

George F Luger

# ARTIFICIAL INTELLIGENCE

Structures and Strategies for  
Complex Problem Solving

Fourth Edition



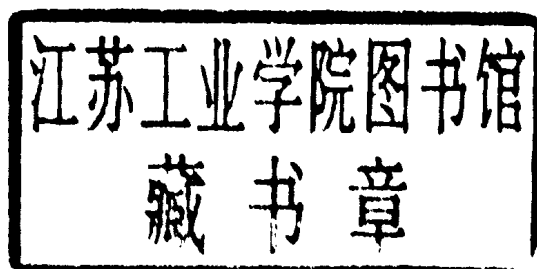
---

# Artificial Intelligence

## Structures and Strategies for Complex Problem Solving

*Fourth Edition*

George F. Luger



---

*An imprint of* **Pearson Education**

Harlow, England · London · New York · Reading, Massachusetts · San Francisco · Toronto · Don Mills, Ontario · Sydney  
Tokyo · Singapore · Hong Kong · Seoul · Taipei · Cape Town · Madrid · Mexico City · Amsterdam · Munich · Paris · Milan

**Pearson Education Limited**  
Edinburgh Gate  
Harlow  
Essex CM20 2JE  
England

and Associated Companies throughout the world

*Visit us on the World Wide Web at:*  
[www.pearsoneduc.com](http://www.pearsoneduc.com)

First published 1989  
**Fourth edition published 2002**

© Pearson Education Limited 2002

The right of George Luger to be identified as author of this work has been asserted by the author in accordance with the Copyright, Designs and Patents Act 1988.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without either the prior written permission of the publisher or a licence permitting restricted copying in the United Kingdom issued by the Copyright Licensing Agency Ltd, 90 Tottenham Court Road, London W1P 0LP.

The programs in this book have been included for their instructional value. They have been tested with care but are not guaranteed for any particular purpose. The publisher does not offer any warranties or representations nor does it accept any liabilities with respect to the programs.

ISBN 0-201-64866-0

**British Library Cataloguing-in-Publication Data**  
A catalogue record for this book is available from the British Library

**Library of Congress Cataloging-in-Publication Data**  
Luger, George F.

Artificial intelligence : structures and strategies for complex problem solving.—4th ed.  
/ George F. Luger.  
p. cm.

Includes bibliographical references and index.

ISBN 0-201-64866-0

1. Artificial intelligence. 2. Knowledge representation (Information theory) 3. Problem solving. 4. PROLOG (Computer program language) 5. LISP (Computer program language) I. Title.

Q335 .L84 2001  
006.3—dc21

2001032715

10 9 8 7 6 5 4  
06 05 04 03 02

Printed and bound in the United States of America

# Artificial Intelligence



We work with leading authors to develop the strongest educational materials in computer science, bringing cutting-edge thinking and best learning practice to a global market.

Under a range of well-known imprints, including Addison-Wesley, we craft high quality print and electronic publications which help readers to understand and apply their content, whether studying or at work.

To find out more about the complete range of our publishing, please visit us on the World Wide Web at: [www.pearsoneduc.com](http://www.pearsoneduc.com)

For my wife, Kathleen, and our children Sarah, David, and Peter.

*Si quid est in me ingenii, iudices . . .*

Cicero, Pro Archia Poeta

GFL

---

# PREFACE

---

*What we have to learn to do  
we learn by doing. . .*

—ARISTOTLE, *Ethics*

---

## Welcome to the Fourth Edition!

---

I was very pleased to be asked to produce a fourth edition of our artificial intelligence book. It is a compliment to the earlier editions, started more than a decade ago, that our approach to AI has been widely accepted. It is also exciting that, as new developments in the field emerge, we are able to present much of it in each new edition. We thank our readers, colleagues, and students for keeping our topics relevant and presentation up to date.

Many sections of the earlier editions have endured remarkably well, including the presentation of logic, search algorithms, knowledge representation, production systems, machine learning, and the programming techniques developed in LISP and PROLOG. These remain central to the practice of artificial intelligence, and required a relatively small effort to bring them up to date. However, several sections, including those on natural language understanding, reinforcement learning, and reasoning under uncertainty, required, and received, extensive reworking. Other topics, such as emergent computation, case-based reasoning, and model-based problem solving, that were treated cursorily in the first editions, have grown sufficiently in importance to merit a more complete discussion. These changes are evidence of the continued vitality of the field of artificial intelligence.

As the scope of the project grew, we were sustained by the support of our publisher, editors, friends, colleagues, and, most of all, by our readers, who have given our work such a long and productive life. We were also sustained by our own excitement at the opportunity afforded: Scientists are rarely encouraged to look up from their own, narrow research interests and chart the larger trajectories of their chosen field. Our publisher and readers have asked us to do just that. We are grateful to them for this opportunity.

