DOCUMENT RESUME

ED 114 116 IR 002 751

AUTHOR	Leal, Ralph A.
TITLE	A Primer on Artificial Intelligence.
PUB DATE	75
NOTE	33p.
EDRS PRICE	MF-$0.76 HC-$1.95 Plus Postage
DESCRIPTORS	*Artificial Intelligence; *Bionics; Computer Programs; *Cybernetics; Information Theory; Intelligence Factors; Machine Translation; *Programing Languages.

ABSTRACT

A survey of literature on recent advances in the field of artificial intelligence provides a comprehensive introduction to this field for the non-technical reader. Important areas covered are: (1) definitions, (2) the brain and thinking, (3) heuristic search, and (4) programing languages used in the research of artificial intelligence. Some examples of artificial intelligence are described. Suggestions for additional reading and a 28-item bibliography are appended. (DS)

Title: A Primer on Artificial Intelligence

Author: Ralph A. Leal
State University of New York - College at Old Westbury

Abstract: This survey of recent advances in the field of artificial intelligence is aimed at the non-technical reader. Its purpose is to present a concise and comprehensive introduction to the field. Areas covered include definitions, the brain and thinking, heuristic search and an overview of programming languages in artificial intelligence research. The appendices include a classical article by Dr. Vannevar Bush entitled "As We May Think" and bibliographical suggestions for additional reading.

Keywords: Artificial Intelligence; AI; Huristic Search; Programming Languages in Artificial Intelligence; Vannevar Bush.

42 pages, appendices, bibliographies

ABSTRACT

Leal, Ralph A. A PRIMER ON ARTIFICIAL INTELLIGENCE. New York, 1974.

Artificial intelligence has been defined as 'the totality of attempts to make and understand machines that are able to perform tasks that, until recently, only human beings could perform, and to perform them with effectiveness and speed comparable to a human.' (Banerji, 1969).

Although the field has many divisions and subdivisions, the most important work can be classified into four areas: (1) Game Playing; (2) Problem Solving; (3) Pattern Recognition, and (4) Semantic Information Processing.

Attempts to understand the basic requirements for the construction of an 'intelligent' machine have led to exploration of the physical functioning of the brain. The most famous theorem in this area is the McCulloch-Pitts Theory of Formal Neural Networks.

An understanding of thought processes is also important to the artificial intelligence field. Nicolai Asomov, of the Institute of Cybernetics of the Academy of Sciences of the USSR, has proposed a general model of a human being which has three types of programs: one for himself - for self-preservation; one for stock - the reproductive instinct, and one for species - the programs of social behavior.

The English mathematician A. M. Turing, a firm believer in the possibility of creating artificial intelligence, has summarized the objections to that possibility in an article entitled 'Can A Machine Think?' Among the objections we find are: (1) the theological; (2) the heads-in-the-sand; (3) the mathematical; (4) the argument from consciousness; (5) the various disabilities argument; (6) Lady Lovelace's argument; (7) the argument from continuity in the nervous system; (8) the argument from informality of behavior, and (9) the argument from extra-sensory perception.

Heuristic search is a mathematical process used to move from an initial state to a desired state using a repertoire of actions when the sequence of steps is not exactly known. Donald Michie believes that this process offers good payoff possibilities for future research in the AI field.

A number of programming languages, characterized by their ability to do list processing and symbol manipulation, have been developed to assist AI researchers. The loci for the creation of these languages have been the Massachusetts Institute of Technology, Stanford Research Institute, and the University of Edinburgh.

Despite the raging debates in favor and against the possibility that 'thinking machines' may be constructed, we must conclude that not all the evidence is yet in, and lean - very tentatively - towards the positive side of the argument.

TABLE OF CONTENTS

Deleted because of copyright restrictions

INTRODUCTION

In 1945, Dr. Vannevar Bush wrote an article in The Atlantic Monthly which has become a classic in the field of information science. In that article, Dr. Bush foresaw a world in which man's quest for knowledge could be tailored to individual needs through the use of 'association trails' stored in each person's home library - the Memex. In the same article, reaching boldly into the future, he wrote:

> In the outside world, all forms of intelligence, whether of sound or sight, have been reduced to the form of varying currents in an electric circuit in order that they may be transmitted. Inside the human frame exactly the same sort of process occurs. Must we always transform to mechanical movements in order to proceed from one electrical phenomenon to another? 1

The September, 1974 issue of Computer Decisions describes research currently underway at the Stanford Research Institute in the following terms:

> Imagine, if you can, a computer installation with none of the input devices you're used to. No card readers, tape readers, crt terminals or teletypewriters. Instead of all this cumbersome and inefficient hardware, you've got a special helmet. You will enter data directly from your brain into the computer. 2

We have chosen the above examples to introduce the subject of our research paper - artificial intelligence - because they are illustrative of the advances which are possible in the various fields of science

---

* 1. Bush, Vannevar. "As We May Think." The Atlantic Monthly, July 1945, p. 108. (See Appendix D.)

