

**ARTIFICIAL  
INTELLIGENCE  
IN BUSINESS**



**EXPERT  
SYSTEMS**

**Paul Harmon  
David King**

# **EXPERT SYSTEMS**

---

## **Artificial Intelligence in Business**

**PAUL HARMON AND DAVID KING**

A WILEY PRESS BOOK

**JOHN WILEY & SONS, INC.**

New York • Chichester • Brisbane • Toronto • Singapore

Publisher: Judy V. Wilson  
Editor: Theron Shreve  
Managing Editor: Katherine Schowalter  
Composition, Design, & Make-Up: Ganis & Harris, Inc.  
Illustrations: Don Sherefkin

*Transcripts of the interaction with MYCIN were produced by the authors using the facilities of SUMEX-AIM. SUMEX-AIM is the Stanford University Medical EXperimental computer exploring Artificial Intelligence applications in Medicine. The assistance of Edward H. Shortliffe and Miriam Bischoff is appreciated.*

Copyright © 1985 by John Wiley & Sons, Inc.

All rights reserved. Published simultaneously in Canada.

Reproduction or translation of any part of this work beyond that permitted by Section 107 or 108 of the 1976 United States Copyright Act without the permission of the copyright owner is unlawful. Requests for permission or further information should be addressed to the Permissions Department, John Wiley & Sons, Inc.

#### Library of Congress Cataloging in Publication Data

Harmon, Paul.

Expert systems.

Bibliography: p.

Includes index.

1. Expert systems (Computer science) 2. Artificial intelligence. 3. Business—Data processing. I. King, David, 1949- . II. Title.  
QA76.9.E96H37 1985 001.53'5 84-27148  
ISBN 0-471-81554-3  
ISBN 0-471-80824-5 (pbk.)

Printed in the United States of America

10 9 8

---

PC-DOS and IBM are registered trademarks of International Business Machines, Inc.

LISP/VM, VM/SP, and YES/MVS are trademarks of International Business Machines, Inc.

Apple is a registered trademark of Apple Computer, Inc.

Macintosh is a trademark licensed to Apple Computer, Inc.

DEC, PDP, DEC-VAX, VMS, VAX LISP, TOPS-20, and VAX are trademarks of Digital Equipment Corporation.

Xerox, Xerox Star, and Xerox 1100's are registered trademarks of Xerox Corporation.

INTERLISP-D is a trademark of Xerox Corporation.

The KEE System is a trademark of IntelliCorp.

Expert-Ease is a trademark of Expert Software International Ltd.

Expert-Ease is a copyright product of Intelligent Terminals Ltd.

Personal Consultant and Explorer are trademarks of Texas Instruments Inc.

MS-DOS is a trademark of Microsoft Corporation.

Lotus and 1-2-3 are trademarks of Lotus Development Corporation.

WordStar is a registered trademark of MicroPro International Corporation.

SRL+ is a trademark of the Carnegie Group Inc.

Symbolics, Symbolics 3600, and MACSYMA are trademarks of Symbolics Inc.

UNIX is a trademark of AT&T Bell Laboratories.

TIMM is a trademark of General Research Corporation.

M.1 is a trademark of Teknowledge Inc.

SeRIES-PC is a trademark of SRI International.

INTELLECT is a trademark of Artificial Intelligence Corporation.

IQLISP is a trademark of Integral Quality.

GCLISP, Golden Common LISP, and K:base are trademarks of Gold Hill Computers, Inc.

Apollo is a trademark of Apollo Computer.

LMI and LAMBDA are trademarks of LISP Machine Inc.

VisiCalc is a registered trademark of VisiCorp.

PROLOG-I is a trademark of Expert Systems, Ltd.

CP/M is a registered trademark of Digital Research, Inc.

INSIGHT Knowledge System is a trademark of Level 5 Research.

TK! Solver is a trademark of Software Arts, Inc.

The Sales Edge is a trademark of Human Edge Software Corporation.

