The Impact of Artificial Intelligence on Society

by
Richard S. Rosenberg

Technical Report 88-20

September 1988

THE IMPACT OF ARTIFICIAL INTELLIGENCE ON SOCIETY

Richard S. Rosenberg
Department of Computer Science
University of British Columbia
Vancouver, B.C., Canada V6T 1W5

ABSTRACT

This paper presents an introduction to a number of social issues which may arise as a result of the diffusion of Artificial Intelligence (AI) applications from the laboratory to the workplace and marketplace. Four such applications are chosen for discussion: expert systems, image processing, robotics, and natural language understanding. These are briefly characterized and possible areas of misuse are explored. Of the many social issues of concern, four are selected for treatment here as representative of other potential problems likely to follow such a powerful technology as AI. These four are work (how much and of what kind), privacy (on which the assault continues), decision-making (by whom and for whose benefit), and social organization (how in a society in which intelligent systems perform so many functions). Finally it is argued that both a major programme of study in this field be launched and that practitioners assume the responsibility to inform the public about their work.

KEYWORDS

artificial intelligence, applications, work, privacy, responsibility

I INTRODUCTION

A few years ago, the question, "What is Artificial Intelligence?" was frequently heard. Clearly that situation does not obtain today. From a somewhat esoteric sub-discipline of Computer Science, AI has emerged as a major force in the universities, the marketplace, and the defense establishment. Many professors have formed companies to market products emerging from their research. Such giant companies as Texas Instruments, Xerox, General Motors, Hewlett-Packard, Schlumberger, and IBM have launched major research and development programs. Industry and the popular media are rife with such terms as expert systems, natural language interfaces, knowledge engineering, image understanding, and intelligent robots. AI has well and truly arrived.[1]

Note that even from its earliest days, some thirty years ago, two paths

have characterized research in AI. One represented the research aim of trying to understand the nature of intelligent behaviour - simply put, the scientific approach. The other is motivated by the attempt to design and build working systems which perform useful tasks normally said to require intelligence - the engineering approach. However, the importance of this distinction should not be overstated. Usually the scientific and engineering approaches co-exist, with some tension, as researchers of both persuasions confront similar problems and influence directions in the field. Nevertheless, in the haste to deliver marketable products, the danger is that science will suffer. For one thing, the effort necessary to solve a wide range of fundamental problems will be diverted to the more immediate needs of commercial enterprises. Of course, the reputation of the field as whole will suffer if exaggerated claims are not realized.

AI is a vital and exciting field and it has attracted many scientists. Research at universities, private laboratories, and government installations is proceeding across many fronts including knowledge representation, reasoning, problem solving, natural language understanding, image processing, expert systems, logic programming and heuristic search, among others. A public discussion of these and other areas was stimulated by an international conference held in Tokyo, in October of 1981, to which about 90 distinguished foreign scientists were invited. At that time Japan announced the initiation of its **Fifth Generation Project**, an ambitious program to develop useful and powerful intelligent computers. Other nations were invited to co-operate, but mindful of Japan's worldwide lead in consumer electronics, they decided to launch their own national and regional projects. The response of the United States, the present world leader in computers in general and AI in particular, is obviously of prime importance. One significant action by the U.S. Department of Defense was the Strategic Computing Initiative (*not Star Wars*) launched in 1983, which has had, and will have, serious implications for AI research and researchers, given the massive amounts of funds to be spent.

All the foregoing represents a prologue (no pun intended) for a preliminary assessment of the implications of successful AI research on society. It should be noted, however, that even partially successful AI systems will have considerable impact. For the purposes of this paper, four areas will be considered: expert systems,image processing, robotics, and natural language understanding. Among the dimensions of this investigation are

such issues as how achievements in AI will affect human self worth especially if one result of the diffusion of intelligent computers is a net reduction in employment. Such a reduction could come about through advanced industrial automation including robots and, in the office, through the realization of a true **Office of the Future.** How will the workplace itself be altered as humans cope with 'intelligent', non-human co-workers? What about the role of such systems in medical treatment, education, financial services, and many other areas of human experience. Other serious issues relate to their use in military decision-making, law enforcement and surveillance, and to the more general concerns of privacy itself.

Any treatment of the impact of AI on society must be situated in the more general concerns of the impact of technology itself. However the nature of AI suggests that as a new technology, its influence may not be merely quantitative, not just more of the same, not just the latest improvement in the continuous progress of the industrial revolution. If its potential is realized, AI may usher in a new age, some of whose features will be discussed below, and it will truly be a qualitative break with the past.

In what follows, a number of issues will be presented which form the basis of a proposed research programme. A necessary component is a description of the current state of AI, at least with respect to immediate commercial applications. This description given in Section 3, briefly characterizes the four areas mentioned above and associates with them some possible social concerns. These concerns are elaborated in the next section under the following categories: work, privacy, decision-making, and social organization. The paper concludes with a call to AI researchers to speak openly and forcefully about their work and its potential for good or ill.

