Science Council of Canada

A workshop on Artificial Intelligence

Sponsored by the
Science Council of Canada

Proceedings 83/1

A WORKSHOP ON ARTIFICIAL INTELLIGENCE

August 1983

Proceedings are published by the Science Council of Canada to share with a wider audience of interested readers the ideas and perspectives brought to Council conferences or seminars by experts in a variety of fields. The papers presented and views expressed in these Proceedings are not necessarily endorsed by the Science Council.

Copies of

A Workshop on Artificial Intelligence
A Proceedings

are available
free of charge from:

The Publications Office
Science Council of Canada
100 Metcalfe Street
Ottawa, Ontario.
K1P 5M1

Des exemplaires du
Compte rendu de l'Atelier sur l'intelligence artificielle
sont disponibles gratuitement

Catalogue No. SS24-19/1983E
ISBN No. 0-662-12716-1

CONTENTS

FOREWORD

Since 1979, the Science Council has been studying the effects of computers and related technologies on industry, education, the workforce and society in general. This major project led to the publication in 1982 of Report 33, Planning Now for an Information Society: Tomorrow is Too Late. The report gave some consideration to the subject of artificial intelligence (AI). However, developments in this area are proceeding extremely quickly, and Council sensed the need to examine the social and economic implications of AI in more detail. Thus, a workshop was convened by Council on 20 and 21 January 1983.

The Ottawa workshop had the effect of acting as a catalyst to bring together members of the AI community from all parts of Canada. For two days, some 60 participants from across the country met in Ottawa to discuss the current state of AI research and consider policies for future development. Those attending included research workers from university, industrial and government laboratories as well as industrialists, policy-makers and venture capitalists.

The Science Council sponsored the workshop and offers these proceedings because, in its view, policy-makers and the public need to become acquainted with the current status of AI, not only in Canada but around the world. Citizens and their representatives should understand how the evolution of this technology will affect employment, the organization of businesses and the structure of society. In light of the new technology, governments will have to reconsider industrial policies and discover the best ways to encourage and accommodate the development of AI. In turn politicians need information on how the new applications will benefit industry. Through the medium of the Ottawa workshop, the Science Council obtained answers to many outstanding questions about the current and future status of AI. This will assist Council in arriving at policy recommendations.

In opening the workshop, Dr. Stuart Smith, Chairman of the Science Council, posed a series of questions: *What is the general state of the institutional infrastructure in which AI is embedded? How can government best assist? How important are grants, projects, the training of personnel and the exchange of ideas?* The meeting then heard a series of overview lectures given by experts in the field. Participants were, of course, speaking with their peers and their talks and discussions ranged from the technical to reviews of topics known to their audience. Without having a general background in the field of AI, much of the material discussed over the two days would be inaccessible to the non-specialist. These proceedings are presented, therefore, as a general, self-contained picture of AI research and development in Canada rather than a collection of abstracts. The report, it is hoped, is faithful to the spirit of the workshop but goes beyond the scope of several of the talks by drawing on additional background material.

A number of people deserve special mention for the success of the workshop and these proceedings. First, Dr. Arthur J. Cordell, science adviser, who thought the idea for an AI workshop to be timely and necessary; Dr. Alfred Stein, Continuum, Toronto, was given the difficult task of setting the agenda for the meeting and preparing the master list of invitees. Finally, Dr. David Peat, consultant, Ottawa, prepared these proceedings which not only capture the flavour and excitment of the workshop, but give the reader an introduction into the area of AI itself.

James M. Gilmour
Director of Research

I. INTRODUCTION

ARTIFICIAL INTELLIGENCE

Some idea of the diversity and rapid evolution of the field can be gained by attempting to define the term "artificial intelligence." The Oxford English Dictionary offers no definition and no topic entry can be found in the Encyclopaedia Britannica (15th edition) or the Von Nostrand's Scientific Encyclopedia (5th edition). The McGraw Hill Dictionary of Scientific and Technical Terms (2nd edition) is more helpful:

> "The property of a machine capable of reason by which it can learn functions normally associated with human intelligence."
>
> Workers in the field offer their own definitions, such as:
>
> "Artificial Intelligence is concerned with understanding the principles of intelligence and building working models of human intelligent behaviour."

and;

> "Artificial Intelligence is the part of computer science concerned with designing intelligent computer systems, that is, systems that exhibit the characteristics we associate with intelligence in human behaviour - understanding language, learning, reasoning, solving problems and so on."

In some of these definitions, however, the computer seems to be very far away:

> "A technical language with which to discipline the imagination."

The concept of AI also involves a general agreement as to the meaning of intelligence itself. Ed Feigenbaum of Stanford University, a noted research worker in the field, offers the following definition:

> "Intelligent action is an act or a decision that is goal-oriented, arrived at by an understandable chain of symbolic analysis and reasoning steps, and is one in which knowledge of the world informs and guides reasoning."

Other workers offer other definitions. Thus, it is probably not too profitable at this stage in the evolution of AI to attempt a single, all-embracing statement for, over its short history, the topic has radically changed in its aims and approaches. In the 1950s, AI was directed towards building an artificial brain that could duplicate a wide range of human functions. Today, research is more concerned with making realistic and practical contributions in a limited number of fields. Yet as computer systems continue to evolve and new theoretical insights are gained, AI may change direction again.

AI today is a truly integrated field of study with its theories, experiments and applications coming from computing, engineering, mathematics, logic, physics, psychology, human engineering, neurology, linguistics and other fields. Each discipline brings with it its own tools and unique approach, and it is not unusual for researchers from several disciplines to be found working side by side on a given research project.

Because there is no single agreed definition of AI nor a common background for AI workers, neither is it always possible to point to computer hardware and say, "This is a piece of AI." A large and expensive computer may carry out its tasks according to conventional mechanical procedures and embody little that could be called AI. However, an inexpensive chess playing game or the barcode reader in a supermarket checkout are both realizations of AI research. On the one hand, an industrial robot represents the frontier of AI research; on the other wrist, there may be a talking watch. Today AI encompasses not so much a single goal but rather a series of diverse and useful applications on the way, with an increasing body of knowledge about abstract systems and insights into the functioning of human intelligence.

How well do these definitions. and discussion help the layreader to become oriented in this rapidly changing field?

In simple terms, AI is an area of research and application involving computers that behave in ways we can recognize as intelligent. An AI computer, along with intelligent beings, must be able to understand something of the world around it and to use this understanding to reason, deduce, estimate and plan. The results of its electronic deliberations will be to respond appropriately to each new situation it encounters.

Intelligent computers therefore gather information about their environment and the world in general through vision, speech recognition, touch and manipulation as well as through the more conventional inputs such as typewriters, touch-sensitive screens, magnetic tape and discs. The computer stores, organizes and integrates the data with other general knowledge that is in its memory banks. Next the computer uses all its background information, together with data specific to its present task, and arrives at a solution to a problem, requests further information or takes some other form of intelligent action. The computer acts through printout, video displays, voice conversation or movement of articulated limbs.

At one time, AI workers sought to embody all these abilities within a single robot or supercomputer. However, the venture proved overly ambitious and, just as life on earth, AI has evolved through a process of specialization. Today AI devices are beginning to fill several of the niches in the informational landscape. These include computer vision, expert systems, machine translation, picture processing, question-answering machines, industrial robots, decision-making systems, office machines and intelligent information banks. Each of these devices employs several principles of AI and these principles and important applications were discussed at the Ottawa Workshop.

Does society need AI? What need is there to build machines that will duplicate human activities?

Unemployment in Canada and many other western nations has become a serious problem, but this topic was not raised. It is difficult to predict the nature of the information society by the end of the twenty-first century, but in the immediate decade the computer will probably act as partner rather than a competitor in the workplace. The reason lies in the computers special abilities and serious limitations. As one participant put it:

> "There are some things we can do so well that we don't even think about them; they're unconscious. These happen to be the very things that computers find hard to do. On the other hand, computers are excellent at making routine calculations and doing the sorts of things we find boring."

At present computers offer no competition in the activities humans enjoy doing and performing well. Computer vision is a good example of this state of affairs. Electronic vision systems are particularly stupid when it comes to recognizing complex objects; their abilities are taxed even by simple, solid objects in a well-designed background. Humans, on the other hand, are marvellously skillful at this and can spot something like a small bird in a landscape or recognize the face of a friend in a moving crowd. In the immediate future computers will certainly not compete with humans in this field of general vision.

However, the computer's "eye" can make accurate measurements of simple shapes it has been trained to see. The computer can rapidly scan a series of photographic plates containing elementary particle tracks and pick out an anomaly or new event. To perform the same task, a human observer would need technical training and the assistance of an accurate measuring apparatus. The human would be far slower, more prone to error and would find the task particularly tedious. In the field of vision, therefore, the computer will act as assistant to relieve humans of tasks they do not wish to perform.

To take another example, the artist's studio or the craftsman's workplace would prove chaotic to the robot, where the need to adapt to each new, creative task would stretch its capacities to the limit. However, the well-designed environment of the conveyor belt is congenial to the industrial robot. Provided a robot can carry out endless repetitions, involving the sort of tasks a human would find "soul destroying," it remains efficient.

Intelligent robots also have the potential to extend human abilities. They can work with the very small (microsurgery and printed circuitry) or the very large (manipulate heavy machinery). They can also work in areas that humans find hazardous such as nuclear plants, chemical works, outer space and the ocean bed. At present, robots and AI computers seem best adapted to carry out the very tasks that humans cannot or have no wish to do. At this stage in the development of AI, it appears that silicon and carbon-based intelligences are in no danger of competition but may eventually coevolve through a process of symbiosis.