Although artificial intelligence, like most engineering disciplines, must justify itself to the world of commerce by providing solutions to practical problems, we entered the

field of AI for the same reasons as many of our colleagues and students: we want to understand and explore the mechanisms of mind that enable intelligent thought and action. We reject the rather provincial notion that intelligence is an exclusive ability of humans, and believe that we can effectively investigate the space of possible intelligences by designing and evaluating intelligent artifacts. Although the course of our careers has given us no cause to change these commitments, we have arrived at a greater appreciation for the scope, complexity, and audacity of this undertaking. In the preface to our earlier editions, we outlined three assertions that we believed distinguished our approach to teaching artificial intelligence. It is reasonable, in writing a preface to this fourth edition, to return to these themes and see how they have endured as our field has grown.

The first of these goals was to “unify the diverse branches of AI through a detailed discussion of its theoretical foundations.” At the time we adopted that goal, it seemed that the main problem was reconciling researchers who emphasized the careful statement and analysis of formal theories of intelligence (the *neats*) with those who believed that intelligence itself was some sort of grand hack that could be best approached in an application-driven, *ad hoc* manner (the *scruffies*). That simple dichotomy has proven far too simple. In contemporary AI, debates between neats and scruffies have given way to dozens of other debates between proponents of physical symbol systems and students of neural networks, between logicians and designers of artificial life forms that evolve in a most illogical manner, between architects of expert systems and case-based reasoners, and finally, between those who believe artificial intelligence has already been achieved and those who believe it will never happen. Our original image of AI as frontier science where outlaws, prospectors, wild-eyed prairie prophets and other dreamers were being slowly tamed by the disciplines of formalism and empiricism has given way to a different metaphor: that of a large, chaotic but mostly peaceful city, where orderly bourgeois neighborhoods draw their vitality from diverse, chaotic, bohemian districts. Over the years we have devoted to the different editions of this book, a compelling picture of the architecture of intelligence has started to emerge from this city’s structure, art, and industry.

Intelligence is too complex to be described by any single theory; instead, researchers are constructing a hierarchy of theories that characterize it at multiple levels of abstraction. At the lowest levels of this hierarchy, neural networks, genetic algorithms and other forms of emergent computation have enabled us to understand the processes of adaptation, perception, embodiment, and interaction with the physical world that must underlie any form of intelligent activity. Through some still partially understood resolution, this chaotic population of blind and primitive actors gives rise to the cooler patterns of logical inference. Working at this higher level, logicians have built on Aristotle’s gift, tracing the outlines of deduction, abduction, induction, truth-maintenance, and countless other modes and manners of reason. Even higher levels of abstraction, designers of expert systems, intelligent agents, and natural language understanding programs have come to recognize the role of social processes in creating, transmitting, and sustaining knowledge. In this fourth edition, we have touched on all levels of this developing hierarchy.

The second commitment we made in the earlier editions was to the central position of “advanced representational formalisms and search techniques” in AI methodology. This is, perhaps, the most controversial aspect of our previous editions and of much early work in AI, with many researchers in emergent computation questioning whether symbolic

reasoning and referential semantics have any role at all in thought. Although the idea of representation as giving names to things has been challenged by the implicit representation provided by the emerging patterns of a neural network or an artificial life, we believe that an understanding of representation and search remains essential to any serious practitioner of artificial intelligence. More importantly, we feel that the skills acquired through the study of representation and search are invaluable tools for analyzing such aspects of non-symbolic AI as the expressive power of a neural network or the progression of candidate problem solutions through the fitness landscape of a genetic algorithm. Comparisons, contrasts, and a critique of the various approaches of modern AI are offered in Chapter 16.

The third commitment we made at the beginning of this book's life cycle, to "place artificial intelligence within the context of empirical science," has remained unchanged. To quote from the preface to the third edition, we continue to believe that AI is not

. . . some strange aberration from the scientific tradition, but . . . part of a general quest for knowledge about, and the understanding of intelligence itself. Furthermore, our AI programming tools, along with the exploratory programming methodology . . . are ideal for exploring an environment. Our tools give us a medium for both understanding and questions. We come to appreciate and know phenomena constructively, that is, by progressive approximation.

Thus we see each design and program as an experiment with nature: we propose a representation, we generate a search algorithm, and then we question the adequacy of our characterization to account for part of the phenomenon of intelligence. And the natural world gives a response to our query. Our experiment can be deconstructed, revised, extended, and run again. Our model can be refined, our understanding extended.