2 AN OVERVIEW OF SOME ISSUES

As discussed in the **INTRODUCTION**, the term AI has now entered the public consciousness. Articles describing computers which think appear regularly in major magazines. In popular computer magazines, AI is a regular feature as languages (Lisp, Prolog), expert system shells (M1,KEE), speech recognition, image processing, and natural language understanding are evaluated. Does this imply that AI, as manifested in working computer

programs has arrived? Of course not. Certainly the mystique of AI, the idea of steady progress towards computers able to perform a wide range of human-like behaviour has been established in the public psyche.

The publicity and the hype are definitely upbeat and expansive and herald a golden age of leisure. All work is to be done by machines, enabling people to realize their potential, free from worry about jobs, rent, and food. Many questions naturally arise, aside from the obvious one about whether or not intelligent computers are really possible.[2] First we face the task of articulating the dimensions along which intelligent systems can be classified as a necessary step to determining where and how they will be used. The more difficult problem of anticipating promising research and development directions is not a major goal and will be considered only in passing. Winner (1986) calls for the development of a philosophy of technology, "to examine critically the nature and significance of artificial aids to human activity" (p. 4). If those artificial aids purport to be intelligent how much more the urgency of the enterprise. Of course, Winner himself is quite skeptical about both the achievements of, and the claims for, AI. He notes that, ".. children have always fantasized that their dolls were alive and talking" (p. 14).

Clearly there are many applications of intelligent systems of direct and unequivocal benefit to society at large and workers in particular. For example dangerous activities in mines, under the sea, in nuclear plants, in the chemical industry, and elsewhere are prime candidates for robots. Less dangerous, but obviously unpleasant, jobs should in the near future be also phased out as exclusively human preserves. There are many ways that this process can take place and if history has anything to teach in this area it is that the protection of workers is not usually the major reason that change takes place. Lessons drawn from the introduction of computer technology into the workplace by David Noble (1984), Harley Shaiken (1984), Phil Kraft (1977), Barbara Garson (1988), and others point out the existence of alternatives and the motives of managers in constraining them.

Of the many issues to consider in this context let me limit the discussion to a preliminary examination of the following ones. I must stress the term, "preliminary" both in the sense of representing a beginning and an attempt to shape a research programme.

1. A realistic evaluation must be attempted of current and near-future prospects for AI applications at home, in the workplace, and in the government.

2. A similar evaluation is necessary of the impact of computer-related technology in the workplace, balancing benefits against perceived problems, including deskilling, monitoring, job loss, restricted promotion paths, breakdown of traditional social organizations in the office, limited entry level opportunities, and health-related concerns.

3. The implications of partially realized intelligent systems in terms of the requirements placed on humans to accommodate to their, the systems', inadequacies must be considered. In the haste to introduce AI into the workplace, pressures may be placed on people to work with systems, which, while advertised as intelligent, are seriously deficient in many areas.

4. Of particular interest is the role of AI in decision-making, whether in financial institutions, in the executive suite, or in diverse military situations such as autonomous land vehicles, pilots' aid, aircraft carrier battle management (all of which are components of the Strategic Computer Initiative launched in 1983) or in the evaluation of possible nuclear attack (either in or out of the context of SDI, the Strategic Defense Initiative).

5. Intelligent systems may find ready application in intelligence activities such as automatic interpretation of tape recordings and the cross-correlation of electronic files. Added to current threats to privacy, the availability of such powerful mechanisms could increase real and anticipated assaults on individual privacy.

6. Futuristic projections of a society without poverty brought about by the extensive diffusion of AI applications have been considered by science fiction writers and futurologists and more recently by AI researchers themselves, including James Albus (1976, 1983) and Nils Nilsson (1984). Even Wassily Leontief (1986), a Noble winner in Economics, has been concerned about such a future. Speculation is interesting but the assumptions underlying the forecasts must be carefully analyzed. Questions to be considered include the following:

What replaces regular work as a necessary part of life?

How is wealth to be distributed if a wage system is no longer operative?
How will the political structure respond?
How will human dignity and self worth be affected if we are no longer defined in part by what we do?

It should be kept in mind that these questions are obviously so difficult to answer, or even to characterize, that only a beginning is made here. However, it is important that they be raised and that a serious discussion be initiated.

3 THE STATE OF AI

Attempting to evaluate an active, diverse, and controversial discipline is a task to be undertaken with some care. Only those areas of AI with an immediate effect will be considered.

Expert Systems

Probably the greatest impact of AI has been in the area of Expert Systems (ES's). In the past few years, ES's have become the most visible and highly publicized product of the entire AI enterprise and indeed they can be characterized as an almost independent discipline. They have been applied in a wide, and growing, number of domains including medicine, education, finance, industry, and the military. In their simplest form, they consist of a body of knowledge, typically quite narrow but necessarily deep, and a reasoning apparatus, or engine, able to use this corpus to perform a range of activities including the ability to answer questions, predict events, offer warnings, or identify meaningful patterns.

Beyond this rather limited overview, work in ES's can be divided into two broad categories - the implementation of specific expert systems employing so called **shells** and research in many diverse areas associated with basic problems in AI such as knowledge representation, reasoning in general and reasoning with respect to uncertainty, knowledge acquisition, testing, verification, explanation, maintenance, and growth, among others. The first category is usually outside AI and involves a cooperative effort in

knowledge acquisition among domain-specific experts, programmers, and others operating within the constraints of a commercially available ES shell. The emphasis on knowledge acquisition dominates because of the necessity of translating informal expertise into the requirements of the knowledge representation formalism of the shell. The inference strategy and problems of uncertain information are accommodated within the given shell and therefore must be accepted by the user.

