Proceedings of the 1986 Winter Simulation Conference
J. Wilson, J. Henriksen, S. Roberts (eds.)

ARTIFICIAL INTELLIGENCE AND SIMULATION: AN INTRODUCTION

R. Greer Lavery
School of Computer Science
Ryerson Polytechnical Institute
Toronto, Ontario, Canada. M5B 2K3

ABSTRACT

Artificial intelligence is one of the most rapidly developing fields in modern applied computer science. It has generated research and applications in several widely varying fields, both within and without the realm of topics considered to be 'core' AI components. At the same time, the science of simulation modelling is also advancing steadily, incorporating new technologies such as graphical output display and more recently, applications of artificial intelligence.

This paper provides a brief description of the major component topics of artificial intelligence. Several references are included to allow readers to examine topics of interest in more depth. The paper also discusses the application of intelligent systems to simulation modelling and describes, with references, some of the development in this area to date.

ARTIFICIAL INTELLIGENCE

Artificial intelligence is not the study of a single topic but rather a collection of related topics. Duda (1979) has stated that: "Historically, AI has both borrowed from and contributed to other closely related disciplines concerned with advanced methods for information processing. Thus, links exist between AI and aspects of such theoretical areas as mathematical logic, operations research, decision theory, information theory, pattern recognition and mathematical linguistics. In addition, research in AI has stimulated important developments in software technology, particularly in the area of advanced programming languages. What distinguishes AI from these related fields, however, is its central concern with all of the mechanisms of intelligence."

This interdisciplinary aspect of artificial intelligence produced many different definitions of AI, most of which are similar in terminology and ideas. Winston (1984a) defines AI as "the study of ideas that enable computers to be intelligent." Hawkinson (1984) defined it as "the science of enabling computers to reason, make judgements and even learn" while Chandor (1970) gives the definition "the use of computers in such a way that they perform operations analogous to the human abilities of learning and decision-taking." Each of these (and in fact most) definitions of AI make reference to computers since computers represent the artificial, i.e. non-human component of AI. The first reference presents four additional definitions of artificial intelligence, one of which however makes no reference whatsoever to computers.

Research and interest in artificial intelligence has expanded in parallel with the development of computers over the past 15-20 years but the true beginning of AI is generally conceded to have been the publication of Alan Turing's classic paper on Computing Machinery and intelligence (1951). The 'imitation game' that he proposed to determine whether or not a machine is intelligent is still used today and is now known as the Turing Test.

Some differences of opinion exist as to which individual topics are the true constituents of the science known as artificial intelligence. A typical, representative summary is given by Gevarter (1985). He lists Modelling and Representation of Knowledge, Heuristic Search, Commonsense Reasoning and Logic and AI Languages and Tools as the basic elements of AI. He also identifies four of the principal artificial intelligence applications as being Expert Systems, Problem Solving and Planning, Natural Language Processing and Computer Vision. A brief discussion of each of these elements and applications forms the portion of this paper devoted to a summary of artificial intelligence.

KNOWLEDGE REPRESENTATION

Since computers are designed and built by humans, to create an intelligent computer it is necessary that the human creators know exactly what knowledge is, how it is stored, accessed, updated, etc. In one very large sense then, AI is really a study of knowledge or intelligence. Considerable AI research is devoted to the study of how humans learn and how they retain, classify and retrieve information. The field of cognitive science is devoted to this area and many books and publications are available on this subject. For example, see Anderson (1983), Barr and Feigenbaum (1982a), Newell (1982) and Simon (1980).

Because of the current widespread interest in computers and computing, much of the public focus and interest in artificial intelligence has been on the artificial component, i.e. the role that computers play in this field. The true emphasis of AI is on intelligence however and while not as glamorous or exciting as computerized teaching models or expert systems, it is the ongoing research into fields such as language comprehension and knowledge representation which will permit even more powerful computer applications to be seen in the future.

HEURISTIC SEARCH TECHNIQUES

In the early stages of its development, AI research was mainly concerned with problem solving. Large decision trees resulted from the sequential decision patterns created (each decision producing a new branch point as new combinations of parameters are created) which meant that the time required to search the tree in some orderly fashion for an acceptable problem solution became very large.

