The social context of artificial intelligence: a guideline and discussion paper

National Research Council of Canada

For the publisher's version, please access the DOI link below./ Pour consulter la version de l'éditeur, utilisez le lien DOI ci-dessous.

https://doi.org/10.4224/23001516

National Research Council Canada
Conseil national de recherches Canada

Associate Committee on Artificial Intelligence
Comité associé de l'intelligence artificielle

The Social Context of Artificial Intelligence

A Guideline and Discussion Paper

June 1989

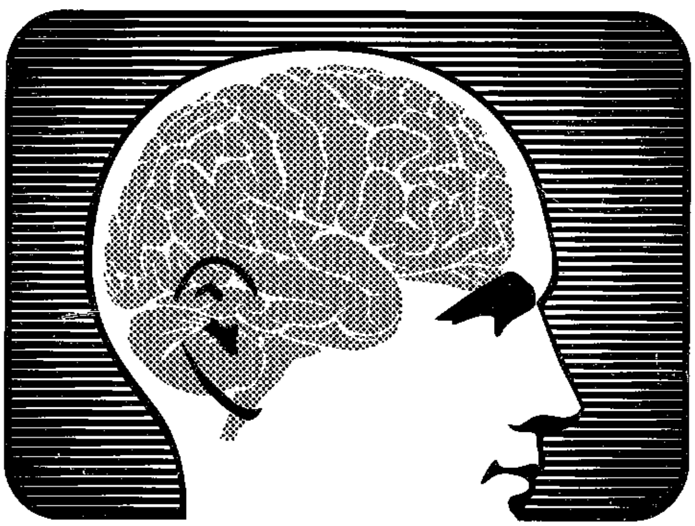

Canada

NRCC No. 30187

Additional copies are available free of charge from:

Editorial Office, Room 301
Division of Electrical Engineering
National Research Council of Canada
Ottawa, Ontario, Canada
K1A 0R6

Des exemplaires supplémentaires gratuits peuvent être obtenus à l'adresse suivante :

Bureau des publications, Pièce 301
Division de génie électrique
Conseil national de recherches du Canada
Ottawa (Ontario) Canada
K1A 0R6

Disponible en français, CNRC n° 30462.

The Social Context of Artificial Intelligence

A Guideline and Discussion Paper

National Research Council of Canada
Associate Committee on Artificial Intelligence

June 1989

Table of Contents

Preface v

I — Introduction 1

II — Social Considerations Arising from Artificial Intelligence 9
- Introduction 11
- Issues and Concerns Arising from Artificial Intelligence 11
- Work 12
- Privacy 12
- Decision Making 13
- Social Organization 14
- Situation Examples and Discussions 14
- A Difference between AI and Data Processing Applications 17
- Regulatory Bodies 18
- An Industry Association for Voluntary Accreditation of AI Products 18
- Codes of Ethics 19
- Validation 19
- Traceability 20
- Documentation 20
- Summary and Conclusions 20

III — Social Considerations Arising from Other Information Technologies 23
- Introduction 25
- Office and Factory Automation 25
- Organizational Changes 26
- Macro and Societal Level Changes 26

IV — A Protocol for the Examination of AI Applications from the Social Point of View 29
- Contexts for a Case Study 31
- The Focal Organizational Context 31

V — Summary and Recommendations 35
- Summary of the Protocol 37
- Additional Issues, Concerns, and Questions to Ask 38
- Summary of Issues and Concerns 41
- Conclusions 41

VI — References 43

Appendix 1 — A Biblioabstract of Selected References on AI in the Social Context 47
- A. Biblioabstract of Selected References 49
- B. Publications of a Continuing or On-Going Nature 56

Appendix 2 — Terms of Reference: NRC Associate Committee on Artificial Intelligence 57
- Subcommittee on the Social Context of Artificial Intelligence 59

Appendix 3 — Organizations and Centres with Activities Related to the Social Context of Artificial Intelligence 61

Appendix 4 — Definition (by Example Publications) of the Diversity, Domain, and Intent of AI 67

Preface

Associate Committees established by the National Research Council Canada have been used for over fifty years on many topics to act as a focal point and forum in Canada for issues of national concern to the engineering and scientific communities. The NRC Associate Committee on Artificial Intelligence (ACAI) was established in April 1987. In its first meetings the ACAI established a sub-committee on the social context of AI, recognizing that artificial intelligence carries with it, not only important technical and economic considerations, but also considerations of a social nature.

Examination of AI in the social context reveals that there are considerations and choices of a social nature related to the way AI may be applied in the workplace, there are social considerations related to the professional responsibility for AI applications and there are large scale macro effects with the potential to affect or influence society as a whole.

The subject of artificial intelligence is bounded on one side by the general subject of information technologies (of which it is a part) and on the other side by the general subject of industrial automation. A very large volume of literature exists on the social aspects of information technologies in general. Similarly, much is available on the social implications of industrial automation. Both of these fields have much longer histories than AI, and it is useful to draw some examples or to point to some of the references available from them. To maintain a forward focus, however, an attempt has been made in this report to restrict most of the considerations presented to the field of artificial intelligence itself.

Near the end of the first meeting of those directly involved in the preparation of this document, one contributor remarked on the limitations of such an endeavour.

"I don't know" he said, "if a group of wisemen ever gathered around the steam engine when it was first invented and if they sought to predict all the social effects that would follow". History has shown us how nearly impossible such an attempt would have been. Nevertheless, in the present instance of artificial intelligence, the contributors to this report believe it is worth something to try.

By producing this report, which includes a protocol for the examination of AI applications from the social point-of-view, it is the hope of those who have contributed to it that our understanding of artificial intelligence will be enhanced, and that the selection and design of AI applications will be undertaken with an improved understanding of the social considerations attached to them in addition to those considerations of a business and technical nature.

Understandably, this report is not intended to provide answers and solutions to the many questions asked or suggested by the situations described. It may provide some degree of guidance, however, by seeking to anticipate them, and presenting them in a balanced and structured manner. The answers or solutions eventually will come from many sources and will only evolve with time and experience. Even the real questions will only evolve in a similar manner.

Since this is a first report, correspondence is invited and will be acknowledged. While it may not be possible to enter into discussion on all communications received, they will be taken into account in any future report or addendum.

Correspondence should be addressed to:

Artificial Intelligence in the Social Context
NRC Associate Committee on Artificial Intelligence
Room 205, Building M-50
National Research Council of Canada
Ottawa, Ontario
K1A OR6

Section I

Introduction

Section I — Introduction

Before one can begin to consider the business, technical or social ramifications of artificial intelligence, it is necessary to have some reasonable definition of what is meant by the term "artificial intelligence".

Artificial intelligence (AI) is not easily defined. It may be considered to be founded on a branch of computer science that concerns the conception, design, development, and application of computer-based systems to perform tasks that require the exercise of intelligence as exemplified by capacity to recognize, reason, decide, and learn. Thus application of AI seeks to incorporate such human-like skills and expertise into computer-based systems. These systems include what are called knowledge-based and expert systems.

It is an essential characteristic of an application of AI that the associated computer system has provision for representing definable elements of knowledge about the context in which the system will be used and about the goals intended to be realized in its use. Further, when the range of conditions of application of the system is appropriately limited, the behaviour exhibited by the system is at least intended to be rationally attributable to the content within it of the representation of its knowledge and goals. These characteristics of AI systems will be illustrated by a variety of examples.

Computer systems which realize artificial intelligence commonly display advanced forms of input and output such as:

- Speech input and voice recognition
- Natural language processing
- Machine vision
- Photo interpretation
- Image understanding
- Advanced forms of sensor-based input
- Voice synthesis and speech output
- Mechanical and manipulation output capability (as in robotics).

Applications of AI include, but are not limited to, the following:

- Natural language processing

- Machine translation

- Expert systems, such as those used in the diagnosis and treatment of disease, monitoring and diagnosing faults in equipment, industrial plants, telecommunications, and electric utility systems

- Operations scheduling and process planning in manufacturing

- Robot and other autonomous systems with vision or other sensing capabilities that permit them to respond to changes in their external environment.

Expert systems are of particular interest because in terms of application they represent a leading edge of AI. They usually exhibit most, but not necessarily all, of the attributes listed in Table I.

A short list of potential application areas is shown in Table II in order to illustrate the breadth and pervasive nature of artificial intelligence. It is necessary to recognize this breadth in order to place the social context of AI in perspective. If the applications of AI become as widespread as some advocates believe, or if even only partially so, its influence will be significant on the daily lives of a large number of people. A purpose of this report is to show in what ways this may happen, and to suggest how the benefits may be maximized while minimizing any undesirable effects.

For example, a set of technical and economic factors that may be used in the selection of expert system applications is shown in Table III. By developing a parallel set of factors for the examination of AI applications from the social point of view (such as those contained in Sections IV and V of this report) an additional set of useful criteria may be brought into play.

This report, in order to meet its purpose of facilitating and enhancing the process of adaptation, must necessarily concentrate or focus on situations where system and technological changes will be needed in order to meet human and social needs, or on situations where social adaptation and organizational change may be needed as adjustments to the new technological capability. While the report must necessarily focus on these issues, it is necessary at the same time to retain a clear image of the benefits which will be obtained from the continued development and application of artificial intelligence. These will be

Table I. Expert System Attributes

Nature of Application and Use

Expert systems are used to provide advice in much the same way that it would be obtained and used if a human expert were available.

Domain expert

In the building of an expert system it is necessary to have available, and to draw on, the knowledge and experience of a person who is already an expert in the subject area. This person is referred to as the domain expert.

Rule based

Rather than being based on procedures written in algorithmic form, many expert systems derive their knowledge from a collection of rules.

Method of knowledge representation

The system knowledge is contained and given by the rules in the knowledge base rather than being represented by procedures given by computer coding as in "conventional" programs.

Symbolic representation

Expert systems tend to work with information in symbolic form rather than being based on calculations performed in numeric form. Use may be made of input, output or rules expressed in a natural language such as English, rather than being expressed numerically.

Search strategies

In seeking a problem solution, expert systems examine, select, and manipulate rules in the knowledge base by means of search strategies such as forward and backward chaining. The portion of the system performing the logic function in an expert system is sometimes known as the inference engine.

Computer language

Expert systems are often written in list processing languages such as LISP, or logic programming languages such as Prolog, rather than in procedurally oriented languages such as Fortran.

Traceability

An important attribute of expert systems is their ability to explain how a given result was obtained. Also known as transparency or explanation facility.

Apprenticeship

There is often a need to subject expert systems to a period of apprenticeship as part of their development because it is harder to verify and validate expert systems than it is for computer programs of a numerical or computational nature.

Designed to grow incrementally

Since the knowledge base is separate from the search logic, it should be possible to add new rules or modify existing rules without having to review or re-organize the original structure.

Table II. Application Areas for AI and Expert Systems (1)

Manufacturing
- Design Synthesis and Optimization
- Process Planning
- Process Parameter Selection, e.g., Robot Arc Welding
- Robotic Vision and Intelligent Robot Systems

Process Control
- Analysis and Reduction of Alarm Conditions
- Defining Process Management Rules
- Modelling the Performance of Best Operators

Materials
- Selection of Best Materials for a Given Design Function
- Selection of Coating Materials and Finishes

Resource Industries
- Evaluation of Exploration Data
- Crop and Herd Management in Agriculture
- Mineral Prospecting
- Mine Planning and Design
- Operations Management

Maintenance
- Diagnosis of Faults in Plant and Office Equipment

Health Care
- Medical Diagnosis
- Determining Optimum Treatment Cycles, e.g., as in Radiation for Cancer

Education
- Intelligent Tutoring
- AI-based Authoring Environments
- Discovery Environments
- Educational Interfaces

Career Guidance
- Advising on Education Paths needed for Employment Desired in Specified Occupations

Transportation
- Air Traffic Control
- Equipment Maintenance
- Inventory Control

Financial Planning
- Advice on Investment Decisions and Planning

Space
- Remote Sensing and Data Interpretation

Environment
- General Weather Forecasting
- Prediction of Hail Storms

Automated Language Translation
Speech Recognition
Natural Language Processing
Military Battle Management

Table III. Criteria for Choosing an Expert System Application Area (2)

- Good Payback
- Well-Bounded Problem
- Real Problem
- Small, Manageable Problem or Application Area
- Stable Proven Technology Available
- Expertise Available

considerable, widespread, and varied, because that is inherent in the nature of AI.

Basically the benefits of AI tend to be derived in the following three ways:

1. Knowledge-based systems tend to capture, represent, and distribute knowledge more uniformly throughout organizations and societal structures. Industrial firms, for example, find that the knowledge and experience gained by long term employees can be retained and distributed throughout the organization in a more uniform and accessible way than was previously possible.

2. Improved decision making is achieved via the quality of the knowledge base and reasoning process employed, and also in the speed of response. Systems are already being developed for a wide variety of applications, particularly in the form of expert systems.

3. As a result of improved and more natural interfaces and means of communication between people, between systems, and between people and computers. The possibility of natural language processing, the possibility of improved methods of machine translation for the many written and spoken languages of the world, the possibilities envisaged from computer recognition of speech and machine vision are all examples of this nature.

Although definitions exist for what is meant by "artificial intelligence", examples constitute one of the best ways by which to gain an understanding of it. Appendix 4 therefore presents brief summaries of papers published in the field of AI. These describe many of the concepts, current developments, and application areas of AI, and illustrate its widespread and pervasive nature. These examples, selected from over 700 recent AI publications, and with emphasis on the practical side, explain much of what AI people plan, propose, and do.

At the same time these illustrate the many sources from which the benefits of AI will be derived, by making the knowledge and experience of experts more widely available, as in expert systems, and by making computers more accessible, as in the many forms of new input/output interface.

By way of illustration, the references in Appendix 4 include the following topics:

Theoretical Concepts

- By studying the mechanisms of the brain, research on neural nets may lead to self-learning systems.
- The influence of language on the ability to reason correctly.
- A comparison of six leading methods for reasoning in the presence of uncertainty and imperfect knowledge.
- Methods for knowledge acquisition by inductive example.
- Why attempts to formalize common sense knowledge have failed.
- A code of conduct for professionalism in AI.
- A probabilistic system for inductive learning and predicting future events.

Systems and Techniques

- The use of AI to extend and enhance conventional information systems.
- How to handle the problem of uncertainty in medical diagnostic systems.
- Variations incurred in expert systems may often be due to the choice of domain expert.
- Establishing company/university joint ventures in AI.
- How to select the correct expert system shell for a given application.
- Human and organizational problems related to expert system applications.
- Difficulties with the maintenance of the knowledge base in an expert system.
- A comparison of how expert system shells deal with uncertainty.
- A practical guide to aid knowledge engineers in the process of interviewing domain experts.
- An expert system for the verification of nuclear test ban treaties by the interpretation of seismic data.

- Techniques for improving the knowledge acquisition process by means of induction.

Applications

- An expert system for advice in the handling of hazardous waste.
- A retrieval system for information about patents.
- A hierarchical expert system for the synthesis of chemical process flowsheets.
- Extensions, by means of expert systems, to the preparation of NC machine tool programs.
- Monitoring the health of rocket engines by integrating vibration analysis and pattern recognition techniques with AI.
- Air traffic control using AI techniques.
- Legal and contractual ramifications of expert systems.
- A system for evaluating project proposals to a government funding program.
- Military and defence applications of expert systems.
- A financial investment assistant for portfolio management.
- A system for screening applications for university admission.
- Providing advice to new and small businesses.
- A system to advice students in career planning.
- Mineral identification by the analysis of the infrared reflectance spectra of rock samples.
- A system for diagnosing faults in telecommunications switching systems.
- A system to assist engineers in setting up problems for structural analysis using the nonlinear finite element method.
- An example of neural network training; driving a vehicle through simulated freeway traffic.
- Expert system applications in the NASA space station.
- The use of expert systems in architecture, engineering, and construction (AEC).

Section II

Social Considerations Arising from Artificial Intelligence

Section II — Social Considerations Arising from Artificial Intelligence

Introduction

There are both micro and macro effects associated with the development and use of almost any technology. Micro effects are those that directly affect the individual worker or user. Macro effects usually evolve more slowly over time and tend to affect large numbers of people indirectly. This often occurs through changes in the societal and economic structure brought about by the technology in question.

For example, in the case of office automation systems, the widespread use of video display units (VDUs) can have a direct effect on the comfort or physical and mental strain of workers who use them for long periods of time, particularly if sufficient attention is not given to ergonomic factors such as lighting, and work station layout in the systems design.

At an intermediate level the use of credit cards, made possible by computer technology, provides a useful illustration. Not only have credit cards obtained widespread acceptance as a convenience to the user, but also they have fostered the establishment of new credit and service companies. These have altered the structure of the financial community. Additionally, the spending habits of people have changed (many in a substantial way), in addition to the simple and direct convenience of the card itself as a means of payment. These are macro effects, brought about by credit cards and computer technology.

As an example of a very wide scale macro effect, some writers see the possibility that thé widespread use of computer systems will gradually alter democracy as we know it today. This is predicted because of a slight but perceived bias in systems design which favours a slow drift towards centralization of power rather than favouring the individual.

The field of industrial automation and computer integrated manufacturing offers another, and perhaps clearer, example of both micro and macro effects. Workers such as welders in the automotive industry have already experienced job displacement, although not necessarily widespread job loss, as a direct effect of robot welding applications. On the other hand, as a macro effect, the plant that makes no technological change, and which suddenly finds itself to be uncompetitive, poses an even greater employment risk. All workers in the company could be adversely affected. Even the standard of living of entire nations can be impacted if their industries somehow fail to keep up with technology and thereby lose competitive position.

In this section, the report will develop and illustrate more fully the subject of social issues, at both the micro and macro level, related to artificial intelligence applications. To place AI in perspective, and to draw from more traditional information technology applications, Section III provides a summary view of social considerations which arise from information technologies other than artificial intelligence. These are closely related, and indeed often impossible to separate from those of AI itself.

Issues and Concerns Arising from Artificial Intelligence

It is difficult to distinguish between the social implications stemming from the use of computers in general, and those arising out of artificial intelligence applications in particular. Most AI applications are still experimental and the social implications will only become clearer when there is a more general use. If its potential is realized, AI may usher in a new age which will truly be a qualitative break with the past. Only the future will tell, but in any event early recognition of the issues involved can serve to ease the transition.

Clearly there are many applications of intelligent systems of direct and unequivocal benefit to society at large and for 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. Other jobs, less dangerous, but obviously unpleasant, should in the near future be also phased out as exclusively human preserves.

AI is a vital and exciting field that has attracted many scientists. Research at universities, in private laboratories and in government installations is proceeding across many fronts. These include knowledge representation, reasoning, problem solving, natural language understanding, image processing, expert systems, logic programming, and heuristic search, among others.

It can be argued that the widespread use of expert systems brings to the fore some very serious issues. By their very nature, expert systems purport to capture, formalize, and disseminate expertise. From a societal 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 expressed in some instances is that the formalization of knowledge as an expert system for some restricted domain can be taken as the representation of the knowledge; on the other hand in many areas of life there is no consensus, no received view and, in the opinion of many, there will never be. Thus a threat possibly posed by the rapid diffusion of expert systems is a limitation of diversity and the imposition of the equivalent of a state religion, with its fixed, dogmatic worldview.

While an advantage of expert systems lies in their ability to disseminate knowledge and make the ability of experts more widely available, difficulties with updating and maintaining the associated knowledge bases is certain in some instances to create a state of frozen and inflexible logic. This could exacerbate the difficulty presented above. Four specific issues may be chosen to illustrate some of the social considerations arising from artificial intelligence technology (3). These involve the influence or inter-relationship between AI and:

- Work
- Decision making
- Privacy
- Societal structure, ethics, and political adaptation

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 and 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, have speculated about a future in which intelligent machines produce the goods and provide many of our services. Without pursuing this topic further, it is noted that serious issues related to income distribution, self-worth, and the basic political organization of society may be involved.

