



---

■

# ARTIFICIAL INTELLIGENCE

*An Applications-oriented  
Approach*

---

■

Daniel Schutzer



VAN NOSTRAND REINHOLD COMPANY  
NEW YORK

---

Copyright © 1987 by Van Nostrand Reinhold Company Inc.

Library of Congress Catalog Card Number 86-13291

ISBN 0-442-28034-3

All rights reserved. No part of this work covered by the copyright hereon may be reproduced or used in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or information storage and retrieval systems—without written permission of the publisher.

Printed in the United States of America

Designed by Beth Tondreau

Van Nostrand Reinhold Company Inc.

115 Fifth Avenue

New York, New York 10003

Van Nostrand Reinhold Company Limited

Molly Millars Lane

Wokingham, Berkshire RG11 2PY, England

Van Nostrand Reinhold

480 La Trobe Street

Melbourne, Victoria 3000, Australia

Macmillan of Canada

Division of Canada Publishing Corporation

164 Commander Boulevard

Agincourt, Ontario M1S 3C7, Canada

16 15 14 13 12 11 10 9 8 7 6 5 4 3 2

**Library of Congress Cataloging-in-Publication Data**

Schutzer, Daniel, 1940–

Artificial intelligence.

Bibliography: p.

Includes index.

I. Artificial intelligence. I. Title.

Q335.S4116 1986 006.3 86-13291

ISBN 0-442-28034-3

---

# ARTIFICIAL INTELLIGENCE

---

# *Preface*

---

**A**lthough artificial intelligence has been studied academically since the late 1950s, the subject has recently generated wider interest because commercial applications now seem to be practical. One major factor in the successful transition of artificial intelligence from academia to industry is the dramatic advances in computer hardware that have occurred in the last two decades. Computer prices and sizes have plummeted, whereas memory capacity and processing speeds have increased to the point that personal microcomputers today possess all the power of the mainframes used by artificial intelligence researchers in the late 1950s and early 1960s. Because artificial intelligence applications tend to be computer-intensive, requiring a great deal of computer resources, the technological advances in the computer industry as a whole have increased the likelihood of successful commercialization.

Another important factor in the likely success of the commercialization of artificial intelligence is the increased maturity of the field itself. Although artificial intelligence is still far from being well understood, certain key unifying concepts and principles have emerged, and several important tools and techniques have been developed in support of these concepts. Much of this book is devoted to an exposition of these unifying concepts and principles and of the supporting tools and techniques that have been developed to facilitate their practical application.

Increased understanding and improved tools have produced many important successful prototypes, and these successes have captured the interest and investment dollars of both industry and government. Numerous projects have been initiated in diverse fields; many of these projects are reviewed herein. Several new companies specialize in artificial intelligence. In addition, artificial intelligence groups have been set up in many established companies, including many Fortune 500 firms. These activities have already led to several commercial successes and offerings, primarily in the area of expert systems and natural language

processing. Concrete examples of these successes are studied in this book.

Since 1983, a private industry consortium has set up the Microelectronics and Computer Technology Corporation (MCC) to develop new advanced computer technologies that apply artificial intelligence (A.I.) technologies. On the government side, the Defense Science Board has ranked artificial intelligence and robotics in the top ten military technologies for the 1980s. Anticipated applications include surveillance; weapons delivery, command, and control; maintenance; and training. The Department of Defense research agency DARPA has initiated an artificial intelligence-based supercomputer project, the Strategic Computing Program, which is intended to provide the necessary computer power for many military applications of artificial intelligence.

Interest in A.I. is worldwide. In June 1982, Japan launched a ten-year program to produce a fifth-generation computer system—a knowledge-information processing system (KIPS) based on A.I. concepts and principles. In early 1983, Japan began a project to develop a next-generation robot capable of advanced, autonomous decision making. Researchers in the robot project draw upon, among other things, results from the fifth-generation computer project. Great Britain has started a national A.I. project called the Alvey Project, and the European Common Market has initiated the ESPRIT program; interest and investment dollars are also evident in such countries as the Soviet Union, France, Germany, Austria, Italy, and Canada.

*Artificial Intelligence* examines these activities and discusses future trends and prospects that might result from them (as well as from complementary advances in computer hardware and software). A revolution in the way we process information is now taking place. This book offers fresh insights into how these directions and trends may affect the reader's field and into how new concepts, principles, tools, and techniques may best be exploited and applied to the reader's own work.