The availability of shells of varying degrees of complexity has resulted in an explosion of ES's. Many companies from the very small to the very large have built ES's to capture some part of their corporate expertise. Systems abound, as described by Feigenbaum and McCorduck (1983) and are increasing at a rapid rate. Magazine articles, newspaper stories, and books abound in describing this growth industry, viz., Winston and Prendergast (1984), Harris and Davis (1986), Bernold and Albers (1985), Pau (1986), to name a few. The applications are seemingly endless but perhaps the following are representative: medical diagnosis, computer configuration, mineral prospecting, oil well analysis, stock market decision-making, financial planning, insurance analysis, electromechanical fault-finding and repair, military analysis and decision-making. More specifically, the list below are taken from a recent conference on expert systems and their applications, called Avignon '88:

- An expert system for the airline gate scheduling problem
- A CAI expert system using an object approach
- ARCHIMEDE: an expert system for teaching geometry
- CAPEX: an audit planning expert system
- CIME: an expert system application in teledetection
- PECUNIA: an expert portfolio advisor
- Multi-expert systems and the evaluation of business risks: the CREDEX system
- SMES: a software modeling expert system
- JURIX: a legal expert system
- OPERA: an expert system to assist design by simulating the assembly of mechanical units
- SAOR: an expert system for the handling of pension funds
- PEOPL: an application of an expert system for the programmed evaluation of personnel
- SPECTRO: an expert system for spectroscopic analysis
- The pilot's associate
- SIAM: an expert system for threat analysis
- SEXTANT: a real-time system for exploring tactics for seaborne planes
- A planning expert system for nuclear power plant maintenance
- ERASME- A road maintenance multi-expert system

- ARPEC: an expert system for designing protocols for clinical tests
- NOSTRADAMUS.1: an expert system for anesthesiology

This list is purposefully long, to suggest that virtually no area has been immune from ES applications and that therefore concern about the possible impact of such ubiquitous systems is warranted. Let me argue that the widespread use of ES's brings to the fore some very serious issues. By their very nature, ES's purport to capture, formalize, and disseminate expertise. From a negative point of view, the effects of this process may include, standardization, homogenization, centralization, legitimization, and a definite sense of authority and control. Part of the concern is that the formalization of knowledge as an ES for some restricted domain can be taken as **the** representation of the knowledge. An analogy might be the legal code, where crime is defined by the existing relevant statutes.

In many areas of life there is no consensus, no received view and in the opinion of many, there will never be. Thus the threat posed by the rapid diffusion of ES's is the limitation of diversity and the imposition of the equivalent of a **state religion**, with its fixed, dogmatic worldview. Further, the objectification of knowledge and experience brings apparent certainty to uncertainty, rules to govern feelings, and regularity where it is unwarranted and perhaps impossible. Some of these fears may seem extreme but underlying them are serious philosophical and social concerns. For example, the by now traditional use of computers to reinforce bureaucratic decisions by appeal to an infallible computer can only be extended by the adoption of even more powerful systems, namely ES's. Of course in the opinion of some critics (Dreyfus and Dreyfus, 1986) ES's are fundamentally limited in the degree to which they can actually capture expertise. As such their use will impose false and possibly harmful constraints in many aspects of human experience. What are the benefits of ES's beyond the immediate economic reasons motivating their rapid adoption?

Indeed there are a number of benefits associated with ES's and these are important in their own right, not merely to offset the possible dangers outlined above. Expertise is a rare commodity and the availability of mechanisms to preserve and disseminate it are to be valued. In such areas as medical emergencies, poison control, and other life-threatening situations, rapid access to specialized knowledge may make the difference between life and death. In remote locations such knowledge may not be

readily available from people, from books, or even from computer databases. Appropriate ES's can substitute for human specialists under some circumstances and who could argue against their use in emergency conditions? Concern may arise, however, when ES's diffuse into more mundane areas.

Let me conclude this section with two additional benefits. Recognizing the uniqueness of certain skills, we can view ES's as a means for preserving them for future generations. Dr. Jack D. Meyers, a renowned internist, has been engaged in a long term process to do exactly this (Pople, 1985). He is firmly convinced that a computer program of a special kind is the only way he can effectively ensure that the knowledge gained over a lifetime can be saved in a useful form. Whether or not the ES is used as a working tool, it can serve another important function, namely, an educational one. An ES can perform a service beyond traditional textbooks in educating the next generation of human experts.

Vision

Computer vision or image understanding is generally concerned with the development of computational models to deal with images, typically accessible from television cameras. Thus a basic problem is given a television picture, determine a (or more than one) real world scene which could have given rise to it. Such a computational model may provide insights into biological vision or be an engineering solution to a given vision problem. Nothing more will be said about the scientific aims of computer vision since our concern is with the impact of applications but it should be kept in mind that vision paradigms emerging from AI may affect scientific methodology in this area.