HISTORY OF AI

A Diverse Introduction

The design of a machine that will duplicate the behaviour of humans has resonances in the myths of Pygmalion's living statue, Paracelsus's homoncule, Rabbi Loew of Prague's Golem and the clockwork mannikins of the eighteenth century. Indeed one of the great arcane secrets of the Middle Ages was the creation of life from inanimate matter. Several of the founders of the science of artificial intelligence have spoken of believing themselves to be the heirs to this tradition.

The first realization that it would be possible to duplicate the functions of human intelligence by machines came about in the 1940s. Early computers were crude by contemporary standards, but their designers realized that their electronic circuits imitated certain processes in the brain. At the time, research workers in cybernetics and neurology predicted that, in principle at least, it should be possible to build a computer that would approximate many of the abilities of the human brain. This electronic intelligence would have an enormous effect not only industrially but also on society as a whole. These visionaries also realized that the computer would be a valuable new tool for exploring that greatest of scientific mysteries - the nature of the human brain.

The first workers in AI were, in fact, exploring a comparison that has often been made between the brain and the most complex technology of a particular period. In the eighteenth century, for example, the brain was supposed to carry out its controlling function on hydrodynamic principles.

Impulses from the brain were directed to the body's organs by means of fluids travelling in narrow tubes. Thus the elaborate play of eighteenth-century fountains became a model for animal behaviour.

In the nineteenth century, the control of new industrial machines by governors attracted scientific attention. Through the study of these governors the principles of "feedback" were discovered and related to homeostasis in animals. By the twentieth century these investigations led to the sciences of cybernetics and general systems theory.

Early in the twentieth century, the brain was modelled on a telephone system with nerves relaying incoming and outgoing signals to a central switchboard. By the 1950s, the switchboard was forgotten and the brain had become a computer in which information is stored and processed.

However, each of these early metaphors is strictly limited in its application. What first appeared to be a promising image, in the end turns out to encompass only a very small part of the brain's total behaviour. The brain, for example, uses the principle of feedback to control certain internal bodily functions, but it also does far more than that. Reflex actions involve the passage of electrical messages to a "switchboard" in the brain or spinal column; but nerves do not operate in the same way as telephone lines and the brain is more than a switchboard.

Today, neurologists realize that the human brain does not operate in the same way as a computer yet the metaphor of "artificial intelligence" is still a valuable one. AI research continues to yield important applications and provides insights into the nature of human vision, language, speech and intelligence. However, history warns us that, in the last analysis, the metaphor of the "computer as brain" may well be replaced by yet another image.

From Dartmouth to Ottawa

The origins of AI as a recognized discipline date from a conference held at Dartmouth College in 1956. There Claude Shannon, Marvin Minsky, John McCarthy, Allan Newell, Herbert Simon and other leading researchers met to define the future directions of AI.

To a large extent, the foundations of the field had already been laid by the time of the Dartmouth Conference. In 1943, for example, Warren McCulloch and Walter Pitts had made a theoretical study of neurons in the brain and deduced that they operated as binary (on-off) switches. The two researchers went on to show that a network of such switches could duplicate all the functions of symbolic logic. This was a surprising and important result for it demonstrated that a collection of simple objects (on-off switches) that also happened to be the building blocks of a computer, were capable of carrying out complex logical operations.

In 1948, Norbert Wiener published Cybernetics in which he drew analogies between the mathematics of servo-mechanisms and the internal behaviour of biological systems. In the same year, Claude Shannon gave a mathematical formulation of information and noise.

By 1950, scientists were taking computer intelligence seriously enough for Alan Turing to propose a test that would determine the extent to which a computer could be judged intelligent. Turing proposed the following: you hold a conversation with some unknown entity by means of written messages. At the end of the conversation you are asked if your unseen partner is human or not. If you judge that a human intelligence was involved in the exchange and the unseen message writer happens to be a computer, then that electronic machine must indeed be judged intelligent. (Turing's test seemed eminently reasonable three decades ago but today we know that the computer doctor and psychiatrist can lull some people into the illusion that they are speaking to a sympathetic human. Is it possible that humans are less intelligent than was once thought!)

In 1952, W. Ross Ashby's Design for a Brain explored more mathematical models of the brain. By 1956, therefore, the theoretical implications of computers, neural networks, information and cybernetics had been much discussed. The future of AI looked promising and the mood of the Dartmouth Conference was highly optimistic. Participants predicted that by 1970 a computer would become a grand chess master, discover significant mathematical theorems, compose music of "classical quality," understand spoken language and provide language translations.

Several research projects were also initiated at Dartmouth which were equally ambitious;

- a complex system of artificial neurons that would begin to function as an artificial brain;
- a robot that would build up an internal picture of its environment;
- a computer program to derive the theorems of the Principia Mathematica;
- a chess playing grandmaster that would also provide significant insights into the nature of human intelligence; and
- a model of the brain's visual cortex.

However, by the early 1960s, the optimistic mood of Dartmouth soured as it became clear that AI involved problems far more complex than anyone had imagined. The various projects proposed in 1956 failed to produce significant results and the difficulties involved were not always clearly understood.

AI had promised much but delivered little of substance with the result that, by the late 1960s, interest in AI had declined. Efforts in machine translation, for example, exposed the subtleties and difficulties inherent in understanding natural language but produced little in the way of commercially attractive translation programs. In 1966, the major US project in machine translation was terminated. Problem solving also ran into serious difficulties when it was discovered that the "solution space" required (the number of possible alternative solutions that the computer would have to investigate), for any fair sized problem, was far too large for existing computers and would involve excessive amounts of computer time. (See "An Aside! Tree Graphs and Combinatorics," page 32.)

By the end of the 1960s, AI had entered a wasteland. Funding for major projects had been reduced or cancelled and it had become difficult to convince university administrators and government policy-makers of the importance of the topic. However, in their darkest hour, AI researchers were also learning from their mistakes. To begin with they realized the error in attempting to build a universal machine intelligence and in tackling problems that were too general in nature. Throughout the 1970s, AI researchers concentrated on well-defined problems and on applications in carefully structured situations.

The 1960s also saw the development of important technical approaches which were to be exploited to advantage in the 1970s. These new techniques can best be illustrated by analogy to humans solving problems. Faced with a difficult problem, a human does not necessarily proceed in a strictly logical fashion and examine every alternative with the same attention. A human makes guesses and uses rules of thumb; a problem may be simplified and broken into smaller parts, or the solution of a similar problem may suggest the strategy for a new approach. Humans also ignore certain details in favour of a "global" approach and attempt to juggle several things at the same time.

AI researchers realized that each problem requires its own particular strategy and can be tackled best with the aid of heuristics or rules of thumb. Computers were designed, therefore, that could do several things at once; in technical terms, carry out parallel processing and time sharing. Also, new programming languages were written that were better adapted to problem solving and AI work. When these insights and technical advances were combined with the more modest goals adopted by AI researchers, the field began a renaissance in the second half of the 1970s. This resurgence was also assisted by advances in the theoretical understanding of problems specifically connected with AI. In linguistics, for example, Noam Chomsky's theory of deep linguistic structures had a considerable effect on machine translations, question-answering machines and the understanding of natural languages.

Even more important in implementing these new advances was the realization that AI systems must have general knowledge of the world. Much of our ability to deal with day-to-day situations comes from the vast store of information knowledge we each hold within us. We know that day follows night, stones fall, shadows move with the sun, vegetables grow in the ground and are purchased at stores. We know about people and society, about beliefs and actions, and about recreation and daily work. Without any of this background information, a general conversation, overheard between two people, would be virtually meaningless and few daily problems could be solved appropriately.

Today, AI workers recognize that computers must also be given such powerful bases of knowledge. Without a knowledge base, even the fastest and largest of machines is powerless to understand the context of a conversation or to solve the simplest practical problems. The programming of speech recognition, for example, relies on the computer's ability to match incoming sounds with the acoustic patterns of sentence fragments it has constructed. However, this matching implies that the computer has anticipated and generated successive fragments of speech. To do so a computer not only has to understand semantics and syntax but also the meaning of a particular sentence within the overall speech and this requires general knowledge concerning the topic of conversation.

In virtually every area of AI today, the computer relies on a knowledge base of facts about the world. This knowledge base has become a key to so many advances in AI that considerable research is being directed towards the best ways of storing and manipulating knowledge.

Advances over the last 5 years can also be traced to evolutions in the computer itself. As more and more computing power is placed on a silicon chip, the computer becomes faster, cheaper, more powerful and more compact. The result is that research groups of even modest size can now have access to considerable computing facilities. Also, problems that once lay outside the capacity of a computer memory or took too much time to solve are now feasible. This process of electronic evolution is still accelerating. As one participant observed:

> "I feel that in the next five years, the champion chess player will be a computer and not a human. This isn't because the computer is any more intelligent, it's just that one order of magnitude increase in speed equals several points in IQ. As the chess machine speeds up, it can work out all the moves and apply various algorithms."

"Intelligent" devices can now be produced in which the computing power is built-in rather than supplied from a distant bank of computers. By the late 1970s, companies had also started to work on the design and manufacture of computers specifically created for AI work and the first applications of AI devices were finding their way into the marketplace. Microprocessors are on sale today that play championship chess, backgammon and checkers. Industrial robots work in factories, machines are voice-operated and barcode readers are found in many supermarkets.