---

# *Acknowledgments*

---

**I** wish to thank my students for their insights and penetrating questions, which helped to shape this book. I would also like to thank my editor, Linda Venator, for her invaluable suggestions and critiques of the book. Finally, I would like to dedicate the book to my loving wife, Myra, without whose patience and support it could not have been written, and to our children, Eric, Richard, and Pamela.

---

# *Contents*

---

PREFACE	vii
ACKNOWLEDGMENTS	ix
<b>1. INTRODUCTION</b>	<b>1</b>
Definition of Artificial Intelligence	1
Key Attributes	2
Historical Background	7
<b>2. THEORIES AND CONCEPTS</b>	<b>13</b>
Knowledge Representation	14
Predicate Calculus	18
Production Rules	23
Frames and Scripts	27
Semantic Networks	31
Dealing with Uncertainty	36
Inference and Deduction	42
Formal Reasoning	44
Forward Chaining and Backward Chaining	52
Frame-based Representations in Complex Reasoning	57
Reasoning with Uncertainty	59
Problem Solving	67
Search Strategies	72
Planning	81
Learning	86
<b>3. TOOLS AND TECHNIQUES</b>	<b>94</b>
LISP	95
Logic Programming and PROLOG	110
Other Languages and Higher-Order Software Tools	122
LISP Machines and Other Special Hardware	123
<b>4. CURRENT APPLICATIONS</b>	<b>130</b>
Expert Systems	131
Machine Perception	142
Natural Language Processing	146



Speech Understanding	164
Vision and Image Understanding	174
Robotics	188
<b>5. SYSTEMS IMPLICATIONS</b>	<b>195</b>
<b>6. FUTURE PROSPECTS</b>	<b>205</b>
Appendix A. Summary of A.I. Research	223
Appendix B. Commercially Offered Expert Systems	232
Appendix C. Commercially Offered Natural Language Systems	247
Appendix D. Speech Synthesis Systems	251
Appendix E. Robot Programming Systems	259
<b>GLOSSARY</b>	<b>263</b>
<b>BIBLIOGRAPHY</b>	<b>266</b>
<b>INDEX</b>	<b>290</b>

---

# 1 ■ Introduction

---

## ■ DEFINITION OF ARTIFICIAL INTELLIGENCE

**A**rtificial intelligence (A.I.) is a field of study concerned with designing and programming machines to accomplish tasks that people accomplish using their intelligence. A.I. also attempts to understand how human beings think, by studying the behavior of machine designs and programs that model current hypotheses and conjectures about some aspect of the human cognitive process. Thus stated, this field of endeavor is almost as old as the human species.

Our understanding of what constitutes intelligence is vague, and definitions of human intelligence are imprecise. The most widely accepted standard measure is the IQ (intelligence quotient) test, but the validity of this test as a true gauge of a person's intrinsic intelligence (potential or achieved) is still hotly contested in many quarters. Consequently, it should not be too surprising to learn that the definition of machine (or artificial) intelligence is equally vague.

Many attempts have been made to define and demonstrate more precisely what is meant by artificial intelligence. Turing (1963) proposed the following test of machine intelligence: if a person engaged in a typewritten discourse with a machine hidden behind a curtain could not determine whether the conversation was with another person or with a machine, the machine could be said to exhibit intelligence. Early A.I. programs addressed this challenge with mixed success. Another attempt to demonstrate machine intelligence involved programming a computer to solve portions of an IQ test. The attempt to define what behavioral characteristics a machine must possess to be considered intelligent, however, is best considered as an evolutionary process; advances in machine intelligence often result in more exact redefinitions of intelligent behavior.

But *Artificial Intelligence* is less concerned with the philosophical question of what constitutes intelligence than with the narrower question of which theories, concepts, tools, techniques, and models produced

from research in this field can be applied to building better computer-based systems (including “smart” robots). To this end, the book concentrates on attributes and features of A.I. that distinguish it from the more conventional algorithmic computer problem-solving methods.

A.I. is a mind-set, a way of looking at and solving problems from a particular point of view. The concepts, methods, and techniques discussed in this book have been developed as a consequence of approaching problems from the relatively heuristic point of view that distinguishes the field of artificial intelligence from other computer science and systems engineering disciplines. The strengths and weaknesses associated with this approach, as well as the methodologies and techniques it has produced, are critically examined so as to enable the reader to recognize attributes and characteristics of a problem that identify it as most suitable for solution through heuristic programming or A.I. processing techniques.