To reduce this search time, heuristic search techniques (often commonly described as 'rules of thumb') were used. Development of programmes to search for heuristic rules to aid in problem-solving is still a significant topic in AI. A good summary of current and recent research in search techniques can be found in Pearl (1984).

COMMONSENSE REASONING AND LOGIC

While considerable progress has been made in modelling or representing human knowledge on computers, the representation of commonsense knowledge has proved much more difficult. Some success appears possible with the modelling of commonsense as low-level reasoning based on a vast amount of experience, but progress is not expected to be rapid in this area.

Logic, however, is ideally suited for computer representation. Since computers represent all information in binary form (0/1, ON/OFF, TRUE/FALSE), logic conditions can be easily expressed in computer code.

It was originally hoped that computational logic would provide a universal problem solving method. While this goal has not been attained, new computer technologies such as LISP machines are providing increased interest in the study of logic and logical inference.

Computational logic, a computational approach to logical reasoning, is divided into two principal parts, propositional logic and predicate logic. As with most of the elements of AI, the study of logic has interrelationships with other AI topics. Moore (1982) gives an excellent example of such interrelationships.

AI LANGUAGES AND TOOLS

Just as specific high-level languages have been developed in computer science for specific types of applications, FORTRAN for scientific calculations, COBOL for business applications, etc., two principal languages have been developed for AI applications. The first, LISP, was developed at MIT in the late 1950's by John McCarthy as a list processing language. It has served as the primary AI programming language in North America since that time. For a description and explanation of this language, the Winston-Horn text (1984b) is recommended. The other language, PROLOG, was developed at the University of Marseille AI laboratory by Colmerauer and Roussel in 1973. PROLOG is the principal AI programming language used outside of North America and was chosen as the programming language for the Japanese 5th generation computer development project. It is currently gaining popularity in North America as well. A popular reference on PROLOG is the book by Clocksin and Mellish (1984).

PROLOG is also beginning to see use as a language for the design of simulation models as described by Adelsberger (1984).

EXPERT SYSTEMS

By far the most visible and publicized product of AI research is the expert system. Feigenbaum (1982) defines an expert system as "an intelligent computer programme that uses knowledge and inference procedures to solve problems that are difficult enough to require significant human expertise for their solution". An expert system consists of a knowledge base of domain facts and heuristics associated with the problem, an inference engine (or control structure) for utilizing the knowledge base in the solution of the problem and a working memory (global data base) for keeping track of the problem status, the input data for the particular problem and the relevant history of what has been done thus far.

Expert systems have been developed and used commercially in many widely varied fields such as medicine, agriculture, mathematics, geology, law, military science, space technology and others. A comprehensive reference on expert systems is provided by Waterman (1985) with additional technical information available in articles by Michaelson et al (1984) and Thompson and Thompson (1984)

PROBLEM-SOLVING AND PLANNING

The need for general-purpose problem-solving techniques to attack problems for which there are no experts has long been recognized. One way of viewing intelligent behaviour is as a problem-solver. Many AI tasks can naturally be viewed this way and most AI programmes draw much of their strength from their problem-solving components. Applications of AI such as image analysis, natural language processing and expert systems all have significant problem-solving components. Approaches to problem-solving such as backward chaining, problem reduction and difference reduction are all being investigated under the general umbrella of AI research. The first example of difference reduction was the General Problem Solver (GPS) programme developed by Newell, Shaw and Simon (1957). It was the first programme to separate its general problem-solving method from knowledge specific to the current problem, a technique fairly common in today's world of inference engines and knowledge bases.

Planning could be described as the design process for selecting and combining individual actions into sequences in order to achieve a desired goal. Planning is essential to problem-solving and Sacerdoti (1979) lists various items of information relative to planning that can be obtained during the search process of a decision tree.

Many models and paradigms for planning systems have been developed. Two typical examples of such systems are KNOBS, described by Engleman, Berg and Bischoff (1979) and ISIS II, described by Fox, Allen and Strohm (1982).