* 2. "Direct Brain-to-Computer Interface On The Way." Computer Decisions, September 1974, p. 13. (See Appendix B.)

* Available in large Libraries

within the span of three decades. What Vannevar Bush referred to as a "... suggestive thought, but it hardly warrants prediction without losing touch with reality and immediateness ..." [3] in 1945 is already being tested in a California laboratory in 1974.

The subject of our study has suffered attacks by a number of respected members of the scientific community, who have put forth powerful arguments to demonstrate that the very concept of artificial intelligence is bizarre and underserving of serious consideration. In the following pages we will attempt to present an unbiased and comprehensive picture of an evolving field.

---

3. Bush, Vannevar, loc. cit.

DEFINITIONS

Donald G. Fink, in his Computers and the Human Mind, defines artificial intelligence in the following manner:

> (a) The ability of machines to organize information into meaningful patterns; ability to recognize, store, recall, and manipulate such patterns in playing games, solving problems, answering questions, etc., and in controlling the actions of other mechanisms; (b) the ability of a machine to adapt to its environment, particularly to respond to patterns of stimulation not explicitly foreseen in its design; (c) the observed performance of such machines as measured by comparison with, or in competition against, human intelligence. 4

Ranan B. Banerji, in his Theory of Problem Solving. An Approach to Artificial Intelligence, puts forth the following definition:

> The field commonly called artificial intelligence may, perhaps, be described as the totality of attempts to make and understand machines that are able to perform tasks that, until recently, only human beings could perform and to perform them with effectiveness and speed comparable to a human. (Italics his) 5

Philip C. Jackson, Jr., in his Introduction to Artificial Intelligence, defines it in these terms:

> "Artificial intelligence" is the ability of machines to do things that people would say require intelligence. 6

---

4. Fink, Donald G. Computers and the Human Mind, p. 225.

5. Banerji, Ranan B. Theory of Problem Solving, p. 1.

6. Jackson, Jr., Philip C. Introduction to Artificial Intelligence, p. 1.

Hubert L. Dreyfus [7] has written that, although the field has many divisions and subdivisions, the most important work can be classified into four areas:

1. Game playing
2. Problem solving
3. Pattern recognition
4. Semantic information processing

The first series of efforts attempt to 'teach' machines to play such games as checkers, chess, poker, etc. with a level of proficiency approaching that of an interested adult. The second set of efforts include training in the solution of mathematical, logical, and other types of problems. The third group, generally using supporting television equipment, seeks to create the ability to discriminate among similar visual or auditory patterns. Finally, in the fourth category we find efforts at machine translation of foreign languages, as well as programs that answer a variety of questions on very clearly-delimited fields.

Professor Frank George, of the Department of Cybernetics at Brunel University, has defined the basic problem of artificial intelligence in these terms:

> The basic problem of artificial intelligence is as to whether or not machines (i.e. arti-

7. Dreyfus, Hubert L. What Computers Can't Do, p. xxxiii.

> ficially constructed systems such as digital computers) can be made to "think." The problem is partly a semantic one, since the word "machine" and the word "think" both present definitional problems, and we could clearly define both terms in such a way that machines could obviously think or could obviously not think as the case may be. 8

The next section of our paper will present a brief analysis of the work which has been done in seeking to link the brain, thought processes, and artificial intelligence.

---

8. George, Frank. "What Is Artificial Intelligence?" Data Processing, September-October 1972, p. 328.

THE BRAIN AND THINKING

In his very penetrating article "The Brain Considered As a Thinking Machine," Dr. William Feindel wrote:

> ... the human brain has certain critical advantages over these electronic devices. It is far more compact. An electronic brain having as many as twelve thousand million thinking units would take up an enormous volume. Even if miniaturized, at the very best an electronic brain would be comparable in size to a very large government grain elevator. The human brain uses a very small amount of electrical energy, something between ten and twenty-five watts. Furthermore, it is portable and comes with a lifetime guarantee! 9

Concern with the brain and thinking goes back to antiquity, but it is only in recent times that scientists have begun to attempt to develop a theory which might be applicable to computers. The most famous theorem in this area is called the McCulloch-Pitts Theory of Formal Neural Networks, and is explained by John von Neumann in his "The General and Logical Theory of Automata" in the following terms:

> The functioning of a formal neural network may be defined by singling out some of the inputs of the entire system and some of its outputs, and then describing what original stimuli on the former are to cause what ultimate stimuli on the latter... Any functioning which can be defined at all logically, strictly, and unambiguously in a finite number of words can also be realized by such a formal neural network. 10

---

9. Feindel, William. "The Brain Considered As a Thinking Machine." in Memory, Learning, and Language, p. 22.

10. von Neumann, John. "The General and Logical Theory of Automata." in The World of Mathematics, pp. 2089-2090.