By far the major application of computer vision is in industry, more specifically, in automated inspection for quality control. Another small but growing area is related to the use of machine vision for the automated guidance of industrial robots. In 1987, vision systems worth approximately $110M were sold in the U.S. with estimates[3] of about $125M for 1988. Most current robots operate essentially in the dark - they are preprogrammed to perform a series of actions, whether welding, painting, or cutting, in an environment required to be appropriately arranged.

Considerable expense is involved in preparing the workplace for robots and this cost must be expended each time the application changes. The incorporation of a vision system and a problem solver with the robot creates a flexible, highly adaptable robot, able to operate in a wide variety of circumstances. Because of the immense difficulty of solving the real-time, unconstrained vision problem, there will be no early, widespread growth of vision-equipped robots. Rather one should expect the gradual appearance of special purpose robots in constrained environments.

Many of the social concerns obviously relate to the impact of robots on both the nature of work and the possibility of unemployment resulting from their use but since these concerns are primarily associated with robots and general industrial automation they will be discussed in the next section. Two other potential issues which may arise are related to the use of vision systems in surveillance and as prosthetic devices. The former is part of the general concern with the erosion of personal privacy in the age of computers. An intelligent vision system could be another component in this assault by permitting images to be correlated with text records.

Prosthetic devices could be important in improving the ability of the visually impaired to function at home and at work. A possible science fiction theme arises involving the use of microprocessor mediated devices to improve biological functions. No controversy should exist over the use of artificial aids but there appears to be some apprehension about the merger of biological tissue with electromechanical systems, the so-called **cyborg** phenomenon. To Weizenbaum (1976, pp. 268-269), certain instances of this integration are "obscene,"

> There are, however, two kinds of computer applications that either ought not be undertaken at all, or, if they are contemplated, should be approached with utmost caution.
>
> The first kind I would call simply obscene. These are ones whose very contemplation ought to give rise to feelings of disgust in every civilized person. The proposal I have mentioned, that an animal's visual system and brain be coupled to computers, is an example. It represents an attack on life itself.

Robotics

The current generation of industrial robots, for the most part, does not incorporate intelligence. As mentioned above, the robot on the assembly line expects to perform its preprogrammed sequence of actions on objects placed precisely within its reach. This requirement both increases the cost and reduces the flexibility of robot applications. A major step to improve the situation is to equip robots with sensory devices and problem solvers to deal with unforeseen circumstances. The primary sensory device is vision, the use of which depends on advances in machine vision. Important requirements are speed, that the scene be analyzed in real time and generality, that complex scenes be amenable to such analysis. Another sense that could be useful is touch in that the end effecters could be made sensitive to hardness, softness, and other features of the objects to be manipulated.

Sensing and analyzing the environment is the first part of the process; a problem solver able to adjust to a wide variety of circumstances is also necessary. Again this component must act in real time to modify the trajectory of the robot to enable it to perform its task. Thus the contribution of AI will be to endow industrial robots with the ability to carry out a wide variety of actions in relatively unconstrained environments, in reasonable time. Of course all will depend on the solution of a number of difficult technical and conceptual problems. Typically the diffusion of AI technology will take place before such solutions are achieved with the result that claims will exceed reality and compromises will have to made in the workplace. There will be an impact in terms of both the number of jobs available and in the nature of these jobs. Clearly, the growing use of robots in conjunction with other aspects of industrial automation, viz., flexible manufacturing, computer integrated manufacturing, will continue the trend of a decreasing proportion of the workforce engaged in manufacturing. The repercussions may well be significant, as many writers have noted.

Although sales of robots are growing in the United States, the rate of growth has not matched most predictions. In fact, in 1985 sales were nearly $600M but fell to $475M the following year[4]. Recovery was marginal to about $480M in 1987 with predicted sales of about $510M for 1988. Thus although the economy has been doing well over the past two or three

years and talk of a post industrial age is common, the sales of robots, surely an indicator of a coming utopia, have not matched such optimism. Nevertheless, the talk has accelerated in volume as the anticipated contributions of AI are factored into the productivity equation. The reasons that robot sales have not soared will not be explored here but are related to capital costs, over- stated claims, unanticipated problems in their introduction, and other market conditions.

From management's point of view, robots will serve to reduce labour costs in a variety of ways - no absenteeism or sickness, improved quality control, reduced training costs, no lengthy grievance procedures, and reduced dependency on the labour market to say nothing of removing unions as a factor in doing business. Furthermore, management will argue that robots are beneficial because they replace workers in dangerous and unpleasant situations such as welding and painting and perhaps will be used more effectively in nuclear radiation environments, in mines, and in undersea activities. But such claims notwithstanding, the primary justification for the large outlays necessary to introduce robots is increased productivity associated with reduced dependency on labour. Thus the fundamental social issues related to industrial automation and robotics are problems in the nature of work encountered by humans in close quarters with sophisticated machines and the potential net loss of large numbers of blue collar jobs, not just jobs displaced. The impact of AI can only accentuate these effects and create additional ones by virtue of accelerating certain trends. Of course ultimate liberation from endless toil may also be a result, as many proponents of AI maintain.