To  
Denny Brown and Steven Hardy

## Photo and Figure Credits

- Figs. 1.1 and 1.3. Adapted from *Knowledge-Based Expert Systems*, by Frederick Hayes-Roth, appearing in *Computer Magazine*, vol. 17, no. 10, pp. 263–273. Copyright © 1984 by IEEE.
- Fig. 4.3 Modified after Collins and Quillian, 1969.
- Fig. 8.1. Modified after a figure originally developed for Teknowledge Inc.
- Fig. 8.2. Reproduced by permission of Expert Systems International. Copyright © 1984: All rights reserved.
- Fig. 8.3. Used with permission of Teknowledge Inc.
- Fig. 8.4. Used with permission of Teknowledge Inc. Photograph by Haruko.
- Fig. 8.5. Used with permission of Teknowledge Inc. Photograph by Haruko.
- Fig. 8.6. Used with permission of Teknowledge Inc. Photograph by Haruko.
- Fig. 8.7. Courtesy of Inference Corporation.
- Fig. 8.8. Copyright © 1983, 1984 by IntelliCorp. All Rights Reserved. Permission to Reprint by IntelliCorp, Menlo Park, California.
- Fig. 8.9. Copyright © 1984 by IntelliCorp. All Rights Reserved. Permission to Reprint by IntelliCorp, Menlo Park, California.
- Fig. 8.10. Copyright © 1984 by IntelliCorp. All Rights Reserved. Permission to Reprint by IntelliCorp, Menlo Park, California.
- Fig. 8.11. Courtesy of American Association for Artificial Intelligence.
- Fig. 9.3. From “Recognition of a Hidden Mineral Deposit by an Artificial Intelligence Program,” Campbell, A. N. et al, *Science*, vol. 217, pp. 927–929, Fig. 4, 3 September 1982.
- Fig. 9.4. From “PUFF: An Expert System for Interpretation of Pulmonary Function Data,” Aikins, J. S., et al., 1982.
- Fig. 10.2. Copyright © 1984 by IntelliGenetics, Inc. All Rights Reserved. Permission to Reprint by IntelliGenetics, Inc. Palo Alto, California.
- Fig. 10.4. Courtesy of General Electric Research and Development Center.
- Fig. 10.5. Used with permission of Teknowledge Inc.
- Fig. 10.6. Used with permission of Teknowledge Inc.
- Fig. 10.7. Adapted from Sholom M. Weiss and Casimir A. Kulikowski, *A Practical Guide to Designing Expert Systems*. Totowa, NJ: Rowman & Allenheld, 1984), p. 109.
- Fig. 13.2. Copyright © 1983, 1984 by IntelliCorp. All Rights Reserved. Permission to Reprint by IntelliCorp, Menlo Park, California.
- Fig. 15.3. Courtesy of American Association for Artificial Intelligence.
- Fig. 15.4. Adapted from Paul Harmon, “The Design of Instructional Materials: A Top-Down Approach,” *Journal of Instructional Development*. Fall 1982, 6, No. 1.

# Foreword

Everyone in business, from trainer to chief executive officer, must daily face problem solving and decision making based on extensive but incomplete, uncertain, and even contradictory data and knowledge. Up until now, the conventional use of the computer has filled rooms with printouts without being of much assistance in these areas. This state of affairs is largely due to a literal-minded approach to the computer.

Well, help is on the way. Paul Harmon and David King have created a fine introduction to the design and use of the expert systems that enlist the computer in a healthy and powerful way to solve the difficult and important problems of business.

“Artificial intelligence” has a forbidding and threatening ring to most of us. The words suggest supplanting, rather than supplementing, and denying, rather than extending, our human capabilities to set goals and directions. But this area of computer science has now fathered the creation of expert systems (also called knowledge systems), and these new systems promise to harvest the remarkable developments of computer science and computers to make all of us more effective.

Furthermore, as this book teaches us, expert systems can be developed that factor experience into decision making in a fashion that is friendly as well as powerful and that gives us comfort and confidence in our actions. A proper expert system can be far more useful and reliable than any expert information source otherwise available to us. Expert systems do not tire or become cranky; they do not bluff but instead tell us the limitation of their knowledge and estimate the uncertainty of their conclusions; they process and distill the experience of many experts and apply it to our problem without bias; they tell us upon our demand what assumptions they are making and

what their line of reasoning is; in short, they add breadth and depth to our reasoning and decision processes.

The authors here introduce us to expert systems over the whole range, from the simpler here-and-now training exercises to prospective systems that will aid the very top levels of a large corporation. They show us how to get started by adopting a limited system for a confined task and observing its costs, its benefits, and how it changes the way we view our tasks and the use of experts. They make a compelling case that the time is ripe to invest in the experience that we shall soon need to move promptly as ever more capable expert systems are developed.

As this development occurs, we shall be applying this experience as we interact with our own experts, with the knowledge engineers who create the expert systems, and with the vendors of tools and systems. This book provides the knowledge that will be indispensable as we decide how and where to start, how to assess the claims and the performance of people and vendors, and how fast to move. There is a great deal of push from the computer scientists, engineers, and ambitious corporations to develop systems rapidly; this book educates us so that we can decide how rapidly to pull expert systems into our businesses.

It is early yet to estimate the magnitude of the contribution expert systems will make to the extension of human capability and to our effectiveness as managers, and it would be more than a little reckless to rank it now along with steam power and electricity. But the contribution will be in that class and will be indeed profound. With the help of this book we can now move into this new world.

Robert L. Sproull

# Preface

Suddenly everyone is talking about artificial intelligence and expert systems. Those who come in contact with the basic ideas behind expert systems immediately begin thinking of ways to improve business by applying this new computer technology to solve some of the persistent and difficult problems they face.

We first became excited about expert systems in 1983. One of us is a psychologist and the other a management consultant who specializes in helping corporations solve human performance problems. We've worked on several projects together, and in early 1983 we found ourselves engaged in helping a new expert systems company develop workshops to explain its technology to businesspeople. The more we've learned about expert systems, the more we have become convinced that expert systems will, in fact, change the way businesses operate. Moreover, expert systems will change the way people think about solving a vast array of different problems.