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. These concerns are primarily associated with the general subject of industrial automation which is discussed in the next section.

Surely intelligent machines can be usefully employed to perform tasks that are 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. In general, knowledge-based systems can be used for applications in the executive suite and on the factory floor. Intelligent aids to information retrieval, decision making, planning, and problem solving are appearing in almost every conceivable area of application. One can only expect that they will continue to improve and be even more widely applied.

Privacy

The potential 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 may have committed crimes as revealed by inconsistencies in their records. The use of AI can only encourage a wider application of such procedures and even the development of new techniques of investigation. Expert systems and AI are being used by police forces in some countries in investigations, including surveillance, and will surely be used in the business world for similar purposes if cost effective.

Governments in many countries and at many levels have enacted laws to deal with the most obvious abuses but various threats remain. The potential problems posed by advances in AI have created new challenges to those concerned with civil liberties in general and privacy in particular but which are likely to be met by a slowly evolving and matching social structure.

Decision Making

This term is 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. 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. The question of the dogmatic aspect of formalized expertise contained within expert systems has already been raised. Indeed, the increasing use of computers must inevitably result in a decrease in individual human decision making. The role of AI will be to accelerate this trend, especially in more critical situations. Of special mention in this respect, the debate over the Strategic Defense Initiative (SDI, or more commonly, Star Wars) has frequently hinged on the question of whether or not the very large software component could perform as required. This system 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. 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 military and defence areas such as 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?

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 individual and collective ability to reason about the world and control our own destiny.

There are a number of benefits associated with expert systems 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. Expert systems can serve as a means of preserving knowledge, in an active and useful way, even for future generations. They can ensure that the knowledge gained over a lifetime by an expert can be saved in a useful form. Expert systems can also perform an educational service, beyond traditional text books, in educating the next generation of experts.

Expert systems may yield substantial benefits in such areas as medical emergencies, poison control, and other life-threatening situations. Rapid access to the specialized knowledge contained in expert systems may make the difference between life and death in these situations. In remote locations such knowledge may not be readily available from people, from books, or even from computer databases. Appropriate expert systems can substitute for human specialists under some circumstances. Who could argue against their use in emergency conditions, but concern may arise, however, when expert systems diffuse into more mundane or less essential areas.

Social Organization

How will society, or better its political institutions, respond to a future in which basic needs of both goods and services are met by machines? Robots and advanced industrial automation are gradually reducing the blue collar workforce. This is taking place largely without AI. Changes are occurring much more slowly within the office but the impact of successful developments in AI, especially in speech understanding, could result in many fewer jobs. Thus a key issue for the future may be what will replace work in most people's lives both as a means to acquire wealth and as a major component in the definition of self-worth.

Since the most widely accepted way to distribute real wealth in society is through wages or salary, and since the envisioned future includes a considerably reduced workforce, two questions emerge: How will people acquire the means necessary to acquire goods and services beyond immediate basic needs and what will replace work as the major activity in most people's lives?

It is not obvious how the social and economic adaptation to such a future will evolve or even that it will. It appears that massive changes in political systems as well as in social organization 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 itself, is so important.

At the same time, 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 have passed since the onset of the Industrial Revolution, yet we still find that unemployment rates are relatively low.

The nature of AI, however, 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 distinctly new age, and ultimately it will truly be a qualitative break with the past. The question then becomes one concerned with the rate of change, and the time available for adaptation.

Situation Examples and Discussions

Four examples have been chosen from many that are possible, to illustrate more clearly and in greater detail some of the social considerations related to the future use of expert systems.

The first example, in realtime process control, if not describing what is already a real situation, describes one that may soon be real. The second example, in the field of medical diagnostics, further develops a number of issues of concern and places them in the realm of direct personal considerations. A third example indicates an instability that may exist if programmed trading in stock market securities reaches the level of widespread practise. A fourth example is the possible initiation of a nuclear war due to an error in computer decision making.

Situation 1 — Realtime Process Control

Assume that a realtime expert system is installed to aid the operator in the control room of a large refinery or petrochemical plant in assigning priority to alarm signals during a major emergency or process upset. Without the expert system, the operator is certain, at these times, to be flooded with alarm signals and notifications of "out-of-limits" conditions, some of which are much more important than others. With the expert system it is hoped that the operator will be in a better position to identify the real cause of the breakdown, and to more quickly determine the corrective action to be taken. The problem, however, is that everything happens fast, in realtime, and the operator may have little choice other than to follow the advice of the expert system.

Who is responsible then if a major upset occurs and the operator follows the expert system's advice but it is erroneous because:

- the expert system rules and analysis were incomplete
- the process upset (or combination thereof) was something never considered and therefore outside the domain of the expert system
- it had actually been developed for a different kind of plant
- there was a logic error in the software
- there was an interface error to parts of the system developed by others
- some sensors had been damaged by the accident and therefore the expert system could not function
- the expert system functioned, but assumed that input data and signals were always complete and correct when in fact they were incomplete or erroneous
- etc.

One could argue that the example provided is not significantly different from traditional control systems with which there is already a wide and generally satisfactory experience, or that it does not differ from expert systems in general. The example is chosen, however, because it does differ in two important ways:

- Expert systems are less deterministic in their performance than sequential computer programs. It is more difficult to predict or test their behaviour and output over the range of input conditions. By their very nature, many expert systems base their results on inferences rather than certainties. Normally the operator may be able to make better judgements and to perform better with the expert system than without it, but there may be no guarantee of this under all conditions.
- In the example given above the expert system operates in realtime and is expected to be used in situations where the human operator has little or no opportunity to cross-check or evaluate its output. He or she will be very dependant on it, and could be forced by circumstances to trust it implicitly. Who, then, bears responsibility to the public for the plant's operation?

This example, from a perhaps hypothetical industrial environment, is sufficient to illustrate the considerations involved. There are parallel considerations for many potential military applications of expert systems or for any artificial intelligence applications which would be expected to operate in realtime, in a variety of situations, where complete pre-testing would be very difficult to arrange.

Situation 2 — Medical Diagnostics

A considerable number of the early and best known expert systems are in the field of medical diagnosis. The reasons for their development can be several fold:

- as aids to teaching the methodology of medical diagnosis,
- to collect and pool expert knowledge in areas where diagnosis is particularly difficult,
- to study a typical application area in order to further develop problem understanding and the AI techniques required for its solution,
- possibly for actual application to diagnosis in clinical situations.

It is in the latter area of real application where potential problems of a social nature may be encountered. To clarify what is meant, and to invoke discussion, a number of different, but possible, scenarios are presented below. These are drawn from potential situations in medical practise, but parallel situations could be found in other expert system application areas, particularly within the professions, such as the provision of legal advice, engineering design, and career guidance to name just a few.

Today many expert systems do not necessarily perform well but exist primarily to illustrate their potential, pending further development. However, just as chess playing programs once performed only at the novice level these now perform, sooner than was expected, at the masters level. We need to look forward already to what may happen if, or when, expert systems begin to equal and then surpass the ability and judgement of those we deem to be experts, and whom we trust, today. We may consider therefore, the following:

- A physician has an expert system available to him to aid in the diagnosis of a patient. At least six possibilities immediately arise:

1. The expert system agrees with his own analysis. The physician proceeds to follow the appropriate treatment procedure with added confidence that it is correct.

2. The expert system does not agree with his analysis. He proceeds to follow his own analysis, which turns out to be right.

3. Same as (2), but it turns out that the expert system was right and the physician was wrong.

4. The expert system does not agree with the physician's analysis. He or she decides to follow the analyses of the expert system, which turns out to be right.

5. Same as (4), but the expert system is wrong.

6. Although the expert system was available, the physician chose not to use it. Several situations arise:

   - This was a good choice, human skill and judgement prevailed and no known problems arose.

   - An error in diagnosis did occur and it was discovered that had the expert system been used and relied upon, the error would not have occurred.

The implications of this latter possibility are even more pronounced if it is considered that the expert system available, by virtue of continued development, may have reached the status of accurately modelling the diagnostic capability of a highly qualified peer group of physicians, exceeding the knowledge and ability of even any single member of that group.

- Yet another series of situations exists if one envisages a situation in a geographically remote region, or in a lesser developed country, where trained physicians are not widely available. The expert system might be marketed and used, even to advantage in the majority of cases, by paramedical personnel, by the patients themselves, or by well-meaning friends acting on their behalf.

Candidates for the assignment of responsibility, especially in instances where errors and problems might have occurred, include:

- the individual physician
- the hospital, clinic or group practise
- the supplier of the expert system (possibly from another country)
- the distributor of the expert system
- the manufacturer of the computer
- the research team that provided the domain expertise and knowledge base
- the developers or supplier of the expert system shell
- the AI team that developed it into a working system for medical diagnosis
- the patient who attempted to use it by his or herself
- the well-intentioned friends of the patient.

It is distinctly possible that in a number of instances the rate of development and application of expert systems, may be limited not by the technology itself but by matters related to legal responsibility. This may be particularly so in the professions. For this reason, some potential applications may only achieve use as teaching aids, rather than being implemented for use in direct practise.

Alternatively, as the technology develops and moves to direct application, the education requirements for professionals deemed responsible for the use of expert systems may rise to a very high and even extraordinary level. This could severely strain the ability of suppliers to provide appropriate program documentation and training. In fact, it may turn out in some instances that practitioners who genuinely understand the program, and who have the ability to use it correctly, are only available at very special centres of expertise.

These centres may often be located outside any particular country, leading in some instances to a loss of national control. A form of "software colonialism" may develop as funds and resources continue to flow through the marketing chain to these sources of expertise and thereby reinforcing centralized development.

On the other hand, this may not necessarily happen. Just as the world once looked almost solely to Versailles for art, or Hollywood for cinema, other developments over time have caused the production capacity for these two examples of

expertise and achievement to become widely distributed. Even though concentration of knowledge may appear to be an initial consequence of expert systems, it may not necessarily continue that way.

Situation 3 — Programmed Trading in the Stock Market

Programmed trading in the stock market is another area where it has already been recognized that regulatory controls may be required. While individual usage is not deemed to be an area for concern at present, multiple simultaneous use creates the conditions for an avalanche effect with rapid, widespread effects of great consequence. Stock markets, such as the New York Stock Exchange, have already given serious consideration to imposing limits on their immediate members for the use of programmed trading. If the technique and practise should become more widespread throughout the financial community at large than it presumably is at present, an even more unstable condition may result than was experienced in October 1987. Fears of a worldwide economic collapse and recession triggered by such action, while it has not happened, are not totally unrealistic.

Many types of systems, both those created by man and by nature, can demonstrate and even suffer from dynamic instabilities. However, except for electric power networks and feedback control systems, few of these, due to the intractability of the problem, have formal and well-developed methodologies for stability analysis.

Situation 4 — The Initiation of Nuclear War

Many computer scientists are concerned with the possibility that a computer error could initiate a nuclear war. The organization known as "Computer Professionals for Social Responsibility" (CPSR) is particularly concerned with this possibility. While many of these concerns stem from questions related to reliability of software in general, particularly the massive software system required in the proposed Strategic Defence Initiative (SDI), some of these concerns could originate from the proposed role of artificial intelligence. An example, in pattern recognition, is the evaluation of sensor data to determine target identification and the overall decision logic used to determine if an apparent event does in fact represent a hostile attack.

A Difference between AI and Data Processing Applications

In general most of the social issues related to the general or conventional use of computers relate also to AI. Thus it is difficult to separate the two. In some instances the role of AI may be to magnify the effect or concern in question, particularly those related to the involvement of computers in decision making as distinct from their use for numerical calculation.

A few examples will illustrate how important it is to ask the right questions in the examination of artificial intelligence applications. An expert system is described in the literature for allocating children to foster homes. This raises the question, "How can the expert system perform this function better than human judgement?", whereas the correct question to ask is more likely, "Does the human–machine combination perform better in this instance than would be the case if decisions were based on human judgement alone?"

Another point to be recognized in the examination of AI applications is that the correct question to be asked in many instances is not just, "Does the human–machine system perform better, on average, than would be achieved with human judgement alone?", but rather "Even in extreme instances does the human–machine system perform better?" In some cases this means that performance must be almost infallible.

For example, in 1988, in a case of mistaken identity, a commercial airliner was shot down and destroyed by missile fire from the naval vessel of another country. The incident occurred "over the horizon", which means that electronic means of detection were clearly involved. At one point in the enquiries that followed this event it was reported that the mistaken identity may have been caused by an error in software employed for pattern recognition and target identification. Later reports indicate that this was not the case, and that the mistake was due to human operator error (4). Had it been due to the automated pattern recognition system, however, this would have fallen into the domain of artificial intelligence, at least according to some definitions.

This illustrates the need, in many AI applications, for extremely careful system validation over a wide range of operational conditions, such that extreme values near boundary conditions are well tested, in addition to average conditions or those central to the normal zone of performance.

A difference with regard to responsibility for public risk clearly exists between data processing applications of computers, with which we have more experience and familiarity, and those involving artificial intelligence and expert systems. In some instances, and according to the interpretations that prevail in various countries, this matter of responsibility may also extend to liability.

For data processing applications it is necessary, when errors occur, that there be a channel for appeal. In many instances the mere availability of a channel for appeal is sufficient. Errors usually occur infrequently, and can be accommodated, for example, if your credit balance is restored correctly next month. In applications where human lives may be involved, or when there is no opportunity to correct the error "next month", something more than "a channel for appeal" and "on-average" performance is clearly required.

Regulatory Bodies

Under somewhat analogous conditions it is not unusual for governments to establish a regulatory or administrative body to ensure that the needs of public safety, or of the individual, are met. Thus in Canada we have the Atomic Energy Control Board (AECB) for nuclear matters, the Canadian Aviation Safety Board, the Canadian Standards Association (CSA), Emergency Preparedness Canada, the Food and Drug Act, the Canadian Human Rights Commission, and many others.

Is a similar, but special, body needed for software and engineering systems reliability in applications of a sensitive nature? This would include applications of artificial intelligence and expert systems as a special section. Given the variety of applications possible, and the difficulty of maintaining expertise in so many diverse areas, it seems reasonable to propose that this function, if needed, should be performed on a decentralized basis, that is, separately, for transportation, civil engineering structures, medical practise, and so forth. This does not eliminate, however, the possible need for a small core group to maintain specialized knowledge and to provide coordination.

Two or three decades ago the Canadian Information Processing Society (CIPS) established an ombudsman on a voluntary basis for complaints concerning data processing applications. As might be expected, it is believed that this function, while performing well, was mainly concerned with appeals concerning billing and accounting systems. In the computer applications of today, and especially of tomorrow, much more is involved.

By way of additional illustration, while we have a facility in Canada that will test sensors for hazards such as electrical safety (the CSA), there is no known body or organization that will test sensors for fitness of purpose. Discussions for the establishment of such a body to test sensors for industrial process control systems (on behalf of their users in the resource and chemical industries) have taken place in the past. Although there was considerable interest and expression of need, these plans and discussions did not come to fruition. In retrospect, the costs and difficulties of maintaining such an organization would be substantial. This is especially true if it is expected to maintain competence over a wide range of devices, although the Scientific Instrument Research Association (SIRA) has maintained at least a partial program of this nature in the U.K. for many years. It would be even more difficult to establish and maintain an analogous body for the testing and validation of advanced sensors and software used in artificial intelligence.

The Canadian Standards Association has announced the availability of four voluntary standards for software quality assurance. In critical applications Q396.1.1 is for the development of software and Q396.1.2 is for "off-the-shelf" software. Correspondingly, in non-critical applications Q396.2.1 is for the development of software and Q396.2.2 is for "off-the-shelf" software. In addition, while use of the guidelines and good practise measures is voluntary, it is also hoped that organizations purchasing the standards will "register" with the CSA and agree to a voluntary audit of the degree to which the recommendations are adopted. The CSA initiative, believed to be the first software quality assurance program in North America, is also being used as a base within the International Standards Organization (ISO) (5).

An Industry Association for Voluntary Accreditation of AI Products

The extravagant claims of some commercial companies in the marketing of their AI products have

created concerns amongst researchers and among the more conservative firms providing AI products and software. In the U.K. it has been proposed that an AI industry association is needed with a code of ethics for advertising in order to protect the reputation of legitimate researchers and firms, to protect the public from claims deemed to be extravagant, and to maintain ethical standards for the AI community. If formed, the association would in effect provide a "Good Housekeeping" seal of approval for the products and work of its members.

Codes of Ethics

Many professions, medicine, law, and engineering for example, include courses on ethics in their university curricula, and maintain a code of ethics for their members following graduation. If it has not already been done in all institutions, it would appear timely and appropriate for subject matter on ethics and professional responsibility to be added to curricula for degrees in computer science. The codes of ethics for the professions might also be examined to determine if they contain clauses that are adequate in the light of today's AI and software technology.

This alone is not enough, and frankly it can be argued that the world is not equipped for what is happening. Not only are the professions less than fully equipped for the technological change that is occurring, but a large amount of computer software, including expert systems and AI applications, will increasingly be written by persons outside the professional societies. This is only natural, and almost certain to increase. Codes of ethics developed within the professions, therefore, do not apply to these situations.

There is also an increasing amount of software in the "public domain", which raises the question, "who is responsible for the integrity of this?"

One could also raise the question as to what is meant by software in the "public domain" or what is the process for placing software in the public domain. This could be discussed at length as a separate treatise.

Validation

Validation is an important, but particularly difficult, aspect of expert systems.

Expert systems, like all other computer programs, can contain undetected and hidden errors. These can arise as errors or "bugs" in the computer coding, or they can have their source in the background analyses and knowledge base. It is particularly difficult to fully test and debug programs that have a large number of program states. The number of branches and possible program states can in fact be extremely large. This poses a combinatorial problem so large that it is impossible to exercise and pre-test all possible conditions before use. This is one reason why some computer programs occasionally run for years before certain errors appear and are only then discovered. In addition, the nature of expert systems is such that it is frequently difficult to identify and pretest the program over the full range of all possible input conditions.

As a further but important point, some expert systems are difficult to test definitively, because the program logic is based on inferences which have only statistical significance, and which therefore are not necessarily conclusive. Unlike engineering programs based on the laws of mechanics or physics, the output of an expert system can carry with it an element of doubt or uncertainty even when the program runs "correctly" in the computer sense. A somewhat similar condition exists with some applications of artificial intelligence other than expert systems, for example, those which use sensor-based input. There is always an element of doubt attached to the performance of the sensor, what it really measures, and how it is interfaced to the computer. This element of doubt, attached to the input, carries through to the program output.

There is an additional doubt with expert systems that the knowledge base may not contain all the necessary rules at any particular time or stage of its development, or that the computerized search through the knowledge base may not have considered all possible paths.

These are some of the reasons why expert systems, unlike other computer applications, are required to serve a period of apprenticeship during which their behaviour is studied in the presence of experts. During this time period confidence in the systems capabilities can gradually rise as positive experience is gained and as some problem areas are corrected. At the same time knowledge is obtained of weaknesses and deficiencies in areas to be avoided until improvements are made.

Traceability

One of the attributes of expert systems is that they should be able to explain to their user how a particular result was obtained. That is, to explain or reveal the logic or reasoning path that was used. Similarly if additional input is requested from the user, the program should be able to explain, on interrogation, why it is needed or how it will be used. This attribute of traceability, or transparency as it is sometimes called, is important to the user in understanding how the program is using the input provided and accordingly how it is reaching its conclusions. This is a particularly important attribute in instances where the user must take professional responsibility for the advice given by the expert system.