Natural Language Understanding

I want to distinguish between the study of language per se and the design and construction of systems to perform specific tasks. The scientific study of language across many dimensions is obviously a legitimate enterprise but the applications of such an endeavour may have serious societal effects. The most common application is the natural language interface, a way to communicate with a computer system, perhaps a database, an ES, or a decision system, by means of a natural language, say English. Of course even the most advanced of such interfaces only approximates English in a rather superficial fashion and here is where

danger lurks. The appearance of understanding conveyed by a natural language interface can be misleading. Geoffrey Pullam (1987, p. 56), in discussing the application of natural language interfaces in the Strategic Computing Initiative, states this concern clearly,

> The inherent and irremediable limitations of computers make them only more unreliable when they run artificial intelligence programs and yet more so if they have natural language front ends. Their reactions to unexpectedly ambiguous utterances, which cannot be guarded against by linguistically untrained users, will always be unpredictable. Natural language systems can only be safe if they are restricted to uses that are relatively benign: off-line, information-garnering uses rather than on-line, system-controlling uses.

Such a possibility falls within the general apprehension that many have about the premature use of AI systems. As mentioned previously, AI applications tend to be used before more general problems are solved, thus giving the impression that they are more competent than is warranted. Users will attribute much greater powers to these systems than the designers either intended or implemented. Dangers abound and it will be difficult to limit the expectations of the average user.

Another application receiving considerable attention is speech understanding. All the problems associated with achieving successful natural language understanding are magnified when the processing must begin with the electrical analogue of the acoustic speech wave. There does exist some confusion in this area because the term speech recognition is also used. However speech recognition applies to a much more restricted domain, namely the recognition of single words or phrases, not unconstrained continuous speech. Many commercial products are available for speech recognition applications such as speaker identification, a limited range of commands, and for use by the physically disabled but the more general problem is still open in spite of ongoing research programs.

During the 1970's a concerted effort was launched by the Advanced Projects Agency (ARPA, now called DARPA), of the U.S. Department of Defense, to develop a speech understanding system within a carefully designed set of design specifications (Newell, 1973). The results were only partially successful but the enterprise provoked the following response from Weizenbaum (1976, pp. 271-272):

> What can it possibly be used for? ... But such listening machines, could they be made, will make monitoring of voice communications very much easier than it is now. Perhaps the only reason that there is very little government surveillance of telephone conversations in many countries of the world is that such surveillance takes so much manpower. Each conversation on a tapped phone must eventually be listened to by a human agent. But speech-recognizing machines could delete all "uninteresting" conversations and present transcripts of only the remaining ones to their masters.

Weizenbaum's remarks were greeted with derision and charges of paranoia, within the AI community but the issue of AI research finding a warm welcome from police and military agencies has not disappeared.

Although Weizenbaum pointed out an application of speech understanding, which many found far fetched, there is clearly a more obvious one with enormous possibilities for dramatically altering a major segment of the working population. An economical, real-time speech understanding system available in every office would substantially affect office employment and bring about the long heralded but currently unrealized **Office of the Future**. What role would many women occupy if the production of letters in the office consisted in speaking into a microphone, via a computer to a laser printer, and subsequently automatically addressed and stamped, or sent electronically, with a copy to file? Such a scenario is not around the corner but its realization would certainly have an incredible effect on work and on privacy (imagine all those documents on-line).

4 ARTIFICIAL INTELLIGENCE AND SOCIETY

Obviously space does not permit a detailed exposition of the many issues associated with the explosive growth of computers but given the revolutionary possibilities of AI, it is worthwhile to suggest those areas, in society, most likely to be affected. Some of these points have been raised in previous sections but they are included here for completeness. The following will focus on work, privacy, decision-making, and social organization.

Work

Of particular concern is the impact of computers on work - both the nature of the job itself and the number of jobs. The relation between technology and work is complicated and operates on many dimensions. The economic imperative to introduce new technology to increase productivity in order to be competitive is alive and well today as it has been since the onset of the Industrial Revolution and before. However, the emergence of AI will bring to the fore the question: Will there be a massive loss of jobs and if so, what kinds of jobs will be available? Various writers, including more recently AI researchers, such as Nilsson (1984), have speculated about a future in which intelligent machines produce the goods and provide many of our services. Of course this is also a theme explored by many science fiction writers from both a utopian and dystopian point of view. Without pursuing this topic, let me note that serious issues of income distribution, self-worth, and the basic political organization of society are involved.

If the foregoing appears to be too speculative and far-fetched then perhaps more immediate concerns of the changing nature of work itself might be considered. The introduction of robots (or more general forms of industrial automation) into the factory or computer networks and associated equipment into the office, has already had, and will most certainly have, a wide impact on workers. Among the problems already identified are deskilling, monitoring, health effects, psychological stress, and the issues of self-worth and dignity. Others are restriction of promotion paths, the breakdown of existing workplace social organization, and a limitation on entry-level jobs. These are discussed in Rosenberg (1986) as well as in many other sources. The incorporation of AI in robots and office automation can only exacerbate these problems, in most cases, but there are clear benefits in others.