NATURAL LANGUAGE PROCESSING

Programming languages such as FORTRAN, PASCAL, COBOL, etc. have long found widespread use because of the relative ease with which they allow humans to communicate with computers. The continuous expansion of computer usage in society however, is producing a still greater need for human-computer communications in areas other than the traditional pair of scientific calculations and business data processing. Applications such as graphics, database creation and querying, facts and rules specification, etc. are being used by an ever-growing spectrum of the population. The fact that the standard high-level programming languages are not designed to facilitate such communication, coupled with the reluctance of many new and naive users to learn a traditional programming language with all its syntax and peculiarities of form, has resulted in a significant research effort in the field of natural language processing, i. e. conversing with a computer in a human language such as English or French. This research is not limited to a single form of communication such as typed instructions but encompasses several media including voice recognition and document reading. All of these topics are included under the general umbrella of artificial intelligence.

If this goal of allowing virtually anyone to use computers and specialized software via natural language communication is to be achieved, it is necessary to have a thorough understanding of how humans communicate. Only then will it be possible to create machines with human-like communication skills. Research in this area is quite active and the rapidly accelerating interest in artificial intelligence is generating and promoting more and more research into language comprehension. More detailed information is available from Barr and Feigenbaum (1982b) and Winograd (1972). Computational linguistics, a science which incorporates elements of linguistics, psychology and philosophy is devoted to the study of how humans communicate and is an area of research central to artificial intelligence.

COMPUTER VISION

Computer vision, i.e. visual perception employing computers, is another active, current area of AI research. Often combined with research into robot applications, computer vision involves the sensing and interpretation of images by machines. It has been estimated that by 1990 one-quarter of all industrial robots will be equipped with some form of vision system. Scores of applications of computer vision systems are available, some currently in use and others waiting for the appropriate technology to be developed. Examples of such applications are object acquisition by robot arms for sorting or packing of items, screening medical images such as chromosome slides, cancer smears and x-ray images, document readers for computer input and automatic guidance of seam welders or cutting tools. Several models of image-understanding paradigms have been developed for analysis and understanding of images, some of which are described by Pratt (1978).

A good example of a computer vision system is the ACRONYM system developed at Stanford University by Brooks et al (1979). The ACRONYM system can also be considered as an example of an expert system. Moreover, since computer vision systems employ elements of knowledge representation, computer vision systems have multiple connections to the topics of artificial intelligence. They are subsequently expected to be a major source of advancement of AI systems over the next few years.

APPLICATIONS OF ARTIFICIAL INTELLIGENCE TO SIMULATION

In the previous sections, many of the component topics of artificial intelligence have been briefly summarized. Applications of AI research are growing exponentially however into fields as diverse as teaching (see Barr and Feingenbaum (1982c) and Sleeman and Brown (1982)) and configuration of commercial computer systems (McDermott 1981). Another highly specialized but steadily increasing use of AI is its application to the field of simulation modelling.

Several potential applications of artificial intelligence to simulation have been proposed in the past two years, research on some of these proposals being already underway while others remain merely ideas, further technological advances being needed before implementation becomes possible. For example, O'Keefe (1986) describes how a natural language interface could help form an 'intelligent front-end' programme for a simulation model which was to be delivered to a client who was an inexperienced user of simulation. Removing the need for the user to be familiar with the simulation language would result in more efficient and probably more productive use of the model as well. The intelligent programme could also be useful if it included some model of the user so that the interface programme could provide various levels of dialogue with the model user depending upon whether the user is experienced or inexperienced in the use and application of simulation models.

In traditional simulation models, the control and data components pertaining to the model are combined in the model structure. This is true whether the model is coded in a simulation language such as GPSS or a general purpose language such as FORTRAN. These two components are provided as distinct, independent entities in an AI programme (the knowledge base and the inference engine) permitting the modeller to alter either component independent of the other. Blending AI techniques with existing simulation techniques would therefore improve simulation procedures by separating these model components.