It should be noted, however, that the presence of this attribute alone, or by itself, does not in all cases completely shift the burden of responsibility for correctness from the program developer to the user. It is of little value, for example, in a realtime expert system sorting out hundreds of alarm signals in seconds during an upset in a large chemical plant to reveal days or weeks after the wrong action was taken that this would have been apparent to the operator at the time if the traceability or transparency capability had been used.

Documentation

Somewhat similar remarks can be made concerning program documentation. If a user is to take full responsibility for an expert system, he must have system documentation available to him that permits him to achieve the level of understanding necessary for this. This poses a special requirement on the user level documentation for expert systems. In a world where even the documentation for many conventional computer programs and software packages does not achieve this level it is a cause of concern.

Not only must the expert system user achieve a very high level of understanding, but also it is necessary to ask if he or she will be able to do so with a reasonable level of effort expended over a reasonable period of time. It is to be expected that the larger and more complex systems will only be adequately understood at special centres of excellence or locations close to their origin and sometimes only by their originators. This is where questions related to knowledge sovereignty and a possible "software colonialism" arise, as mentioned elsewhere in this report.

If a deep and substantial knowledge is required by the capable user of a large system, the knowledge and education requirements for high level expert system developers is even greater. It is conceivable that a lifetime of hard work and institutional support may be necessary in the future in many instances.

The statement that good program documentation is necessary seems to imply that good documentation alone will solve the problem. This is not the case. It can take years of experience and close association with a computer program to fully understand the implications of its output. Expert systems are an example of this, and as expert systems increase in their depth of reasoning, this learning time for the user will increase. Good documentation is necessary, but it is only a starting point.

Summary and Conclusions

While codes of ethics, professional responsibility, and regulatory bodies are useful mechanisms, they are by no means likely to be sufficient. Computer software, including expert systems, will soon be almost as common as people. It will originate from a wide variety of sources, and from people in many countries. Only a small percentage of these sources will fall within those covered by professional codes of ethics, in terms of either software generation or use. Thus, it is fair to say that, if some form of regulation on supply or use is needed, it will be extremely difficult to establish and achieve.

Caveat emptor, let the buyer beware, might appear as another alternative, but is equally nonviable. Software is already being given away as an inducement to the purchase of other goods and services, and this practise will increase as software becomes more and more a low cost commodity item. In many instances the buyer (or acceptor) of software will lack the skills and knowledge for any in-depth evaluation or validation of software he acquires.

Should some form of control become necessary, there is virtually no workable method by which it could be implemented that would not permit circumvention by those who might wish to do so. It appears that in most instances this could be quite

readily accomplished. In the meantime, the current system rests on:

- the integrity of software developer(s), wherever they may be;
- the codes of ethics of professional societies and other bodies, in instances where they apply;
- regulatory codes for applications already recognized to be sensitive to public safety, e.g., nuclear power plant control;
- caveat emptor;
- education and awareness of the general public.

Section III

Social Considerations Arising from Other Information Technologies

Section III — Social Considerations Arising from Other Information Technologies

Introduction

Some information technologies, for example, office automation and manufacturing automation, are more fully developed and more widely used than artificial intelligence. AI may or may not represent a sharp departure from these, and may or may not represent a sharp departure or breakpoint in the history of technology development. While many feel it does, and some feel it may not, there is agreement that the societal issues related to artificial intelligence, even though there may be some that are new in magnitude or in kind, are at least similar to those experienced in the related information technology fields. Accordingly, it is useful to draw experience from these fields and present it in a manner useful to those concerned with the examination of AI applications.

Office and Factory Automation

Office automation based on computer and communication techniques can introduce impacts on the individual, in the organization, and at the society level. A study, from the social impact point of view, of the Canadian Office Automation field trials has produced a summary of impact at these three levels, as shown in Table 4 (6).

Most of these impacts apply also to industrial and factory automation, which is an area with additional or greater concerns regarding competitive position and the avoidance of job loss at the firm and organizational level. This concern extends even to the national level (7).

In summary, in both office and factory automation, the concerns and questions raised can be used as a checklist for systems design from the social point of view. In most instances they apply equally well to AI applications. They can be grouped as follows, and include:

Individual Level

1. A first consideration is the ergonomic effects on those who continue to work at the same job, but as users of the new system after it is installed. What is the effect of the system on the user's job? Is the user's job enhanced or degraded? Do the users have more or less time for other aspects of their jobs? Do the users make better or worse use of their time? Is their performance improved or degraded? Is their job more or less tedious? Is their job more or less difficult? Will adequate training on the new system be provided, both before and after it is installed? Will the workers and users have an opportunity to identify preferences and participate in the design of the system?

Ergonomic factors include, for example, musculoskeleton and mental strain for those who work with video display units (VDUs) for long periods of time. Extensive studies have shown that these problems are frequently due to improper work place design. Loss of job satisfaction can occur also in instances where there is utter simplicity and lack of challenge in the new job. In essence, is the job enhanced or degraded by introduction of the system? This is a key question.

The potential also exists in many systems for increased surveillance of employee work output, quality, and work habits. This can include the counting of units of production, quality levels, and even to counting key strokes and break periods for operators of word processing units (8). The Canadian CAD/CAM Council anticipated this potential development of increased monitoring and also machine pacing. They recommended in the 1980 report "Strategy for Survival" that such monitoring should, as a cautionary step, only be implemented with the knowledge and understanding of the workers involved and that machine pacing of the individual should be avoided in the system design (9). These recommendations recognize, for example, that where the monitoring is for the purpose of establishing group incentive pay rates the employees are unlikely to fear, and even to welcome, the system in question.

On a more general level a new factor, termed "techno-stress", is being experienced particularly in Computer Integrated Manufacturing systems, not only by system operators such as CAD operators, but also by technical, engineering, and managerial personnel as they strive to keep abreast of the technology, to keep the company competitive,, and to cope with the constant technological and organizational change (10). The pressure to work faster, to maintain ever tightening schedules, and at the same time to keep on top of the system can be unrelenting.

2. For those whose jobs are displaced by the system, but who continue to be employed, the change in job, skills, and the likely need for retraining are paramount concerns.

3. Direct reductions in labour due to technology are usually handled by natural attrition and a reduction in the hiring rate. This does therefore contribute to unemployment (as well as to job protection) but seldom on a one-to-one basis. For those displaced to the point of unemployment, retraining, achieving new job entry skill levels, and maintaining a sense of personal self-worth are priority concerns, in addition to the obvious need and concern for economic income.

Organizational Changes

At the organizational level, within the firm, changes are occurring due to both office and factory automation. There are changes in the form and structure of management, and changes in the role of middle management. Advanced information systems are generally reported to be leading to the introduction of broader, flatter management structures. These have fewer layers and increased responsibility for the individual manager, who must be more technically knowledgeable than before, because he or she now interacts more directly with the system and less directly with people. In the words of one case study report, which clearly illustrates the trend.

> When I first started we had a department with a lot of high grade people looking after 25% of the inventory, and at one point in time we had probably 10 or 12 people looking after $25 million. Now we have 3 people looking after $125 million (10).

Automation may cause an increase, in shift work, or it may cause a decrease. As the amount of capital investment per employee rises, there is a natural desire on the part of its owners to run the plant on a second shift, third shift, and even on a continual basis in order to shorten the time required for production runs, and to reach objectives for return on investment. For these reasons, as investment in information and automation systems rises, there will be a general trend towards more shift work. For example, it is almost impossible today to conceive of a chemical plant or petroleum refinery that could run competitively one shift per day, or even be started and stopped daily in order to run one shift. With the introduction of Computer Integrated Manufacturing, manufacturing plants will also tend to become more like that. As a first order effect, this will cause an increase in shift work, which most people regard as undesirable.

Off-setting this is the trend to advanced automation systems that are capable of unattended operation during second and third shifts. These "lights out" factory automation systems, of the highest order in the technical sense, lead to a decrease in shift work. They are being introduced substantially for this reason, particularly in Japan, where many workers dislike shift work. The day shift clears away all the production of the night before, attends to the normal activities of the organization, and prepares the material, tooling, and information system for the next period of unattended operation which will take place when they go home. Not a bad state of affairs in many ways.

Labour unions are also concerning themselves with the organizational changes induced by modern information and automation systems. Not only are they concerned with the job protection and retraining needs of their members, but as the percentage of "blue collar" and direct workers decreases it portrays an environment of reduced membership and power for the unions themselves. This is a fundamental technologically induced change, even though nearly all unions today have positive views on the need for new technology itself, when properly introduced.

Macro and Societal Level Changes

The macro and societal effects of the information technologies, including industrial automation, are largely structural and economic. These are effects which are evidenced at the national and international level and which are not readily attributed to the individual or to any specific organization. Over time they can be very strong and forceful, as for example, we witness the plight of the third world countries who somehow missed out on the first industrial revolution beginning some two hundred years ago, and which seems to deepen and continue.

We see, therefore, that the impacts of the information technologies, and industrial automation in particular, go far beyond the company bottom line. The performance of a nation in the

adoption, or non-adoption, of technology has a direct bearing on its economic and social well-being and standard of living for both current and future generations (11, 12).

Exhibit 2-1. Levels of Impact of Office Automation and Associated Major Issues

1. INDIVIDUAL LEVEL
   - A. Control
     - i. Monitoring
     - ii. Pacing
     - iii. Autonomy: decision making and responsibility
     - iv. Accountability
     - v. Worker participation
   - B. Nature of Work
     - i. Increased stress
     - ii. De-skilling
     - iii. Job interest/monotony
   - C. Job Meaningfulness
   - D. Career Development
     - i. Appropriate pay benefits
     - ii. Job security
     - iii. Opportunity to learn and develop new skills
     - iv. Opportunity to advance
   - E. Social Needs
     - i. Communications networks
     - ii. Balance between work and home
     - iii. Provision of a group structure
     - iv. Privacy

2. THE ORGANIZATIONAL LEVEL
   - A. Management and Organizational Control
     - i. Changed authority structures
     - ii. New forms of management
     - iii. Changed role of middle level management
     - iv. The use of measurement systems
   - B. Organizational Structure
     - i. Removal of levels
     - ii. Two-tiered effect
     - iii. More contract work
     - iv. Changed importance of units
   - C. Changed Hours and Place of Work
     - i. Shift work
     - ii. Off-hours work
     - iii. Cottage industry work
     - iv. Satellite work places
     - v. Work from the home
   - D. Career Pathing
     - i. Managerial/professional employees
     - ii. Clerical workers
   - E. Work Design

3. MACRO LEVEL — SOCIETAL LEVEL (Major economic and social concerns)
   - A. Employment Impacts
     - i. Effects on jobs
       - layoffs
       - jobless growth
       - job displacement
     - ii. Effects on occupations
       - occupational shifts
       - occupational segregation
     - iii. Effects on skills requirements
       - rising skills requirements
       - emphasis on computer literacy
       - de-skilling
     - iv. Increased job mobility
     - v. Dual labour market
   - B. Changed Social Patterns
     - i. Reduced work time
     - ii. Changing meaning of work
     - iii. Work in off-hours and work from the home
     - iv. Transportation/telecommunications trade-offs

Section IV

A Protocol for the Examination of AI Applications from the Social Point of View

Section IV — A Protocol for the Examination of AI Applications from the Social Point of View

The purpose of this protocol on the bringing into use of a system employing artificial intelligence is to alert to the sociotechnical dimensions of doing so the persons who:

- define the ground for the intended uses of the system,
- create the system,
- bring it into use,
- operate it,
- have their tasks influenced by its use,
- evaluate its worth.

The importance of reflecting on these sociotechnical dimensions lies in the reality that systems of artificial intelligence have the potential to carry out significant tasks of perception, reasoning, decision making, and taking action that have characteristically been performed by persons. This potential raises new and profound aspects of questions about the relation of technological change to work, organization, and culture. Here the word culture is taken to mean the whole evolving social ecology which forms the basis on which the members of a society interpret their experiences and shape their relations with one another and their environment into a way of life. This ecology is multi-dimensional and encompasses the social character of science, engineering, technology, economics, law, politics, the arts, morality, and religion.

Contexts for a Case Study

In the uses of artificial intelligence, as indeed of all technology, a fundamental cultural question can be put in two extreme ways.

- First, are persons to be viewed as being required (by whom under what authority) to adapt to technological artifacts?

- Second, are artifacts to be viewed as being brought into use to complement, preserve, and assist the expression of (whose) human talents?

Since technological change can and must be considered from the differing perspectives of individuals, groups, organizations, and indeed society there is clearly a spectrum of answers lying between the limiting questions as stated. For this reason it is imperative in case studies that ecological questions of intent, expectations, and consequence be considered at a number of different contextual levels which are selected by consciously shifting the focus of attention by a process of what may be called "cognitive zooming".

The Focal Organizational Context

The organizational setting within which an application of artificial intelligence is embedded constitutes a focal context from which attention can be directed inward to operational structure and can be directed outward to various social contexts. The embedding of a use of artificial intelligence is an innovation in the organization. This innovation represents:

- the set of all considerations and actions associated with the bringing into practical operational use the system embedded in the organizational setting,

- the consequences to the organization and its people.

Just as the innovation of an industrial product such as a new material, device, or system does not typically follow a linear sequential process from basic science, to applied science, to engineering design, to production, and to market distribution, the innovation of an application of artificial intelligence in an organization will not proceed simply from prior specification of intended use, to system design, to implementation and evaluation. The interactions between those who require, specify, design, construct, install, document, operate, manage, and evaluate will typically be complex. These interactions will reflect the evolving interplay of social and economic purposes expressed through operational practices within and outside the focal organization setting.

*IDENTIFY

- with respect to the selection of an application for embedding, identify at both the corporate and relevant social levels distinctive contextual characteristics.

For example, a private-sector industry has the underlying corporate purpose of making profit to provide a reasonable return on investment to the owners. This objective is carried out within an environment of corporate law and regulation, which together with socioeconomic circumstances, local, regional, and global, establish a context of risk and competitive opportunity for the responsible achievement of profits.

*IDENTIFY

- the distinctive socioeconomic circumstances of the organizational setting (such features as the form of the corporate entity, its relation to organized labour, to the local and wider environment of competition and regulation, to local social, political, and ecological conditions).

- within the organizational setting identify the motivation for bringing into use the application. Is there a perceived problem to be addressed? Is the instance part of a corporate strategy for technological change? Is there an internal champion as a driving force? Is there persuasion from outside suppliers? Is there peer pressure from the industrial sector?

* IDENTIFY

- within the concrete situation of application the proposed grounds of justification for the necessary expenditure and stress of organizational change? In an industrial setting these elements will characteristically have to do with technical feasibility, profitability, and quality of human relations. Factors that may be pertinent include potential for increased productivity, improved product quality, improved working conditions and job satisfaction, reduced downtime, enhanced flexibility of production and organizational responsiveness, reduced costs or energy, materials, training, labour, and management.

To achieve reasonable clarity of response to the foregoing question of justification it is necessary to examine how the application of artificial intelligence will be embedded into and change the existing organizational context.

*IDENTIFY

- the organizational changes necessary to accommodate the system. These changes will include redefinition and reassignment of operational tasks and the associated communications network and structure of accountability. Such redefinitions will account for job creation, job displacement, machine and process replacement, and identify as enhancements or reductions significant shifts in the scope of both operational and managerial tasks. These definitions of organizational change will also identify anticipated changes in the degree of centralization/decentralization of managerial control and of worker participation.

Since current systems of artificial intelligence, for example in the form of expert systems, have sharply limited and specialized domains in which they can perceive, question, reason, make decisions, and take actions in the face of uncertainty, it is essential that the relation of its operation to human oversight be understood by all persons responsible for the operations and consequences of the employment of the system. This requirement raises novel issues of task definition, corporate accountability, and legal liability.

*IDENTIFY

- the organizational and operational contexts and characteristics of the system in terms of the functions of perception, questioning, being questioned, reasoning and explaining reasoning, decision making, and taking action through effectors.

- the boundaries or limits of the domain of expertise represented in the system and the manner in which human oversight and accountability for the use of the system is to take place in the organizational context. The designers, suppliers, and the domain experts employed have a distinctive responsibility to document and explain to the using organization "the microworld" within which the system is able to function with the intended reliability.

These identifications will serve to define the extent to which the system operates to fulfil complete tasks or functions in an advisory or enabling manner with respect to the human operation and management of the organization.

There is substantial evidence to show that technological and therefore organizational change is most effectively carried out and evolved when those persons participating in and affected by the change are kept informed of the contextual elements which their experience, skills, and talents enable them to grasp and potentially contribute to.

*IDENTIFY

- the practices to be followed in the organizational context:

- for keeping informed from the earliest stage of system choice, and "employing" the capabilities of those who will participate in and become part of the innovation,

- to inform and assist in a timely manner persons whose customary jobs will cease to exist,

- to document evolutionary changes in design and operation,

- to provide training in the use of the system.

Whether or not systems of artificial intelligence will realize the expectations held out for them in enhancing the performance of the organizations in which they become embedded as innovations will become gradually known only through a holistic process of evaluation against as many contextual elements of the embedding process as can reasonably be examined.

Systems of artificial intelligence represent a unique phase of technological change in which the consequences of use will gradually reverberate more deeply into the form and ethos of the organizational setting than has any previous stage of technology. Sensitivity to questions of change of organizational ethos is an essential dimension in the evaluation of consequences of use. As the characteristics of those organizational settings undergo change, the consequences will reverberate into the whole social ecology of our culture. The background documents for this protocol consider key dimensions of such reverberation.

*IDENTIFY

- the practices intended to be followed to assess the contribution (or otherwise) of the system to the performance and overall welfare of the organization using technical, organizational, financial, and socially germane norms. These practices will include taking note of changes in ethos reflected, for example, in internal social networks among participants (workers and managers), and seek to understand changing attitudes to the place and significance of work in an organizational context wherein the exercise of perception, reasoning, and decision making is distributed among persons and machines.

The latter assessment will require accumulation of understanding from the larger context of sets of similar applications characteristic of the development of the field of artificial intelligence and a sensitivity to the evolving social ecology of technological change.

Section V

Summary and Recommendations

Section V — Summary and Recommendations

Applications of artificial intelligence should not be pre-judged, out of context. This means that they can only be examined within their own organizational and unique setting. Accordingly, this report, while recommending that planned or existing applications can usefully be examined from the social as well as from the business and technical point of view, does not purport to pre-judge applications, but does present a protocol by which such examination can be made. Central to the protocol is the belief that technological systems can have influence or impact at the direct level on users and workers and former workers; at the intermediate level on organizations and; at the macro level on nations and human society at large. While artificial intelligence may or may not represent a significant departure from past forms of technological development (and information systems in particular), many of the criteria for evaluation derived by experience from these earlier forms, apply also to artificial intelligence. It appears likely in the minds of many that artificial intelligence carries with it some concerns that are new in kind, or that are at least of heightened magnitude from those of the past. For example, a fundamental and singular thing about AI is that it is changing the way in which the process of reasoning and decision making is proportioned between people and machines. Questions related to the professional responsibility for expert systems, and perhaps also for responsibility and liability in the legal sense particularly when mistakes occur, arise from these sources.

As with any technology, there are both pros and cons associated with it, and particularly with regard to the way in which it is used. Among those experienced with artificial intelligence, the following have been noted by some organizations to date:

Pros

1. increased job satisfaction
2. quality of work improved
3. VDUs (video display units) are less noisy and more flexible than the mechanical devices they replace
4. greater response and ability to be on top of affairs
5. new systems are user friendly, simpler to learn
6. computer solution is often cheaper than mechanical means

Cons

1. job security may be threatened
2. physical strain, mental strain, and possible health hazards associated with the use of VDUs
3. fear of innovations, loss of value of "experience"
4. fear of change, inability to cope with non-permanence
5. technology introduction requires planning, education, and attention to details
6. need for more flexible, more educated mind-set
7. loss of need for "specialist knowledge"
8. shortage of specialists
9. utter simplicity — job de-skilling and lack of challenge.