Mortimer Taube [11] quotes Warren Weaver's explanation of the theorem by stating that it is theoretically possible to construct a robot which could deduce any legitimate conclusion from a finite set of premises, provided that the construction included a set of regenerative loops.

Arbib,[12] who has been strongly influenced by the ideas of Warren McCulloch, analyzes the problem from two distinct angles: what he has termed the artificial intelligence approach, and what he calls the brain theory approach. He distinguishes between the two approaches by stating that the former is based on carrying out operations on data which is passively stored, while the latter emphasizes the parallel activity of a multitude of operations. This second approach is referred to as "... distributed action-oriented computation in layered somatotopically organized machines." [13]

Nicolai M. Asomov, of the Institute of Cybernetics of the Academy of Sciences of the USSR, is the author of an article entitled "Simulation of Thinking Processes," which appeared in the book Purposive Systems. [14] Asomov believes that simulation or modeling of thinking is necessary for studying human cognition, for creating artificial intelligence, and for developing models of human interaction.

---

11. Taube, Mortimer Computers and Common Sense, p. 19

12. Arbib, Michael A. The Metaphorical Brain. New York, Wiley-Interscience, 1972.

13. Arbib, op. cit., p. 5.

14. Asomov, Nicolai M. "Simulation of Thinking Processes." In Purposive Systems. New York, Spartan Books, 1968.

He proposes a general model of a human being which can be represented by an automaton having three types of programs; one for himself - for self-preservation; one for stock - the reproductive instinct, and one for species - the programs of social behavior. He goes on to state that six programs determine human behavior. They are:

1. Program of bodily feeling - long perception and processing of information from the body.
2. Perception of outer influence and recognition of models of meaning and qualities of the surrounding world.
3. Program of action - imparting of energy and information to the surrounding world.
4. Program of speech.
5. Program of consciousness.
6. Program of creativity and work. [15]

Clearly, if an automaton could be constructed and programmed with the above list of attributes, we could categorically state that we have achieved 'artificial intelligence.' Unfortunately, we are far from having arrived at that point.

In his principle of complication in reproduction, von Neumann felt that below a certain point, if automata could reproduce, they would only be able to make less complicated machines than themselves. Above that point, and at least conceptually, it would be possible for automata to create higher entities. [16]

---

15. Asomov, op. cit., p.41.

16. von Neumann, op. cit., p. 2098.

Herbert Anschutz, in an article entitled "Prospects for the Development of the Psychocybernetics of Intelligent Behavior," also found in Purposive Systems, writes that cyberneticists make use of three basic disciplines: information theory, the theory of control, and automaton theory, and that knowledge about structure can be largely replaced by knowledge about function, "... about the behavior of the black box over an abstract set of states." [17]

He further states, "... we have found that every data-processing system, including the human brain, is subject to a generalized combination of information theory and the theory of automata." [18]

We have discerned two principal problems connected with the area of the brain and thought processes. First, despite intensive research, little is known of the way in which the brain functions in physical terms. Second, although it is possible to postulate a series of requirements for the creation of artificial intelligence, the bridge between concept and construction has not been established, except for some minimal examples which we will examine below.

In our next section we will consider a series of arguments presented by A.M. Turing in negative response to the question "Can Machines Think?"

---

17. Anschutz, Herbert, in Purposive Systems, p. 113.

18. Anschutz, op. cit., p. 117.

CAN MACHINES THINK ?

A. M. Turing, the British mathematician and logician, must be credited with two significant contributions to our field, both of which are found in his article "Can A Machine Think?" [19]

The first contribution, which has come to be known as 'Turing's test,' is a positive one. The author suggests that one way in which artificial intelligence could be tested world be through what he calls 'The Imitation Game.' The object of the game, which is played by three individuals: a man (a); a woman (b), and an interrogator (c) who may be of either sex, is to enable the interrogator to determine the sex of either of the other two players. The rules of the game call for (a) to try to get (c) to make a wrong identification, and for (b) to try to get (c) to make a right identification. (a) is free to lie, but (b) must always tell the truth. The players are physically separated, so (c) cannot gain any insights into the sex of the players except through their printed responses. Turing suggests that if a machine could be designed and programmed to successfully play the role of (a), it could be credited with possessing what he defines as artificial intelligence.

Turing's second contribution, although couched in negative terms, is also a positive one. As a believer in artificial intelligence, Turing sought to marshall the arguments which had been put forth against the possibility of such an occurrence. We have chosen to summarize these arguments.

---

19. Turing, A. M. "Can a Machine Think?" in The World of Mathematics, pp. 2099-2123.

The theological objection. This argument states that thinking is a function of man's immortal soul, given to him by God. Therefore, no animal or inanimate object can think.

The 'heads in the sand' objection. Inasmuch as the consequences of machines thinking are too dreadful, let us not consider the possibility and fervently hope that they never will.