Surely intelligent machines can perform tasks undesirable for people, for example, both underground and undersea mining, dealing with hazardous wastes, welding and spray painting, and handling noxious gases. Of greater significance is the intellectual benefit of intelligent systems for improving efficiency in every aspect of human endeavour. We have already discussed ES's but, in general, knowledge-based systems can be used in the executive suite and on the factory floor. Intelligent aids to information retrieval, decision-making, planning, and problem-solving are appearing and will

continue to improve. For those holding jobs which require, at their core, the ability to make decisions, AI augmented systems offer both hope and despair. Hope exists in the potential power of the new systems to "amplify intelligence," to offer to the mind what motors have offered to muscles. Despair lurks in the threat to human autonomy, to the very essence of what makes us human - our ability to reason about the world and control our own destiny. Even unrealized AI may project such a threat.

Privacy

The growing threat to individual privacy represented by the increasing use of computer databases has long been recognized by civil libertarians as well as the general public. The existence of private databases containing employment, credit, cable, fuel, and medical records, among others, and public databases storing tax, census, education, and voters' records, creates the possibility for abuse because of the ease of accessing these records under a variety of search conditions. Such terms as computer profiling and computer matching have become quite common recently, as well as controversial. The former relates to the attempt to predict behaviour, of a potentially criminal kind, by defining a profile which can then be searched for in existing files, thus identifying individuals "likely" to exhibit such behaviour. Computer matching has already been employed in many situations to cross index files in order to determine whether or not individuals have committed crimes as revealed by inconsistencies in their records. For example, an examination of property records might reveal that Mr. X has received considerable income on the sale of a piece of land while at the same time collecting unemployment insurance, as determined from an examination of payments listed in the files.

Both of these practices have aroused considerable debate as they tend to presume guilt *a priori* for some individuals, who are to be identified through a search for confirming evidence in the computer matching case and subsequently placed under surveillance in the profile case. Such actions seem to be a violation of presumed innocence and may place individuals under jeopardy. The use of AI can only encourage wider applications of such procedures and even the development of new techniques of investigation. ES's and AI are being used by police forces in investigations, including surveillance, and will surely be used in the business world for similar

purposes if cost-effective. Of course genuine criminals may be apprehended by means of this technology and this is surely worthwhile. But it must be remembered that personal privacy has been under assault ever since records have been kept. This assault has intensified since the computer and associated databases have become readily available. Governments in many countries and at many levels have enacted laws to deal with the most obvious abuses but various threats remain and the potential problems posed by advances in AI have created new challenges to those concerned with civil liberties in general and privacy in particular.

Decision-Making

This term will be used here in an all-encompassing sense to cover activities regularly carried out by individuals, companies, institutions, and governments, involving the assimilation of information, its organization, and finally its employment based on experience, special knowledge, theory, and perhaps, even intuition. It hardly needs remarking that every aspect of life involves decisions, whether made by the individual or made for him, or her, by others. As such, decision-making represents a fundamental component of human existence and threats to human autonomy, however couched in friendly terms, are of serious concern. In the discussion on ES's, the question of the dogmatic aspect of formalized expertise was raised and it is to this topic that I wish to return.

The science fiction version of this possibility has become a popular theme, with the negative aspects explored by such authors as Jack Williamson (1963, 1981) and Jack Chalker (1986). Briefly, the premise is that massive computer systems have been given, or even take powers, in order to prevent humans from destroying themselves. Unfortunately, for humans, the system interprets its mission so literally that they are sufficiently constrained that life loses all its challenge. It becomes impossible to exercise even a modicum of meaningful independent behaviour. That such a future fascinates science fiction writers does not make it in any way inevitable, of course, but it does raise some interesting questions. Indeed, the increasing use of computers must inevitably result in a decrease in human decision making and the role of AI will be to accelerate this trend, especially in more critical situations.

In this respect, the debate over the Strategic Defense Initiative (SDI, or more commonly, Star Wars) has frequently turned on the question of whether or not the very large software component could perform as required. This system, expected to be an order of magnitude larger than Unix, will monitor information gathering devices, assimilate the information, *decide* on a response, coordinate the response, and continue with these activities until the end. AI is certainly expected to play a role in this system as one of its harshest critics, Parnas (1985), has noted in a dissenting commentary. Note that the issue of computer decision making does not begin with AI but rather that the reliance on AI may exacerbate the potential problem in a fundamental way.

This reliance on AI is being pursued in other areas as well including the development of the autonomous land vehicle and the battleship management system. The former can be thought of as an armoured, autonomous vehicle able to navigate the battlefield, avoid obstacles, and report on conditions. Its successful development requires advanced image understanding, problem solving, and decision-making. The battleship system is seen as an intelligent aid to a naval commander, under engagement, who must deal with many simultaneous events. Here is a system, upon which a commander will rely in dangerous situations, perhaps putting his men in some jeopardy. But then would they be better off without the benefits possibly available from sophisticated software? A chilling, partial answer to this question occurred in 1988 when an Iranian passenger airplane was shot down in the Gulf of Arabia by a U.S. ship boasting the most sophisticated electronics equipment available (including computers of course), yet unable to distinguish this plane from a much smaller fighter aircraft and unable to determine that the plane was not descending in attack mode.