Certain software tools and special hardware are available today to assist in applying these concepts, methods, and techniques; obtaining them and getting more information about them are discussed in later chapters. The current status and maturity of these tools and techniques are assessed, and past and existing examples of where this technology has been applied are reviewed. Finally, some forecasts of likely futures for this technology are explored.

## ■ KEY ATTRIBUTES

More than anything else, artificial intelligence represents a particular problem-solving mind-set that differs from more traditional algorithmic-oriented approaches. Much of the difference is traceable to the nature of the problems A.I. researchers tackle. A.I. research has concentrated on solving problems that are demonstrably solvable by human beings but for which no currently well-formulated, computationally feasible methodology exists. This class of problems includes:

- How we learn
- How we play games
- How we communicate with one another
- How we perceive (see, hear, speak, listen, write)
- How we create

Thus, A.I. has been concerned with human mental activities that are among the least well understood.

The approaches taken to solve these problems can best be characterized as informal, symbolic, and conceptual rather than as quantitative. Solutions are generally nonoptimal and are reached principally through approximate rules of thumb (for example, if  $A$  and  $B$ , then  $C$ ), and through logical or plausible inferences based on these rules (for example, if  $A$  implies  $B$  and  $B$  implies  $C$ , then  $A$  implies  $C$ ), rather than through calculation or step-by-step procedures.

Because A.I. has developed largely in the research community and because the methodologies needed to solve problems of interest to this group are poorly understood, A.I. has been largely a trial-and-error science. Consequently, the techniques developed have been designed with change and evolution in mind. The methods employed tend to be nonprocedural, flexible, and adaptive in nature; additional knowledge (facts and rules) can be added, deleted, or modified easily without system design modification. Furthermore, A.I. system architecture usually facilitates evaluating, examining, and analyzing proposed changes before their permanent addition.

The knowledge is often expressed in an object-oriented representation (where data, procedures, and rules are combined into conceptual units and communicate with each other through messages), rather than a system of separate data, separate procedure-oriented knowledge representations. One reason the problems under investigation are not amenable to well-understood quantitative methodologies is because the data are dirty. Available information is often noisy, full of uncertainty and errors, incomplete, and often inconsistent.

Many researchers attracted to A.I. are cognitive psychologists by training or inclination who are interested in understanding how human beings think. They observe and study people in various problem-solving situations and attempt to formulate theories and hypotheses at the macro level (not at the level of neural connections) to explain observed behavior. They develop computer models that test these hypotheses by attempting to duplicate the observed behaviors. These models and theories have strongly influenced the A.I. field, inspiring many A.I. concepts and techniques.

Of course, many other A.I. researchers are not concerned with whether or not their approach faithfully models the way human beings solve problems; they are concerned only that their programs work. These researchers are justified in seeking more novel, machine-unique solutions by the fact that (at least for the current crop of computer hardware) the computer and the human brain differ in several basic ways. The human brain has about 40 billion neurons (a neuron can be thought of

as representing approximately 1 byte of information), whereas today's computers typically have main memories ranging from about 2 million bytes (2 Mbytes) for personal microcomputers to several hundred million bytes for large mainframes. Since we believe that decision making, learning, and other "intelligence-oriented" functions use only a comparatively small percentage of the brain's total capacity—typically 10 to 30 percent—the equivalent of only about 10 billion bytes (10 Gbytes) of neural memory probably are available to a human being for intelligence-oriented functions. This, however, represents a far greater capacity than that of today's computers. Moreover, each neuron has from 1,000 to 10,000 inputs and outputs, with over 100 trillion interconnections. By contrast, today's computer components are relatively sparsely interconnected, with no more than 4 inputs per logic gate. On the other hand, neural impulses typically travel at a speed of about 10 miles per hour, whereas electrons in an electronic circuit can travel at the speed of light. Finally, the brain's neurons generally fire on a majority threshold basis, whereas computer components work on a binary logic basis.

The final assessment of computer problem solving and the human paradigm has not yet been made. A case supporting the notion that computers should not necessarily emulate human problem-solving techniques is the current crop of chess-playing machines. These are specially designed machines that can rapidly anticipate many more moves than most people can, but they possess little strategic planning capability. They certainly do not play chess the way people do, yet they perform reasonably well against human opponents. They do, however, have recognizable shortcomings, and researchers are attempting to supplement current chess-playing machine designs with human-style approaches, including strategic planning, the setting of intermediate goals, and associated tactics.