The mathematical objection. There are a number of results of mathematical logic which can be used to show that there are limitations to the powers of discrete-state machines. The best known of these objections is known as Goedel's theorem,[20] and it shows that in any sufficiently powerful logical system statements can be formulated which can neither be proved not disproved within the system, unless it can be shown that the system itself is inconsistent.

The argument from consciousness. This argument states that a machine cannot be said to think until it composes a sonnet or writes a piece of music and is conscious of the fact that it has done so.

The argument from various disabilities. The form of this argument allows for machines to perform a number of 'thinking' tasks, but suggests that there is one or more tasks that a machine could not possibly perform.

Lady Lovelace's argument. The most detailed information available about Babbage's Analytical Engine comes from a memoir written by Lady Lovelace. In this memoir, a classic in the field, Lady Lovelace states

---

20. Nagel, Ernest and Newman, James, R. "Goedel's Proof." in <u>The World of Mathematics</u>, pp. 1668-1695.

that the machine can do whatever we know how to order it to perform, no more, no less.

The argument from continuity in the nervous system. This argument states that the nervous system is a continuous-, rather than a discrete-state machine. Therefore, no machine could hope to reproduce the workings of the nervous system.

The argument from informality of behavior. It is not possible to predict each and everyone of man's activities. Therefore, it is inconceivable to create a machine which could faithfully reproduce man's behavior.

The argument from extra-sensory perception. Since we must acknowledge the existence of such unexplained phenomena as telepathy, clairvoyance, precognition, and psycho-kinesis, we cannot possibly expect machines to account for these things.

Clearly, the arguments presented by Turing constitute a powerful array of logic against the possibility of artificial intelligence. The reader is urged to seek out the reference article to see how the author demolishes each and every one of the arguments described above.

Donald Michie is a member of the faculty of the Department of Machine Intelligence and Perception at the University of Edinburgh, one of the three leading places in the field of artificial intelligence. At a meeting of the Mathematical Society of the University in 1970, Michie explained the development of a new technology, such as artificial intelligence, in terms of three evolutionary stages, each of which overlaps or merges with the others over time.

The first stage of technological development is the phase of ad hoc innovation, exemplified by the Wright brothers in the field of aeronautics or Watt in the development of steam power. The second phase is the phase of formalization, that phase at which a new science is born, thermodynamics in the wake of the steam engineers, or aerodynamics in the wake of the fliers, or information and switching theories in the wake of Marconi and Bell. The final phase is one of stabilization and systematic development, when the researchers have moved on to newer fields and the practitioners have arrived.[21]

Michie believed that machine intelligence, as a whole was in headlong transit between phase one and phase two.

In the same article, Michie describes a class of problems which are defined in the following terms:

Given: an initial state
a definition of a desired state
a repertoir of actions,

Find: a sequence of actions which will transform the initial state into a desired state

---

21. Michie, Donald. "Heuristic Search." The Computer Journal, February 1971, pp. 96-102.

Michie believes that in the future, heuristic search studies, at least within the context of artificial intelligence, are likely to pay increased attention to devising automatic methods of imposing classifications on problem spaces. The object of this activity would be to subdivide the spaces into regions, and to have a different operator-selection regime be associated with each region.[22]

---

22. Ibid.

> For more than a decade, the list processing and symbol-manipulation languages -- such as COMIT, IPL, LISP, SLIP (Bobrow and Raphael 1964) -- have been the media for almost all AI achievements. Although the effectiveness of research with these languages has improved dramatically due primarily to greatly expanded memory sizes and new interactive debugging facilities, the languages themselves have remained remarkably stable. 23

The article from which the above quote was extracted provides a very useful summary of some of the computer languages which have been utilized in AI research. We have reproduced capsule descriptions of these languages for the sake of increasing reader awareness in a not-very-common field.

ABSET is a programming language based on sets which was developed at the University of Aberdeen, during 1971.

ECL is an extensible language system developed at Harvard in 1972 for use with automatic programming. It offers a combination of agreeable language sources, an interpreter for list structure program representation and various levels of compiler.

LISP-70 was developed in 1973, and is based on the LISP system. Although the full system is not yet available, a prototype called MLISP2 has been used in pattern matching and automatic rule maintenance.

PLANNER and MICRO-PLANNER, a subset, were written in LISP in 1972, at the MIT AI laboratory. MICRO-PLANNER runs only on the PDP-10 at MIT, but transfer to other LISP systems for experimental use is being accomplished.

---

23. Bobrow, Daniel and Raphael, Bertram. "New Programming Languages for Artificial Intelligence Research." Computing Surveys, September 1974, pp. 153-174.

POPLER 1.5 is a programmatic language implemented as an extension of POP-2, a system developed at the University of Edinburgh for research in artificial intelligence.

QLISP/INTERLISP are languages which were developed directly from QA4. They eliminate the slow utility functions and lack of debugging tools of that language. QLISP was piggy-backed unto INTERLISP, and both were developed at the Stanford Research Institute.

