific to AI? What are the characteristic features of successful AI applications? What relations has AI to psychology, linguistics, and computer science? What can be said about the future of AI?

Knowledge Representation (2 LE + 1 EX)

Significance of knowledge representation, basics, declarative representation, semantic nets, logical representation, procedural representation, frames, representation of vague knowledge, default reasoning, semantic primitives, pattern matching, knowledge representation languages, FRL, KRL, KLONE, RLL, flavors.

Problem Solving, Search, and Planning (2 LE)

Using production systems for the representation of problems, problem solving as search in graphs, algorithms for directional and best-first search, search in AND/OR-graphs to solve decomposable problems, search algorithms for the choice of moves in game trees, problem solving with time-variant data, planning in connection with heuristic search, hierarchical planning, belief revision systems.

Logic in AI (Formal Methods in AI) (2 LE)

Modal logic, deontics, time logic, knowledge logic, action logic, multivalued logic, nonmonotonic logic, combinatorial logic, application of logic to AI.

Formal Methods in AI (2 SE)

Formal definition of problems and their resolution, generation of strategies in games, rewriting rule systems, lambda calculus.

Theorem Proving (2 LE)

Skolemization in the predicate calculus, transformation into clauses; Herbrand interpretations, semantic trees, theorem of Herbrand; method of Davis and Putnam, completeness of the method; the resolution principle; unification and algorithms for unification; completeness of the resolution; refinement of the resolution principle − linear resolution; paramodulation.

Mathematical Logic (two units, each 2 LE)

Logic with quantors − language, interpretations, structures, models, calculi, deductions, deduction theorem, prenexe normal form, Skolem normal form, dual Skolem normal form, realizability, refutation, independence of axioms, specific constants, "Herbrandisierung", "Dualherbrandisierung", "Herbrandsches Praenexformeltheorem", "Herbrandsches Reduktionstheorem", "Herbrandsches Refutationstheorem", the resolution principle, introduction of functors and predicates, completeness theorem of Goedel, theorems of Loewenheim−Skolem, theorems of Lindstroem, theorems of Trachtenbrot.

Fuzzy Inferencing (1 LE)

Real-world knowledge, incompleteness, fuzziness, inconsistency, fuzzy facts, fuzzy inferences, logic, probability theories, fuzzy set theory, possibility theory, other theories of fuzziness, applications in medicine.

Learning and Knowledge Acquisition (1 LE)

Machine learning, learning by instruction, learning by example, learning by analogy, learning simple concepts, learning multiple concepts, multistep learning; automated knowledge acquisition, construction and modification of knowledge bases.

Functional Programming and its Application to AI (2 LE + 1 EX)

Concepts, development and significance of functional programming; introduction to functional programming with examples in LISP; differences between AI programming and traditional programming; application of functional programming to various subfields of AI.

PROLOG and Logic Programming (2 LE + 1 EX)

Introduction to the programming language PROLOG based on Clocksin–Mellish, presentation of examples taken from Coelho–Pereira, lambda calculus, additional concepts for logic programming, comparison of the predicate calculus with PROLOG.

Expert Systems (2 LE + 1 EX)

Range of expert systems; application and characteristic examples of expert systems; structure of expert systems – knowledge base, inference component, dialogue component, explanation component, planning component, knowledge acquisition component; architecture of expert systems – methods of knowledge representation, reasoning, inference and control structures; construction of expert systems, knowledge engineering.

Artificial Intelligence in Medicine (2 SE)

Introduction to specific AI methods used in medicine via special AIM systems: decision support for diagnosis, therapy, and prognosis; representation of medical knowledge: symptoms, diseases, and diagnostic and therapeutic procedures, hierarchies, classification models, causal models; inference methods, generation of hypotheses, trigger, evaluation methods, incompleteness and uncertainty; system control, consideration of risks and costs; explanation of the underlying reasoning.

Natural Language Understanding (2 LE + 1 EX)

Problems encountered when processing natural language, linguistics, phrase structure grammar, transformation grammar, analysis of text, syntactic parsing, semantic parsing, knowledge representation,

reasoning, scripts, plans, generation of utterances in natural language, dialog strategies, applications in medicine.