We have consulted with many large corporations and several small businesses. In each case, we worked with a client to help solve some practical business problems. Occasionally, we encountered problems that simply could not be solved with existing technologies. More often, we have faced problems that were so difficult or complex that they could only be partially solved. As we learned about expert systems, we realized that this new technology will make it possible to develop quick, pragmatic solutions to a wide range of problems that currently defy effective solutions.

We have written this book for executives, middle managers, computer systems personnel, and corporate trainers. We are convinced that each of these groups of people needs to know about expert systems. In order to address such a diverse audience, we have

included some material that will be of interest to one group and not the other. We have avoided unnecessary computer and business jargon. Where we introduce technical terms we define them with clear, nontechnical examples.

The following suggestions should help you maximize the useful information that you gain from this book.

Everyone should read Chapter 1. It provides a general introduction to the subject and defines key terms.

If you are an executive or a middle-level manager, you should probably skip to Section Four and review how expert systems are likely to be used in the near future. Then, depending on how much depth you want, you can look at Section Three for some ideas on the steps involved in actually developing a system, or at Section Two for information on the costs involved in purchasing system building tools or for examples of systems that have already been built.

If you are involved in computer systems development or implementation, you will probably want to read the chapters in the order in which they are presented. If you already know about the theoretical underpinnings of AI and expert systems, you might skip Section One and only return to it after you have looked at Section Two. Our reference list, explained below, will be especially useful in locating primary sources.

If you are involved in training or are particularly interested in applying expert systems to the practical problems of improving employee productivity, we recommend that you read Chapters 3, 7, 8, 11, and then Section Four. This route through the book emphasizes the development of "small knowledge systems." After taking this path you will probably

want to go back and read Section One more carefully, but that is something you can decide a little later.

The book does develop concepts in a systematic manner and thus, if you skip around, you may encounter terms that you don't know. We have included a glossary as Appendix A that should help readers who choose an unconventional route.

To make the book flow as smoothly as possible, we have avoided footnotes and references. Annotated citations are collected in Appendix C. Information about companies is located in Appendix B. We have limited our references to commonly available books or magazine articles. This is consistent with our goal of providing a broad overview of expert systems and their business applications. For readers needing more technical depth, we identify books in our reference list that provide citations to the extensive technical and scholarly literature of artificial intelligence.

Our understanding of applied artificial intelligence has been shaped by a number of people. Our professional and academic interests in cognitive science

provided the foundation. Our first professional contact with expert systems occurred when we assisted Teknowledge Inc. with the development of training programs. Many people at Teknowledge have given us advice and shared their ideas with us, and we gratefully acknowledge their influence.

We have talked to many other people—more than we can list—who have helped us to clarify our ideas. A few individuals stand out and we'd like to acknowledge them: Thomas F. Gilbert, Frederick Hayes-Roth, Cliff Hollander, Peter E. Hart, Tom Kehler, Claudia Mazzetti, and Karl M. Wiig.

We also thank James Eilers, who helped prepare the manuscript and graphics for the book, and Theron R. Shreve and the many other helpful people at John Wiley & Sons. Finally, of course, we must express our appreciation to our families for tolerating our obsession with this project.

Paul Harmon  
David King



# Contents

Foreword	ix
Preface	xi
1. <b>Introduction</b> Artificial Intelligence • Knowledge-Based Expert Systems • Conventional Programming Versus Knowledge Engineering • The Expert Systems Market	1
<b>Section One Basic Concepts and Techniques</b>	13
2. <b>MYCIN</b> MYCIN's Expertise • A Consultation with MYCIN • How Good an Expert Is MYCIN? • Summary	15
3. <b>Human Problem Solving</b> Human Information Processing • The Production System as a Processing Model • Problem Solving • Varieties of Knowledge • The Nature of Expertise • Summary	22
4. <b>Representing Knowledge</b> An Informal Look at a Knowledge Base • Strategies for Representing Knowledge • Semantic Networks • Object-Attribute-Value Triplets • Rules • Frames • Representing Facts and Relationships Using Logic • Summary	34
5. <b>Drawing Inferences</b> Inference • Control • The Future of Representation and Inference • Summary	49
6. <b>MYCIN Revisited</b> Where Is the Expertise in MYCIN?	61
<b>Section Two Languages, Tools, and Systems</b>	77
7. <b>Languages and Tools</b> Levels of Software • The Language-Tool Continuum • AI Languages and Environments • Knowledge Engineering Tools	79
8. <b>Commercial Tools</b> Categories of Tools • Consultation Paradigms • Cost Considerations • Training and User Interface Issues • Evaluating Knowledge Engineering Tools • Commercially Available Tools • Small System Building Tools • ES/P ADVISOR • Expert-Ease • INSIGHT Knowledge System • M.1 • Personal Consultant • Additional Small Tools • Large, Narrow System Building Tools • EXPERT • KES—Knowledge Engineering System • OPS5 • S.1 • TIMM—The Intelligent Machine Model • Large, Hybrid System Building Tools • ART—The Automated Reasoning Tool • KEE—The Knowledge Engineering Environment • LOOPS • Additional Large Tools • Summary	92