Thanks to advances in computer design and new understandings about AI, like knowledge bases, new programming languages, strategies and heuristics, the field has emerged from its wasteland. Today AI is a fast-moving field that, in certain countries, is being heavily funded. Although AI researchers are no longer willing to make the generalized optimistic predictions of the Dartmouth Conference, they can at least point to successful applications and promising areas of study.

II. TOPICS IN AI RESEARCH

AI is a particularly integrated field of study where research workers from many different disciplines pool their knowledge and experience to solve given problems. Although AI research today is divided into a number of specific areas, each particular advance has implications for the whole field and embodies principles from several different areas of AI.

For clarity of presentation, this paper is divided into different sections, each dealing with a particular field of AI. It must be borne in mind, however, that these divisions are to some extent arbitrary for the interests of various AI workers overlap and common problems occur in different fields.

TOPIC I: KNOWLEDGE AND THE COMPUTER

One of the most significant fields of AI research today is that of knowledge representation and bases. Knowledge bases are ordered collections of facts about the world that are used in fundamental work on expert systems, problem solving, computer vision, natural language understanding and speech recognition. The topic in itself also represents an important advance in how the computer's abilities are used. As one participant put it:

> "When the computer moves from processing numbers to actually dealing with knowledge itself, then it has made a tremendous evolutionary leap."

During the Ottawa Workshop, there were several talks and discussions on how knowledge can best be represented and manipulated by a computer. The participants also debated which programming languages are most appropriate for knowledge representation and how knowledge bases can be used in question-answering machines and expert systems.

Expert Systems

Expert systems represent the most visible success of AI today. They are of obvious practical importance to businesses and to professionals and a clear demonstration to non-experts, of how computers can solve problems in an intelligent fashion.

A definition of this important new area was given by one speaker:

> "Expert systems are computer systems which attempt to duplicate or approach the performance of an expert in his field."

Expert systems are designed to capitalize on the knowledge of experts in certain fields. When such systems are provided with background information on a particular case or situation, they are able, for example, to make a medical diagnosis, predict a chemical structure or suggest the location of a mineral deposit.

Some of life's problems can be solved by the mechanical application of simple rules, for example, setting up a small balance sheet or tracing the fault in a simple machine. These problems can also be solved by a well-written computer program and are not the province of expert systems.

Dr. Douglas Skuce of the University of Ottawa explained that the ideal field for an expert system is one in which the problems involved are more subtle and cannot be directly solved with a limited number of rules. When doctors make a diagnosis, for example, they use rules of thumb, intuition and judgements based on partial information, and combine all this with a lifetime of accumulated knowledge and experience of similar cases. The expert system attempts to formalize this tacit knowledge in a way that can be used by the computer.

A perfect expert system not only gives accurate solutions to complex problems but explains how it arrived at its conclusions. It learns from experience, restructures new knowledge and detects internal inconsistencies in its knowledge base. The perfect system also realizes what is relevant to a problem and what can be neglected and, like any good expert, it is aware of its own limitations. Although commercially available expert systems incorporate several of these features, they all have limitations and none has yet reached the level of a true human expert.

During the late 1970s and early 1980s, a number of expert systems were designed and marketed, several of which are explained in Table I. The systems PUFF, MYCIN and CADUCEUS have a high success rate when it comes to the diagnosis of a limited range of illnesses, DENDRAL and CHRYSALIS are able to interpret mass spectra and electron-density maps respectively, PROSPECTOR indicates the location of minerals and has already correctly predicted the location of a molybdenum deposit. Other expert systems are used in the design of electronic equipment to aid computer programmers and as mathematical assistants. Expert systems also form the basis of some computer-aided learning systems.

Table I - A Selection of Expert Systems

System	Function
PUFF	Diagnoses and recommends treatment of pulmonary dysfunction.
PROSPECTOR	Aids in the exploration of nine important minerals.
DENDRAL	Suggests organic molecular structures based on mass spectrograph data.
MYCIN	Diagnoses bacterial infections and recommends antibodies.
VM	Interprets clinical data on a patient's breathing.
AM	Discovers mathematical concepts.
EL	Analyses electronic circuits and determines voltages and currents.
GUIDON	Uses the MYCIN and PUFF expert systems as a basis for teaching medical diagnosis.
ABSTRIPS	Devises plans for a robot to move objects between rooms.
MOLGEN	Designs experiments in molecular genetics.
CHRYSALIS	Interprets the electron-density maps of protein molecules.

Some expert systems operate at a high level of knowledge and performance. These systems act as consultants for professionals such as doctors and engineers and, in effect, make the experience and knowledge of a world expert available in many different locations. Other expert systems operate at the level of assistants or apprentices and remove some of the burdens from human professionals.

Notwithstanding their successes, current expert systems are limited in their areas of expertise and, in some cases, do not offer satisfactory explanations as to how they reach a conclusion. Also, current expert systems have the irritating habit of "crashing" when they reach the boundaries of the field they are programmed to cover. Rather than indicating that they have reached their limits, these systems continue to offer solutions with confidence even when these solutions are absurd.

A major barrier to the development of new expert systems lies in the time taken to extract knowledge rules and heuristics from a human expert and the problems of incorporating these rules into a computer program. At present a knowledge engineer must work with the human expert for many months or years. During this period, the engineer attempts to learn the different strategies and approaches that the expert employs and converts these into an expert program.

Human experts are not always clear as to how they use their intuition and experience and the engineer may have to begin with general cases and develop new insights by presenting the expert with special cases, exceptions to the rule and instances where a rule must be broken. One possibility for removing this labour-intensive bottleneck is to have the expert system itself extract the knowledge. This can be done, for example, by generating a large number of cases and presenting them to the expert for consideration. (The field of knowledge representation, itself, is an active one.)

Expert systems of the future will have to handle far more rules and facts than present systems. They will be structured to deal with incomplete or imprecise knowledge, they will accommodate knowledge from several different sources and operate in more flexible ways. The research in these cases was discussed by John Mylopoulis, Zenon Pylyshyn and Douglas Skuce.*

There was general agreement at the workshop that the field of expert systems should be pursued in Canada. Some research work is already being done in this field which appears to be ideal for further exploitation. For example, development of the PROSPECTOR system involved Canadian expertise and information on Canadian resources, yet the actual system was designed in California. Its development costs were not large, $2 million. The Canadian AI community should be involved in the production of similar expert systems, in particular systems which would assist in the development and exploitation of natural resources.

Knowledge Representation and Logic Programming

Knowledge representation deals with the ways in which facts and information are structured and manipulated in the computer and how they can be used to make inferences and deductions. It is essential in the design of an expert system and other AI applications where general world knowledge is needed, and involves some of the conceptually most difficult and exciting issues in AI.

In the 1960s, research workers still had as their goals general purpose systems that would play games at a high level and solve complex problems, but such systems were rule based, an approach that was soon found to be limited. By the late 1970s, a different, knowledge base approach was attempted in which computers were given access to general facts about

---

* See p. 69.

particular domains. The success of such attempts depends, of course, on having a proper representation of knowledge within the computer.

Knowledge representation depends critically upon choosing a programming language that is sufficiently precise to make correct deductions yet, at the same time, is flexible enough to contain all the relationships that exist between a wide range of knowledge facts. Natural languages, for example, have considerable flexibility but lack precision. Mathematics and conventional programming languages have precision but are too narrow in their range. A new level of languages, therefore, had to be created for knowledge representations; these were based upon features from mathematics, natural and programming languages. John Mylopoulis of the University of Toronto explained that three basic approaches have been used in the development of such knowledge representation languages.

The first approach makes use of the rules of symbolic logic. For example, statements of the form:

"If X is a student, then X is a person,"

are used to structure the knowledge base. If the computer is later supplied with the particular fact that "John is X" it will be able to deduce that John is also a person. By stringing sets of logical relationships together ("If.... then...," "Either... or..." and so on), it becomes possible to make powerful inferences from a particular given set of facts. Deductions made with this approach are at present wide but are limited to those of traditional symbolic logic. To be of more practical use, the knowledge representation system must also deal with partial or incomplete knowledge, heuristics, beliefs, orders involving time, metaphors, actions, plausability and so on. Research projects are attempting to incorporate some of these more flexible features.

A second approach makes use of the relationships that exist between objects and facts. John of the previous example may not only be a student but also the husband of Mary, the father of David, owner of a dog and a red car, taller than Henry and older than Paul. Indeed, a complex web of relationships interconnect John with other people and objects in the knowledge base. By laying down such relational networks, the knowledge base develops a particularly rich structure, far superior to that currently available in a logic-based representation. However, it is difficult to make new inferences using this relational approach.

A third attempt at representation is "process driven." The programming language makes use of the input of a fact to trigger a series of processes within the knowledge base. The fact "John is a student" may trigger off a process that results in the deduction "John is a person" or "John is the husband of Mary." Representations using processes are currently fashionable amongst the designers of expert systems but, as with relational networks, they are limited in their reasoning potential.

Dr. Myopoulis explained how contemporary research is being directed to a new generation of languages that use features from all three of these approaches. Randy Goebel, of the University of Waterloo, discussed PROLOG, one of these new, logic-based languages. PROLOG is proving particularly popular amongst the AI community and has been chosen by the Japanese for use in their fifth-generation computer. Doug Skuce, who had earlier talked about expert systems mentioned the LESK language that is designed for knowledge acquisition from an expert in a "user friendly" way.