SAIL is a language project of the Stanford Research Institute AI Project. It is considered to be one of the most stable, debugged, and heavily used of the languages in the article. It runs on a PDP-10 under the DEC 10-50 monitor.

SMALLTALK is a language developed in 1973 which uses the notion of classes as an interrogative mechanism. The language enables a user to obtain an answer from an instance of a class without knowing whether the data is stored or procedurally developed.

The languages described in the article are characterized by their ability to do list processing and/or symbol manipulation, and have been developed to aid AI researchers in their efforts to enable computers to perform more and more complex tasks.

SOME EXAMPLES OF AI

This section refers the reader to a number of written descriptions of experiments in which computers were made to play games, solve problems, translate sentences, or recognize patterns. However, it must be understood that the written word is a poor substitute for the excitement which pervades a computer installation when a program designed to accomplish one or another of the above tasks is being tested for the first time, or when it finally accomplishes the task it was designed to carry out.

A general type of problem which AI researchers have attempted to solve is the situation-space problem. Briefly described, a problem of this type "... consists of an initial situation, a set of possible situations, and a set of possible actions, together with a specification of how the various situations can be produced from each other by different actions, and the specification of a final, desired situation, or goal." [24]

One of the best programs for solving situation-space problems is the General Problem Solver (GPS) program of Newell, Shaw, Simon, and Ernst. The program uses descriptions of objects and operators - situations and actions - and using a means-ends analysis technique (a name given by the authors) proceeds as follows:

1. Evaluate the difference between the current situation and the goal.
2. Find an operator that typically lowers the type of difference found in step 1.

---

24. Jackson, op. cit., p. 72.

3. Check to see if the operator found in step 2 can be applied to the current situation; if it can, then apply it, else determine a situation required for the application of that operator, and establish it as a new (sub) goal; they go to step 1. 25

The program has been applied to many different simple problems, such as the Missionary-Cannibals Problem, the Three-Coins Problem, and the Tower of Hanoi Problem. However, when the program was given the Seven Bridges of Konigsberg Problem, shown to be unsolvable by Euler in 1736, it tried again and again until it gave up.

Another computer program which has achieved a modicum of fame is the checker-playing program written by Dr. A. L. Samuels, of the International Business Machines Corporation. Samuels taught the computer how to play checkers and to continually improve its game until "... In August 1962, the computer took on Mr. Robert W. Nealy, 'a former Connecticut checkers champion, and one of the nation's foremost players,' and Mr. Nealy lost." 26

Other games for which programs have been written, with varying degrees of success, are Chess, Nim, Go, Tic-Tac-Toe, and Bridg-It. Many of these games are available commercially, and it is worthwhile for the interested reader who wants to explore the program implications to acquire a manual version and to seek to plot the probabilistic alternatives of the various moves.

Nilsson suggests that problem solving has two basic elements:

---

25. Jackson, *op. cit.*, p. 76.

26. Cote, Jr., Alfred J. *The Search for the Robots*, p. 199.

representation and search. The first element deals with the way in which a problem is formulated, and Nilsson states

> The problem of representation is common to any problem-solving approach, but unfortunately research in artificial intelligence has not yet produced any general automatic method for skillfully formulating problems. 27

The second element - search - deals with the steps used in arriving at the solution of a given problem.

The possibility that a machine might be able to perceive its surroundings has interested AI researchers for some time. One of the considerations in constructing a robot capable of surviving in the real world is that it must have the capability to analyze its environment and to take the necessary steps to avoid mishaps. The field of pattern perception encompasses visual patterns, sound patterns, symbol patterns, and even reasoning patterns. However, not all of these segments have been explored.

Among the exercises which computers, usually with a television camera hook-up, have been able to perform, we find one in which the computer is shown a series of aerial photographs of bridges running over rivers and bridges running over roads, and each type is identified for the computer. Upon recall, the computer has scored better than 80% in properly identifying the subtle differences.

One of the classical efforts in the area of perception was undertaken by Frank Rosenblatt, of the Cornell University Aeronautical

---

27. Nilsson, Nils J. Problem-Solving Methods in Artificial Intelligence, p. 8.

Laboratory, was a machine called the Perceptron. The first Perceptron was designed with the ability to recognize the letters of the alphabet. The commercial application of the idea in the form of OCR (optical character recognition) will not escape the reader.

Semantic information processing is probably the area of artificial intelligence in which most experimentation has been carried out. In his book Semantic Information Processing, Marvin Minsky provides detailed analyses of a number of programs which have been successfully developed in the field.

The first program discussed in Minky's work is Bertram Raphael's doctoral thesis. It describes the development of a program named SIR: Semantic Information Retrieval, in which "... an attempt (was made) to build a memory structure that converts the information it receives into a systematic, efficient representation ... The system ...works by understanding the statements when they are made, consolidating this understanding by adding to or modifying the network." [28]

Daniel G. Bobrow is the author of a program called STUDENT, which "... accepts as input a comfortable but restricted subset of English which can be used to express a wide variety of algebra story problems." [29]

Space limitations preclude a more extensive analyses of the myriad of programs which have been written in support of AI efforts.