9.	<b>Early Systems</b> DENDRAL • MACSYMA • HEARSAY I and II • INTERNIST/CADUCEUS • PROSPECTOR • PUFF	134
10.	<b>Recent Systems</b> XCON (R1), XSEL • GENESIS • DELTA/CATS-1 • DRILLING ADVISOR • A Microprocessor-Based Electrophoresis Interpreter • Summary	155
<b>Section Three</b>	<b>Developing Expert Systems</b>	175
11.	<b>Building a Small Knowledge System</b> The Role of Small Systems • Building a Small System • Summary	177
12.	<b>How Large Knowledge Systems Are Developed</b> Knowledge Engineering • Knowledge System Development • Phase I: Selection of an Appropriate Problem • Phase II: Development of a Prototype System • Phase III: Development of a Complete Expert System • Phase IV: Evaluation of the System • Phase V: Integration of the System • Phase VI: Maintenance of the System	195
<b>Section Four</b>	<b>The Expert Systems Market</b>	209
13.	<b>The Market for Knowledge Engineering</b> What Needs Will Drive the Knowledge Engineering Market? • How Will the Knowledge Engineering Market Be Organized? • How Will Businesses Use Knowledge Systems? • The Macroenvironment • Capitalization of AI • Summary	211
14.	<b>Knowledge Systems in the Next Five Years</b> Types of Companies • Types of Systems • Large, Hybrid Systems • Large, Narrow Systems • Professional Workstations • Small, Specialized Systems • Smart Generic Software • Systems in Hardware • Other Applications	226
15.	<b>Expert Systems for Training</b> Education Versus Training • STEAMER • DEBUGGY • GUIDON • DELTA • PUFF • Other Training Applications • Training Managers and Salespeople	236
16.	<b>Preparing for the Knowledge Systems Revolution</b> Acquiring More Knowledge • Seminars and Consultants • Experimenting with Small Systems • Hiring Outside Consultants • Developing an In-House Knowledge Engineering Group • A Final Word	249
	<b>Appendixes</b>	255
	Appendix A. Glossary	257
	Appendix B. Companies Active in Expert Systems and AI-Related Research	268
	Appendix C. References	274
	Index	279

# 1.

---

## Introduction

This book is about expert systems and how they will change the world of business. In the next 10 years, the concepts and techniques described in this book will revolutionize what we do with computers. Some computers will still process data, crunch numbers, and perform all the tasks we expect computers to handle today. New software and new computers, however, will soon be available.

Since World War II, computer scientists have tried to develop techniques that would allow computers to act more like humans. The entire research effort, including decision-making systems, robotic devices, and various approaches to computer speech, is usually called artificial intelligence (AI). Most AI efforts remain in the research labs. A collection of AI techniques that enables computers to assist people in analyzing problems and making decisions, called knowledge-based expert systems, however, has recently proved its value, and numerous commercial applications are now underway. Expert systems are being developed to assist managers with complex planning and scheduling tasks, diagnose diseases, locate mineral deposits, configure complex computer hardware, and aid mechanics in troubleshooting locomotive problems.

As recently as 1980, expert systems research was still confined to a few university research laboratories. Today, the United States, Japan, England, and the European Economic Community are all in the process of launching major research programs to develop and implement expert systems in the near future. Many Fortune 500 corporations are assembling AI departments, venture capitalists are rushing to invest in entrepreneurial expert systems companies, and expert systems technology is well on its

way to commercial success.

Expert systems will change the way businesses operate by altering the way people think about solving problems. This new technology will make it possible to develop quick, pragmatic answers for a wide range of problems that currently defy effective solutions.

Consider, for example, the problem of designing a program to train salespeople to develop highly structured financial proposals when the services and the packaging options are constantly changing. You can teach the general skills, but the specific products and techniques will be out of date before the course is even completed. Imagine, instead, that you could design a workshop that would teach the salespeople the general skills and then teach them to use a personal computer-based program that interacts with the salesperson by asking questions and then recommending appropriate options. It would be as if each salesperson could talk with a senior salesperson who had all the latest information whenever he or she wanted advice. Imagine, further, that the computer program that provided this advice was so modularized and so user-friendly that it could be quickly updated by product specialists rather than computer programmers. Thus, the program would always represent the latest products and packaging strategies. This sort of scenario is on the verge of becoming a reality.

Expert systems technology will also help America solve its productivity problems. It will help businesses reorganize themselves into more efficient and effective organizations. It will do this by helping individuals solve problems more quickly and efficiently than they can today.

Along with expert systems, executives will soon

have powerful computerized workstations that will enable them to handle more information in more complex ways than is now considered possible. Using these workstations, managers will be able to monitor more activities and personnel while simultaneously increasing the quality and quantity of decisions that they make.

Individuals responsible for operating complex equipment such as chemical plants and nuclear reactors will soon be assisted by expert systems that will monitor the equipment, anticipate problems, and make intelligent suggestions to the operators.

Training will also be revolutionized by the introduction of expert systems. Most companies are struggling to teach employees to use new procedures and to understand and explain new products. Skills that are now difficult to teach will become easy once we equip the employees with smart programs that will assist them in performing the tasks. Moreover, the concepts underlying expert systems will change the way we think about the tasks that people perform. In the near future, when analysts study jobs, they will be able to specify exactly what knowledge is and is not necessary to perform the jobs with a precision that is currently impossible.