Also, as described by Shannon et al (1985), the use of a knowledge base rather than a standard base allows the incorporation of different types of information such as rules, judgements, intuition and experience (heuristic knowledge). Such a system is more flexible than the traditional data base system with its rigidly defined data structures and permits AI systems to provide solutions to some problems not solvable by traditional data base information systems. Application of these systems to simulation modelling should allow more detailed and more flexible models to be created.

Modellers know that there are several different forms possible for the model of a given system, particularly if there are different goals or objectives for the simulation. Each model is unique in that it deals with a specific aspect of the investigation and at the same time it often models only a part of the total system. One of the expectations of expert simulation systems is that they will be able to support an organized base of models and domain specific knowledge whose partial perspectives can be integrated to achieve a comprehensive whole.

The ultimate objective of a simulation analysis dictates both the structure of the system model and the nature of the experiment to be run. The appropriateness of a model to the overall objective is usually what determines how useful a simulation study is to an organization. The expertise to make judgements concerning the level of detail to incorporate in a model requires considerable experience with both simulation techniques and the problem domain. This, then represents another potentially valuable way in which artificial intelligence can enhance the design and application of simulation models.

Two other applications among the many documented are automatic selection of simulation output data for analysis as described by Reddy et al (1985) and automatic control of display data from a simulation analysis as described by McCandless (1986).

The infusion of artificial intelligence procedures to the world of simulation should not be cause for alarm among simulationists, foreseeing their replacement by expert systems. Rather they should expect to become even more valuable as the application and analysis of the sophisticated simulation models of the future permit forecasting and decision-taking of increased reliability.

REFERENCES

"A Workshop in Artificial Intelligence." Proceedings, Publications Office, Science Council of Canada, Ottawa, 1983.

Adelsberger, H.H. (Nov. 1984). "PROLOG as a Simulation Language", Proceedings, 1984 Winter Simulation Conference, pp 501-504.

Anderson, J.R. (1983). The Architecture of Cognition. Harvard University Press, Cambridge, Mass.

Barr, A. and Feigenbaum, E. (1982). The Handbook of Artificial Intelligence. Kaufman Publishing Co., Los Altos, CA., Vol. 1, Ch. III.

IBID, Vol. 1, Chapter IV.

IBID, Vol. 2, Chapter IX.

Brooks, R., Greiner, R. and Binford, T.O. (1979) "The ACRONYM Model Based Vision System". Proceedings of the International Joint Conference on Artificial Intelligence. Vol. 6, pp 105-113.

Chandor, A.A. (1970). A Dictionary of Computers. Penguin Books, Markham, Ontario, p 36.

Clocksin, W.F. and Mellish, C.S. (1984). Programming in PROLOG, 2nd edition, Springer-Verlag Publishing Co. N.Y.

Duda, R.O. (1979). "State of Technology in Artificial Intelligence". Research Directions in Software Technology, P. Wenger editor, MIT Press, Cambridge, Mass., pp 729-749.

Engleman, C., Berg, C.H. and Bischoff, M. (1979). "KNOBS: An Experimental Knowledge-Based Tactical Air Mission Planning System and a Rule-Based Aircraft Identification Simulation Facility." Proceedings IJCAI-79, pp 247-249.

Feigenbaum, E.A. (1982). "Knowledge Engineering for the 1980's", Dep't of Computer Science, Stanford University, Stanford, Cal.

Fox, M.S., Allen, B. and Strohm, G.A. (Aug. 1982). "Job-Shop Scheduling: An Investigation in Constraint-Directed Reasoning". Proceedings of the Second Annual National Conference on Artificial Intelligence, Carnegie-Mellon University, pp 155-158.

Gevarter, W.A. (1985). Intelligent Machines. Prentice-Hall Pub. Co. Englewood Cliffs, N.J.

Hawkinson, L. (Feb. 1984). "High-Level Languages Based on LISP Aid Expert System Design", Computer Technology Review, Westword Productions Inc., Vol. 4, #1.