In summary, although A.I. has much in common with other computer science disciplines, it differs from more conventional computer science areas in the following respects:

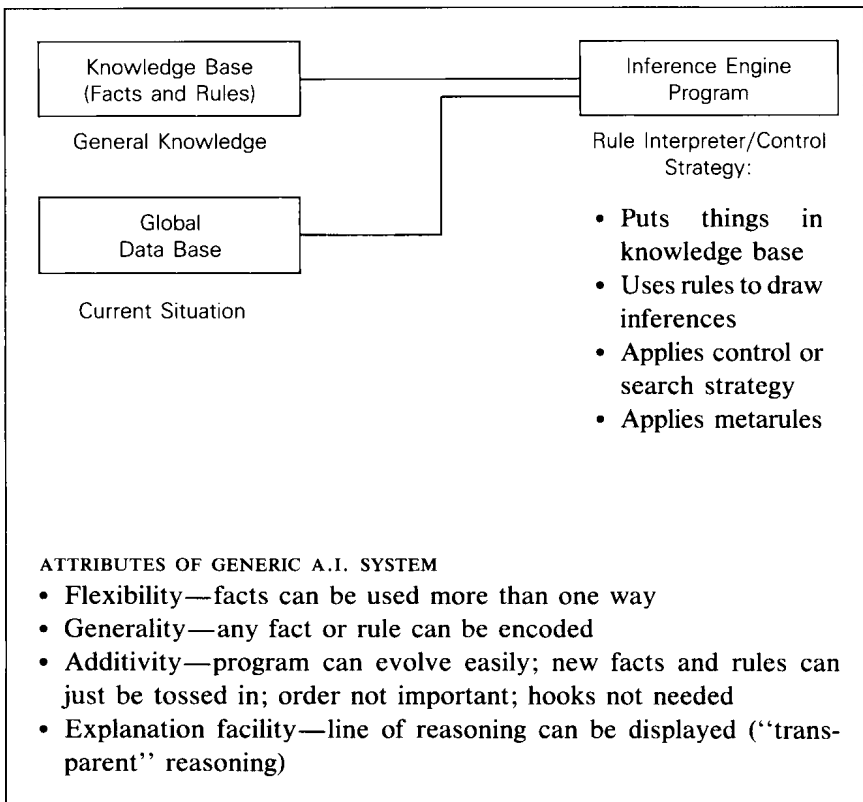
1. Viewpoint (plausible and logical reasoning, instead of quantitative calculation)
2. Subject matter (mental activity—very knowledge-intensive)
3. Tolerance for errors and imprecise data
4. Symbolic manipulation (instead of numeric orientation)
5. Evolutionary design principles (nonprocedural, anticipating addition and change)
6. Knowledge-based design

7. Inference and deduction capabilities (has a line of reasoning and can explain itself)
8. Heuristic or approximate problem-solving approach

Of course, any of these individual attributes can be found in more conventional computer science disciplines; for example, compilers do symbolic processing. Taken together, however, these attributes characterize a unique, human-style approach to problem solving that differs substantially from the mainstream computational approach.

### A GENERIC A.I. SYSTEM

Typical A.I. system architecture is pictured in figure 1-1. It includes a knowledge base (in which the problem-domain knowledge is explicitly represented as a collection of facts, rules, and relationships about the



**1-1.** Generic A.I. system.

problem-domain objects and concepts), a global data base (which contains assertions about the current case/problem under consideration, expressing the current state of knowledge about the particular problem being solved), and an inference engine program (which represents the control or procedural knowledge that updates the global data base, feeds information to the global data base, determines how to use the rules in the knowledge base to draw conclusions, and establishes the sequence in which different rules in the knowledge base are brought to bear in solving the problem).