Fundamental to almost all the considerations posed is a changing balance in how decision making is conducted and performed. This is a change in balance between the portion delegated to machines, and the portion reserved for people as human beings. This changing human/machine symbiosis is viewed by many as a fundamental point of departure in the development of technology. Since the beginning of the industrial revolution, large changes have been witnessed in the division of labour between people and machines. We now, potentially, have the beginning of a new era with a significant change in the division of reasoning and decision making between people and the computer. The outcomes of this are likely to be profound.

Summary of the Protocol

1. Are persons required to adapt to the artifacts of the system, and if so, under what authority?
2. Will the system, as brought into use, preserve and assist the expression of human talents? If so, whose talents will be so preserved and assisted?
3. Examination of the system should identify:
   - The characteristic purposes of the organization at both the corporate and relevant social level
   - The distinctive socio-economic circumstances of the organizational setting (i.e., form of the

corporate entity, relationship to organized labour, competitive environment, regulatory environment, etc.)

- The motivation(s) for introducing the application. (e.g., is it to solve a perceived problem, is there an internal champion or external advocate such as a supplier?)
- The justification(s) for the application, the expenditures necessary, and the organizational changes necessary.
- How the system will be embedded into the existing organizational context, and how it will change that context.
- The organizational changes necessary to accommodate the system. This will include job creation, job displacement, machine and process replacement, the identification of enhancements, redefinition and reassignment of operational tasks, changes in the communications network, changes in accountability, changes in the degree of centralization of managerial control and of worker participation.
- How the system will function in the face of uncertainty and the degree of human oversight required.
- The organizational operational contexts and characteristics of the system in terms of the functions of perception, questioning, being questioned, reasoning, explaining reasoning, decision making, and taking action through physical effectors.
- The boundaries or limits of the domain of expertise represented in the system.
- The manner in which human oversight and accountability for use of the system is to take place in the organizational context.
- "The microworld" within which the system is able to function with the anticipated reliability. This will serve to identify the extent to which the system operates to fulfil complete tasks or functions in an advisory or enabling manner with respect to the human operation and management of the organization.
- The practices to be followed in the organizational context:
  - for "employing" the capabilities of those who will participate in and become part of the innovation and for keeping them informed from the earliest stage of system choice
  - to inform and assist in a timely manner persons whose customary jobs will cease to exist
  - to document evolutionary changes in design and operation
  - to provide training in the use of the system
- To the extent possible, how the consequences of use of the system will reverberate into the form and ethos of the organizational setting and eventually into the culture of society itself.
- The practices intended to be followed to assess the contribution (or otherwise) of the system to the performance and overall welfare of the organization using technical, organizational, financial, and socially germane norms. These practices will include taking note of changes in ethos reflected, for example, in internal, social networks among participants (workers and managers), and seek to understand changing attitudes to the place and significance of work in an organizational context wherein the exercise of perception, reasoning, and decision making is distributed among persons and machines.
- The latter assessment will require accumulation of understanding from the larger context of sets of similar applications characteristic of the development of the field of artificial intelligence and a sensitivity to the evolving social ecology of technological change.

Additional Issues, Concerns, and Questions to Ask

A list of issues and concerns to be considered when applying an AI system and/or when reviewing a case study should also include:

How and why has the application been chosen?

- Different reasons for implementing an AI system will result in very different social consequences. For example, was the application chosen to reduce manpower, to increase productivity, to free up manpower for other activities, or for training purposes?
- Was there user involvement in the choice and design of the application? Did the AI system benefit from the user's understanding, advice and creative abilities?

How does the AI System fit into operations?

- Does it address a problem that is recognized as a problem by the users?
- Are there organizational changes necessary to accommodate the system and if so how will they be addressed?
- Will implementation of the system introduce or cause changes in the communication patterns within the organization?
- Will there be job displacement? Will there be job creation?
- Will there be productivity gains? How, and to whom, are the benefits of these to be distributed?
- When and how is the technology introduced to staff (both to the domain experts and users)?
- What are the initial expectations about the technology prior to implementation? How is the technology to be evaluated after implementation?
- What level of the organization is most directly effected by the technology? management? workers?
- What is the effect of the system on the user's job? Is the user's job enhanced or degraded? Do the users have more or less time for other aspects of their jobs? Do the users make better or worse use of their time? Is their performance improved or degraded? Is their job made more of less tedious? Is their job made more or less difficult?
- What are the costs and benefits as seen by the domain expert involved in building the system?
- What are the costs and benefits as seen by the user of the system?
- What will the consequences be in terms of the centralization or decentralization of the organization or in terms of changes in management; increase/decrease in flexibility/rigidity of responses?
- Is the user interface well designed?

Recognizing that the system, when dealing with uncertainty, may give erroneous or incorrect output in the form of advice, conclusions, or direct action, has the responsibility for this been clearly defined?

- Is the responsibility chain back to the originators of the system components clearly defined? For example, are the domain experts, knowledge engineers, system engineers, and computer scientists involved identified or identifiable by name?
- Does use of the system involve a risk to the safety and health of the public at large, or to individual members of the public? Could use of the system result in personal damage to a group or to an individual?
- How is responsibility for use of the system allocated and shared between persons or bodies such as the originators of the system components, the practitioner who is operating the system, the end user on whose behalf it is being utilized? Where necessary, is this allocation of responsibility clear in the legal sense?
- If all, or part of, the system input is sensor based, how is the system protected against false, erroneous, or incomplete input or errors in the meaning and interpretation of the input? How are such conditions alerted to the user? Is the system "fail safe" from the user point of view if the system is suddenly interrupted and made unavailable?
- Does the domain of the system fall within the domain of one of the professions, such as engineering, medicine, or law? Does its use comply with the regulations and the codes of ethics of the professions concerned?
- Does the domain of the system lie within the bounds of other regulatory bodies such as the transportation act, food and drug act, environmental regulations, health and safety, the handling of hazardous wastes, etc.? If so, has compliance with those regulations been considered and assured?

- Recognizing that use of the system may improve performance "on-average", but not necessarily in all circumstances, what provision exists to screen out or minimize those instances in which degraded results will occur?

- In instances where human intervention should, or must, be applied to the system output, will the operator or user of the system have reasonable time in which to apply this intervention? This is particularly important in the case of real-time systems.

- In instances where false or undesirable information or actions may result from use of the system, does a reasonable channel of appeal exist for those who may be affected?

Documentation
- Is documentation complete and available to the user, permitting him to fully understand and comprehend the "micro-world" within which the program operates?

- Does this documentation clearly define the limits and limitations of the program?

- Can the user reach the level of understanding necessary to use the program safely and properly with a reasonable amount of study and effort over a reasonable period of time? What level of prior education and experience does this assume?

- Can the system be accessed and used by persons without this level of capability?

Validation
- Is the system validation considered complete, or is the system for approved experimental use only? What special safeguards does this require?

- Has the system validation considered all necessary possibilities?

- What is the source of domain knowledge or analysis as assumed and used in the system? Is this domain knowledge or analysis description included in the documentation? Is this in terms the user can understand? Are the assumptions and limitations clearly defined?

Traceability
- Can the system readily and adequately explain how any given result was obtained, and identify the information which was used or not used?

- Can the system readily and adequately explain how any item of input normally required, or requested, will be used?

Updating
- What provision has been made for reviewing the system, keeping it current and up to date? How are users informed of these changes?

Privacy
- Does normal use of the system affect the privacy of any individual? Could this result through abnormal use, or in combination with some other system?

National Control
- Does use of the system remain within the national boundaries of the country of origin? In cases where usage is considered outside such boundaries, does this result in a loss of national control in the social or economic sense?

Ethical
- Are there concerns or issues of an ethical nature associated with the system or its use?

The above are issues and concerns, largely expressed in the form of questions that should be explored when considering new AI technology or when reviewing an existing application of the technology.

The concerns and the ways in which they are expressed may be different in subtle but important ways from the larger population of interest, i.e., those groups implementing AI technology into their organizations. It may be noted that procedures and methods for generating questionnaires designed to elicit the concerns and issues from the population of interest are available that avoid the pitfall of assuming that the evaluators are representative of the population of interest.

Summary of Issues and Concerns

1. A realistic examination should be attempted of current and future prospects for AI applications at home, in the workplace, and in the government. The discussion and protocol for

such examination contained in this report is intended to encourage and facilitate this process.

2. Such examination would include 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 system 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.

5. Intelligent systems may find ready application in intelligence activities such as the 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 writers and more recently by AI researchers themselves. 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.

Conclusions

- Artificial intelligence has a great potential to contribute to the well-being of humanity in both the economic and social sense.

- Research in artificial intelligence may contribute, in consort with parallel research by other disciplines, to the understanding of the human mind and even to the understanding of human behaviour.

- Applications of artificial intelligence have the potential to greatly extend the boundaries for current computer applications into regions involving recognition, reasoning, problem solving, and learning. Applications include natural language processing, machine translation, diagnostic systems, and autonomous systems, including robots, with sensing and logic capabilities permitting them to respond to changes in their external environment.

- The range of application of expert systems, which is one branch of artificial intelligence, is particularly broad, reaching into almost every field of human endeavour in business, industry, education, health care, recreation, the arts, defence, and war.

- Aside from the direct applications of artificial intelligence, research in artificial intelligence may contribute to the understanding of how people think, which includes logically, creatively, and irrationally, and to how people learn. The understanding gained may lead to improvements or enhancements of these processes.

- At the same time, artificial intelligence may alter the way in which man thinks of himself. Some see man's exposure to his weaknesses as a beneficial process. With others there is a concern that by apparently reducing thinking to a mechanistic process, the public perception may be that people are regarded as "nothing but clockwork". This may have "indirect effects on self esteem and social relations [that] could be destructive to many of our most deeply held values" (13).

- The most distinctive and singular thing about AI is that the process of reasoning used to make decisions becomes fundamentally altered in the way it is distributed between people and machines.

- Because people may be directly affected by artificial intelligence applications, and by expert systems in particular, either individually or in some instances in large numbers, there is a special need to assure the quality, verification, safety, and reliability in a wide variety of applications.

- At the present time there does not appear to be any specific mechanism for this in Canada, or elsewhere, other than the professional responsibility of individuals and the codes of ethics that the professions have developed for their members.

- The diversity of applications for artificial intelligence suggests that this decentralized approach to responsibility may be the best to follow, but with special attention drawn to the responsibilities involved and an examination of the codes of ethics, where available, to ensure that these new conditions are adequately represented.

- To assist in the examination of artificial intelligence applications in the social context, a protocol for the examination of case histories or proposed applications has been developed and is presented herein.

- Many other new technologies are currently being introduced besides artificial intelligence. Some of these, such as advanced manufacturing systems are likely to precede artificial intelligence in terms of a very broad scale implementation, and with very widespread social and economic effects. The final impacts of artificial intelligence, whatever they may be, will occur in the presence of these parallel activities, some of which are known, and also in the light of other events which are not known or predictable.

The Industrial Revolution, Revisited

Computers **Thinking**

Relax. Don't Worry. Machines will do the Work.

Section VI

References

Section VI — References

1. Scrimgeour, J., Expert Systems — A Review of their Potential and Application in Canada, *Canadian Conference on Industrial Computer Systems*, May 1986 (principal source).*

2. Peacocke, R., Zlatin, D., Artificial Intelligence at Bell-Northern, *Artificial Intelligence (Expert Systems) Symposium*, NRC, January 27, 1987 (with adaptation).

3. Rosenberg, R., The Impact of Artificial Intelligence on Society, Department of Computer Science, University of British Columbia (manuscript submitted for publication, August 1988).*

4. *IEEE Spectrum*, September 1988 p. 3.

5. CSA Guidelines Address Quality Assurance for Software Developers, *Direct Access*, January 27, 1989.

6. Booth, P.J., Plowright, T.I., *A Study on the Social Impacts of Office Automation*, Department of Communications, Ottawa, Ont. December 1982.

7. Scrimgeour, J., *Towards the Factory of the Future, an Index and Guide to CAD/CAM Technology*, NRC report 24261, March 1985.*

8. Garson, B., *The Electronic Sweatshop: How Computers are Transforming the Office of the Future into the Factory of the Past*, Simon and Schuster, 1988, 288 pp. hardcover.*

9. *CAD/CAM Technology Advancement Council, Strategy for Survival*, Department of Industry, Trade & Commerce, Ottawa, September 1980.*

10. Robertson, D., Wareham, U., *Computer Automation and Technological Change: Northern Telecom*, National Automobile, Aerospace and Agricultural Implement Workers Union of Canada (CAW Canada). January 1988.*

11. Scrimgeour, J., Computer Integrated Manufacturing, *CSME Bulletin*, January 1989.

12. Chisholm, A.W.J., The Social Effects of Intelligent Manufacturing Systems, *Robotics and Computer Integrated Manufacturing (G.B.)* Vol. 3, No. 2, pp. 157–164, 1987.*

13. Boden, M.A., Artificial Intelligence. *The Oxford Companion to the Mind*, 1987.

*Reviewed in the Biblioabstracts Section, Appendix 1.

Appendix 1

A Biblioabstract of Selected References on AI in the Social Context

Appendix 1 — A Biblioabstract of Selected References on AI in the Social Context

A. Biblioabstract of Selected References

H.H. Brune. The Social Implications of Information Processing. Information and Management 1 (1978). pp. 14–156. North-Holland Publishing Company.

The author, with the federal Department of Communications, presents his view that serious long-term social and economic dislocations are likely if Canada fails to align overall social and economic policies with the powerful on-going transition towards an information-based economy and society. Survey data from 1030 households are presented to show that, contrary to popular belief, privacy is not a major issue in the minds of the public, but that fears of unemployment created by automation, errors due to the misuse of data, and depersonalization caused by poor and unresponsive systems design are the major sources of concern.* The author argues that system designers must accept a larger share of responsibility for these factors, while premature legislative measures should be avoided.

* Note — In a later presentation to the Ottawa CIPS/ACM chapter December 1978, the same author included additional survey data in which information processors (i.e., not households) placed the above first three concerns last and in exactly reverse order. The fair information reporting and privacy issues were placed in first and second place. Obviously there is a difference of opinion between the public on one hand, and computer professionals on the other.

G.D. Mackintosh. Software Liability: Are Expert Systems Any Different. Canadian Artificial Intelligence, April 1988. pp. 1–14.

Referring to two incidents where human lives were lost through software malfunctions (one of them dealing with software that originated in Canada), the author raises the question "How does the legal system currently deal with liability for software failure?" It is explained that much depends on whether software is a product or service because the legal system treats product and service liability differently. For expert systems the issues are even more complex because they are becoming learning systems. A "veteran" neural network, for example, may differ substantially from its original state. In a dispute a supplier could argue that since a system was well behaved when it left the factory, it is the client who must have "brought it up" badly. Special considerations are cited with regard to professional systems in which expert systems may eventually out-perform their human counterparts. It is concluded that, even in their infancy, expert systems possess characteristics that set them apart from traditional software and which confuse their legal standing.

C. Bryson. AI Engineering Issues Provoke Debate. Canadian Artificial Intelligence, April 1988. p. 12.

This short article reports on discussion by AI experts at the 1987 CIPS Edmonton annual conference following a presentation by Star Wars reliability critic D. Parnas on the subject of "Why engineers should not use artificial intelligence". The author of this paper is reported to object, not so much to AI as a field of research, but to those who claim that it is a developed technology that should be put immediately to practical engineering use. In rebuttal, experts attributed these claims to overzealous entrepreneurs rather than bona fide researchers and defended AI and computer science on the grounds that there needs to be a more careful use of terminology and a deeper presentation of the issues.

M. Kirby. Legal and Ethical Issues in AI. The Computer Law and Security Report, 1987. Vol. 2, Part 6, pp. 4–7.

Comments on the current strength and rate of development of AI in Victoria, Australia, and states that the eventual goal of all research on AI is to develop computer systems which surpass human capabilities in reasoning, problem solving, sensory analysis, and environmental manipulation. Although some AI commentators do not expect this goal to be achieved within 50 years, the fundamental issue is whether human intelligence and society will be the master or the slave of these developments. The paper discusses the implications for peace and world survival, noting that if expert systems take control of defence preparedness, they may put human life and health at risk on a very large scale if ever they go wrong. As an ethical question, surrendering the entire future of civilization to artificial intelligence is deemed not morally acceptable, at least at this stage.

R.M. Restak. The Brain — The Last Frontier. Warner Books 1979. July 1980.

Presents a review of brain research in recent years, with emphasis on the combination of studies in the behavioral sciences with studies in the brain sciences. This combination is known as psychobiology. According to some authorities, brain research is considered to be the most promising area of research likely to produce award winners for the Nobel Prize by the year 2000.

It could be argued additionally that many aspects of AI research are closely linked to this and may therefore make contributions, in important and crucial areas (which this book proposes), such as the behavioral causes for acts of terrorism.

M.S. Blumenthal. Programmable Automation and the Workplace. IEEE Technology and Society Magazine, March 1985. pp. 10–15.

The author, with the U.S. Congress, Office of Technology Assessment, examines the impact of programmable automation on persons in the workplace. Emphasis is given to the displacement of labour and to the effects of robots which are seen to be concentrated within a few industries and not cataclysmic. It is emphasized that displacement incurs change, but does not necessarily translate into unemployment since this depends on other factors. Demand for technical personnel is expected to increase, and for craftworkers to fall.

C.C. Gotlieb. Artificial Intelligence: A Perspective and a Dilemma. CIPS Review, September/October 1985. pp. 126–128.

Artificial intelligence is defined as the machine simulation of intelligent, i.e., human, behaviour. The widespread interest in this field by computer science students, by researchers, and for use in defence systems as one application area is underscored. The danger foreseen is that in spite of successful demonstrations of AI techniques in specific or narrow domains, "true" intelligence that can adapt to a variety of situations is as far away as ever. Too much is expected too soon. A period of disillusionment and backlash in terms of support for AI may therefore result. The dilemma seen for serious AI researchers is how to take advantage of the support currently available for AI, without promising more than can possibly be achieved.

J. Liebowitz. Common Fallacies About Expert Systems. ACM Computers and Society, Vol. 16, No. 4, pp. 28–33.

Discusses common fallacies concerning expert systems as commonly perceived by the layperson, namely, the myths that expert systems do not make mistakes, that they can learn from their mistakes, that they are deep reasoning systems, that they are easy to build, that individuals can be replaced by expert systems, and others. Suggestions are given for laypersons to become better informed about expert systems by personal study and by taking courses.

J. Scrimgeour. Expert Systems — A Review of Their Potential and Application in Canada. Proceedings, Canadian Conference on Industrial Computer Systems, Canadian Industrial Computer Society (CICS), May 1986.

Following review of an early computer aided engineering program for design synthesis of large synchronous motors that had many, but not necessarily all, of the attributes of modern expert systems, the paper identifies a number of organizations involved in expert systems technology development and lists a spectrum of application areas where expert systems could be applied. It is suggested that tools and software to aid in the construction of expert systems will soon be widely available. By comparison the domain knowledge which is necessary in organized form for application development will be relatively scarce, and that the wisdom to select the most needed applications will be particularly hard to find.

M. Newman. Professionals and Expert Systems. A Meeting of the Minds. ACM Computers and Society, Vol. 18, No. 3, July 1988. pp. 1–27.

While the speed and capacity of computers to store vast amounts of data account for many applications and their current impact in the work place, the writer contends that artificial intelligence and expert systems will directly affect the professions which, until recently, have largely escaped the computer revolution. Questions related to who is responsible for mistakes, the programmer or professional (if he exists), and the establishment of standards of conduct are deemed crucial. Examples are discussed in the legal, accounting, and medical fields.