There was some discussion at the meeting on the merits of various languages but one participant observed:

> "The problem with all these languages is that people act like apostles and push the language they believe in. The field's being spoiled by so many unjustified claims."

Another suggestion was to avoid a totally unified approach in a computer system and use one language for knowledge representation and another for control.

The concepts of "Frames" and "Scripts" were also mentioned in relation to knowledge representation. "Frames" were devised at MIT by Marvin Minski and involve stereotypes of knowledge or clusters of facts that can be manipulated as a block. "Scripts" were developed to cover common situations in which humans often find themselves involved. A well-known example of a Script concerns the type of action that a customer in a restaurant might experience. By means of this Script, the computer becomes aware of a range of human experience that includes goals, beliefs and plans besides hard factual knowledge.

Present systems of knowledge representation are limited to several hundred facts and the workshop participants agreed that a whole new range of applications will open up once these systems are able to handle thousands or tens of thousands of facts. The barrier to such an increase appears to lie with the knowledge representation language itself. Some speakers felt that a significant breakthrough in representation is needed before knowledge bases increase in size by an order of magnitude. However, to give an idea of the scale of the problem involved, Myopoulis pointed out that some of the issues of knowledge representation involve questions that have preoccupied philosophers for many generations without solution.

Question-Answering Systems

Computer buffs aside, most of us would prefer to talk to a computer in our own everyday language. A person who uses the computer as a tool in an office, home or factory is generally uninterested in the way in which the program is written or how the knowledge is structured. Such a person does not wish to learn a programming language or to spend several hours with a system manual. The ideal situation would be to ask the computer a direct question and receive an answer in everyday language. Based on this answer, the user might then choose to enter into a more detailed conversation with the computer and ask for further facts, explanations and predictions.

The goal of question-answering sytems is to produce a "user friendly" computer that will hold an intelligent and productive conversation with a human user. The conversation itself could be held using speech or typed questions and answers in ordinary language. Expert systems are particularly appropriate candidates for augmentation by a question-answering system. As one participant put it:

> "What we want is the sort of medical expert system where a worried mother can lift up the phone and explain her child's symptoms. The computer would tell her not to worry or suggest she brings the child in to see the doctor."

Ray Perrault pointed out that question-answering systems must have capabilities that lie far beyond the understanding of natural language - a difficult enough task in itself. Such systems would have to deal with the general flow of a conversation, they would have to understand its context, to refer back to earlier answers and to deal with ambiguous questions.

Current systems are stretched to the limit by even the most simple conversations. A question-answering system may, for example, be able to deal with:

> "What Canadian manufacturing company's profits were higher than average in 1980?"

But what will it make of the follow-up?

> "What do they make?"

Here the ambiguity of the word "they" can only be resolved if the computer remembers and understands the previous question and answer. The word "make" is even more difficult for it could refer to profits, income or products.

The first generation of such systems had their language-generating programs integrated with rules for the semantics and context of the conversation. They were fast and versatile, but were incapable of dealing with subtleties of language; also, any change to the data base involved modifying the whole system. A second generation of machines had data bases that were independent of syntax generation. This step allowed for a richer use of language and more acceptable replies but required very large computers and more computing time.

At present, commercial systems are unable to hold connected conversations where any subtlety or dependence on context is involved. However, some research systems that are able to make deductions during the conversation have greater versatility. If such systems are perfected they will have considerable commercial possibilities.

Dr. Perrault pointed out that, as with so much AI work, the key to the problem is to limit the system's domain of application. Within a well-defined domain, question-answering machines can be effective and useful, but a problem arises in discovering the best way to limit the domain of a question-answering system. For example, even with a very limited data base, an involved question still requires considerable computing power to produce a satisfactory reply. One useful line of research would be to discover guidelines about the most profitable ways to limit the domain of a question-answering system.

An Aside! Tree Graphs and Combinatorics

Although none of the speakers specifically addressed the topic of tree graphs or solution searches, the issue was implicit in much of the workshop discussion. Tree graphs and the various strategies that are employed to search through the branches of a solution tree are ever present problems in AI research. The concepts involved can be readily illustrated by reference to a game of chess.

Suppose Black is to move at a particular point in the game. Black considers the prudence of a given move and envisions the possible countermoves that White could make. Each of White's countermoves will in turn open up new possibilities to Black; some of them will strengthen Black's position and others lead to danger. Black therefore considers not only the immediate move but the possible responses to each of White's countermoves and so on. The implications of a single move can therefore be represented by branches on a tree:

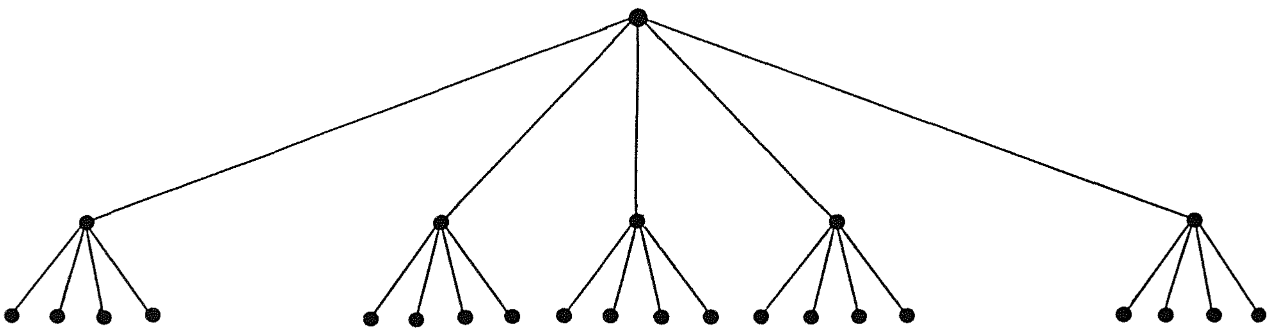

It is clear that the further ahead Black plans the game, the larger and larger become the number of possible moves about which Black has to think. If Black plans far enough ahead, then the number of moves involved become astronomically large. Within all of these possibilities, Black hopes to trace down a branch to a possible checkmate of White. Black also bears in mind the possible sub-branches which, if taken by White, may cause Black to lose the initiative. The state of the game can therefore be considered by scanning the implications of each node, by searching the branches both laterally and in depth. Finally, a particular strategy which involves zigzagging down a branch from node to node will be chosen.

Many problems in AI are analogous to the chess game for each particular move or action gives rise to a range of possibilities that fan out like twigs on a branch. In some AI situations, the process is performed in reverse. Diagnosis of a disease for example starts with the twigs (symptoms) and back-tracks to a particular disease (one branch). However, whichever way the tree graph is read, AI researchers always come up against the "combinatorial explosion", which is so-called because as the computer moves from level to level within a tree graph, the number of possibilities that must be considered increases dramatically.

Tree graphs crop up in AI applications including vision, speech recognition and problem solving. In all cases the computer reaches its limit when it no longer has the time or the memory capacity to search through so many branches. Some idea of the magnitude of the task can be gained by referring back to the chess game. If the fastest computer available were to calculate all possible configurations in a 50-move chess game, it would take it longer than the age of the universe!

Human chess players do not get caught in the combinatorial explosion; yet are still able to play long games of chess and solve complex problems. This is because humans have the advantage of drawing on experience, anticipating certain overall patterns, adapting strategies for given situations and so on. In chess, a human player may learn to concentrate on the centre of the board, recognize an opponent's overall strategy and condense a particular series of moves using mental shorthand.

AI researchers are similarly learning how to beat the combinatorial explosion by means of heuristics, generating and testing tentative solutions, breaking down a problem into sub-problems and searching both from the top and bottom of a tree. Also, new approaches to programming favour parallel processing in which several branches are searched simultaneously.

The difficulties of searching amongst the branches of a tree graph is one of the major computational barriers in many fields of AI. Researchers now recognize that a breakthrough in AI will come about once the computer has self-knowledge about the task and can perform a tree search in an intelligent fashion.

Conclusions

It was clear from the discussion that expert systems and knowledge representation is a particularly active and promising field. Several expert systems have already been marketed and are generating considerable interest within the business world. It is clear that as new systems are developed, they will have a readymade sales interest.

Several of the participants argued that the design of expert systems was a particularly promising field for Canada, especially in the field of resource management. The development of a new system does not require a large AI team or a major outlay of capital. The knowledge and experience to design and produce such systems is available in Canada, and such a product would have a promising worldwide market.

TOPIC II: NATURAL LANGUAGE AND MACHINE TRANSLATION

Computer understanding of natural language is of considerable importance in AI applications. Question-answering systems depend upon the ability to understand and generate sentences in natural language. Speech synthesis and recognition rely similarly on the ability of computers to manipulate and understand natural language. Also, an application of particular relevance to Canada is that of machine translation.

Language translation is an expensive, labour-intensive activity and, in a bilingual country like Canada, it represents a significant cost that has to be added to the overhead of many projects. The possibility of automated machine translation is, therefore, of national importance.

Early attempts at computer translations were crude and had a tendency to produce absurdities. A much-quoted example involved the translation of the English sentence "The spirit is willing, but the flesh is weak" into Russian and back into English. The translation produced by the computer read "The wine is agreeable but the meat is spoiled." There are a number of outstanding problems involved in language translation and this example illustrates one of the most serious, namely, the problem of meaning. Words carry many levels of meaning, as for example the word "flesh". The Oxford English Dictionary contains two pages of definitions for this word which involve subtle shades of meaning in regard to its common usage as a noun and such curious usages for the verb as, "to reward a hawk or hound" and "to plunge a weapon into the flesh." Some of these meanings are made clear by the sentence in which the word occurs, others require a knowledge of the general context of the passage or even information that lies outside the text itself, for example, general knowledge about the world and the way people act. Such considerations become of supreme importance where computer translations are concerned.