In short, the whole business environment should become much more rational. More information will be gathered, synthesized, and put into useful form more rapidly than has ever before been possible. Individuals responsible for managing businesses will be in a position to utilize this information. In effect, middle managers of large corporations will be equipped with a staff of 15 to 20 automated experts who will always be available to answer questions and give advice about problems those managers face.

The introduction of expert systems will also prove exciting to experts and professionals. These systems will help experts define problems and determine what knowledge is available to solve problems in ways they have never considered before. As expert systems are built, experts will be freed to focus on the more difficult aspects of their specialty. This, in turn, will result in solutions to new problems, and the range of problems that experts can solve will widen.

## **ARTIFICIAL INTELLIGENCE**

As World War II ended, separate groups of British and American scientists were working to develop what we would now call a computer. Each group wanted to create an electronic machine that could be guided by a stored program of directions and made to carry out complex numerical computations. The principal British scientist, Alan Turing, argued that such a general-purpose machine, once developed, would have many different uses. Reflecting his knowledge of the accomplishments of formal logic in the years before the war, Turing argued that the fundamental instructions given to such a machine ought to be based on logical operators, such as "and," "or," and "not." One could then use such very general operators to assemble the more specialized numerical operators needed for arithmetic calculations. Moreover, programs based on logical operators would be capable of manipulating any type of symbolic material that one might want to work with, including statements in ordinary language.

The American scientists, being more practical, knew that the machine was going to be expensive to build. In addition, they assumed that they would not build very many of them. And since they were confident that they were building a machine that would do only arithmetic calculations, they decided against using logical operators and chose instead to use numerical operators, such as "+," "-", and ">." This decision, which the British subsequently followed as well, resulted in large computers that are essentially very fast calculating machines. In spite of the great proliferation of computers since 1946, this decision always seemed like a reasonable one to most people involved with computers. Until very recently.

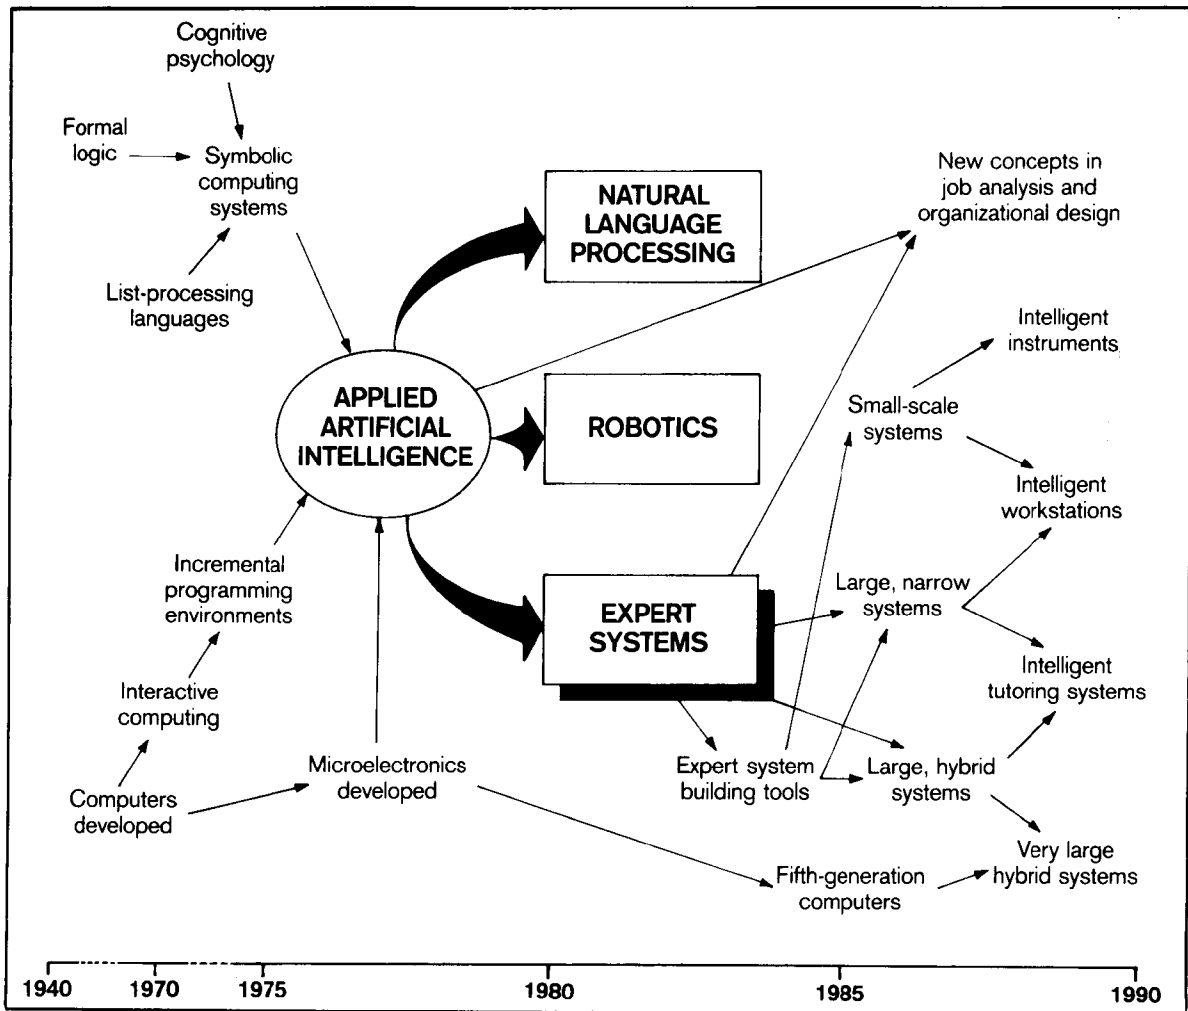
In spite of the fact that computers were built as numerical processors, a small group of computer scientists continued to explore the ability of computers to manipulate non-numerical symbols. Simultaneously, psychologists concerned with human problem solving sought to develop computer programs that would simulate human behavior. Over the years, individuals concerned with both symbolic processing and human problem solving have formed that inter-