H.H. Rosenbrock. Robots and People. Control Systems Centre Report No. 592. Fourth Hartley Lecture, November 24, 1981, University of Manchester, Institute of Science and Technology (UMIST).

Having identified control engineering as the "fifth primary technology" alongside civil, mechanical, electrical, and chemical engineering, the writer then focuses on robotics and the social effects of automation. A clear need is seen for relieving men and women of the need to perform work that is unavoidably dangerous, or excessively noisy, dirty or strenuous to the point where it imposes severe stress. On the other hand, automation appears to have created or multiplied the number of jobs or tasks that are trivial and repetitive. Rather than breaking jobs to be automated into fragments, and leaving only those pieces which cannot be automated to men and women, an alternative is proposed in which the tasks requiring the greatest level of human ability and skill would be reserved for people. This is a reversal of roles, in which man continues to be master rather than become servant of the machine. A flexible manufacturing cell which follows these principals is currently being designed at UMIST.

J. Scrimgeour. Towards the Factory of the Future. An Index and Guide to CAD/CAM Technology. NRC report 24261. March 1985.

Chapter 15 of this extended guide to manufacturing and automation technologies increasingly utilized in the factory of the future contains biblioabstracts of thirty references from INSPEC and the CAD/CAM Newsletter dealing with the social implications of CAD/CAM (now known widely as CIM), Ergonomics and Human Engineering. The effects of automation are seen as direct effects on employed workers in the workplace, direct effects on the structure and competitiveness of companies and indirect effects such as social costs on the unemployed in both the industrialized and lesser developed countries. References deal primarily with management, skills changes, retraining, and structural adjustments.

C.S. Lewis. Computerized Manufacturing Automation, Education and the Workplace. U.S. Congress, Office of Technology Assessment, Washington, DC. OTA-CIT-235, 1984, 471 pp. $14.00(U.S.).

A comprehensive review of computerized manufacturing automation is presented with historical perspectives and alternate approaches that might be adopted in the U.S.A. Emphasis is placed on the economy, education, work environment, and technology development. A stronger role for the federal government is advocated.

CAD/CAM Technology Advancement Council. Strategy for Survival. Department of Industry, Trade and Commerce, Ottawa, Ont. September 1980.

Contains discussion of issues and recommendations to Canadian industry, educational institutions, and government pertaining to the adoption of computer-integrated design and manufacturing systems. Early adoption of these technologies is urged if Canadian industry is to remain competitive and the Canadian standard of living is not to be eroded. A special section in the report draws attention to the social implications, including the direct ergonomic effects on workers, plus the midterm and macro effects such as unemployment and the international balance of payments if a competitive manufacturing industry is not maintained. Ultimately, new social structures for the distribution of wealth may be required as automation continues and the number of opportunities to work, in the normal sense, decreases.

A.W.J. Chisholm. The Social Effects of Intelligent Manufacturing Systems. Robotics and Computer Integrated Manufacturing (G.B.), Vol. 3, No. 2, pp. 157–164, 1987.

The discussion in this paper focuses exclusively on the external or "macro" social effects of advanced manufacturing systems while deliberately omitting the micro, internal social effects.

The intelligent manufacturing systems of the future will be highly productive and effective, but they may have undesirable social effects as well as beneficial ones. Possible motivations for developing them are therefore examined against a background of the long, continuing resistance to the use of innovations in production machinery since the start of the industrial revolution. The claim that advanced computer-integrated manufacturing technology has a specially beneficial effect on the prosperity of nations is critically examined, as well as the nature of the beneficial synergism between manufacturing and services in creating wealth. A sombre view is taken of international competition which seems to be leading to an "ultimate" competitive race. The continuing substantial reductions expected in manufacturing employment may be much greater than any compensating employment increase in the service sector, owing to the application of the new technologies in this sector which is only now getting under way. Because full employment in traditional form is now widely seen to be an unattainable goal, a social revolution is required in attitudes to employment, work, and its value. The paper ends with a proposal to consider ways of creating new social and economic structures in order to make this social revolution a practical reality, before social catastrophe occurs. This includes replacement of the "employment" ethic with a "work or usefulness" ethic embracing voluntary unpaid as well as paid work, a reduction of the working life, and other measures.

A. Bundy and Clutterbuck. Raising the Standards of AI Products. IJCAI-85, pp. 1289–1294.

S. Turkle. The Second Self: Computers and the Human Spirit. New York: Simon and Schuster, 1984.

The impact on children, primarily, of computers, by an author with training in psychology. A highly regarded book.

New Technologies in the 1990s: A Socio-economic Strategy. Conclusions and Recommendations of a Group of Experts OECD, Paris. 1988. 23 pp.

This report contains the policy conclusions and recommendations from a full report to be issued later by a group of independent experts working under appointment to the OECD. The starting point of the group's analysis is that technological change is a social process and therefore a need exists for a broad consensus about the impact of new technologies. The report underscores the employment-creating potential of new technologies, particularly the information technologies, while stressing the need for organizational and social adaptation as a dynamic ongoing process if the expected benefits are to be obtained. Disequilibria are likely to result in time and space when patterns of job creation and job destruction do not match in time or geographically. Adaptation and new institutional arrangements are needed by and for workers, managers, labor, organizations, and governments. The need is stressed for education and training throughout life, an active society in which all members can contribute (but with variations in the concept of full employment), the creation of new job opportunities, expanding world markets, and support to the Third World Countries such as Brazil, Mexico, China, and India, who may be next to join the ranks of the newly industrialized countries (NICs). In summary, "Technological innovation and social change must be seen as an integrated process and should be managed as such. The diffusion and management of innovation must be a part of that technological policy."

D. Robertson and J. Wareham. Computer Automation and Technological Change: Northern Telecom. National Automobile, Aerospace and Agricultural Implement Workers Union of Canada (CAW Canada). January 1988.

Northern Telecom is one of the world's largest manufacturers of digital switching telecommunications products and a leader in electronics, computer, and software technology with emphasis on the latter. The 194-page document presents a case study report on the impact of technological change at Northern Telecom conducted under the CAW's Technology Project. Because Northern Telecom is highly advanced in computer-based automation, the study will be like a look into the future for most other firms and organizations. The report describes a wide variety of both office and factory automation projects at three Canadian locations with details on management objectives, system functions, and worker's reactions, both positive and negative, in this highly competitive industry. Throughout the report there is emphasis on fierce competitive pressure, the compelling need for constant innovation and productivity improvement, and the closely associated techno-stress which affects both managers and workers. The current period is seen as a watershed in the history of manufacturing. The most frequently used word is "change". Products are different. Offices are different. Factories are different. Jobs are different and organizational structures are different. It would be an oversimplification to abstract further from this excellent and nearly book-length report other than to underscore the sense

of urgency throughout, particularly with regards to the three key concepts of TQC, time, quality, and costs.

C. Freeman (and other speakers). Information Technology and Employment Conference. IBM International Education Centre, La Hulpe, Belgium, June 27–28, 1985, IBM European International Services Company, Bruselles.

Proceeding from many studies in the mid-1970s on the employment effects of the chip and microprocessor, many of which were overly pessimistic, the author now sees a growing number of studies in which professional economists have become involved. Empirical evidence now is recognized to indicate that even though technology-induced gains in labour productivity may have led to job destruction at the enterprise level and in certain sectors, these have been compensated at the macroeconomic level by the growth of demand and by the emergence of entirely new opportunities for employment elsewhere in the economy. However, time lags and employment diversion prevent an immediate and perfect balance on a local basis. Some investigators believe that since the service sector is now also directly affected by technical change, and will no longer be able therefore to act as a compensatory area of growth, that the industrialized countries are consequently facing a prolonged period of mass unemployment.

Information technology is viewed by the principal author and his co-workers as a new techno-economic paradigm effecting the management and control of production and service systems throughout the economy. It is regarded as a technological revolution requiring structural and social adaptation and with the potential to unleash a new era of high economic growth in the OECD economies.

R.S. Rosenberg. The Impact of Artificial Intelligence on Society. Manuscript submitted for publication in AI and Society.

This paper presents an introduction to the social issues related to AI applications by discussing four typical applications (expert systems, image processing, robotics, and natural language understanding). Four social issues are chosen for discussion and illustration: the impact of AI (and automation) on work, i.e., the nature of future work and employment levels; the role of machine implemented decision making; a potential threat to individual privacy from the increased use of computer data bases, computer profiling and computer matching; and finally, questions of social organization as society and its political institutions adapt to a future in which something else may need to replace work as a means of wealth distribution and as the basis for determining the self-worth of the individual.

T. Winograd and F. Flores. Understanding Computers and Cognition: A New Foundation for Design. Norwood, NJ: Ablex Publishing, 1986.

This book represents a departure from Winograd's previously long-held beliefs about AI. He (and Flores) take a holistic approach to AI and its impact on society.

T.D. Sterling. Democracy in an Information Society. Information Society, Vol. 4, Number 1/2. 1986. Crane Russak & Company. (The author is at the School of Computing Science, Simon Fraser University, B.C.)

The theme paper by Sterling in this 143-page issue of "The Information Society" discusses factors and implications of the growing use of computers and information. His paper is followed by comments from nine other authorities plus a rejoinder by Sterling. Understandably many issues are discussed, ranging from whether information technologies may undermine democratic political forms, alter the balance of political power, lead to increased centralization, provide new information processing capabilities for individuals or small groups, provide conveniences such as library notices and welfare checks or, by extension, turn into surveillance systems. The possible destabilizing effect of large scale technological unemployment on the politics of democratic countries is considered. All in all, Sterling sees computer-based technologies as tending to have a slowly evolving anti-democratic political impact, but that it is pre-mature by far to fear the demise of the democratic order.

T.D. Sterling. Analysis and Reanalysis of Shared Scientific Data. Annals of the American Academy Shared Scientific Data, 495, January 1988.

Knowing that enormous strides have been made in the ability to record, store, retrieve, manipulate, and communicate data, this investigative paper considers instances in which this may, or may not, lead to an increase in the use of shared data files amongst scientific investigators. Instances where the use of shared data has been found widely acceptable include the public use of data tapes containing health, census data, or of the comprehensive bibliographic files used in literature searches. However, the sharing of data between individual investigators is much less common. Resistance to the granting of access includes the possible loss by the holder to monetary, political, or psychological reward, a fear of disagreement or of interpretation of the data, and even the potential exposure of extreme bias or fraud. Technical problems also exist, such as the cost and the time required to become adequately familiar with how the data are stored, how they need to be accessed, and how they can be manipulated. Data sharing between scientists depends on individual person-to-person arrangements and is considered unlikely to flourish on a broad scale.

B. Garson. The Electronic Sweatshop: How Computers are Transforming the Office of the Future into the Factory of the Past. Simon and Schuster, 1988. 288 pp.

From a series of interviews and anecdotal accounts with employees in many organizations, ranging from McDonalds Hamburgers to stockbrokers at Merrill Lynch, the writer see a gradual and continuing change of work practices involving electronic pacing, performance monitoring, surveillance, and standardization of output in the service industries and even in the professions. This is predicted to gradually embrace managerial positions and eventually executive levels as well. As an example, reservation clerks at American Airlines have had the clean-up time between the end of one call and the beginning of the next call automatically assigned to them reduced to thirteen seconds. Employees at a Canadian airline using the same software have only been reduced to sixteen seconds in their latest union contract, and their union flatly opposes individual monitoring. Electronic monitoring is reported to be widely practised in a majority of organizations where networked computers or terminals are installed. Even in the military, computers first introduced for record keeping and ordinance supply logistics have become command and control systems for measurement and strategic decision making at locations far removed from the scene of battle. Many examples are cited, and while the writer believes it still possible to choose the computer but not the sweatshop, this choice is deemed unlikely to happen.

S.J. Frank. What AI Practitioners Should Know About the Law. Part 1. AI Magazine, Spring 1988. pp. 6–75.

The three main methods used for the protection of intellectual property are reviewed according to law and practise in the U.S.A. These are copyright, patents, and trade secrets. The review discusses the applicability of these to computer software in general, and even more specifically to AI, which may entail special considerations. Since they are idea-driven, knowledge-based systems seem to be poor candidates for copyright protection. Mathematical formula or algorithms tend to be excluded from the subject matter accepted for patentability and ideas themselves are stated to never be patentable subject matter. For AI this tends to leave trade secret protection as the remaining method. This is considered unfortunate because it is the one method that does not involve disclosure, and therefore tends to inhibit the sharing of ideas deemed necessary for progress.

S.J. Frank. What AI Practitioners Should Know About the Law. Part 2. AI Magazine, Summer 1988. pp. 109–114.

The second part of this two-part article discusses two additional questions: first, the issue of tort liability as it applied to assigning financial compensation to victims of physical, economic or emotional injury, and secondly the use of computers in the court room. Tort liability tends to be different for products than it is for services. The liability for products tends to be stricter because product sellers are regarded as more able to spread a loss over their larger number of customers. Because of the close liaison that normally exists between expert system developers and users, they are likely to be regarded as being closer to the provision of a service than as a (mass) marketed product. In cases of use in the professions, such as a diagnostic system used by a physician, the physician is likely to be deemed responsible for any professional misjudgment, for otherwise the expert system could be claimed to be performing in the profession, which is not allowed. What will happen when expert systems exceed the ability of human practitioners is not clear. The admissibility of computers and expert systems as expert witnesses in court hearings is also discussed, but deemed unlikely since they would be unable to withstand convincing realtime cross examination.

M.R. La Chat. Artificial Intelligence and Ethics: An Exercise in the Moral Imagination. The AI Magazine, Summer 1986. pp. 70–79.

The possibility of constructing a personal AI raises many ethical and religious questions that have only been dealt with seriously by imaginative works of fiction; they have largely been ignored by technical experts and by philosophical and theological ethicists. Arguing that a personal AI is possible in principle and that its accomplishment could be adjudicated by the Turing test, the article suggests some of the moral issues involved in AI experimentation by comparing them to issues in medical experimentation. Finally, the article asks questions about the capacities and possibilities of such an artifact for making moral decisions. It is suggested that much a priori ethical thinking is necessary and that such a project cannot only stimulate our moral imaginations, but can also tell us much about our moral thinking and pedagogy, whether or not it is ever accomplished in fact.

J. Weizenbaum. Not Without Us. Computers & Society, Vol. 16. 1986.

In this abbreviated English translation of a talk given in Germany by a noted computer scientist from the U.S.A., a call is made to computer professionals, individually and collectively, to recognize that they have the obligation and ability to halt the development of weapons systems in general, and in particular, nuclear weapons systems capable of mass destruction. The rational that researchers cannot control the applications to which their discoveries will be put and therefore are not responsible for the result, or the rational "If I do not do it, someone else will", are not considered as acceptable.

A. Borning. Computer System Reliability and Nuclear War. Communications of the ACM, February 1987. Vol. 30, No. 2, pp. 112–131.

How dependent should society be on computer systems and computer decision making? What are the cost-benefit trade-offs between the advantages of computerization (greater efficiency, speed, precision, and so forth), and the jeopardy we are in when critical computer systems break down or otherwise fail to meet our intentions? Given the devastating consequences of nuclear war, it is appropriate to look at current and planned uses of computers in nuclear weapons command and control systems, and to examine whether these systems can fulfil their intended roles. The author explores these questions and issues, emphasizing errors of data interpretation that have already occurred in the past, problems associated with correct systems specification, design, and testing and the decreasing decision and responses times associated, for example, with a "launch-on-warning" policy or with the short flight times of submarine launched missiles.

M.M. Waldrop. A Question of Responsibility. The AI Magazine, Spring 1987. pp. 29–36.

The writer sees robots, not just as automated manipulators according to their current industrial definition, but as including all information systems and networks intended to serve humanity. Starting from the three well-known Laws of Robotics, postulated as engineering safeguards by Asimov, he shows how even these can involve decision conflict. As intelligent and decision-making systems proliferate and become embedded in many applications (just as electric motors are today), they will frequently carry with them a value structure or set of assumptions not necessarily apparent at first evaluation. This is most sharply illustrated by the "launch-on-warning" strategy of U.S. defence policy.

D. Bellin and G. Chapman (editors). Computers in Battle — Will They Work? Harcourt Brace Jovanovich Publishers, Boston 1987. 362 pp. ISBN 0-15-121232-5.

This publication of eleven papers and essays has been largely collected under the auspices of the Computer Professionals for Social Responsibility (CPSR) in view of the possibility that computer scientists may now be at the same threshold with regard to the impact of their work on society that nuclear physicists occupied in the 1940s. Contributions provide an overview on the use of computers in battle, concerns for system reliability, the role of software, and the near impossibility of achieving total systems reliability in large scale systems, especially where evolutionary development and verification are precluded by the nature of the application. The book concludes with a list and description of organizations of scientists, such as the CPSR, who are concerned about the arms race, and others for the evaluation of technology.

T. Athanasiou. High-Tech Politics: The Case of Artificial Intelligence. Socialist Review, No. 92, 1987. pp. 7–37.

This paper reviews artificial intelligence and knowledge systems from the social impact point of view. The author seeks to place realistic but knowledgeable bounds on what computers can and cannot do. Knowledge systems are defined more as "competent systems" rather than "expert systems". They are seen as continuing to be special purpose systems. Their use is seen as restricted to clearly bounded domains where it is possible to capture human skill in the form of declarative rules since they lack any common sense or any general sense of the real world. Concern is expressed, however, with regard to their increasing use. This may nurture bureaucracy and encourage standardization in decision making as well as influencing other outcomes of a social nature. Concerns are expressed, from a safety point of view, with regard to the use of knowledge-based systems for realtime application in process control, as well as for the use of AI in military decision making where human beings may be "out of the loop".

Z.W. Pylyshyn. Computers, Knowledge & the Human Mind. In Creativity and Liberal Learning, D.G. Tureck (ed). Ablex Publisher. 1987.

The intellectual significance of recent developments now occurring in computer science suggest that history will someday record that many classical problems of philosophy and psychology were transformed by a new notion of process, namely that of symbolic or computational process.

Just as the first industrial revolution was founded on the dual principles of division of labour and the interchangeability of parts, it is believed that fundamental changes will result from the basic characteristics of computers and recent developments in artificial intelligence. The versatility (plasticity) of computers makes it possible to handle complex problems through a hierarchical systems design called layering. The building up of layers allows specialists to contribute, while the suppression of detail permits ordered growth and a means for organizing complexity. Since artificial intelligence machines exhibit behaviour which is dependent on the knowledge encoded in them, they therefore represent a historical departure from any previous non-living object.

B. Publications of a Continuing or On-Going Nature

Computers and Society, Quarterly newsletter of the Special Interest Group on Computers and Society (SIGCAS) of the ACM.

Technology and Society, A quarterly magazine published by the IEEE Society on Social Implications of Technology (SSIT).

AI Magazine, Quarterly magazine of the American Association for Artificial Intelligence.

AI & Society, The Journal of Human and Machine Intelligence", Karamjit S. Gill (ed.), SEAKE Centre, Faculty of Information Technology, Brighton Polytechnic, Moulescoomb, Sussex BN2 4GJ, UK.

Appendix 2

Terms of Reference: NRC Associate Committee on Artificial Intelligence

Appendix 2 — Terms of Reference: NRC Associate Committee on Artificial Intelligence

Subcommittee on the Social Context of Artificial Intelligence

Objectives and Activities

1. To conduct a brief literature search and prepare a bibliography of references pertinent to the scope of the sub-committee.
2. To develop a brief report or working paper to identify typical issues and areas of concern.
3. To develop a short list of sub-committee members and a larger network of interested individuals with knowledge and expertise in the subject area.
4. To hold at least one meeting or consultative workshop to review material available, areas of concern, available expertise, and develop guidelines for the avoidance of problem areas.