Pierre Isabelle and Laurent Bourbeau described the approach of the TAUM group in Montréal and the various attempts that had been made to improve machine translation. The field had been influenced during the 1970s by the theories of Noam Chomsky who proposed the existence of "deep linguistic structures" located in the brain. According to Chomsky, language arises from language-independent processing that is genetically inherited by all humans. The particular language that is learned in childhood is coupled to this deep linguistic structure so that not only speech but also reading and writing become possible.

Chomsky's theories suggested to the AI community that a language independent structure or "interlingua" could be used as the basis for computer translations. Sentences in one language would be transformed into symbolic forms of "interlingua" and then transformed into the target language. In practice, however, such a translation program proved particularly difficult to carry out so TAUM took a modified route.

The Montréal-based TAUM project was formed in 1965 from a group of linguists, translators and computer scientists to develop a practical computer system for English to French translation. TAUM decided on an approach using "pivot languages"; a modified form of deep linguistic structure. Each English sentence is first analyzed for syntax, semantics and word meanings and then translated into "pivot English." This intermediate step has the effect of exposing the full linguistic structure of the English sentence. Using an internal dictionary and rules for transformation, the sentence in pivot English is then converted into pivot French. Finally, the TAUM program translates the sentence from pivot French into the French language.

The TAUM program is not sufficiently general to make translations of books and newspapers so the group decided to apply it to what are called "sub-languages"; that is, situations, such as technical manuals or command languages, in which syntaxes are restricted and vocabularies are limited. One of the first applications of the TAUM program to such a sub-language was the translation of weather reports. TAUM-METRO has translated some 2.5 million words each year with only 20 per cent of the sentences requiring revision. This success was followed with TAUM-AVIATION for the translation of maintenance manuals.

An advantage of the TAUM system is the independence of its software from internal grammars and dictionaries. It is, therefore, possible to apply TAUM to different technical sub-languages or, indeed, to languages other than French and English without the need to revise the whole system.

The cost of a TAUM machine translation was $0.083 per word as compared with a human translation of $0.145 per word. However, because all texts were subject to human revision, this raised the cost of a TAUM-produced and translator-supervised text to $0.183 per word. On the basis of these figures, the TAUM attempt did not appear to be commercially competitive. TAUM's contract to perform machine translations was terminated by the Department of the Secretary of State and the research group disbanded in September 1981.

The failure of the TAUM translation team to survive was regretted by many present at the workshop. However, an announcement was made by a member of the Secretary of State's Department that a study would be funded to look into the feasibility of computer translations.

James Mackie from MITEL pointed out how useful computer translations could be when they are used with restricted sub-languages. MITEL was making use of a commercially available system to translate its telecommunications documentation into a variety of languages. The original text is written in clear, simple English using a restricted syntax and vocabulary.

By keeping the limits of the computer's translation program in mind when the original text is written, the machine is able to produce acceptable translations at an economical cost.

TOPIC III: VISION AND SPEECH

The fields of expert systems, knowledge representation and the understanding of natural language belong, loosely speaking, to the "intellectual" side of the computer. The electronic machine must also be given senses and the ability to gather information about its environment and to use this information to take action. As one participant said:

> "So much of this workshop is about thinking, yet a lot of the actual work in AI today is about getting information from the world and then altering it."

The earliest computers interacted with the outside world by means of punched cards, punched and magnetic tape, discs, drums and typed input and output. Today, computers have the rudiments of vision, understand simple speech, sense the position of surrounding objects by means of robot limbs and interact with humans via sensitive screens and tablets. The two major areas of computer vision and speech were discussed at the workshop.

Computer Vision

So much of what we know about the world is gathered by our eyes that it is natural for AI researchers to attempt to give vision to their computers. It is a task, however, that proves to be one of the most difficult in the whole AI field.

Some of the earliest experiments in computer vision were carried out at the SRI Laboratories in California between 1969 and 1971. Using television cameras as eyes, SHAKEY I and II were able to find their way through a series of interconnected, block-filled rooms.

As scientists attempted to extend the abilities of computer vision, problems emerged that were of surprising complexity. Our own vision works so well that we are never conscious of the complicated processes involved. We scan a landscape and, faster than thought, spot the sudden movement of a rabbit. We walk through a jostling crowd and recognize the face of a friend last seen a decade ago. We run downstairs and notice a mark on the wall that was not there the day before. As one participant put it:

> "Our visual system is truly astounding but it is also the result of a billion years of selective evolution. A billion years which has produced eyes, optic nerves and a visual cortex specifically designed to perform tasks geared to survival."

There does not exist even a tentative general theory of vision, and set beside the visual system of a human, SHAKEY and the vision systems that came after it are particularly crude. Contemporary research has moved away from computer vision towards more restricted problems and to vision systems that work in limited, structured environments. Research topics include image processing, picture restoration, scene analysis and a host of other specialized sub-divisions.

One of the most notable successes of computer vision is in areas where human seeing has no special ability; that is, the fields of image enhancement, picture correction and reconstruction, image compression and tomography. Here the computer's "eye" makes no attempt to read or understand a photograph but treats it as digital information that is to be processed. Computers can remove defects in a photograph, correct colour values and generally improve a two-dimensional image.

Picture processing is now a matter of routine with images taken from planetary probes, earth resource satellites and medical apparatus; it is even used in commercial video production for special effects. Computer tomography is applied in medical diagnosis to obtain three-dimensional information of the body from a series of flat-slice images. Image compression permits the storage of a larger number of visual images in a computer memory or the transmission of images with a reduced flow of data.

Computer systems can also be used for pattern recognition where a limited number of simple shapes are displayed against a uniform background. Applications include the reading of barcodes on merchandise, counting particles, inspecting electronic circuits and scanning blood cells. The computer has the advantage of being fast, efficient and tireless. Its vision can also be used to measure the dimensions of simple changes directly and accurately, an ability not shared by the human eye.

Whereas the computer may be able to improve the image of a blurred photograph, it is generally beyond its abilities to pick out and identify even the most general solid objects that are represented on that two-dimensional surface. It is only in restricted situations, such as photographs of the earth's surface taken from a given height, that computer vision has, at present, any chance of making a commercial contribution.

Humans know how to read photographs because they possess general knowledge about how things appear in the real three-dimensional world. Humans are able to decode the flat image because they understand clues as to scale and distance, they know how shadows form, they associate certain textures with grass, water and rock, they recognize a given object from many angles and they know how solid objects are mapped onto a two-dimensional surface by the camera.

The computer can only begin to decypher a photograph when it too is supplied with additional information about the scene. One way of doing this is to use the computer vision system as an assistant to a human expert. T. Kasvand of the National Research Council and A.G. Fabbri of the Geological Survey of Canada developed GIAPP to place image-processing skills at the disposal of geologists who wish to pick out and compare details from a series of maps.

Professor Mackworth from the University of British Columbia described his own group's work on the interpretation of satellite pictures. The computer is able to "read" the image only after being supplied with information on illumination, ground cover and surface orientation. Another experimental program, MAPEEZE, is able to reconstruct the scene represented by an aerial photograph. A human operator provides the computer with a rough sketch map of roads, bridges, rivers and mountains in the area. MAPEEZE then uses that information to produce a full colour map based on the photograph.

When it comes to general vision of full three-dimensional scenes, the computer is even more limited. One AI goal is an industrial robot whose eyes can be used for inspection, general control and the direction of manipulating arms. Such robots are, at present, only effective in highly structured environments involving a limited number of different objects. However, by carefully arranging a scene, the robot can be used to advantage. In an ideal vision situation, a limited number of well-defined objects are carefully arranged and illuminated. The simple expedient of painting a target on an object will markedly increase a robot's vision ability. One participant defended the robots' limited vision:

> "You've got to remember that the workplace has evolved for the convenience of humans and not robots. What may seem congenial to us is confusing and complicated to a robot. They're going to work best in the sort of standardized surroundings that humans find boring, mechanical and repetitive."

Mackworth felt that Canadian research into computer vision is of high quality but, without additional support, it will rapidly fall behind the rest of the world. One possibility for the future is the transfer of high-quality visual systems onto silicon chips. In this way, new vision systems will be fast in operation and cheap to produce.

Whereas other applications of AI, such as machine translation, may have to compete on the grounds of cost effectiveness with humans, this is not necessarily true with computer vision. A particularly attractive area for development is vision for robots to be used in outer space, ocean beds, atomic reactors, chemical works and other environments that are hostile to humans.

Speech Recognition and Synthesis

It is advantageous to interrogate computers using ordinary speech, for computers that can engage in verbal conversations with their users will have a host of applications from improved question-answering systems to aids to the handicapped and machines directed by verbal commands. However, as with other fields of AI, the full development of computer speech also depends on having electronic representations of general knowledge about the world.

Dr. M.L. Blostein explained the work of the Bell Northern Research - INRS Speech Communication Research Group and began by playing an example of computer-generated speech. True synthetic speech involves the translation of a text into acoustic signals. At present synthetic speech is intelligible although, in the case of microprocessor systems, considerable improvement is still needed particularly where long sentences are involved. An alternative approach is to construct speech out of pre-recorded, electronically compressed segments but this requires the capacity of a whole IC chip (integrated circuit chip).