In A.I. systems, program control is generally not a predefined, step-by-step procedure in which order is important. It is more of a trial-and-error procedure in which searches are made of a space of candidate solutions, and heuristics are used to prune the combinational growth that occurs in most complex real-world problem searches. The inference engine applies the control or search strategy; it determines when to apply which rules against what part of the data base to produce an output or to reach a goal or conclusion. This strategy is often expressed by heuristic rules of thumb that are pattern-invoked, triggered by the specifics of the problem state, called metarules. This approach is characterized by:

- *Flexibility*: facts can be used in more than one way.
- *Generality*: any fact or rule can be encoded.
- *Additivity*: the program can evolve easily, with new facts and rules added randomly. The order is not critical, and hooks are not needed. The program tolerates incomplete answers, occasional errors, misses, and less than optimum results.
- *Transparent reasoning*: often these programs are run interactively and provide, upon request, an explanation in human language of the line of reasoning used—how a particular conclusion was reached or why additional information was requested.

Because the knowledge base is usually an explicit (declarative) representation of the domain knowledge, it is maintained in a form that makes it accessible for more than one kind of use. For example, the knowledge base might contain a description of the various pieces in the game of chess, their legal moves, a few strategic concepts such as the importance of controlling the center of the board, and a list of successful opening moves. The global data base contains the evolving state of the particular chess game being played. The knowledge base can be used by the inference engine to determine the next move to make. It can also be used to answer questions about the game being played and the current situation.

The size and quality of the knowledge base is important in determining system performance. Generally, the greater the knowledge, the higher the quality of the system's problem-solving ability and the better the performance of the system. Of course, the larger the number of rules to consider, the more computer processing time and memory are required. This processing overhead can sometimes be reduced through a clever control or search strategy, much as a human expert can be efficient in determining which of the many facts and rules to apply to a particular problem situation.

The inference engine is generally a separate subsystem that interprets the contents of the knowledge base for one or more purposes and applies problem-domain-specific procedural knowledge to focus and limit what would otherwise be a large search.

Artificial intelligence encompasses many different ideas and disciplines. Another way to characterize A.I. is by the fields of study it encompasses. The January 1982 Computing reviews classification system by Samet and Ralston defines A.I. as including:

- Expert systems
- Game playing
- Automatic programming
- Deduction and theorem proving
- Knowledge representation
- Learning
- Natural language processing
- Problem solving
- Robotics
- Vision and scene understanding

## ■ HISTORICAL BACKGROUND

Artificial intelligence research has gone on almost since the introduction of the digital computer. It is based on the idea that all intelligent activity can be formalized and described by some sort of computable function—that is, that all human intelligence can, in principle, be mimicked by a digital computer. In a sense, A.I. carries the banner first raised by the ancient Greeks, who believed in the fundamental rationality of human beings.

Work in A.I. began in earnest in the 1950s, with great expectations that important successes would be achieved quickly. Industry anticipated early commercialization. Most of the seminal work during this period



was done at Rand-Carnegie (by researchers that included Alan Newell, Herb Simon, and J.C. Shaw) and at MIT (by researchers that included Marvin Minsky and John McCarthy). The initial problems undertaken included games and theorem proving. The timeline below identifies the succession of developments.

**1950**—Claude Shannon designed a chess-playing program that introduced the “game tree” concept.

**1955**—Alan Newell, J.C. Shaw, and Herb Simon showed that chess playing was analogous to the problem of proving theorems in symbolic logic.

**1956**—A Dartmouth workshop, organized by John McCarthy, coined the term *artificial intelligence*.

**1957**—Newell and his associates developed their Logic Theorist program, which proved theorems in symbolic logic using heuristic rather than exhaustive techniques. They invented IPL (Information Processing Language), the first list-processing language, to assist them in this work.

**1958**—Newell and his associates developed a chess program in IPL. They also developed the famous General Problem Solver (GPS), which served as an executive for managing the general searches used to solve problems in elementary logic, chess, high-school algebra word problems, and question-answering systems.

**1958**—John McCarthy developed programs capable of question-answering and developed the computer language LISP. LISP is currently the language of choice for most A.I. researchers.

**1959**—McCarthy worked on systems capable of deductive reasoning and information retrieval.

Research at this time centered on developing and understanding general problem-solving strategies, known as “weak methods.” The problems proved much more difficult, and the general problem-solving approaches too weak to allow their solution on the existing computer hardware within reasonable time constraints.

By the early 1960s, few practical applications of A.I. had emerged, industry’s interest had dissipated, and general disillusionment with A.I. research prevailed. Research continued in a few academic centers, namely Carnegie-Mellon, Stanford, MIT, and SRI (Stanford Research Institute). During this period, research began to move in a different direction. It addressed problems of much narrower scope and augmented the general problem-solving methodology (the weak methods) with very specific knowledge. The developments of the 1960s are listed below.