---

28. Minsky, Marvin, ed. Semantic Information Processing, p. 4.
29. Minsky, op. cit, p. 135.

CONCLUSIONS

In his evaluation of computer applications in the field of artificial intelligence, Donald G. Fink writes

> We find that almost all the computer applications fail to qualify as "intelligent" under (a) or (b) of the definition (see p. 3) - useful, yes, beyond the fondest dreams of Eckert and Mauchly. Computers save time and tedium almost beyond comprehension, but intelligent, in the sense of being able to answer questions not explicitly foreseen? Not proved for the great majority of them. 30

Stanley Jaki, in his very penetrating analysis of the relationship between Brain, Mind and Computers, states

> The stage has been reached where the question "Can machines think?" is viewed as a projective test of personality. Those answering "yes" are described as "self-confident humanists..." Those who refuse to admit the possibility of thinking machines are called "doubters, pessimists... 31

Hubert L. Dreyfus, the most outspoken of the critics of artificial intelligence says

> ... my general thesis will be that the field of artificial intelligence exhibits a recurring pattern: early, dramatic success followed by sudden unexpected difficulties. 32

And then there was HAL, the computer in 2001: A Space Odyssey ...

---

30. Fink, op. cit., pp. 225-226.

31. Jaki, Stanley L. Brain, Mind and Computers, pp. 255-256.

32. Dreyfus, op. cit., p. xxxiii.

On a more positive vein, Jackson writes,

> ... it may be possible for artificial intelligence to greatly reduce the amount of human labor necessary to operate the economy of the world ... Computers and AI research may play an important part in helping to overcome the food, population, housing, and other crises that currently grip the earth ... 33

The following lines from Locksley Hall, by Alfred, Lord Tennyson, are representative of this author's feelings towards his subject:

> For I dipt into the future, far as human eye could see
> Saw the Vision of the world, and all the wonder that would be.

33. Jackson, *op. cit.*, p. 398.

LC Science Tracer Bullet

Reference Section, Science and Technology Division
Library of Congress, 10 First Street, S.E., Washington, D.C. 20540

ARTIFICIAL INTELLIGENCE
Compiled by B. I. Kravitt

TB 73-4 January 1973

SCOPE: Includes material on the technical and philosophical aspects of the question: Can a machine think? The interdisciplinary nature of the topic is reflected by its appearance in the literature of the engineering sciences, mathematics, linguistics, psychology and the biological sciences.

An introduction to the topic appears in:

Davis, Ruth M. Artificial intelligence. In Encyclopedia of library and information science. v. 1. New York, M. Dekker, 1968. p. 627-632. Z1006.E57 Vertical file*

SUBJECT HEADINGS under which books on artificial intelligence can be located in the LC card catalogs include the following:

ARTIFICIAL INTELLIGENCE (Highly relevant)
CYBERNETICS (Highly relevant)
BIONICS (Relevant)
CONSCIOUS AUTOMATA (Relevant)
SELF-ORGANIZING SYSTEMS (Relevant)
INFORMATION THEORY (More general)
MACHINE TRANSLATING (Related)
PATTERN RECOGNITION (Related)
PERCEPTRONS (Related)

BASIC TEXTS:

Dreyfus, Hubert L. What computers can't do: a critique of artificial reason. New York, Harper and Row, 1972. 259 p. Q335.D74 1972*

Findler, N. V., and Bernard Meltzer, eds. Artificial intelligence and heuristic programming. New York, American Elsevier, 1971. 327 p. Q335.A787

*Available in the Science Reading Room Collection
**Available in the Science Reading Room Microform Collection

Feigenbaum, Edward, and Julian Feldman. Computers and thought. New York, McGraw-Hill, 1963. 535 p. Q335.5.F4*

Jaki, Stanley L. Brain, mind and computers. New York, Herder and Herder, 1969. 267 p. Q335.J3*

Taube, Mortimer. Computers and common sense. New York, Columbia University Press, 1961. 136 p. Q310.T33

ADDITIONAL TEXTS:

Wiener, Norbert. Cybernetics, or control and communication in the animal and the machine. New York, J. Wiley, 1948. 194 p. Q175.W6516*

______________. The human use of human beings; cybernetics and society. Boston, Houghton Mifflin Co., 1950. 241 p. Q175.W6517*

Other books including material on engineering cybernetics are shelved under the following LC call numbers: TA166-168 and TJ212-225.