To produce good synthetic speech, the computer must be able to analyze each sentence and apply rules of inflection, punctuation and intonation as well as the characteristics of an individual speaker. A major difficulty in the development of computer speech is the lack of full understanding of how human speech is produced.

When it comes to automatic speech recognition, the problem is even more complex. At present a single IC chip has the capacity to recognize a small vocabulary of isolated words spoken by a trained speaker. To extend this to carefully spoken, connected speech requires a series of circuit boards and even with this computing power, the vocabulary and syntax must be highly restricted. Full conversational speech with untrained speakers is at present beyond the capacity of a computer. It is difficult to give a time frame for this latter development. Some laboratory systems work well as experimental devices but as soon as they are put into the field, they look less attractive. As a participant said:

> "It's not that easy to speak to one of these machines. We don't really know how to get people to 'live' with a limited vocabulary and syntax. Possibly command languages will be the most profitable area to explore."

Speech-recognition systems work by matching the incoming sound against acoustic patterns generated in the computer, and, if the two patterns happen to match then the word or phrase is recognized. This process requires the computer to constantly anticipate what is being said, given the overall context and earlier parts of the conversation. Hence, the computer is constantly making sentence hypotheses based on the rules of natural language and some general knowledge about the subject under discussion.

Improvements in speech recognition involve knowledge representation, strategies for constructing more effective hypotheses, searching amongst hypotheses, pattern matching and so on. Also, more work must be directed to understanding how human speech is produced.

Dr. de Mori of Concordia University spoke of his own research into the computer recognition of continuous speech. His approach employs an integrated system to extract acoustic clues and generate syllabic and lexical hypotheses. It operates by processing the various sub-tasks in parallel.

Dr. Blostein argued that significant research in this field could only be carried out by an organization with considerable computing power at its disposal. He suggested that successful systems should be transferred immediately to silicon chips. Cheaply produced speech-recognition systems could then be used by smaller research groups and for field trials in homes and offices. The results of these trials and laboratory assessments would lead to further improvements in design.

A good speech-recognition system probably lies 5 to 10 years into the future. For the present, however, commercial systems have a use in well-defined situations and with limited vocabularies and syntax. One obvious field of application is "command languages" where machines are controlled by a series of limited voice commands. Although the voice-operated typewriter lies far in the future, research could be carried out now on how people will react when they find that they can hold conversations with a machine.

TOPIC IV: APPLICATIONS AND EFFECTS

The Office of the Future

Much has been written about the office of the future and how the new technologies will transform the functions within an office. Up to now, the major transformation has resulted from word processors and, in the near future, their influence will be even further extended as they are given the capability to make decisions, produce graphs and other documents and are integrated with electronic mail services.

Tom Carey of the University of Guelph pointed out that the research and development is badly needed in the field of user-system interactions. Office equipment of the future will make use of natural language, speech recognition and multi-mode interfaces. (In multi-mode interfaces several "clues" are given by the user, for example, gestures can be picked up by means of a gyroscope attached to the arm.)

The general user of office equipment prefers to learn to use a new system in ways that are as painless as possible. In place of complicated instructions and bulky manuals, the system itself should guide the user through its various functions at an individually designed pace. Instead of error messages, an electronic device could offer help and instruction. Executives who find themselves lost or confused when faced with a computer are unlikely to be impressed by a bluntly worded error message.

At present, it is not at all clear how this new technology will change the office. For example, when word processors were first marketed, it was assumed that their major function would be to act as time-saving devices. In fact, they do not so much save time as allow writers to produce a much better first draft. One function of the new machines may be to make the skills of an expert more readily available throughout an office, but, at present, no one is at all clear if the new technology will reduce the number of office jobs or create more.

Part of the problem arises because the functioning of the office is not well understood. On paper, for example, an organization may function through hierarchical lines of command but there will also exist a more subtle, informal structure within that same organization. This informal system can be seen during a coffee break. The way in which new office machines interact not only with the formal office structure, but with this informal sub-structure, will be of particular importance.

Long-Range Planning and Decision Making

Computers have traditionally been used by scientists to solve problems that can be rigorously formulated. However, once the element of AI is added, the computer becomes capable of more subtle tasks such as decision making, value judgements and planning for the future.

In arriving at a policy decision, a human expert takes into consideration the potential behaviour of the physical world (resources, economics, environment, etc.) together with the various goals and values that society supports. This process demands, to begin with, an understanding of the complex interactions that exist between, for example, the marketplace, global economics, farming, mining, manufacturing and employment. Also, policy-makers are sensitive to the values society sets upon certain ways of life. They must, therefore, take into account beliefs and goals together with the possible implications of the particular decision that is to be made. Thus, if all or part of this task were to be carried out by computer, it would be of immense importance.

George Strobel of the Université de Montréal discussed the ways in which intelligent computers can assist in decision making and long-range planning. One approach is to use the computer as an experimental tool for investigating the consequence of a particular policy decision. Economists and systems theorists have constructed dynamic models that represent the evolution of complex interrelated systems such as the environment and the economy. These mathematical models consist of sets of coupled differential equations together with initial conditions, adjustable parameters and so on.

By programming dynamical models on the computer, a human expert can experiment with the effects and long-range implications of different scenarios, policy decisions and world trends. The computer in this case acts as a tool to decision making, whereby a number of alternative policies can be explored through a process of trial and error. The assessment of the results, in terms of values to society or detriment to the environment, is still, however, taken by human experts.

A more advanced form of computer planning not only models the economy and the environment, but the processes of decision making itself. In the GRIPS system at the Université de Montréal, the computer is supplied with information on the relative importance of needs, values and goals together with various indicators such as pollution, life expectancy, existence of war, etc. By combining these with a model of the environment and the economy, the results of a particular scenario can be evaluated by the computer.

The GRIPS system offers a complete modelling of human decision making where complex issues are involved. The computer works out the environmental and economic implications of a particular policy, searches within this scenario for the various indicators of quality of life and produces an overall evaluation of the result. The GRIPS system also features "automatic programming," it accommodates "fuzzy" input and supplies the results of its analysis and planning in terms of graphs and tables.

There was little time available at the workshop to discuss the far-reaching implications of systems like GRIPS, but the possibility of evaluating policies, making decisions and carrying out long-range planning by computer should be of importance to industry and governments.

HARDWARE AND THEORY

Computer Hardware

Much of the discussion at the workshop focussed on the practical tasks that AI systems are called upon to perform and on the software that should be developed. Professor T. Pietrzykowski of Acadia University provided a welcome contrast by addressing the problems of computer design and manufacture. AI systems today are required to carry out pattern matching, extensive tree graph searches, specialized numerical calculations and memory operations as well as parallel processing. Yet the computers available are still manufactured according to a philosophy used in their ancestors of the 1940s and 1950s.

Von Neuman's original design called for a centralized control and data flow but this is particularly inappropriate for the type of computer used in AI work. Today, there is good reason why computer design, or architecture, should change radically. For example, the central processor once represented the major cost of a computer but in some current machines, it is about one-tenth the cost of a memory unit. Novel designs could use tens or hundreds of processors to direct and control the flow of data through various parts of the machine. Areas of the computer could be dedicated to special operations or hard-wired for vision and speech analysis, intelligent memories could direct the flow of data and special modules would provide a well-designed interface with the user.

Such radical changes in computer architecture are not beyond the capabilities of the Canadian AI community. Pietrzykowski did not believe that the research and development involved would require massive grants or large teams and could be handled by relatively small groups. Rather than attempting to build a general purpose fifth-generation machine, such a group could concentrate on producing an intelligent computer for a specific task. AI computers could be designed for robotics, vision, speech recognition, data base fronts or expert systems.

A computer manufacturer at the workshop agreed that the design and manufacture of such a machine was well within the capabilities of the Canadian AI fraternity. Such a development would allow Canada to leapfrog over conventional computer technology and secure a position amongst the next generation of computer manufacturers.

Role of Theory

As a complement to Dr. Pietrzykowski's talk on hardware, Ray Reiter of the University of British Columbia spoke on the role of theory in AI work. AI theory has recently experienced a renaissance in the fields of logic programming, expert systems, metatheory, new forms of reasoning and "smart data bases." In the area of "metatheory," contemporary research is being directed towards "knowledge about knowledge" or intelligent control over the computer's internal processes. This also involves the ability to deal with incomplete knowledge and belief systems. Considerable theoretical work is also involved in the development of logic-based languages for dealing with knowledge about the world and for making inferences from this knowledge.

Dr. Reiter pointed out that these theoretical studies relate to profound philosophical problems that had exercised thinkers for the last 2000 years. He had no reason to be confident that such questions as the "meaning of meaning" or "belief," would be resolved in the immediate future. Indeed, the word of John McCarthy, one of the founders of AI, are worth recalling:

> "One can expect true artificial intelligence 'somewhere between four and four hundred years' in the future."

GENERAL SUMMARY

It became clear from the talks and discussions held in Ottawa that Canadian researchers are represented in every important field of AI. (There were no specific discussions on robotics but the National Research Council is building an active group to study robotics and the applications of AI to manufacturing technology. Companies like SPAR Aerospace, with CANADARM, are also developing the required mechanical skills.)

Although Canada has not developed an expert system of its own, the community has all the necessary expertise to design and produce a commercially attractive system. Machine translation is another field where skills, a ready market and national priorities coincide. Considerable advances have also been made in the areas of long-range planning and decision making using AI computers. The field may now be at a stage where it could be of definite practical use to policy-makers and planners.