Adopted: February 16, 1988

Members

Dr. R. De Mori, CRIM & McGill University, Montreal, P.Q.*

Professor C. Gotlieb, University of Toronto, Toronto, Ont.†

Dr. J. Ham, University of Toronto, Toronto, Ont.‡

K. Hayes, Canadian Labour Congress, Ottawa, Ont.†

Dr. C. Lajeunesse, Centre de recherche informatique de Montréal (CRIM), Montreal, P.Q.†

Dr. A. Mackworth, University of British Columbia, Vancouver, B.C.†

A. Mayman, National Research Council of Canada, Ottawa, Ont.†

Professor R. Rosenberg, University of British Columbia, Vancouver, B.C.†

J. Scrimgeour, National Research Council of Canada, Ottawa, Ont.†

G.F. Sekely, Canadian Pacific, Toronto, Ont.†

Dr. B. Smith, Acquired Intelligence, Victoria, B.C.†

G. Thomas, Price Waterhouse, Ottawa, Ont.*

Professor W. Vanderburg, University of Toronto, Toronto, Ont.†

---

* As of October 1988.
† As of May 1988.
‡ Chairman.

Appendix 3

Organizations and Centres with Activities Related to the Social Context of Artificial Intelligence

Appendix 3 — Organizations and Centres with Activities Related to the Social Context of AI

IFAC Technical Committee on the Social Effects of Automation

The IFAC TC on the Social Effects of Automation with 100 members is the largest IFAC technical committee, and becoming also more active. For example, 50 people from 15 countries attended their workshop in Karlsruhe, Germany, in September 1986, dealing with the skills needed for work in flexible manufacturing systems, in flexible assembly systems, and in computer-aided design.

In addition to special conferences and workshops, the committee arranges for sessions on the social effects of automation to be included among the 115 sessions and 500 papers which are typical of the IFAC Congress, held every three years.

The current chairman of the IFAC TC on the Social Effects of Automation is:

Dr. T. Martin
Kernforschungszentrum
Karlsruhe
Postfach 3640
D-7500 Karlsruhe 1
Federal Republic of Germany

Swiss National Research Program

This program focusses on the "Humanization of Work and Technological Development" and the training of informatics students in the minor subject of labour science at the Zurich Technical University.

ref. Martin, Ulich, Warnecke, IFAC 87

Profile 21 Project — The Role of the Manufacturing Professional in the year 2000

Papers were given on this topic at the 1988 Autofact Conference October 31 – November 2 in Chicago. This included results of a 15-month study commissioned by the Society of Manufacturing Engineers and undertaken by A.T. Kearney, Inc. in Chicago to identify what new skills and education the manufacturing engineer will require by year 2000, what the typical manufacturing workplace will be like, how global competitiveness will affect work, life styles, etc.

Socioeconomic Aspects of Robotics and Integrated Manufacturing (SEARIM)

SEARIM is a special interest group of faculty and students involving manufacturing engineers, computer engineers, and social scientists. SEARIM's research focus is (1984) on the factors that tend to retard the implementation rate of new technology, the enhancement of job creation and remedies for workers displaced by technology. SEARIM is currently (1984) involved in a human development case study at the Ford Motor Company's Dearborn plant. SEARIM is a unit of the Program in Urban, Technological and Environmental Planning. Inquiries should be directed to:

The SEARIM Project
Program in Urban, Technological, and Environment Planning
The University of Michigan
506 East Liberty
Room 218, Carver Building
Ann Arbor, MI 48104
Telephone: (313) 763-4190

Control Systems Centre

University of Manchester Institute of Science and Technology (UMIST)
P.O. Box 88
Manchester, M60 IQD
England

Dr. H.H. Rosenbrock, closely associated with this institute, is a strong advocate of the principal that the information and automation technologies should be applied in ways which will enhance the creative abilities of the user and not, for example, used to deskill the operator by reducing him/her to the role of machine tender.

Centre for Society, Technology and Values

Centre for Society, Technology and Values
University of Waterloo
Waterloo, Ontario
N2L 3G1

e.g., See Social Issues Conference, Canadian Artificial Intelligence, July 1988, pp. 29–31.

Centre for Technology and Social Development

Dr. W. Vanderburg, Director
Centre for Technology and Social Development
Faculty of Applied Science & Engineering
Rosebrugh Building
University of Toronto
Toronto, Ontario
M5S 1A4

Netherlands Organization for Applied Scientific Research (TNO)

ref. Ir. J.L. Remmerswaal

Reported to be continuing the leadership of the CIRP (International Institution for Production Engineering Research) interdisciplinary research program on the computer-integrated factory initiated by Dr. Eugene Merchant and previously lead by Professor A.W.J. Chisholm, University of Salford, U.K.

Office of Technology Assessment, U.S. Congress, Washington, DC

Although no specific reference has come to our attention to date, this organization is almost certain to have issued some evaluation reports on AI and expert systems, and to be studying the field in an on-going manner.

Centre for Social and Economic Issues

L.G. Tonatzky
Director, Centre for Social and Economic Issues
Industrial Technology Institute
Ann Arbor, Michigan

Author of the paper Learning to Use Computers That Think, which stresses the need for training in CIM.

Systems Evaluation Centre Admiral Management Services

Camberley, U.K.

The centre operates a software testing and evaluation service for users and suppliers of systems where safety, security, or high quality are considered to be crucial. Examples of safety critical applications given are the control of air traffic, nuclear power stations, fly-by-wire aircraft, military systems, health care, process, and manufacturing plants.

National Automobile, Aerospace and Agricultural Implement Workers Union

205 Placer Court
North York, Willowdale
Ontario
M2H 3H9

Under their CAW Technology Project, a case study report has been developed on Computer Automation and Technological Change: Northern Telecom. For a review of this report see the preceding appendix of reference publications.

Technology and Work Program

Charles Richardson
Director, Technology and Work Program
University of Lowell
Lowell, MA 01854

This organization has prepared case study reports on the introduction of new technologies and their effects on the work force at two locations as follows:

- The Boston Herald (newspaper)
- The G.E. Factory of the Future, Lynn, MA

Computer Professionals for Social Responsibility (CPSR)

CPSR is an organization of computer professionals committed to informing both policy-makers and the public about the appropriate uses and potential misuses of computer technology. Ongoing activities include studies of the Strategic Defense Initiative, the Strategic Computing Initiative, challenges to civil liberties, and the funding of computer science research.

Contact:

CPSR, Inc.
P.O. Box 717
Palo Alto, CA 94301
Telephone (415) 322-3778

TC9 of IFIP — Relationship between Computers and Societies

Technical Committee number 9 (TC9) of the International Federation of Information Processing Societies is concerned on a continuing basis with the relationship between computers and societies.

The Canadian contact established by CIPS for TC9 is:

Professor A. Clement
Department of Computer Science and Mathematics
York University
North York, Ontario
M3J 1T3
Telephone (416) 736-5232

COMP.RISKS

COMP.RISKS is a news group concerned with computer risks moderated by P. Neumann of Stanford Research Institute. It typically has 2–3 items of interest per week. COMP.RISKS may be accessed on most computer networks, including ARPANET, Internet, Usenet, Bitnet etc.

On those networks or installations that will accept the command rn for "Read News" this may be accessed as rn COMP. RISKS. In case of difficulty consult with your local computer network support group for advice on procedures for access.

Similar information is published quarterly in hard-copy form in ACM "Software Engineering Notes", which is an informal publication of the ACM Special Interest Group on Software Engineering (SIGSOFT), available from the ACM Order Department, P.O. Box 64145, Baltimore, MD 21264. To submit material, camera ready or on-line, contact:

NEWSLETTER EDITOR:
Peter G. Neumann
SRI International, BN168
Menlo Park, CA 94025
Telephone (415) 859-2375
(Newman@CSL.SRI.COM)

Canadian Workplace Automation Research Centre

The centre, which operates under an advisory board formed by the federal Minister of Communications, operates primarily to provide leadership in computerized office systems, to synthesize the needs of users, to become a focal point for information exchange in the field of workplace automation, and to foster cooperation between experts and client groups. A catalog of study reports for the 1986–1988 time period, available in French or English, is available. Contact:

Canadian Workplace Automation Research Centre
1575 Chomedey Boulevard
Laval, Quebec
H7V 2X2
Telephone (514) 682-3400

or

Information Services
Department of Communications
300 Slater Street
Ottawa, Ontario
K1A 0C8
Telephone (613) 990-9400

Stanford Research International (SRI), Palo Alto, CA

Xerox Palo Alto Research Center (PARC)

Individuals at Universities and Educational Institutions in Canada

- University of British Columbia (R. Rosenberg (computer science), Ed Levy (philosophy), Joe Smith (law))
- Simon Fraser University (T. Sterling, Margaret Benston, computing science)
- York University (Andrew Clement, computer science)
- Queen's University (Elia Zureik, sociology)
- University of Waterloo (Robin Cohen, computer science)
- Ontario Institute for Studies in Education (Professor David Olsen, literacy, competence, expertise)
- University of Toronto (C. Gotlieb, computer science, Joanne Marshall, Faculty of Library and Information Science)
- University of Toronto AI Group (Professors Hector Levesque, John Mylopoulos, Ray Reiter, John Tsotsos, and Graeme Hirst)

All have shown interest in social aspects of AI on different occasions, and the work of this committee has been discussed with some of them.

Individuals at Universities in the United States*

- Stanford University (Terry Winograd, computer science)
- University of Washington (Alan Borning, computer science)
- City University of New York (Abbe Mowshowitz, computer science).
- University of California at Irvine, CA 92717 (Robert Kling, Department of Information and Computer Science)
- University of Southern California (Dr. Ann Majchrzak, Institute of Safety and Systems Management)
- Lindenwood College, St. Charles, MO (J.F. Pooley, Mathematics and Computer Science Department)

Unions

Ontario Federation of Labour
15 Gervais Drive
Don Mills, Ontario
M3C 1Y8

Canadian Union of Public Employees
(Marc Bélanger)
21 Florence Street
Ottawa, Ontario
K2P 0W6

Canadian Labour Congress
2841 Riverside Drive
Ottawa, Ontario
K1V 8X7

Quality of Working Life Centre

Hans von Beinum
Ontario Ministry of Labour
400 University Avenue, 15th Floor
Toronto, Ontario
M7A 1T7

ACM special Interest Group on Computers and Society (SIGCAS)

ACM
11 West 42nd Street
New York, NY 100366

Westminister Institute for Ethics and Human Values

London, Ontario
(Abbyan Lynch)

Centre for Professional and Applied Ethics

University of Manitoba
(Arthur Schafer)

Registered Nurses Association of Ontario

(Margaret Keating, Diana Dick)

* A recent survey of the ACM Special Interest Group on Computers and Society (SIGCAS] which has approximately 800 members, showed that 37% of members were also members of the Computer Professionals for Social Responsibility (CPSR) organization. 37% were also identified as being members of the IEEE Computer Society, thus indicating that the degree of interlocking and cross communication is quite substantial. The SIGCAS survey identified "ethics related to users and professionals, develop curricula" as clearly having the highest priority for attention.

Appendix 4

Definition (by Example Publications) of the Diversity, Domain, and Intent of AI

Appendix 4 — Definition (by Example Publications) of the Diversity, Domain, and Intent of AI

The field or domain of artificial intelligence is hard to define, partly because it is pervasive and can enter into or become a part of many other fields of endeavour, yet its presence and contribution has a distinct and unique flavor.

This is illustrated by the attached brief list of references.

If the definition of physics is "Physics is What Physicists Do", then in the final analysis one may also have to say, "AI is whatever you make it", or "AI is what AI people do".

All of the references which follow have been listed by INSPEC in their publication "Key Abstracts in Artificial Intelligence". While formal definitions of "Artificial Intelligence" are useful, one of the best ways to gain a more complete appreciation of what it involves, is by gaining familiarity with AI projects and publications on a world wide basis. The following pages, reprinted with permission from Canadian Artificial Intelligence January 1989, permit an initial understanding of this nature, and are a regular feature of that publication.

World Watch

World Watch on AI Applications and Development is sponsored by the National Research Council's (NRC) Associate committee on AI. Based on the information provided in the abstracts, the references provided have been selected by the secretariat of the NRC Associate Committee on AI as a representative sample of interest and value to Canadian industry. Abstracts provided are reprinted from "Key Abstracts in Artificial Intelligence" with permission from INSPEC. INSPEC is widely recognized as the leading English-language database covering the published information in the field of physics, electronics and computing. Information contained in the INSPEC services is collected on an international basis from over 4,000 Journals and 1,000 Conference Proceedings. INSPEC is a division of the Institution of Electrical Engineers, Station House, Nightingale Road, Hitchin, Herts, UK. All INSPEC's products and services are available in North America from the INSPEC Dept. IEEE Service Centre, 445 Hoes Lane, P. O. Box 1311, Piscataway, NJ 08855-1331, USA.

Persons wishing to obtain copies of references cited should contact their nearest technical library or the Canada Institute for Scientific and Technical Information (CISTI), NRC, Building M-55, Montreal Road, Ottawa, Ontario K2A 0S2, (Phone: (613) 993-1585, Telex: 053-03115). For on-line ordering, CAN/OLE users may use the CAN/DOC command. Envoy users type "COMPOSE CISTI."

1.0 THEORETICAL ASPECTS

1254 The Perceptron and AI: a New Old Way Forward?
D. Ellison
Dundee Coll. of Technology, UK
Comput. Bull. (UK), vol. 4, pt. 1, pp 27-30, March 1988

During the early days of artificial intelligence, Perceptrons were studied intensively. These "brain models" were loosely based on neural mechanisms and there was great hope that using the serial (Von Neumann architecture) digital computers of the day large leaps forward in AI would be made. Unfortunately, they went out of favour after several elementary limitations of Perceptrons were discovered. There is a growing resurgence in the "neural" based approach; as a result of improved experimental techniques, more is known about the detailed structure and functioning of neurons and brains; parallel computers are becoming generally available; and with the VLSI technology, it is feasible to put something akin to simple neural nets on chips at reasonable cost; and, finally, the mathematical theories of learning and distributed memory and processing have improved. (18 refs.)

1262 Expressiveness and Tractability in Knowledge Representation and Reasoning
H.J. Levesque
Dept. of Computer Science, U. of Toronto
Toronto, Ontario, Canada
R. J. Brachman
Computational Intelligence (Canada), vol.3, no.2
pp 78-93, May 1987

A fundamental computational limit on automated reasoning and its effect on knowledge representation is examined. Basically, the problem is that it can be more difficult to reason correctly with one representational language than with another and, moreover, that this difficulty increases dramatically as the expressive power of the language increases. This leads to a tradeoff between the expressiveness of a representational language and its computational tractability. It is shown that this tradeoff can be seen to underlie the differences among a number of existing representational formalisms, in addition to motivating many of the current research issues in knowledge representation. (37 refs.)

1299 Expert System Models for Inference with Imperfect Knowledge: A Comparative Study
A. Goicoechea, A.P. Sage, D. A. Schum
Sch. of Inf. Techno. & Eng., George Mason U.
Fairfax, Virginia, USA
Proceedings of the 1987 International Conference on Systems, Man, and Cybernetics (Cat. No. 87CH2503-1), Alexandria, VA, USA, 20-23 Oct 1987 (New York, NY, USA: IEEE 1987), pp 559-63 vol.2

A detailed comparative study is presented of six leading methods for reasoning: Bayes' rule, Dempster-Shafer theory, fuzzy set theory, the MYCIN model, Cohen's system of inductive probabilities, and a class of nonmonotonic reasoning methods. Each method is presented and discussed in terms of theoretical content, a detailed numerical example, and a list of strengths and limitations. The same numerical example is addressed by each method, so that the assumptions and computational requirements that are specific to each method can be highlighted. Guidelines are offered to assist in the selection of the method that is most appropriate for a particular problem. (12 refs.)

1337 Knowledge Acquisition by Inductive Learning from Examples
J. Selbig
Dept. of AI, Central Inst. of Cybern. & Inf. Processes
Berlin, East Germany
Analogical and Inductive Inference. International Workshop All '86 Proceedings, Wendisch-Rietz, East Germany, 6-10 Oct 1986 (Berlin, West Germany: Springer-Verlag 1987), pp 145-63

Before the author describes his approach to the problem of learning the action part of IF (pattern) THEN DO (action) rules, he gives a survey of the problem of knowledge acquisition for expert systems. In connection with this work, he focuses on automatic knowledge acquisition by learning methods. (26 refs.)

1509 A Selected Artificial Intelligence Bibliography for Operations Researchers
B. Jaumard
Rutgers U., New Brunswick, New Jersey, USA
Peng Si Ow, B. Simeone
Ann. Oper. Res. (Switzerland), vol. 12, no. 1-4, pp 1-50, 1988

The authors have compiled a selected, classified, and annotated artificial intelligence bibliography specifically addressed to an operations research audience. The bibliography includes approximately 450 references from the areas of search (including heuristics and games), automatic deduction (including theorem proving, logic programming, and logical aspects of databases), planning, learning, and knowledge-based systems (with numerous specific applications to management, engineering, science, medicine, and other fields). They have also added a general references section, as well as a special section on artificial intelligence/operations research interfaces. (457 refs.)

1527 Comments on "A Critique of Pure Reason", and Reply
D. McDermott, J.F. Allen, H.A. Kautz, D.G. Bobrow, M.J. Stefik, K.A. Bowen, R.J. Brachman, E. Charniak, J. de Kleer, J. Doyle, K.D. Forbus, P.J. Hayes, C. Hewitt, G.E. Hinton, J.R. Hobbs, D. Israel, R. Kowalski, J. McCarthy, V. Lifschitz, R.C. Moore, N.J. Nilsson, A. Pentland, D. Pole, R. Reiter, S.J. Rosenschein, L.K. Schubert, B.C. Smith, M.E. Stickel, W.M. Tyson, R. Waldinger, T. Winograd, W.A. Woods
Computational Intelligence (Canada), vol.3, no. 3, pp 161-237, Aug 1987

McDermott (ibid., vol.3, no.3, p.151-60, 1987) argued that attempts to formalize commonsense knowledge using first-order logic had failed and that the logicist approach was based on a false premise. The other authors, except Woods, defend logicism and criticize McDermott's views in various ways in 26 papers. McDermott's reply is in two parts; a clarification of his earlier arguments, and a discussion of the nature of knowledge representation and whether AI programs differ from any others in this regard. Finally, Woods argues that logic was designed for other types of knowledge representation, and that logicists try in effect to use a good tool for an inappropriate job. He then develops an account of procedural semantics, arguing that it is a better foundation for commonsense knowledge formalism.

1774 A Code of Professionalism for AI
B. Whitby
AISB Q. (UK), no. 64, pp 9-10, Spring 1988

As AI becomes important as a commercial product, a change in attitudes to professionalism may be required. Commercial AI cannot simply adopt the informal, experimental nature of academic AI. Given the need to attempt some more towards professionalism, the publication of a code of conduct seems an easy and effective first step. Obviously to be of value a code should be acceptable to the majority of AI practitioners — a difficult objective, given the diverse nature of the field. The publication of a code of conduct should be a catalyst in the move towards professionalism in AI. It is only after a code is acceptable to those who will actually have to work within it that it will have any real meaning. Before that point it is merely a discussion document. (3 refs.)