disciplinary subfield of computer science called artificial intelligence (AI). AI researchers are concerned with developing computer systems that produce results that we would normally associate with human intelligence.

About 15 years ago a number of corporations thought that some of the research coming out of the AI laboratories would prove useful in business. Several companies set up AI groups to develop practical applications. By and large, these efforts failed because

AI programs were too costly to develop, were too slow, and didn't produce sufficiently practical results. AI programs were simply too complex to run on the computers that existed at the time. However, AI researchers continued to work in the universities and made steady theoretical progress. Meanwhile, the development of microelectronics technology resulted in a new generation of faster, more powerful, and relatively inexpensive computers. Today AI has once again emerged from the laboratories. Existing compu-

**Figure 1.1** The evolution of expert systems.



ter hardware, combined with significant theoretical advances in AI, has resulted in a technology whose time has come.

AI can be subdivided into three relatively independent research areas. (See Figure 1.1.) One group of AI researchers is concerned primarily with develop-

ing computer programs that can read, speak, or understand language as people use it in everyday conversation. This type of programming is commonly referred to as natural language processing. Another group of AI scientists is concerned with developing smart robots. They are especially concerned with

**Table 1.1** An Overview of the Key Events in the History of Artificial Intelligence

Period	Key Events
Pre-World War II roots	Formal logic Cognitive psychology
The postwar years, 1945–1954  <i>Pre-AI</i>	Computers developed H. Simon, <i>Administrative Behavior</i> N. Wiener, <i>Cybernetics</i> A. M. Turing, "Computing Machinery and Intelligence" Macy Conferences on Cybernetics
The formative years, 1955–1960  <i>The initiation of AI research</i>	Growing availability of computers Information Processing Language I (IPL-I) The Dartmouth Summer Seminar on AI, 1956 General Problem Solver (GPS) Information processing psychology
The years of development and redirection, 1961–1970  <i>The search for general problem solvers</i>	A. Newell and H. Simon, <i>Human Problem Solving</i> LISP Heuristics Satisficing Robotics Chess programs DENDRAL (Stanford)
The years of specialization and success, 1971–1980  <i>The discovery of knowledge- based systems</i>	MYCIN (Stanford) HEARSAY II (Carnegie-Mellon) MACSYMA (MIT) Knowledge engineering EMYCIN (Stanford) GUIDON (Stanford) PROLOG Herbert Simon—Nobel Prize
The rush to applications, 1981–  <i>International competition and commercial ventures</i>	PROSPECTOR (SRI) Japan's Fifth-Generation Project E. Feigenbaum and P. McCorduck, <i>The Fifth Generation</i> U.S.'s Microelectronics & Computer Technology Corp. (MCC) INTELLECT (A.I.C.) Various corporate and entrepreneurial AI companies

how to develop visual and tactile programs that will allow robots to observe the ongoing changes that take place as they move around in an environment. A third branch of AI research is concerned with developing programs that use symbolic knowledge to simulate the behavior of human experts. Table 1.1 gives a brief overview of the key developments in AI.

## **KNOWLEDGE-BASED EXPERT SYSTEMS**

Professor Edward Feigenbaum of Stanford University, one of the leading researchers in expert systems, has defined an expert system as:

. . . an intelligent computer program that uses knowledge and inference procedures to solve problems that are difficult enough to require significant human expertise for their solution. Knowledge necessary to perform at such a level, plus the inference procedures used, can be thought of as a model of the expertise of the best practitioners of the field.

The knowledge of an expert system consists of facts and heuristics. The “facts” constitute a body of information that is widely shared, publicly available, and generally agreed upon by experts in a field. The “heuristics” are mostly private, little-discussed rules of good judgment (rules of plausible reasoning, rules of good guessing) that characterize expert-level decision making in the field. The performance level of an expert system is primarily a function of the size and the quality of a knowledge base it possesses.

Feigenbaum calls those who build knowledge-based expert systems “knowledge engineers” and refers to their technology as “knowledge engineering.” Early systems were usually called “expert systems,” but most knowledge engineers now refer to their systems as “knowledge systems.”