Social Organization

How will society, or better its political institutions, respond to a future in which basic needs, both goods and services, are met by machines? Robots and advanced industrial automation are gradually reducing the blue collar workforce and this is taking place largely without AI. Changes are occurring much more slowly within the office but as we have noted above, the impact of successful developments in AI, especially in speech understanding, can result in many fewer jobs. Thus the issue, for the future

may be what will replace work in most people's lives both as a means to acquire money and as a major component in the definition of self worth. However desirable the anticipated high-tech future is, and this may itself be debatable, the means for achieving it are rarely spelled out, as most visions of utopia neglect to describe how society will make the torturous trip from the current world to the promised land.

Since the most important way to distribute the real wealth in society is through wages or salary (ignoring stocks, bonds, real estate, etc.) and the envisioned future includes a considerably reduced workforce, two questions emerge: How do people acquire the means necessary to acquire goods and services beyond immediate basic needs and what replaces work, with all its trappings, in most people's lives? Nilsson (1984), Albus (1976, 1983), and Leontieff (1986) suggest partial answers to these questions. For Leontieff one response is that the average work week will have to be reduced so that people will work less but not at lower salaries, thus maintaining their earning power by sharing in the increased wealth produced by the advanced technology. For Albus, the new technology, represented by robots, will, by some undefined process, be owned, in part, by the very workers that they displace thereby providing for these workers a share of the wealth that the robots earn. Albus has referred to his vision as **People's Capitalism** (1976), which others may refer to as socialism , or even communism (recall the idea of the ownership of the means of production by the workers.)

Nilsson, a major figure in the AI community, paints an enticing vision of a future free from drudgery, made possible by AI. Removing the need to work as the primary means to satisfy wants will permit people to realize their potential by doing what they really want to do. Thus the apparent psychological need to work will prove illusory as technology liberates society from such a mundane requirement. Nilsson offers answers to the question of how wealth is to be distributed by quoting several authors, including Albus but does not see it as a serious problem. Except for some temporary problems during the transition period, the benefits will far outweigh the difficulties. The vision is almost purely utopian.

I would argue, and have argued, that crucial problems have been swept aside in the enthusiasm of describing the wonders to be brought about by AI. It is not clear that any of the proposed schemes for distributing the wealth resulting from AI will indeed work given that the current political system

in operation in the most technically advanced countries is capitalism, unfettered free enterprise. How societies are supposed to move from such an economic system to a system in which accumulated wealth is distributed to individuals independent of their direct responsibility for earning this wealth remains a mystery. Up to now the earnings from natural resources such as minerals, oil, coal, lumber, fish, and the land itself, accrue only to those who extract it, fish for it, or farm it, with a rather small portion returning to the state in the form of license fees and royalties. In fact if the state received a greater portion, entrepreneurs would be discouraged from searching for new resources. Based on this experience, what reason is there to believe that new forms of wealth achieved by the use of new technologies, such as AI, will be made universally available?

It is not all obvious how the technologically-induced utopia will come about or even that it will. Massive changes in political systems as well as social organizations will be necessary. Work and money are just part of the equation. Autonomy, self respect, and civil liberties are others. None of these are gifts bestowed by a benevolent state, especially one which is the product of major technological innovations. It is for this reason that the process involved in moving towards a new society is so crucial and that an awareness and realistic understanding of how technology operates, perhaps the philosophy of technology, referred to earlier, is so important.

5 CONCLUSIONS

These will be brief. This paper is one small step, among many, to try to understand how AI, representative of other new and far-reaching technologies such as genetic engineering, will manifest itself in the future. If this effort were to be classified under the category of futurology, the major point would be lost. It is important to keep in mind that the relations among economics, politics, social organization, and technology are deep and intricate and it would be foolish to trivialize them by predicting a future based on unfettered technology. If AI is seen as a natural and expected continuation of the historical evolution of technology, then there is no reason to expect its effect to be substantially discontinuous with the past. Some two hundred years since the onset of the Industrial Revolution, we find that unemployment rates are relatively low.

The contribution of AI would have to be revolutionary, in the true meaning of this word, to transform the world in a way which would result in a major decrease in real employment.

Concern with the impact of a given form of technology should not be left, in general, to "outsiders." I do not mean, of course, that practitioners should be regularly involved in the criticism business but rather that as part of their professional responsibility, both to themselves and to society, they should be prepared to speak out about their work and to inform the public about possible implications. An excellent review of many important social issues related to AI can be found in Athanasiou (1987). I also recommend Weizenbaum's (1986) impassioned call to the computer science community, in general, and to the AI sub-community, in particular, to recognize its responsibility by considering carefully the outcome of its work. Furthermore he argues that computer professionals must be aware of how dependent society is on their expertise for future developments, especially military applications.

For those who would respond that the duty of professionals is to their clients and that their personal views should be kept private, I would point out that the potentially far reaching implications of AI demand that its researchers take an active role in initiating a public debate on the issues, in order to educate society at large. Such a demand is not unique to computer science; note that at many law schools, in the U.S., there is currently a movement afoot, called Critical Legal Studies (CLS), which seeks to instill social responsibility within the educational process. For example, in a brief statement, describing the goals of the movement, we read the following[5]:

> Although CLS advocates are clearer about what they oppose than propose, they favor manipulating existing law to guarantee economic equality and to eliminate distinctions based on class,race, and sex. In short they'd dispense with most legal precedents.