The technically difficult areas of computer vision and speech recognition are also well represented in Canada. Although the final solutions to many of the problems involved lie in the unpredictable future, there are many important applications to be made along the way. Already aspects of computer vision and picture processing are being used in science, medicine and industry and year by year, more applications are expected.

In theoretical AI, Canadians continue to work on some of the most fundamental problems. The depth of the issues involved transcends a single discipline such as AI and embraces fundamental questions in philosophy. These problems may, by their very nature, lie beyond solution even though the human race is involved in rephrasing and clarifying them from century to century. This process is a valuable part of the world's intellectual heritage, and it may also have the side effect of stimulating work in more practical fields.

Some consideration was also given to the problems of computer design and hardware manufacture. Canada is rapidly developing skills and experience in the microelectronics industry and there are good arguments to support a national initiative in the production of AI computers. There is plenty of room for competition and innovation in this area and Canada's best strategy may be to design intelligent electronic devices for specific fields of application.

III. A POLICY FOR AI

THE ISSUES

What role should governments play in AI research?

Should Canada build a special centre for AI work?

What are the national priorities for AI research and development?

How soon will AI applications reach the marketplace?

Questions like these were debated at special sessions, question periods and coffee breaks during the two-day workshop. Whereas participants were in general agreement that AI is an important and promising field, they were less clear as to which policies and strategies would foster this work in Canada. High-quality work is at present being carried out by a number of smallish groups but unless a rapid increase in funding and support takes place, Canada will soon trail behind Japan, Europe and the United States in development and applications.

Research Funding

During the first day of the workshop, a senior government representative led a discussion on the possible roles that government could have with respect to AI research. He pointed out that one of the most important stimuli for research in Canada is the Natural Sciences and Engineering Research Council (NSERC) grant scheme. These grants are designed to support long-term research and personnel development and could have a considerable effect on AI work.

Where applications are of ultimate benefit to society, governments can also fund research through strategic grants. Some government departments can participate in a more direct form of support when applications coincide with a particular mandate. For example, the Department of Energy, Mines and Resources supports remote sensing for natural resources. The National Research Council maintains an interest in robotics and computer vision for industrial applications, the Department of the Secretary of State encourages machine translations, and the Department of Communications has a tradition of development in communications technology and office equipment. In all these cases, a government department may fund research teams or engage directly in research and development.

The Workshop was asked to consider to what extent should government take the lead in Canadian AI. Should it give direction through active policies and research, or should it act as a coordinator and help to foster communication between interested groups? In the lively discussion that followed, research workers from the universities made it clear that they felt, themselves, poorly served by NSERC. They pointed out that if AI is to flourish in Canada, then some improvement in funding is needed rapidly.

Several suggestions were offered as to why NSERC support is so low. To begin with AI is a new and poorly defined topic that embraces research workers in computing, engineering, mathematics, physics, psychology and linguistics. Often a grant application will "fall between two stools" or be considered by a committee that is unsympathetic to the field. Also, it appears that no active AI researcher sits on an NSERC committee, and no committee is particularly disposed towards the field of AI.

A further difficulty is that AI research normally involves a team effort and NSERC may not look favourably on research teams. Also NSERC grants are designed for long-term research and some authorities question if most AI profits are of a long-term duration. The workshop argued strongly that NSERC should develop more understanding of the nature, significance and future of AI research. Workers in the field should lobby NSERC members, representatives of the AI community should sit on NSERC committees and possibly a special NSERC committee should devote its attention to the field and declare AI a high-priority topic.

Government: Leader or Adviser?

A senior government representative asked whether the federal government should take a strong lead in AI research or act as adviser and coordinator. Many participants felt that the government's role should be an active one in stimulating R&D. They pointed to the $48 million that will be spent between 1982 and 1984 on the Japanese Fifth Generation Project (a national program in Japan involving the design and production of an AI computer), and a $120 and $200 million budget is projected for 1985-1988. The response of the United Kingdom to this announcement had been rapid and resulted in the establishment of the British Program for Advanced Information Technology with a budget in excess of $600 million over the next 5 years. The United States for its part continues to maintain several major industrial and university AI research laboratories. By contrast, the Canadian government appears to have done next to nothing.

Many of the participants argued with the utmost vigour that the Canadian government must be made aware of the importance of funding AI research before it is too late. One suggestion was that the Science Council should act as an advocate for the AI community and, in particular, should make an analysis of the economic benefits of AI and its future applications.

Another speaker pointed out that the Canadian government is lagging behind the general public in its understanding of these issues. He felt that through exposure to the various media, the public had become well aware of the importance of AI and the radical changes that were taking place in Japan and other countries. The public also realizes that Canada is failing to participate in this challenge.

There was general agreement that the AI community should lobby government at all levels to publicize its message. An important factor in raising the nation's consciousness about AI will be the Canadian Society for the Computational Study of Intelligence.

Leadership: Tomorrow is Too Late

The Workshop participants, therefore, felt strongly that the government should move rapidly to demonstrate its support of AI and not wait until proven applications or breakthroughs emerge. One speaker pointed to the way in which a new technology can dominate the marketplace and gave the example of how Japanese manufacturers have captured 98 per cent of the video recorder market. This did not come about through any dramatic leap in technology or fundamental Japanese advance. Rather video recorders evolved through a gradual process of improvement. Many other industrial nations had the capabilities and know-how to make such advances but, in the case of Japan, the project was vigorously supported from its inception. The consequence of the Japanese confidence in the future of its product was that a totally new market was captured and held.

The same dynamics may apply to the various areas of AI. Rather than involving a single dramatic breakthrough, the field will probably advance through a series of applications and improvements each of which represents a modest but definite advance. If Canada is to have a role in this process and a slice of an enormous potential market, then the government must demonstrate a strong and long-term commitment to R&D. If this support does not come about quickly, then Canadian companies and research workers will lose their positions in the field and will rapidly fall behind the other technological nations.

An AI Research Centre

If the government does decide to give financial support to the AI community, how best should this funding be used? One solution is to create an AI Centre. Canada cannot boast any major AI laboratories that are privately funded. Most of the expertise is concentrated in university departments, in small groups that tend to be underbudgeted and overwhelmed with the need to train graduate students. Although the research carried out in these groups is often of high quality, the groups themselves are spread too thinly across the country. What is needed is to bring research workers together into large groups where a "critical mass" can be created for stimulating, rapidly advancing R&D.

One possibility for larger research teams is to have a special NSERC committee fund two or three main research centres. Another scenario is to create a distributed laboratory in which smaller groups can share resources and ideas. AI researchers need access to state-of-the-art computing facilities. At present, this involves a capital investment of between $50 and $150 thousand per researcher. A major research centre, or distribution centre, would provide such computing power and the infrastructure and funding for long-term software development.

Not everyone agreed that the poverty of university research is the major problem. Some felt that the industrial sector is in most need of government support and suggested that an AI centre for industrial applications should be built.*

A further implication of the present inadequate funding is that it leaves talented graduate students with little option other than to move to foreign (generally US) laboratories. This is not only a matter of job vacancies but of the presence of stimulating and active groups where promising young minds can feel challenged. The creation of special AI research centres in Canada would have the effect of protecting home-grown talent and attracting some of the best research workers on the international market.

If such an AI research centre is built in Canada, there are good arguments that applied research should occupy a strong position within its organization. Historically it has been found that the best advances and

---

* See p. 65.

the most stimulating environment exists when an active applied group interacts under the same roof with basic research.

Several speakers made the point that such research centres function best when they maintain strong links with the universities. Possibly university research workers could arrange joint appointments at the proposed centre and funding could be provided for "partial release" of AI researchers. The majority of AI work is performed in departments of computing science where the staff are seriously overworked. By providing partial release, selected faculty would be freed from teaching and administration to spend more time on research or to take a joint appointment at an AI centre.

AI and Technology Transfer

On the second day of the workshop, an active debate was held on the problem of technology and applied research. Several speakers from the industrial sector pointed out that the government had provisions to fund both university research and product development but that there is no funding for applied research outside a university setting. Several industrialists warned that the position of applied research within the whole field of AI is particularly weak. Applied research would be strengthened not only by an influx of funding but also by establishing more efficient communication between all sectors of the AI community.

A major problem involves technology transfer. Not every exciting breakthrough at the laboratory bench reaches the level of product development. The TAUM group, which was disbanded through lack of support, is an example of how a valuable new technology may become lost forever.

Experience shows that it is not sufficient for a new process or device to be commercially attractive, provision for technology transfer must also be built into a project right from the start. Too many university researchers do not think about this transfer until it is too late. On the other hand, research from the laboratory bench often appears abstract to people in business, and there may be problems in convincing an industrialist of the significance of an advance.

There was a general plea that AI groups in the university should develop a greater degree of sophistication about the business world and plan carefully for the implications and applications of work they are doing.

There were several proposals as to how the link between the left brain (product development) and the right brain (basic research) of AI could be strengthened. One suggestion was to set up a centre to concentrate on applied research for the industrial sector.

Applied research is an expensive investment which most firms in the high technology field cannot afford. Where research is carried out and a new topic is investigated, each firm generally finds itself retracing the same steps before it reaches a frontier in technology. This duplication of effort is a costly and wasteful business but companies for their part like to keep research results to themselves. Companies with their own laboratories do not like to be involved with other groups but prefer their projects to be secret.