1793 OBSERVER: a Probabilistic Learning System for Ordered Events
K.C.C. Chan
Dept. of Syst. Design Eng., U. of Waterloo, Ont., Canada
A.K.C. Wong, D.K.Y. Chiu
Pattern Recognition, 4th International Conference. Proceedings, Cambridge, UK, 28-30 March 1988 (Berlin, West Germany: Springer-Verlag 1988), pp 507-16

Given a sequence of observed events which are ordered with respect to time or position and are described by the coexistence of several discrete-valued attributes that are assumed to be generated by a random process, the inductive prediction problem is to find the probabilistic patterns that characterize the random process, thereby allowing future events to be predicted. This paper presents a probabilistic inference technique for solving such a problem. Based on it, a learning program called OBSERVER has been implemented. OBSERVER can learn, inductively and without supervision, even if some observed events could be erroneous, occasionally missing, or subject to certain degrees of uncertainty. It is able to reveal the patterns and regularities inherent in a sequence of observed events and cannot only specify, in a clearly defined way, the happenings in the past but also gain insight for prediction. The proposed technique can be applied to solve different problems in artificial intelligence and pattern recognition where decisions concerning the future have to be made.

2.0 SYSTEMS AND TECHNIQUES

1358 Systems Analysis Techniques for the Implementation of Expert Systems
D. R. Martinez
Arco Oil & Gas Co., Dallas, Texas, USA
M. G. Sobol
Inf. Softw. Technol. (UK), vol.30, no. 2, pp 81-8, March 1988

Knowledge-based expert systems incorporate human expert knowledge with the use of computer systems. Today, the wide availability of expert systems shells allows the knowledge engineer to implement specific rules for a desired application. The final design of the artificial intelligence system is the outcome of a detailed study in an organized fashion. The application of system analysis tools for designing knowledge-based expert systems is presented. The paper illustrates the application of the tools with a simplified example drawn from the oil and gas exploration business. The use of a systematic approach in designing expert systems should help the knowledge engineer clearly identify the facts and rules representative of the acquired human knowledge. (15 refs.)

1360 Trends in Artifical Intelligence
I/S Anal. (USA), vol. 26, no. 2, pp 1-12, Feb 1988

Until recently, AI applications were generally thought to require special hardware and software and work best in stand-alone applications. Recently, however, that belief has been reversed. Now, many of the benefits of AI are seen to come from using it to extend and enhance conventional information systems. The report looks at two areas in AI that are likely to become integral parts of many information systems — expert systems and natural language processing systems. In addition, tools used to develop expert systems may foretell features of future, all-purpose programming environments. (14 refs.)

1379 Management of Uncertainty in a Medical Expert System
D.L. Hudson, M.E. Cohen
California U., San Francisco, California, USA
International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems, Paris, France, 30 June-4 July 1986, Berlin, West Germany: Springer-Verlag 1987, pp 283-93

The use of uncertainty in a rule-based expert system for the analysis of chest pain is discussed. The system, EMERGE, has been evaluated retrospectively and prospectively and has been found to perform extremely well. The original system has been altered to handle degrees of presence of symptoms and variable contribution of antecedents. It also utilizes a logical construct which generalizes traditional AND/OR logic. (26 refs.)

1391 Toward Considering Psychological Measurement Issues when Developing Expert Systems
L.Adelman
Dept. of Inf. Syst. & Syst. Eng., George Mason U.
Fairfax, Virginia, USA
Proceedings of the 1987 International Conference on Systems, Man, and Cybernetics (Cat. No. 87CH2503-1), Alexandria, VA, USA, 20-23 Oct 1987 (New York, NY, USA: IEEE 1987), pp 1044-8 Vol.3

There are four possible sources (or determinants) of the quality of different components of the knowledge base contained in an expert system: the domain experts, the knowledge engineers, the methods used by the knowledge engineers when working with the domain experts, and the problem setting and structure used to facilitate knowledge elicitation. A study has been undertaken to assess the extent to which system validity is affected by variability of any of one of these. Toward this end, the results of reanalyzing the data from an experiment varying domain experts, knowledge engineers, and elicitation methods when developing multi-attributed hierarchies is presented. No significant effects were obtained for elicitation method or knowledge engineer, and the results suggest that domain experts were the largest source of variation. (12 refs.)

1415 Experiences with OPS5 in the Development of Expert Systems at Bayer
F. Biegler-Konig
Bayer AG, Leverkusen, West Germany
Expert Systems in Production Engineering. Proceedings of the International Workshop, Spa, Belgium, 18-22 Aug 1986, (Berlin, West Germany: Springer-Verlag 1987), pp 157-61

Techniques of AI are supposed to be one of the most important technologies of the future. Great interest has been shown in the use of expert systems, but much effort must still be invested to get more experience in their development and applications. At

Bayer some studies were made about the advantages of the application of expert systems in chemical engineering. Many areas of possible applications were located, especially for small, compact and easy to use systems with connections to conventional software (Fortran, graphics). Small systems only need a limited amount of development effort and are therefore better suited to gain experiences and assess the power of the new technology. (4 refs.)

1549 A Company/University Joint Venture to Build a Knowledge-Based System
J.R. Weitzel, K.R. Andrews
South Carolina U., Columbia, South Carolina, USA
Manage. Inf. Syst. Q. (USA), vol. 12, no. 1, pp 23-34, March 1988

A joint venture between a university-based research institute and a health insurance company to build a knowledge-based system to perform medical review of health insurance claims is described. The article examines the impact of the differing cultures of the company and the university on the cohesion in the joint knowledge engineering group. It also examines the current literature on the development cycle for building knowledge-based systems as the framework for analyzing the events in the project, particularly the influence of the claims review task on system design. From another perspective, it examines participant roles in terms of shifts of attention among domain knowledge, knowledge representation, system performance, and the kinds of skills needed to improve the evolving system. The conclusion includes a series of recommendations that may assist other companies and universities setting up similar joint ventures. (16 refs.)

1553 Elements of Expert System Shells
M. Fontana, J. Zeimetz
Periscope Co. Inc., Atlanta, Georgia, USA
PC Tech J. (USA), vol. 6, no. 5, pp 63-5, May 1988

To build the appropriate expert system for a given problem, the developer must first select the proper tool, known as an expert system shell. Deciding which shell is best means knowing which questions to ask. To help ask the right questions, the authors present criteria for evaluating shells. In particular, they discuss knowledge representation methods, inference engines and tracing.

1557 The KREME Knowledge Editing Environment
G. Abrett, M.H. Burstein
BBN Labs., Cambridge, Massachusetts, USA
Int. J. Man-Mach. Stud. (UK), vol. 27, no.2, pp 103-26, Aug 1987). *AAAI Workshop on Knowledge Acquisition for Knowledge-Based Systems*, Banff, Alberta, Canada, Nov 1986

One of the major bottlenecks in large-scale expert-system development is the problem of knowledge acquisition: the construction, maintenance, and testing of large knowledge bases. This paper provides an overview of the current state of development of the Knowledge Representation Editing and Modeling Environment (KREME). It is an extensible experimental environment for developing and editing large knowledge bases in a variety of representation styles. It provides tools for effective viewing and browsing in each kind of representational base, automatic consistency checking, macro-editing facilities to reduce the burden of large-scale knowledge-base revision and some experimental automatic generalization and acquisition facilities. (22 refs.)

1571 Making Friends of Man and Machine [Expert Systems]
S. Ottley
Syst. Int. (UK), vol. 16, no. 4., pp 23-5, April 1988

The author warns of some human and organizational problems to consider when implementing an expert system, from capturing an expert's knowledge to gaining user acceptance.

1577 Expert Systems Maintenance: Post-implementation Support
J.E. Caviedes
North American Philips Corp.
Briarcliff Manor, New York, USA
IEEE MONTECH '87 Conferences: COMPINT
(Cat. No. 87CH2518-9), Montréal, Que., Canada, 9-12 Nov 1987 (New York, NY, USA: IEEE 1987), pp 21-4

Knowledge-base maintenance is a growing concern within the community of expert systems developers. One of the factors limiting the initial rush to apply expert systems technology is the difficulty of post-delivery maintenance. Since the quality of an expert system depends more on the knowledge engineering skills of its developers than on the implementation toolset, it is almost impossible to predict its maintainability. It is argued that maintenance tools can be developed in parallel with the expert system if it has a mature and stable architecture. (10 refs.)

1578 The Software Engineering of Systems with Expert Components
W.B. Frekes, C.J. Fox
AT&T Bell Labs., Holmdel, New Jersey, USA
Proceedings of the Twenty-First Annual Hawaii International Conference on System Sciences. Vol. II. Software Track (Cat. No. 88TH0212-1, Kailu-Kona, HI, USA, 5-8 Jan 1988 (Washington, DC, USA: IEEE Comput. Soc. Press 1988), pp 48-53

Integrating expert system components into production software can be difficult because expert system development environment are typically incompatible with traditional software engineering technology. In an effort to deal with this problem, the authors are developing CEST, a C expert system toolset. CEST is a library of inference engines implemented as C functions callable from C programs, and a workbench of knowledge engineering support tools. CEST allows easy integration of expert system components into C-based software systems and provides knowledge-engineering support tools analogous to traditional software-engineering support tools. (27 refs.)

1584 A Comparison of the Manipulation of Certainty Factors by Individuals and Expert System Shells
D. Kopcso, L. Pipino
Babson Coll., Babson Park, Massachusetts, USA
W. Rybolt
Proceedings of the Twenty-First Annual Hawaii International Conference on System Sciences. Vol. III. Decision Support and Knowledge Based Systems Track (Cat. No. iiTH0213 9), Kailua-Kona, HI, USA, 5-8 Jan 1988 (Washington, DC, USA: IEEE Comput. Soc. Press 1988), pp 181-8

The treatment of uncertainty in expert system shells is addressed, starting with a review of the modeling of uncertainty by expert system shells. An experiment to replicate earlier work investigating the manner in which individuals manipulate certainty factors in comparison to commercial shells is discussed. Comparisons are made among seven commercial shells, both personal-computer (PC)-based and mainframe-based, and individuals. A significant difference between individuals and shells themselves is indicated. Some implications for both expert system and decision-support-system methodologies are discussed. (25 refs.)

1829 Knowledge Harvesting: a Practical Guide to Interviewing
M. Davies
Software A & E, Chichester, UK
S. Hakiel
Expert Syst. (UK), vol. 4, no.1, pp 42-50, Feb 1988

It is up to the knowledge engineer not only to do his best to get

the relevant information from the domain expert, but also to make the expert's part of the interaction as pleasant and as easy as possible, and to generate and maintain in him as much interest and enthusiasm as possible. The aim of this article is to provide some guidelines on how to do all this, culled from the experience of the author and the reported views of other workers in the area. Because this is intended as a practical guide rather than a learned treatise, it is orientated primarily towards quick reference. Each section covers one major part of the knowledge harvesting task; within each section, the text is principally organized into notes. It is recommended that the reader make a quick pass through the whole document before starting a knowledge harvesting exercise, and refer to specific parts in more detail as and when required during the course of the exercise. This report does not go into the absolute merits of various approaches. (3 refs.)

1851 Artificial Intelligence: Current Technologies and Tools
W.E. Bracker, Jr.; L.C. Bracker
Tech. Res. Associates, Tucson, Arizona, USA
B.R. Konsynski
Proceedings of the Seventh Annual Conference and Exposition: Computer Graphics'86. Anaheim, CA, USA, 11-15 May 1986 (Fairfax, VA, USA; Nat. Comput. Graphics Asoc. 1986), pp 37-47, vol.3

AI has many subfields. Among them are: computer vision, voice-recognition, natural language, robotics and expert systems. The authors provide an introduction to expert systems, discuss languages used in expert system development, consider current technologies in support of AI, and present via a comparison matrix, AI workstations. (83 refs.)

1859 SEA — an Expert System for Nuclear Test Ban Treaty Verification
C.L. Mason, R.R. Johnson
Dept. of Appl. Sci., California U.
Livermore, California, USA
R.M. Searfus, D. Lager
Proceedings of the Australia Joint Artificial Intelligence Conference — AI '87, Sydney, NSW, Australia, 2-4 Nov 1987 (Sydney, NSW, Australia: U. of Sydney 1987), pp 11-25

Presents an expert system that interprets seismic data from Norway's regional seismic array, NORESS, for underground nuclear weapons test ban treaty verification. Three important aspects of the expert system are that it emulates the problem solving behaviour of the human seismic analyst, it acts as an assistant to the human analyst by automatically interpreting and presenting events for review, and it enables the analyst to interactively query the system's chain of reasoning and manually perform an interpretation. The general problem of seismic interpretation is described. The expert system is presented in terms of problem solving strategy, representation structures and user interface elements. (12 refs.)

1868 Induction, Knowledge and Expert Systems
J.R. Quinian
Key Centre for Advanced Comput. Sci.
New South Wales Institute of Technology
Sydney, New South Wales, Australia
Proceedings of the Australian Joint Artificial Intelligence Conference — AI'87, Sydney, NSW, Australia, 2-4 Nov 1987 (Sydney, NSW, Australia:U. of Sydney 1987), pp 223-40

Expert systems technology is a cornerstone of applied artificial intelligence that has many potential benefits for the manufacturing and service industries. The principal barrier to the more widespread use of expert systems is the difficulty of assembling and debugging the requisite collections of expert knowledge. Induction, the process of generalizing from examples, provides a proven means of expediting the knowledge acquisition process. this paper sketches the key ideas, supported by a case study, and discusses experimental techniques for improving the intelligibility of the induced knowledge. (29 refs.)

3.0 APPLICATIONS

1419 REGWASTE: An Expert System for Regulating Hazardous Wastes
G. Anandalingam
Dept. of Syst., Pennsylvania U.
Philadelphia, Pennsylvania, USA
Proceedings of the 1987 International Conference on Systems, Man, and Cybernetics (Cat. No. 87CH2503-1), Alexandria, VA, USA, 20-23 Oct 1987 (New York, NY, USA: IEEE 1987), pp 634-9, vol. 2

The author reports on an expert system called REGWASTE, which has been implemented to assist the US Environmental Protection Agency in regulating hazardous waste. The expert system simplifies a complex decision problem, provides alternatives arranged according to likelihood, and advises the user on further information needed to finalize the regulation. (6 refs.)

1421 A "Neural" Network that Learns to Play Backgammon
G. Tesauro
Center for Complex Syst. Res., Illinois U.
Champaign, Illinois, USA
T. J. Sejnowski
1987 IEEE Conference on Neural Information Processing Systems — Natural and Synthetic Abstracts of Papers
(Cat. No. 87CH2386-1), Denver, CO, USA, 8-12 Nov 1987 (New York, NY, USA: IEEE 1987), pp 28

Summary form only given. A class of connectionist networks that have learned to play backgammon at an intermediate-to-advanced level is described. The networks were trained by a supervised learning procedure on a large set of sample positions evaluated by a human expert. In actual match play against humans and conventional computer programs, the networks demonstrate substantial ability to generalize on the basis of expert knowledge. The study touches on some of the most important issues in network learning theory, including the development of efficient coding schemes and training procedures, scaling, generalization, the use of real-valued inputs and outputs, and techniques for escaping from local minima. Practical applications in games and other domains are also discussed.

1425 Artificial Intelligence: Current Trends Expanding
J. K. Carter
Office (USA), vol. 107, no. 1, pp 82, Jan 1988

The field of AI has experienced its share of exciting trends and developments with research in areas such as neural networks. Of major importance to outside research labs is progress towards integration of AI-based system functions into conventional data processing and office system environments. The trend toward integration of AI and conventional information processing is allowing computerized business systems to expand beyond clerical tasks into more knowledge-intensive applications. Because of their ability to support a decision-making process from inception to conclusion, expert systems, integrated with conventional data processing systems add significant value to a business by leveraging the valuable time of professional personnel.

1448 Applying AI Techniques for Patent Information Retrieval
D. Vermeir, E. Laenens, J. Dierick
Dept. of Math. & Comput. Sci., Antwerp U.
Wilrijk, Belgium
World Pat. Inf. (USA), vol. 10, no.1, pp 26-36, 1988

Presents an overview of recent developments in software technology, especially information retrieval and expert systems. Particular consideration is given to the possible applications in the area of user-friendly access to patent information systems. A proposal for an expert system, that could act as a knowledgeable intermediary between the end user with no information retrieval experience and the various host systems, is described. (16 refs.)

1449 Overview of Artificial Intelligence and Expert Systems for Information Professionals
D.T. Hawkins
AT&S Bell Labs, Murray Hill, New Jersey, USA
Online '87 Conference Proceedings,
Anaheim, CA, USA, 20-22 Oct 1987
(Weston, CT, USA: Online 1987), pp 94-6

Interest in AI and expert systems is growing rapidly as AI-based systems move out of the research laboratory and into the marketplace. AI will be a major research effort in the coming years, with information retrieval a prime application. As interest in AI and its related areas grows, online searchers and other information professionals need to understand the terminology and basic principles so they can intelligently evaluate the new systems that are appearing. The author introduces AI as it relates to information retrieval and online systems. It defines and explains the terms commonly used in AI, describes some AI programming languages and hardware, and briefly discusses some well known prototype systems. Applications of AI for information retrieval are then described, and finally, some commercial software packages now beginning to appear on the market are introduced. (5 refs.)

1457 A Prototype Expert System for Synthesizing Chemical Process Flowsheets
R. L. Kirkwood, M.H.Lockhe, J.M. Douglas
Dept. of Chem, Eng., Massachusetts U.
Amherst, Massachusetts, USA
Comput. Chem. Eng. (UK), vol. 12, no. 4, pp 329-43, Apr 1988

PIP (process invention procedure) is an hierarchical expert system for the synthesis of chemical process flowsheets. It uses a combination of qualitative and quantitative knowledge, arranged in a hierarchical structure. The heuristics are used to select the unit operations, to identify the interconnections between these units, to identify the dominant design variables are to identify the process alternatives at each level of the hierarchy, while the quantitative models are used to calculate process flows, the equipment sizes and costs the raw material and utility costs and the process profitability at each level. A hybrid, expert system control architecture was developed for PIP that allows these two types of knowledge-bases to interact in such a way that the heuristic rules "fire" the appropriate subroutines used to evaluate the flowsheet. PIP attempts to invent a flowsheet using a depth-first strategy, where one of the goals is to see if there is some reason why none of the alternatives will ever be profitable. Thus, it attempts to complete a design before any alternatives are considered. If profitable operation is observed over some range of the design variables at a particular level, the PIP proceeds to the next level and adds more detail to the flowsheet. However, if the process is not profitable, PIP examines the alternatives, starting with the ones found at the earliest level, and if no profitable alternative can be found the design project is terminated. Hence, PIP allows a design engineer to invest initial chemical flowsheet structures rapidly, to estimate the optimum design conditions for this flowsheet, to identify possible process alternatives and to quickly screen these alternatives. (33 refs.)

1458 EMPS — Making Automated NC Programming Feasible
FMS Mag. (UK), vol. 5, no.1, p.25-6, Jan 1988

Automation of the NC programming function is the goal of a new project initiated by Computer Aided Manufacturing-International. The system will be called EMPS (expert manufacturing programming system). So far, the project has developed a comprehensive design for an automated processor. Using techniques such as solid modeling geometric tolerancing, expert systems, features recognition, and volume decomposition, the processor will automatically determine the volumes of material to be removed, the appropriate NC machine fixturing requirements, cutting tool assemblies, cutter paths, and machinability data (feed and speed rates). The challenges lie in three areas: task decomposition, work element generation and tool path generation.

1490 An Approach to Integrating Expert System Components into Production Software
W. B. Frakes, C.J. Fox
AT&T Bell Labs., Holmdel, New Jersey, USA
Proceedings 1987 Fall Joint Computer Conference — Exploring Technology: Today and Tomorrow
(Cat. No. 87CH2468-7), Dallas, TX, USA, 25-29 Oct 1987
(Washington, DC, USA: IEEE Comput. Soc. Press 1987),
pp 50-6

The authors describe work on CEST (C expert system tools), an expert system function library and workbench for the Unix/C environment that has been developed to incorporate expert systems techniques into production software. CEST is being designed to consist of a library of inference engines implemented as C functions and callable from C programs, and a workbench of knowledge engineering support tools for building, analyzing, and maintaining knowledge bases. The authors describe the first tool in CEST, called AVIEN, which is a backward chaining inference engine implemented as a set of C functions. (21 refs.)