Surely, a relatively new discipline, such as computer science, can establish for itself the practice of encouraging its members to speak openly and forcefully about their work.

ACKNOWLEDGMENTS

I would like to mention the help of Sheryl Adam, Wolfgang Bibel, and Alan Mackworth in attempting to make this paper more readable. They are not responsible for the opinions expressed, which are of course mine.

NOTES

1. See Andrew Pollack, "Setbacks for Artificial Intelligence," **The New York Times**, March 4, 1988, First Business Page, p.32. Pollack reports that a number of leading companies in ES development tools, Lisp machines, and ES applications lost money and instituted layoffs in 1987. Although he notes that AI has failed to live up to its promise, the major reason for these difficulties is poor financial management rather than poor products.

2. This important question will not be answered here (or anywhere else for that matter) for the issue at hand is to try to understand how the gradual diffusion of this most advanced of all technologies will affect society. On the other hand if the critics, Dreyfus (1972) and Dreyfus and Dreyfus (1986), are to be believed, computers will never be intelligent and all questions revert to the traditional study of the impact of technology in general on society.

3. **Electronics**, January 7,1988, p.87.

4. **Electronics**, January 8,1987, p. 67; January 7,1988, p.88.

5. See Leslie Helm and Lawrence J. Tell, "The Radical Rumblings Shaking Up Law Schools," **Business Week**, June 6, 1988, pp. 116.

REFERENCES

Athanasiou, Tom (1987), "High-Tech Politics: The Case of Artificial Intelligence," **Socialist Review**, (17:2), No. 92, March/April, pp. 6-37.

Albus, James S. (1976), **People's Capitalism**, College Park,MD: New World

Books.

_____ (1983) "The Robot Revolution: An Interview with James Albus, **Communications of the ACM** (26:3) pp. 179-180.

Avignon' 88. Eighth International Workshop on Expert Systems and their Applications, Avignon, May 30-June 3, 1988.

Bernold, Thomas and Albers, Gunter (eds.) (1985), **Artificial Intelligence: Towards Practical Application**, Amsterdam: Elsevier Science Publishers.

Chalker, Jack L. (1986), **Lords of the Middle Dark**, New York: Ballantine.

Dreyfus, Hubert L. (1972), **What Computers Can't Do**, New York: Harper and Row.

Dreyfus, Hubert L. and Dreyfus, Stuart E. (1986), **Mind Over Machine: The Power of Human Intuition and Expertise in the Era of the Computer**, New York: The Free Press.

Feigenbaum,Edward A. and McCorduck,Pamela (1983), **The Fifth Generation: Artificial Intelligence and Japan's Challenge to the World**, Reading, MA: Addison-Wesley.

Garson, Barbara (1988), **The Electronic Sweatshop: How Computers Are Transforming the Office of the Future into the Factory of the Past**, New York: Simon and Schuster.

Harris, Larry R. and Davis, Dwight B. (1986), **Artificial Intelligence Enters the Marketplace**, New York: Bantam.

Kraft, Philip (1977), **Programmers and Managers: The Routinization of Computer Programming in the United States**, New York: Springer-Verlag.

Leontief, Wassily and Duchin, Faye (1986), **The Future Impact of Automation on Workers**, New York: Oxford University Press.

Newell,A., et al (1973), **Speech Understanding Systems, Final Report of a Study Group**, New York: American Elsevier Publishers.

Nilsson, Nils (1984), "Artificial Intelligence, Employment and Income," **THE AI MAGAZINE**, Summer, pp. 5-14.

Noble, David F. (1984), **Forces of Production**, New York: Alfred A. Knopf.

Parnas, David Lorge (1985), "Software Aspects of Strategic Defense Systems," **American Scientist**, September-October, pp. 432-440.

Pau, L.F. (ed.) (1986), **Artificial Intelligence in Economics and Management**, Amsterdam: Elsevier Science Publishers.

Pople Jr., Harry E. (1985), "CADUCEUS: An Experimental Expert System for Medical Diagnosis," In Winston and Prendergast (1985), pp.67-80.

Pullam, Geoffrey (1987), "Natural Language Interfaces and Strategic

Computing," **AI & Society** (1:1), July-September, pp.47-58.
Rosenberg, Richard S. (1986), **Computers and the Information Society**, New York: John Wiley & Sons.
Shaiken, Harley (1984), **Work Transformed: Automation and Labor in the Computer Age**, New York: Holt, Rinehart and Winston.
Weizenbaum, Joseph (1976), **Computer Power and Human Reason**, San Francisco: W.H.Freeman.
_____ (1986), " Not Without Us," **Computers and Society**, (16:2,3) Summer/Fall .
Williamson, Jack (1963), **The Humanoids**, New York: Lancer, (originally published in 1949).
_____ (1981), **The Humanoid Touch**, New York: Bantam.
Winner, Langdon (1986), **The Whale and the Reactor**, Chicago: The University of Chicago Press.
Winston, Patrick H. and Prendergast, Karen A. (eds.) (1984), **The AI Business: Commercial Uses of Artificial Intelligence**, Cambridge,MA: The MIT Press.