HANDBOOKS, ENCYCLOPEDIAS and DICTIONARIES which contain information on artificial intelligence include:

Encyclopaedia of cybernetics. Translated from the German Lexikon der Kybernetik by G. Gilbertson. Manchester, Manchester U. P.; New York, Barnes and Noble, 1968. Q310.L413 1968*

Meetham, A. R. Encyclopaedia of linguistics, information, and control. Oxford, New York, Pergamon Press, 1969. 718 p. Q360.M35 1969*

BIBLIOGRAPHIES:

Alum, Fazlul. Cybernetics: automation, computers, control, ergonomics, information theory, and machine translation: a subject guide. Provisional ed. London, New Science Publications, 1968. 110 p. Z7405.C9A4*

Campbell, Alan. Bionics and biocybernetics. Washington, Aerospace Technology Division, Library of Congress, 1968. 366 p. (ATD report 68-77-108-4) Z664.23.A2

Dutton, John M., and William H. Starbuck. Computer simulation models of human behavior: a history of an intellectual technology. In Institute of Electrical and Electronics Engineers. Systems, Man and Cybernetics Group. IEEE transactions on systems, man and cybernetics, v. SMC-1, Apr. 1971: 128-171. Q300.I43

Gibbs, Katye M. Bionics and related research; a report bibliography. Arlington, Va., Armed Services Technical Information Agency, 1963. 177 p. AD 294 150 Vertical file*

Sweitzer, Dorothy I. Biological and artificial intelligence. Pasadena, Jet Propulsion Laboratory, California Institute of Technology, 1960. 185 p. (Astronautics information literature search, no. 254) Z7405.A7S9 Vertical file*

____________. Supplement. Pasadena, Jet Propulsion Laboratory, California Institute of Technology, 1961. 160 p. Z7405.A7S9 Suppl. Vertical file*

SELECTED TECHNICAL REPORTS:

Illinois. University. Biological Computer Laboratory. Accomplishment summary, 1971-1972. Urbana, 1972. 275 p. AD 744 009**

Massachusetts Institute of Technology. Artificial Intelligence Laboratory. Description and theoretical analysis (using schemata) of planner: a language for proving theorems and manipulating models in a robot. Cambridge, Mass., 1972. 403 p. AD 744 620**

Stanford University. Computer Science Department. Artificial Intelligence Project. Project technical report. Stanford, Calif., 1971. 1 v. AD 724 867**

System Development Corporation, Santa Monica, Calif. Problem solving and learning by man-machine teams. Santa Monica, 1971. 229 p. AD 729 070**

STATE-OF-THE-ART REVIEWS and CONFERENCE PROCEEDINGS:

American Society for Cybernetics. Purposive systems; proceedings of the first annual symposium of the American Society for Cybernetics. New York, Spartan Books, 1969, c1968. 179 p. Q300.A45

American Society for Cybernetics. Cybernetics and the management of large systems; proceedings of the second annual symposium of the American Society for Cybernetics. New York, Spartan Books, 1969. 264 p. Q295.A45

American Society for Cybernetics. Cybernetics, simulation, and conflict resolution; proceedings of the third annual symposium of the American Society for Cybernetics. New York, Spartan Books, 1971. 249 p. HM136.A58

American Society for Cybernetics. Cybernetics, artificial intelligence, and ecology; proceedings of the fourth annual symposium of the American Society for Cybernetics. New York, Spartan Books, 1972. 333 p. Q335.A444

International Joint Conference on Artificial Intelligence, 1st, Washington, D.C., 1969. Proceedings. Bedford? Mass., 1969. 715 p. Q335.5.I55 1969

Machine intelligence. Proceedings of the Machine Intelligence Workshop. 1st, 1967- New York, American Elsevier. Q335.M27

JOURNAL ARTICLES and other literature on artificial intelligence are indexed primarily in the following abstracting and indexing services:

Computer Abstracts (1957-) Z6654.C17C64*
See: Artificial Intelligence (Section 4) and index

Computer & Control Abstracts (1966-)*
See: Artificial Intelligence (Section 61.40)

Computer and Information Systems (1962-) QA76.I46*
See: Artificial Intelligence (Section CA2) and index

Engineering Index (1884-) Z5851.E62*
See: Cybernetics

Quarterly Bibliography of Computers and Data Processing (1971-) QA76.Q3*
See: Artificial Intelligence

Science Citation Index (1961-) Z7401.Z87*
See especially Permuterm Subject Index (1966-) for entries under Artificial Intelligence

Other indexes, listed here, should be used for an exhaustive search. Only a limited return can be expected for the time spent. Directions are generally given in the front of each issue.

Computing Reviews (1960-) QA76.C5854*
Cybernetics Abstracts (1964-) Q300.R413*

REPORTS and other types of literature are indexed in the following guides:

Scientific and Technical Aerospace Reports (1963-) TL500.S35*
See: Artificial Intelligence
Cybernetics

Note: Consult reference librarian for location of abstracting and indexing services in the Science Reading Room.