The first systems were built by interviewing a recognized human expert and attempting to capture that expert’s knowledge, hence the term “expert systems.” Recently, however, several systems have been built that contain knowledge of a difficult decision-making situation that is quite useful, but hardly the equivalent of a human expert. To avoid suggesting that all systems built by means of knowledge engineering techniques capture the knowledge of a

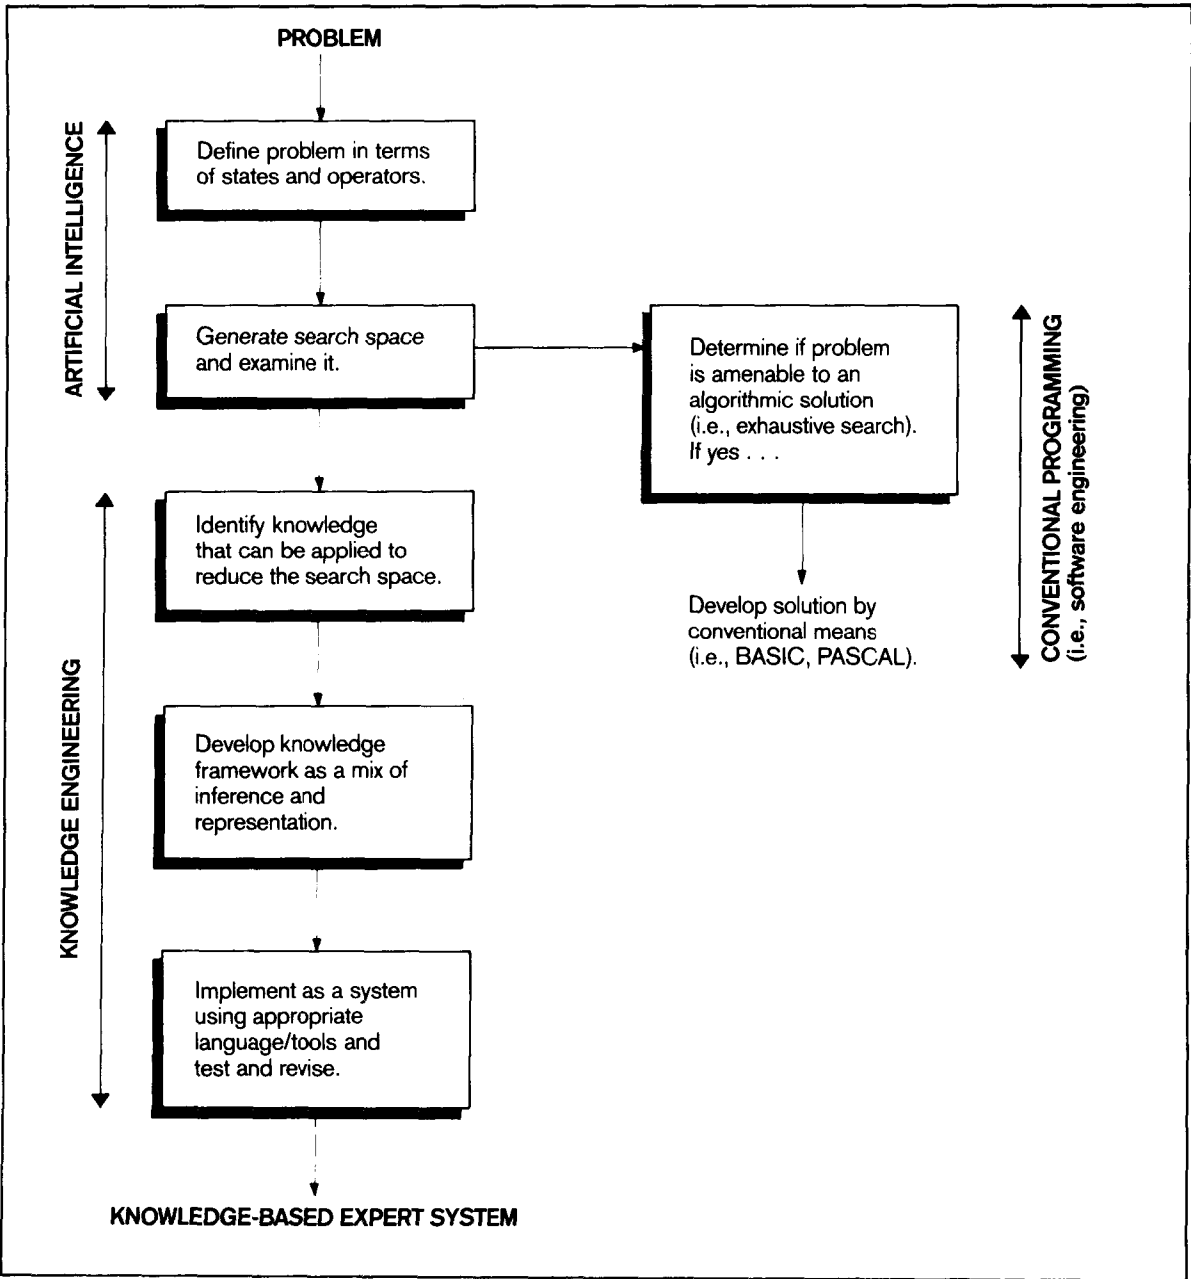
human expert, “knowledge systems” is rapidly becoming the preferred name. Throughout this book we use “expert systems” and “knowledge systems” as synonyms, but we generally refer to large systems as “expert systems” and smaller systems as “knowledge systems.”