An applied research centre may be one way around this problem of high investment costs and the duplication of results. It would be jointly funded by government and industry with close connections to the universities. Research on common problems in AI development would be available to Canadian industry as a whole. Also, there would be provision for individually funded research, contracted by a particular company, and conducted under secure and secret conditions.

Regardless of how the research organization is set up, there was a general feeling that it should not be completely under government control and should foster traffic between basic research and product development.

Some concern was expressed that by concentrating so much applied research in a single location, the necessary spirit of competition would be lacking. On the other hand, there is plenty of competition from Japan, Europe and the United States. Participants also discussed individual rewards for AI research and several industrialists felt that a successful research worker should have a tangible share in any advance.

POLICY RECOMMENDATIONS

The policy recommendations that emerged from the two-day Ottawa workshop are:

1) The government should become aware of the nature and future of AI research, including economic benefits and future applications. It should also realize the need for immediate response to this rapidly moving field.

2) The government should give firm and long-term leadership to the whole AI community.

3) Special fields should be identified that correspond to national needs, international markets, current AI expertise and an already existing industrial infrastructure. Expert systems for resource development, language translation and AI related to robotics are examples of such fields.

4) An NSERC committee should be set up to fund AI research in the universities as a special priority area.

5) One or more research institutes should be established to develop teams with the necessary critical mass for rapid innovation and development. These centres would provide smaller groups with access to facilities and superior computers.

6) Special attention should be given to applied research and the transfer of technology. An applied research institute, jointly funded by government and industry, should be established.

7) The design and manufacture of AI computers and devices for specific applications should be undertaken in Canada.

8) The Canadian AI community should establish an electronic information network among its members. Also, the Science Council, NRC or some other organization should sponsor workshops to increase communication between all sectors of AI research, development and commercialization in Canada.

A WORKSHOP ON ARTIFICIAL INTELLIGENCE

sponsored by

The Science Council of Canada
100 Metcalfe Street, Ottawa

20 January

9:00 - 9:30	Welcoming Address Dr. Stuart Smith, Chairman Science Council of Canada
9:30 - 9:45	**Introduction** Mr. Jorge Miedzinski, Chairman of Meeting and Deputy Executive Director Science Council of Canada
9:45 - 10:15	**Artificial Intelligence: A Perspective** Professor Zenon Pylyshyn, Director Centre for Cognitive Studies University of Western Ontario
10:15 - 10:45	Coffee
10:45 - 11:15	**Knowledge Representation** Professor John Mylopoulis Department of Computer Science University of Toronto
11:15 - 11:45	**Progress in Automatic Translation and Natural Language Processing** Pierre Isabelle and Laurent Bourbeau (TAUM)

11:45 - 12:15	**Logic Programming** Professor Randy Goebel Department of Computer Science University of Waterloo
12:30 - 13:30	Lunch
13:30 - 14:00	**Computer Vision** Professor Allan Mackworth Department of Computer Science University of British Columbia
14:00 - 14:30	**Question-Answering Systems** Professor Ray Perrault Department of Computer Science University of Toronto
14:30 - 15:00	**Expert Systems** Professor Douglas Skuce Department of Computer Science University of Ottawa
15:00 - 15:30	Coffee
15:30 - 16:15	**The Role of Government and Government Laboratories in Artificial Intelligence** Keith Glegg, V.P. Industry National Research Council
16:15 - 17:30	**Discussion: Directions for AI Research in Canada. What should be done and how to do it?**

21 January

9:00 - 9:30	**Speech Recognition and Synthesis at Bell Northern Research Laboratories, Montréal** Dr. Maier Blostein
10:00 - 10:30	**AI Applied to Long-Range Planning and Decision Making** Professeur George Strobel Département de psychologie Université de Montréal
10:30 - 10:45	Coffee
10:45 - 11:15	**AI and Office Automation** Professor Tom Carey Department of Computer Science University of Guelph

11:15 - 12:30	**Discussion. Applications. How soon? How Good?**
12:30 - 13:30	Lunch
13:30 - 14:00	**Hardware for Artificial Intelligence** Professor Thomaz Petrokowski Department of Computer Science Acadia University Wolfville, Nova Scotia
14:00 - 14:30	**The Role of Theory** Professor Raymond Reiter Department of Computer Science University of British Columbia
14:30 - 14:45 p.m.	Coffee
14:45 - 16:45 p.m.	**Artificial Intelligence: A National Priority? A Canadian Fifth Generation System: Industry and Government Involvement? Where will the Money Come From?**
16:45 - 17:00 p.m.	Closing Remarks Mr. Jorge Miedzinski

LIST OF PARTICIPANTS

Lewis Auerbach
CRTC
Hull, Québec

John W. Ball
GEAC Computers
Markham, Ontario

Robert Baser
Department of Communications
Ottawa, Ontario

Marc Beauchamp
President
Novacap Investments
Montréal, Québec

Laurent Bourbeau
Bureau des Traductions
Direction de la Technologie
Secrétariat d'État
Hull, Québec

David Burden
Convergent Technologies
Ottawa, Ontario

Professor Tom Carey
University of Guelph
Guelph, Ontario

Professor Nick Cercone
Department of Computer Science
Simon Fraser University
Burnaby, British Columbia

Donald Chisholm
President
Innovation & Development
Northern Telecom Limited
Mississauga, Ontario

William Coderre
National Research Council
Ottawa, Ontario

John Collins
AES Data
Montréal, Québec

Arthur J. Cordell
Science Council of Canada
Ottawa, Ontario

Alistair Davidson
Laventhol and Horwath
Toronto, Ontario

Nick Dawes
Systemhouse Ltd.
Ottawa, Ontario

Denzil Doyle
Kanata, Ontario

Professor E.W. Elcock
Department of Computer Science
University of Western Ontario
London, Ontario

Nardin Ghahary
HiTech Canada
Ottawa, Ontario

Professor Randy Goebel
Department of Computer Science
University of Waterloo
Waterloo, Ontario

Takashi Gomi
Gomi A.I. Systems
Kanata, Ontario

Gerald Grant
Woods Gordon
Toronto, Ontario

Pierre Isabelle
Montréal, Québec

Kenneth Iverson
I.P. Sharp Associates Ltd.
Toronto, Ontario

A. Kasvand
National Research Council
Ottawa, Ontario

Professor Roger Kaye
Department of Computer Science
Carleton University
Ottawa, Ontario

Professeur Richard Kittredge
Département de linguistique
et philologie
Université de Montréal
Montréal, Québec

Maurice L'Abbé
Science Council of Canada
Ottawa, Ontario

Professor Wilf Lalonde
Department of Computer Science
Carleton University
Ottawa, Ontario

J. James Mackie
Vice President
Mitel Corporation
Kanata, Ontario

Professor Allan Mackworth
Department of Computer Science
University of British Columbia
Vancouver, British Columbia

J. Maksym
Defence Research Establishment
Atlantic
Dartmouth, Nova Scotia

Gwynneth Martin
Treasury Board
Ottawa, Ontario

David McCart
SB Capital
Toronto, Ontario

Murdoch McKinnon
CAE Electronics Ltd.
Montréal, Québec

P. Mermelstein
Bell Northern Research
Montréal, Québec

J. Miedzinski
Science Council of Canada
Ottawa, Ontario

J. de Montingy
Convergent Technologies
Ottawa, Ontario

Renato de Mori
Associate Professor
Computer Science Department
Concordia University
Montréal, Québec

Margaret Munro
Southam News
Ottawa, Ontario

Professor John Mylopoulis
Department of Computer Science
University of Toronto
Toronto, Ontario

Nobuyuki Otsu
(NRC Visitor)
Senior Research Officer
Mathematical Engineering Section
Inormation Science Division
Electrotechnical Laboratory
Ottawa, Ontario

Professor Ray Perrault
Department of Computer Science
University of Toronto
Toronto, Ontario

Professor Thomaz Pietrzykowski
School of Computer Science
Acadia University
Wolfville, Nova Scotia

Birendra Prasada
Bell-Northern Research
Verdun, Québec

John Pugh
School of Computer Science
Carleton University
Ottawa, Ontario

Kenneth Pulfer
Vice-President
Laboratories
National Research Council
Ottawa, Ontario

Professor Zenon Pylyshyn
Department of Computer Science
University of Western Ontario
London, Ontario

Professor Raymond Reiter
Department of Computer Science
University of British Columbia
Vancouver, British Columbia

Judith Scott
Gandalf Technologies Inc.
Manotick, Ontario

Robert W. Shaw
Public Works Canada
Ottawa, Ontario

Professor Douglas Skuce
Department of Computer Science
University of Ottawa
Ottawa, Ontario

Alfred Stein
Continuum
Toronto, Ontario

Bob Steklasa
Financial Post
Ottawa, Ontario

David Stewart-Patterson
Globe and Mail
Ottawa, Ontario

Professeur M.G. Strobel
Département de Psychologie
Université de Montréal
Montréal, Québec

G.D. Sutton
President
Canadian Enterprises
Development Corporation Ltd.
Toronto, Ontario

Marla Tenenbaum
Elliott & Page
Toronto, Ontario

Andrej Tenne-Sens
Department of Communications
Ottawa, Ontario

John Tsotsos
Professor
Department of Computer Science
University of Toronto
Toronto, Ontario

George Young
Industry, Trade and Commerce
Ottawa, Ontario

Professor S.W. Zucker
Computer Vision & Robotics Lab
Department of Electrical Engineering
McGill University
Montréal, Québec