1474 Rocket Engine Health Monitoring System (HMS) via an Embedded Expert System (EES)
J. Pooley, T. Homsley, W. Tech
SPARTA Inc., Huntsville, Alabama, USA
J. Jones, P. Lewallen
Proceedings of the 1987 International Conference on Systems, Man. and Cybernetics (Cat. No. 87CH2503-1),
Alexandria, VA, USA, 20-23 Oct 1987 (New York, NY,
USA: IEEE 1987), pp 1127-32 vol. 3

The SPARTA embedded expert system (SEES) is an intelligent system that directs the analysis of rocket engine maintenance requirements by placing confidence factors on possible engine status and then recommending a course of action to an engineer or engine controller. This technique can prevent catastrophic failures or costly rocket engine downtime because of false alarms. Further, the SEES has potential as an onboard flight monitor for reusable rocket engine systems. The SEES methodology synergistically integrates vibration analysis, pattern recognition, and communications theory techniques with an artificial intelligence technique. Results from processing real data are presented. (5 refs.)

1465 Air Traffic Control using AI Techniques
M. H. Davies
Sist. Autom. (Italy), vol. 33, no. 287, pp 1292-302,
Dec 1987, in Italian

Describes a project carried out by Systems Designers Ltd. in collaboration with RSRE (Royal Signals and Radar Establishment) which has already resulted in a prototype air traffic control system using AI techniques. The process of command and control is considered and the requirements of the air traffic control operation described in detail. An expert system type aid, ATC Aid, is proposed, with a description of its architecture, mode of operation, and the way it interacts with the human controller. Verification and validation of the system are considered. It is concluded that the feasibility of an AI approach to air traffic control has been demonstrated. (10 refs.)

1500 Extensions and Modifications to the Prolog Language in View of Actual Application Programming
Expert Systems in Production Engineering. Proceedings of the International Workshop, Spa, Belgium, 18-22 Aug 1986
(Berlin, West Germany: Springer-Verlag 1987), pp 145-56

The initial work that the authors have done has proved that expert system technology should be regarded as an extension of conventional data processing. With these techniques it has become possible to solve problems which could not be solved before by conventional methods. As with all new technologies, the techniques of AI and especially expert systems are based on an initial theory developed over recent years. This theory was mainly based on a general view of the possible thematics derived by mathematicians and informatic engineers. In view of the recent work carried out by the authors, they find that a number of extensions and modifications have to be executed to make the use of the Prolog language more efficient in the hands of the users.

1609 Expert Systems and the Law
A. Bloch
Micro Syst. (France), no. 84, pp 203-4, March 1988, In French

Legal ramifications of the use of expert systems are considered with particular reference to the two sections: the knowledge base, and the inferential mechanism. Contractual details regarding construction of specific expert systems are being developed with the emphasis on turnkey contracts.

1612 The Alvey DHSS Demonstrator Project: Applying Intelligent Knowledge-based Systems to Social Security
G.N. Gilbert
Dept. of Sociology, Surrey U., Guildford, UK
Inf. Age (UK), vol. 10, no. 2, pp 113-15, April 1988

Prototype decision support systems being developed under the Alvey Program are outlined. They are designed for possible use by the UK Department of Health and Social Security and cover the assessment of benefit claims in local offices, the formulation of legislation by policy-makers, and the provision of advice and assistance to claimants. An attempt is being made to utilize the strengths of both user and system in the design. (6 refs.)

1618 QUARTZ: an Intelligent Assistant for the Analysis and Evaluation of Project Proposals
K. Dalkir, J. Muzard
Canadian Workplace Autom. Res. Centre, Laval, Que.
IEEE MONTECH '87 Conferences: COMPINT
(Cat. No. 87CH2518-9), Montréal, Que., Canada,
9-12 Nov 1987 (New York, NY, USA: IEEE 1987), pp 208-10

QUARTZ is an intelligent decision-support system that will assist Canadian government analysis in their evaluation of project proposals submitted for funding under the economic regional development assistance program for Quebec. The authors primarily focus on the knowledge engineering methodology used to develop QUARTZ. In particular, the knowledge acquisition, knowledge representation, and knowledge validation stages of system development are discussed, both in terms of lessons learned and future research directions to be undertaken. (7 refs.)

1621 Potential Defense Applications of Expert Systems
V. Shah
Southwest Texas State U., San Marcos, Texas, USA
G.D. Buckner
IEEE Aerosp. Electron. Syst. Mag. (USA), vol.3, no.2,
pp 15-21, Feb 1988

The authors provide an overview of expert systems and how they may effect the development of future defense applications. Military uses of computers are outlined, and expert systems fundamentals are described. AI research and development efforts by the military are examined, and potential military applications are discussed. Expert systems efforts at NASA, by the US Air Force, and for the Strategic Defence Initiative are considered. (19 refs.)

1649 A Financial Investment Assistant
K. Kandt, P. Yuenger
Teknowledge Fed. Syst., Thousands Oaks, California, USA
Proceedings of the Twenty-First Annual Hawaii International Conference on System Sciences. Vol. III.
Decision Support and knowledge Based Systems Track
(Cat. No.88TH0213 9), Kailua-Kona, HI, USA, 5-8 Jan 1988
(Washington, DC, USA: IEEE Comput. Soc. Press 1988),
pp 510-17

The analysis of financial markets is a time-consuming complex and error-prone process. The system described which is still under development, is an attempt to improve this process by partially automating the acquisition, analysis, and selection of financial market investments. The ultimate goal is to fully automate this activity. The current approach is to use technical analysis and fundamental analysis to determine when to buy, sell, or hold various instruments, and AI techniques to select a portfolio of stocks and/or options based on the goals of the system user. The tool uses a dynamic interface that is reactive to human interaction. (25 refs.)

1671 Expert Systems for University Admissions
J.S. Edwards, J.L. Bader
Aston U., Birmingham, UK
J. Oper. Res. Soc. (UK), vol. 39, no. 1, pp 33-40, Jan 1988

Describes the construction of an expert system to help the admissions tutor for a university degree in business and management which receives some 2000 applications for entry each year, using the SAGE shell. What originally began as a "demonstrator project" is shown to be of practical value, in terms of both producing a usable expert system and clarifying and questioning the selection criteria used by the admissions tutor. A particular conclusion emerging from this work which may be relevant to many expert systems applications is that the ethical and practical considerations dictated that some questions involving judgement could not be delegated by the admissions tutor to the clerical staff. It thus became necessary to develop two versions of the expert system, one a full "administrations-tutor system", the other a more limited version for day-to-day use. (9 refs.)

1701 Expert System Tools Emerge from the Technology of AI
F.J. Bartos
Control Eng. (USA), vol. 34, no. 14, pp 36-8, Dec 1987

AI is a technology that works to make computers behave more like humans in executing and solving problems. AI is fundamentally different from conventional programming. It is a more symbolic than numeric process, employing heuristic or rule of thumb procedures over explicit ones to arrive at solutions. It is meant to be easily modified or expanded and can live with less than optimum results — even accepting an incorrect answer or two. It is emphasized that AI, in the form of its expert systems, is starting to go to work for process and control engineers.

1705 Expert System for Wire Cutting EDM, Based on Pulse Classification and Thermal Modeling
W. Dekeyser, R. Snoeys, M. Jennes
Inst. voor Werktuigkunde, Heverlee, Belgium
Robot. Comput.-Integr. Manuf. (UK), vol. 4, no. 1-2, pp 219-24, 1988, (Manufacturing Science, Technology and Systems of the Future, Ljubljana, Yugoslavia, 12-14 Sept 1985)

Although wire electric discharge machine (EDM) has become a fully competitive machining technique during the last decade, the process performance is limited by some typical restrictions. In order to improve machining performance, a multi-disciplinary expert system was developed. It mainly consists of two parts: an observation procedure and a decision procedure. The observation of the wire EDM process is carried out by means of a powerful electronic device, called the EDM pulse discriminating system. This device delivers "on-line" a "chaos" of information of the process. In the decision phase, the incoming data are analyzed and patterns are detected by means of a computer and eventually the parameter settings of the EDM machine are changed in order to obtain more optimal working conditions. This decision procedure is partially based on the results of a mathematical thermal model of the wire, enabling one to predict the influence of some parameter changes. Further, a technological database is permanently accessible for the control strategy program. Machining speed can be increased without enhancing the risk of wire rupture, because in the case of an increased danger of rupture, appropriate action is taken by the control strategy. The paper gives some results on how the thermal model may predict some thermal overload of the wire. It also describes how frequently observed breakage at sharp corners of the wire path may be dealt with. It further illustrates the actions taken by the expert system when various types of disturbance occur. The consideration of practical experience with a thermal model yields an increased level of machine autonomy ensuring a performance level obtainable only with a skilled operator. (7 refs.)

1710 Glass Annealing Process Simulation Using Expert System: a Glass Industry Application of Artificial Intelligence
R.A. Herrod, J.W. Rickel
Texas Instrum. Inc., Dallas, Texas, USA
T. Garland
IEEE Trans. Ind. Appl. (USA), vol. 24, no. 1, pt. 1, pp 43-8, Jan-Feb 1988

How an AI system was used for simulation of a glass annealing process is discussed. The system consists of two parts: a planner that uses rules-of-thumb (in the form of an expert system) to determine control settings for the glass processing, and a simulator that uses control settings to derive a glass temperature curve. The two subsystems work together to produce the necessary Lehr control settings to anneal the glass product. A short review of expert system technology is given, followed by a discussion of the project and lessons learned. (4 refs.)

1726 RATIONALE: Developing Expert Systems that Reason by Explaining
S. Abu-Hakima
Dept. of Syst. & Comput. Eng., Carleton U., Ottawa, Ontario
F. Oppacher
IEEE MONTECH '87 Conferences: COMPINT
(Cat. No. 87CH2518-9) Montreal, Que., Canada, 9-12 Nov 1987 (New York, NY, USA: IEEE 1987), pp 13-16

Explanations generated by current expert systems are often terse, include bookkeeping information irrelevant to most users, and are isolated from the reasoning processes of the expert system. A methodology is presented for building knowledge-based systems that reason by explaining. The methodology is the basis for RATIONALE, a tool that is implemented in Prolog and ties together the processes of knowledge acquisition, expert system construction, and explanation of system reasoning. RATIONALE does this by using domain knowledge together with explicitly represented strategies. Reasoning proceeds by constructing a hypothesis tree having a root that contains the most general diagnosis of the expert system. Guided by a focusing algorithm, the tree branches into more specific hypotheses that explain the more detailed symptoms provided by the user. As the expert system is used, the hypothesis tree also serves as a dynamically generated explanation tree. (22 refs)

1733 A Knowledge-based Approach for Real-time Systems Debugging
J.P. Tsai, K.Y. Fang, V.R.K. Thalla, H. Gandhi
Dept. of Electr. Eng. & Comput. Sci., Illinois U. Chicago, Illinois, USA
Proceedings of the Twenty-First Annual Hawaii International Conference on System Sciences. Vol. II. Software Track (Cat. No. iiTH0212-1), Kailua-Kona, HI, USA, 5-8 Jan 1988 (Washington, DC, USA: IEEE Comput. Soc. Press 1988), pp 533-40

The authors consider a method of testing unpredictable sequences called program execution monitoring. In this method, a snapshot of the system events and happenings are recorded in a non-interfering manner, i.e. without corrupting the critical timing requirements of the system. The traces thus collected are analyzed to isolate the bug in an off-line mode. The examination of voluminous traces is a laborious and tedious task that requires a high degree of expertise. It is shown how this expertise can be encapsulated in a knowledge-based system and the examination process automated to at least localize the fault and answer questions about its reasoning.

1885 Victorian Business Assistance Referral System: Development and Early Use
M. Frazer
Dept. of Ind., Technol. & Resources Melbourne, Victoria, Australia
Proceedings of the Australian Joint Artificial Intelligence Conference — AI '87, Sydney, NSW, Australia, 2-4 Nov 1987 (Sydney, NSW, Australia: U. of Sydney 1987), pp 44-60

An expert system to provide advice on assistance to new or small businesses in use at a number of sites in Melbourne. The system has been developed by BBJ Computers International Pty. Ltd of South Melbourne, for the Victorian Government's Department of Industry Technology and Resources. The paper discusses three matters: the Department's motivation in commissioning the system; the planning and development of the system; and early experience of its use. The presentation is from the point of view of the Department as the client, rather than from the point of view of the technical developers of the system. (3 refs.)

1887 Expert Systems for Business Applications
J. Liebowitz
Dept. of Manage. Sci., George Washington U. Washington, DC, USA
Appl. Artif. Intell. (USA), vol. 1, no. 4, pp 307-13, 1987

Expert systems are becoming more prevalent in financial and business applications. With venture capital pouring into the expert systems area, there is greater interest by commercial firms. Companies are either building in-house expert systems capabilities or contracting out to expert system developers to construct expert systems for business applications. This paper presents much of the work being done in developing expert systems in business and discusses fertile areas for constructing more business expert systems. (10 refs.)

1897 Expert Process Planning System with Solid Model Interface
S. Joshi, N.N. Vissa, Tien-Chien Chang
Sch. of Ind. Eng., Purdue U., W. Lafayette, Indiana, USA
Int. J. Prod. Res. (UK), vol. 26, no. 5, pp 863-85, May 1988

Presents an integrated hierarchical framework of a process planning system with a CAD interface. The objective of the project discussed is to integrate design with process planning using AI techniques. The development of a CAD interface is discussed with respect to automated feature recognition,

determination of tool approach direction, and deciding the precedence relationship between the features. Sample results from the CAD interface are presented. The expert system for the process planning module is discussed with the part representation and knowledge base, and the plan generation procedure. The module uses hierarchically organized frames for both part representation and the knowledge base. (40 refs.)

1899 Knowledge Based Simulation Techniques for Manufacturing
R.E. Shannon
Dept. of Ind. Eng., Texas A&M U. College Station, Texas, USA
Int. J. Prod. Res. (UK), vol. 26, no.5, pp 953-73, May 1988

The art and science of simulating complex manufacturing systems is rapidly changing. A great deal of attention is being devoted to the possibilities of bringing AI and expert systems technology into simulation methodology. Such systems will hopefully allow models to be quickly developed, validated and run with as much of the necessary expertise as possible built into the software. This paper addresses: the motivation and need for developing such systems; the nature of such systems; the potential benefits of this technology over existing approaches; and the current state-of-the-art as it applies to simulation. (60 refs.)

1909 The Expert System for Career Planning: ACES
S. Geffin, T. Burges, B. Furht
Dept. of Electr. & Comput. Eng., Miami U. Coral Gables, Florida, USA
Microcomput. Appl. (USA), vol. 6, no.3, pp 71-7, 1987

The design of a knowledge-based student advisor system implemented on a microcomputer is discussed. The student advisor system is intended to assist graduating students, majoring in either electrical, computer engineering, or computer science in deciding upon an area of specialization. The system can also prescribe a list of graduate schools which match the student's qualifications, location desired and the area of specialization decided by the system. The programming tool used to develop the system was the Texas Instrument Personal Consultant run on an IBM/AT microcomputer. (11 refs.)

1926 GOLD: an Expert System for Mineral Identification from Reflectance Spectra
R. Lister
Basser Dept. of Comp. Sci., Sydney U., NSW, Australia
K. Ali, R. Buda, C. Horsfall, W. Buntine
Proceedings of the Australian Joint Artificial Intelligence Conference — AI '87, Sydney, NSW, Australia), 2-4 Nov 1987 (Sydney, NSW, Australia: U. of Sydney 1987), pp 29-43

GOLD is an expert system for identifying minerals, some of which are indicative of fine-grained sub-surface gold deposits. The expert system identifies minerals by searching for characteristic features in the infrared reflectance spectra of rock samples. The system is being built using the Prospector expert system shell which represents knowledge using an inference net and propagates uncertainty using subjective Bayesian reasoning. A prototype has been installed at CSIRO and is currently on trial. GOLD's performance compares favourably with that of domain experts. The system is being extended to deal with spectra produced by samples containing more than one mineral. (11 refs.)

1932 Application of Artificial Intelligence in Telecommunications
K.J. Macleish, S. Thiedke, D. Vennergrund
GTE Commun. Syst., Phoenix, Arizona, USA
GLOBECOM Tokyo '87. IEEE/IECE Global Telecommunications conference 1987. Conference Record (Cat. No. 87CH2520-5), Tokyo, Japan, 15-18 Nov 1987 (New York, NY, USA: IEEE 1987), pp 329-33, vol. 1

The authors examine the ways in which AI techniques are used now and will be used in the future to assist in many aspects of telecommunications systems. The current use of AI consists primarily of expert systems designed to aid in diagnosing complex equipment. The authors briefly examine such a system, NEMESYS, which diagnoses intermittent faults in the GTD-5 EAX digital central office switch, and focus on why this project required AI techniques. Furthermore, the authors explore how other AI techniques, such as natural language processing, distributed AI, and speech recognition, will be used along with expert systems to assist both operations and services in what they call the network of the future. (10 refs.)

1939 Knowledge Based Maintenance in Networks
D. Peacocke, S. Rabie
Bell-Northern Res., Ottawa, Ontario, Canada
GLOBECOM Tokyo '87 IEEE/IECE Global Telecommunications Conference 1987. Conference Record (Cat. No. 87CH2520-5), Tokyo, Japan, 15-18 Nov 1987 (New York, NY, USA: IEEE 1987), pp 1833-8, vol. 3

The maintenance advisor for DMS-100 is considered. MAD is an interactive expert system for helping operating company personnel perform maintenance on the DMS-100 family of digital switches. A prototype is described that is based on using product documentation and experience obtained from craftspeople and field technical support to create the knowledge base for an expert system that diagnoses and chooses corrective actions for DMS-100 problems. Its operation is described, and steps for further evolution of the system in a network maintenance role are outlined. (4 refs.)

1947 ESA: Expert Structural Analysis for Engineers
B.W.R. Forde, S.F. Stiemer
Dept. of Civil Eng., UBC, Vancouver, British Columbia
Comput. Struct. (UK), vol. 29, no. 1, pp 171-4, 1988

The finite element method of analysis is possibly the most popular tool available for the numerical solution of complex problems in engineering. A consistent algorithmic approach to formulation, discretization, and solution procedures for this method has resulted in the wide-spread development of extremely powerful computer software. This power has its price: expertise is required for efficient analysis. A knowledge-based expert system called expert structural analysis (ESA) is being developed which will allow average structural engineers to skillfully use a nonlinear finite element analysis program. The paper examines the nature of the algorithms employed in the analysis program and introduces the principle issues associated with human and artificial expertise. A case study is used to demonstrate the application of a prototype expert system. (9 refs.)

1949 Teaching Artificial Neural Systems to Drive: Manual Training Techniques for Autonomous Systems
J.F. Shepanski, S.A. Macy
TRW Inc., Redondo Beach, California, USA
vol.848, pp 286-93, 1988, *Intelligent Robots and Computer Vision.* Sixth in a Series, Cambridge, MA, USA, 2-6 Nov 1987

The authors have developed a methodology for manually training autonomous control systems based on artificial neural systems (ANS). In applications where the rule set governing an expert's decisions is difficult to formulate, ANS can be used to extract rules by associating the information as expert receives with the actions he takes. Properly constructed networks imitate rules of behaviour that permit them to function autonomously when they are trained on the spanning set of possible situations. This training can be provided manually, either under the direct supervision of a system trainer, or indirectly using a background mode where the network assimilates training data as the expert forms his day-to-day tasks. To demonstrate these methods the authors have trained an ANS network to drive a vehicle through simulated freeway traffic. (5 refs.)