Computers and Natural Language (4 SE)

Comparison and discussion of different methods for processing natural language, adaptation of selected methods to the special requirements of German, implementation of these methods, rating the methods using acquired practical experience.

Scene Analysis (2 LE)

Data, preprocessing, color, detection of edges and corners, analysis of regions, structural analysis; representation of characteristic features of a scene, segmentation, transformation 2D–2.5D, 2D–3D, stereo; single scene interpretation, generation and verification of hypotheses concerning objects, scene interpretation, hypotheses of movements, analysis of movements.

Pattern Recognition (2 LE + 2 SE)

Basics of the theory of pattern recognition, relation of pattern recognition to other subfields of AI; finding patterns, preliminaries to pattern matching, "Glaettungsverfahren", using filters, going along lines, detecting edges; selection of variables (feature extraction) – statistical and syntactical methods; statistical classification – parametric methods, nonparametric methods, learning systems; syntactical classification – pictorial grammars, stochastic grammars.

Robotics (2 LE)

Definition and description of robots used in industry – definition of the notion "Industrie-Roboter" (Industrial robot); concept of the industrial robot – kinematics, types of construction of robots, power, control, measuring systems, effectors, sensors; application areas; peripheral devices for industrial robots – conveyor devices, article recognition; reasons for using industrial robots – economy, better quality, greater flexibility, no need for trained workers, innovation, danger of losing market share, more-human working conditions, law; examples.

Laboratory Exercises on Robotics (4 LU)

Solving simple problems using a robot that works according to the Skara principle, development and debugging of programs in an AML (Automatic Manufacturing Language) for the control of robots and peripheral devices.

Automatic Programming (1 LE)

Problem description, logical programming, automatic program synthesis, knowledge-based program synthesis, automatic program transformation, program synthesis using examples, automatic program verification, systems for automatic programming.

Symbolic Mathematics (2 LE)

Algorithms based on number theory, algorithms for symbol manipulation, programming languages for symbol manipulation, canonical forms of formulas, new algorithms for symbolic integration, integration viewed as a decision problem, systems for formula manipulation – MACSYMA, REDUCE, SCHOONSHIP, SAC-ALDES, MuMath.

Intelligent CAI and Tutorial Systems (1 LE)

Types of ICAI systems – monitors for problem solving, trainer, instructor, advisor; components of ICAI systems – expert, model of the trainee, tutor system; application areas – symbolic integration, failure analysis for electronic equipment, interpretation of NMR data, medical diagnosis, games, debugging of programs, advising the user of complex systems, user-friendly interfaces; examples of ICAI systems; special deduction methods.

Cognitive Science (1 LE)

The computer metaphor (the computer viewed as a model of human information processing) as a point in common between artificial intelligence, psychology, philosophy, and linguistics; performance models, exploration models, structural models; research in intelligence; results; is a new form of human self-understanding required?

Seminar in Neuroinformatics (2 SE)

The general structure and function of the brain, information processing in neural nets, switch elements, output regulating elements, models of the cerebellum; the feedback loop, perception, higher senses, three-dimensional frequency analysis in the visual system, comparison of visual sensation to hearing; differentiation between algorithmic and nonalgorithmic thinking processes, control problems, specialization of the two hemispheres (speech, etc.), memory and recall.

Seminar in AI (2 SE)

Working in special subfields of AI — the student gains knowledge of the problems, features, and ranges of particular AI systems by preparing a report and giving a talk.

Practical Exercises in AI (4 PR)

The student learns the techniques and methods of AI programming by designing and implementing his/her own small AI system.

Impacts of AI (1 SE)

The responsibility of the scientist (politician, businessman, ...) for his own actions. What impacts can be foreseen by considering existing AI programs? What impacts will be generated by future developments? What kinds of impacts exist? Is it possible to differentiate between "good" and "bad" scientific projects? Would such a differentiation lead to action? AI as "inhuman intelligence"?

Name Index

Subject Index