Government Reports Index (formerly U.S. Government Research and Development Reports Index) (1946-) Z7405.R4U513*
See: Artificial Intelligence
Cybernetics

Rand Corporation. Selected Rand Abstracts (1963-) AS36.R284*
See: Artificial Intelligence
Cybernetics

PRIMARY JOURNALS that often contain articles relevant to artificial intelligence are:

Artificial intelligence Q335.A785
Association for Computing Machinery. Journal QA76.A77
IEEE transactions on information theory Q350.I2
Institute of Electrical and Electronics Engineers. IEEE transactions on computers TK7885.A1F2
Institute of Electrical and Electronics Engineers. IEEE transactions on systems science and cybernetics (1965-1970) TA168.I22
Institute of Electrical and Electronics Engineers. Systems, Man and Cybernetics Group. IEEE transactions on systems, man and cybernetics (1971-) Q300.I43

SELECTED MATERIALS available in the Science Reading Room vertical file include:

The age of thinking machines. Business week, no. 2200, Oct. 30, 1971: 90-92+

Hamming, Richard W. The computer and the intellectual frontier. Computers and automation, v. 21, June 1972: 25-29+

Kugel, Herbert C., Computer intelligence and general software development. Software age, v. 4, Oct./Nov. 1970: 14-15.

Thring, M. R. Robots in home and industry. In American Society for Cybernetics. Cybernetics, artificial intelligence, and ecology; proceedings of the fourth annual symposium of the American Society for Cybernetics. New York, Spartan Books, 1972. p. 253-268.

Turing, A. M. Intelligent machinery. In Machine intelligence 5; proceedings of the machine intelligence workshop. New York, American Elsevier, 1970. p. 3-23. Q335.M27

BIBLIOGRAPHY

Arbib, Michael A. The Metaphorical Brain. An Introduction to Cybernetics as Artificial Intelligence and Brain Theory. New York, Wiley-Interscience, 1972.

Banerji, Ranan B. Theory of Problem Solving. An Approach to Artificial Intelligence. New York, American Elsevier Publishing Company, Inc., 1969.

Barrow, H. G., et. al. "Tokyo-Edinburgh Dialogue on Robots in Artificial Intelligence." The Computer Journal, volume 14, February 1971, pp. 91-95.

Bobrow, Daniel and Raphael, Bertram. "New Programming Languages for Artificial Intelligence Research." Computing Surveys, volume 6, September 1974, pp. 153-174.

Bush, Vannevar. "As We May Think." The Atlantic Monthly, volume 176, July 1945, pp. 101-108.

Cote, Jr., Alfred J. The Search for the Robots. New York, Basic Books, Inc., Publishers, 1967.

Davis, Ruth M. "Artificial Intelligence." in Encyclopedia of Library and Information Science. Volume 1. New York, Marcel Dekker, Inc., 1968.

"Direct Brain-to-Computer Interface on the Way." Computer Decisions, September 1974, p. 13.

Dreyfus, Hubert L. What Computers Can't Do. A Critique of Artificial Reason. New York, Harper & Row, Publishers, 1972.

Feindel, William, ed. Memory, Learning, and Language. Toronto, University of Toronto Press, 1960.

Fink, Donald G. Computers and the Human Mind. An Introduction to Artificial Intelligence. Garden City, Anchor Books, 1966.

Frank, George. "What is Artificial Intelligence?" Data Processing, volume 14, September-October 1972, pp. 328-330.

Jackson, Jr., Philip C. Introduction to Artificial Intelligence. New York, Petrocelli Books, 1974.

Jaki, Stanley L. Brain, Mind and Computers. New York, Herder and Herder, 1969.

Krevitt, B. I., comp. "Artificial Intelligence." LC Science Tracer Bullet, TB 73-4, January 1973.

Licklider, J. C. R. Libraries of the Future. Cambridge, The M.I.T. Press, 1965.

Michie, Donald. "Heuristic Search." The Computer Journal, volume 14, February 1971, pp. 96-102.

Minsky, Marvin L. "Artificial Intelligence." Scientific American, September 1966, pp. 246-252+.

________. Semantic Information Processing. Cambridge, the M.I.T. Press, 1968.

Nagel, Ernest and Newman, James R. "Goedel's Proof." in The World of Mathematics. Volume 3, pp. 1668-1695.

Newman, James R., ed. The World of Mathematics. New York, Simon and Schuster, 1956.

Nillson, Nils J. Problem-Solving Methods in Artificial Intelligence. New York, McGraw-Hill Book Company, Inc., 1971.

Shannon, Claude E. "A Chess-Playing Machine." in The World of Mathematics. Volume 4, pp. 2124-2133.

Taube, Mortimer. Computers and Common Sense. New York, McGraw-Hill Book Company, Inc., 1961.

Turing, A. M. "Can A Machine Think?" in The World of Mathematics. Volume 4, pp. 2099-2123.

Von Foerster, Heinz, et. al. eds. Purposive Systems. Proceedings of the First Annual Symposium of the American Society for Cybernetics. New York, Spartan Books, 1968.

Von Neumann, John. The Computer and the Brain. New Haven, Yale University Press, 1958.

________. "The General and Logical Theory of Automata." in The World of Mathematics. Volume 4, pp. 2070-2098.