AI is a research field concerned primarily with studying problem solving in the abstract. Knowledge engineers, on the other hand, focus on replicating the behavior of a specific expert when he or she is engaged in solving a narrowly defined problem. (See Figure 1.2.) Many regard this shift from the study of generic problem-solving techniques to a focus on building systems that contain large amounts of specific knowledge about a particular problem as the major conceptual breakthrough in AI in the last 15 years.

In subsequent chapters we shall develop each of these concepts in considerable detail. For the moment, it is sufficient to note that AI’s contribution to knowledge engineering lies in its insights into how to analyze problems and develop general search strategies to use in solving problems. Knowledge engineers are concerned with identifying the specific knowledge that an expert uses in solving a problem. Initially, the knowledge engineer studies a human expert and determines what facts and rules-of-thumb the expert employs. Then the knowledge engineer determines the inference strategy that the expert uses in an actual problem-solving situation. Finally, the knowledge engineer develops a system that uses similar knowledge and inference strategies to simulate the expert’s behavior.

If a program is to function like a human expert, it must be able to do the things that human experts commonly do. For example, experts consult with others to help solve problems. Thus, most knowledge systems ask questions, explain their reasoning if asked, and justify their conclusions. Moreover, they typically do this in language that the user can easily understand. They allow the user to skip questions, and most can function even when the user provides incomplete or uncertain data. In other words, knowledge systems interact with a user in pretty much the same way that a human consultant does.

Figure 1.2 The different concerns of AI and knowledge engineering.





Expert systems are knowledge-intensive computer programs. They contain lots of knowledge about their specialty. They use rules-of-thumb, or *heuristics*, to focus on the key aspects of particular problems and to manipulate symbolic descriptions in order to reason about the knowledge they are given. They often consider a number of competing hypotheses simultaneously, and they frequently make tentative recommendations or assign weights to alternatives. The best expert systems can solve difficult problems, within a very narrow domain, as well as or better than human experts can.

All this is not to suggest that most of today's expert systems are as good as human experts. The technology is new and just beginning to be applied to tough commercial problems. Today's knowledge systems are confined to well-circumscribed tasks. They are not able to reason broadly over a field of expertise. They cannot reason from axioms or general theories. They do not learn and, thus, they are limited to using the specific facts and heuristics that they were "taught" by a human expert. They lack common sense, they cannot reason by analogy, and their performance deteriorates rapidly when problems extend beyond the narrow task they were designed to perform.

It's reasonable, for example, to consider developing a small expert system to assist a manager in the analysis of a specific type of cash flow problem or to help managers decide how to respond to a particular type of employee error. In most companies, these problems are probably sufficiently well-defined and adequately constrained so as to result in a very useful small system. One would not, however, want to try to develop a system to help a manager analyze a potential legal problem. Most legal problems tend to be poorly defined and broad-ranging. Moreover, they typically involve large amounts of common sense and reasoning by analogy—things that existing expert systems are unable to do.

On the other hand, knowledge systems do not display biased judgments, nor do they jump to conclusions and then seek to maintain those conclusions in the face of disconfirming evidence. They do not have "bad days"; they always attend to details, and they always systematically consider all of the possible

alternatives. The best of them, equipped with thousands of heuristic rules, are able to perform their specialized tasks better than a human specialist. And, as we shall explain in subsequent chapters, new concepts and techniques that are now being introduced are certain to result in knowledge systems that are more flexible and powerful than those to which we have just alluded.

## **CONVENTIONAL PROGRAMMING VERSUS KNOWLEDGE ENGINEERING**

Conventional programming techniques have been used to create the large data processing systems we commonly associate with computers. These systems are capable of collecting and processing large volumes of data. They process this data by means of complex algorithms. Algorithms are simply step-by-step procedures that guarantee that the right conclusion will be reached when the correct data have been entered. For example, each evening all of the data regarding all of the changes in every account at your local bank are fed into a computer. Then a very complex algorithm works through the data, making additions and subtractions, calculating the proper fees, and finally arriving at the bank's overall balance for the day. The numbers differ each day, but they are always processed in the same way, and they always result in a predetermined conclusion—the bank's overall balance.

Today's large computer systems automate complex and time-consuming clerical tasks that previously would have required hundreds or thousands of clerks. Indeed, prior to 1950, the word "computer" referred to a human who made calculations according to a set procedure. The U.S. Census Bureau, for example, was staffed by a small army of computers who slowly and systematically processed the data that census takers had gathered. The first internally programmed electronic machines designed to do similar tasks were called "electronic computers." Once electronic computers managed to supplant all their human competitors, of course, they acquired the shorter name.