McCandless, T.P. "PDP Mechanisms for Intelligent Display Control". Proceedings of the 1986 Conference on Intelligent Simulation Environments, P. Luker and H. Adelsberger editors, The Society for Computer Simulation, San Diego, CA. pp 87-91.

McDermott, J., (1981). "R1: The Formative Years", AI Magazine, Vol. 2, No. 2, pp 21-29.

Michaelson, R.H., Michie, D. and Boulanger, A. (April 1984). "The Technology of Expert Systems". BYTE, Vol. 10, #4, pp 303-312.

Moore, R.C. (Aug. 1982). "The Role of Logic in Knowledge Representation and Commonsense Reasoning". Proceedings of the National Conference on Artificial Intelligence, Pittsburgh, PA., Kaufman Oub. Co., Los Altos, CA. pp 428-433.

Newell, A. (1982). "The Knowledge Level". Artificial Intelligence, Vol. 18, pp 87-127.

Newell, A., Shaw, J.C. and Simon, H.A. (1959). "Report on a General Problem-Solving Programme". Proceedings of the International Conference on Information Processing (ICIP), Paris, Unesco House Publishers.

O'Keefe, R. (Jan. 1986). "Simulation and Expert Systems - A Taxonomy and Some Examples". Simulation, Vol. 46, #1, pp 10-16.

Pearl, J. (editor) (1984). Search and Heuristics. Elsevier Pub. Co., N.Y.

Pratt, W.K. (1978). Digital Image Processing. J. Wiley and Sons, N.Y.

Reddy, Y.V., Fox, M.S. and Husain, N. "Automating the Analysis of Simulations in KBS", Proceedings of the 1985 Conference on AI, Graphics and Simulation, the Society for Computer Simulation, San Diego, CA., pp 34-40.

Sacerdoti, E.D. (1979). "Problem-Solving Tactics". AI Magazine, Vol. 2, #1, pp 7-14.

Shannon, R.E., Mayer, R. and Adelsberger, H. (June 1985). "Expert Systems and Simulation". Simulation, Vol. 44, #6, pp 275-284.

Simon, H.A. (Jan.-Mar. 1980). "Cognitive Science: The Newest Science of the Artificial" Cognitive Science, Vol. 4, #1, pp 33-46.

Sleeman, D. and Brown, J.S. editors. (1982). Intelligent Tutoring Systems. Academic Press, New York, N.Y.

Thompson, B.A. and Thompson, W.A. (April 1984) "Inside an Expert System". BYTE, Vol. 10, #4, pp 315-330.

Turing, A.M. (Oct. 1950). "Computing Machinery and Intelligence". MIND, Vol. 59, pp 433-460.

Waterman, D.A. (1985). A Guide to Expert Systems, Addison-Wesley Publishing Co., Reading, Mass.

Winograd, T. (1972). Understanding Natural Language , Academic Press, New York.

Winograd, T. (Jan.-Mar. 1980). "What Does It Mean to Understand Language?". Cognitive Science, Vol. 4, #1, pp 209-241.

Winston, P.H. (1984). Artificial Intelligence. Addison-Wesley Publishing Co. Reading, Mass. p 1.

Winston, P.H. and Horn, B.K.P. (1984). LISP, 2nd edition, Addison-Wesley Publishing Co., Reading Mass.

AUTHOR'S BIOGRAPHY

GREER LAVERY is a professor and the Director of the School of Computer Science at Ryerson Polytechnical Institute, Toronto, Canada. He received a B.A.Sc. in Chemical Engineering from the University of Toronto in 1963, an M.A.Sc. in Chemical Engineering from the University of Waterloo in 1967, an M.A.Sc. in Management Sciences from the University of Waterloo in 1977 and is currently engaged in doctoral studies at the Ontario Institute for Studies in Education, Toronto. He is the author of a textbook, "Programming with System/370 Assembler Language" and a number of papers on simulation. He is a member of APEO and SCS.

R. Greer Lavery, P. Eng.
School of Computer Science
Ryerson Polytechnical Institute
350 Victoria Street
Toronto, Ontario, Canada. M5B 2K3
(416) 979-5063