## New with This Edition

---

I, George Luger, am the sole author of the fourth edition. Although Bill Stubblefield has moved on to new areas and challenges in computing, his mark will remain on the present and any further editions of this book. In fact this book has always been the product of my efforts as Professor of Computer Science at the University of New Mexico together with those of my professional colleagues, graduate students, and friends: the members of the UNM artificial intelligence community, as well as of the many readers that have e-mailed me comments, corrections, and suggestions. The book will continue this way, and to reflect this community effort, I will continue using the prepositions *we* and *us* when presenting material. Individual debts in the preparation for this fourth edition are listed in the *acknowledgement* section of this preface.

We revised many sections of this book to recognize the growing importance of agent-based problem solving as an approach to AI technology. In discussions of the foundations of AI we recognize intelligence as physically embodied and situated in a natural and social world context. Apropos of this, we present in Chapter 6 the evolution of AI representational schemes from associative and early logic-based, through weak and strong method approaches, including connectionist and evolutionary/emergent models, to situated and social approaches to AI. Chapter 16 contains a critique of each of these paradigms.

In creating this fourth edition, we considered all topics presented earlier and brought them into a modern perspective. In particular, we added a reinforcement learning section to Chapter 9. Algorithms for reinforcement learning, taking cues from an environment to establish a policy for state change, including *temporal difference* and *Q-learning*, are presented.

Besides our previous analysis of data-driven and goal-driven rule-based systems, Chapter 7 now contains case-based and model-based reasoning, including examples from the NASA space program. The chapter includes a section on the strengths and weaknesses of each of these approaches to knowledge-intensive problem solving.

Chapter 8 describes reasoning with uncertain or incomplete information. A number of important approaches to this problem are presented, including Bayesian reasoning, belief networks, the Dempster–Shafer model, and the Stanford certainty factor algebra. Techniques for truth maintenance in nonmonotonic situations are also presented, as well as reasoning with minimal models and logic-based abduction. We conclude the chapter with an in-depth presentation of Bayesian Belief Networks and the clique-tree algorithm for propagating confidences through a belief network in the context of new evidence.

Chapter 13 presents issues in natural language understanding, including a section on stochastic models for language comprehension. The presentation includes Markov models, CART trees, mutual information clustering, and statistic-based parsing. The chapter closes with several examples, including the applications of text mining and text summarization techniques to the WWW.

Finally, in a revised Chapter 16, we return to the deeper questions of the nature of intelligence and the possibility of intelligent machines. We comment on the AI endeavour from the perspectives of philosophy, psychology, and neuro-physiology.

## The Contents

---

Chapter 1 introduces artificial intelligence, beginning with a brief history of attempts to understand mind and intelligence in philosophy, psychology, and other areas of research. In an important sense, AI is an old science, tracing its roots back at least to Aristotle. An appreciation of this background is essential for an understanding of the issues addressed in modern research. We also present an overview of some of the important application areas in AI. Our goal in Chapter 1 is to provide both background and a motivation for the theory and applications that follow.

Chapters 2, 3, 4, and 5 (Part II) introduce the research tools for AI problem solving. These include the predicate calculus language to describe the essential features of a problem domain (Chapter 2), search to reason about these descriptions (Chapter 3) and the algorithms and data structures used to implement search. In Chapters 4 and 5, we discuss the essential role of heuristics in focusing and constraining search-based problem solving. We also present a number of architectures, including the blackboard and production system, for building these search algorithms.

Chapters 6, 7, and 8 make up Part III of the book: representations for AI and knowledge-intensive problem solving. In Chapter 6 we present the evolving story of AI

representational schemes. We begin with a discussion of semantic networks and extend this model to include conceptual dependency theory, frames, and scripts. We then present an in-depth examination of a particular formalism, conceptual graphs, emphasizing the epistemological issues involved in representing knowledge and showing how these issues are addressed in a modern representation language. In Chapter 13, we show how conceptual graphs can be used to implement a natural language database front end. We conclude Chapter 6 with more modern approaches to representation, including Copycat and agent-oriented architectures.

Chapter 7 presents the rule-based expert system along with case-based and model-based reasoning systems, including examples from the NASA space program. These approaches to problem solving are presented as a natural evolution of the material in the first five chapters: using a production system of predicate calculus expressions to orchestrate a graph search. We end with an analysis of the strengths and weaknesses of each of these approaches to knowledge-intensive problem solving.

Chapter 8 presents models for reasoning with uncertainty as well as the use of unreliable information. We discuss Bayesian models, belief networks, Dempster–Shafer, causal models, and the Stanford certainty algebra for reasoning in uncertain situations. Chapter 8 also contains algorithms for truth maintenance, reasoning with minimum models, logic-based abduction, and the clique-tree algorithm for Bayesian belief networks.

Part IV, Chapters 9 through 11, offers an extensive presentation of issues in machine learning. In Chapter 9 we offer a detailed look at algorithms for symbol-based learning, a fruitful area of research spawning a number of different problems and solution approaches. These learning algorithms vary in their goals, the training data considered, their learning strategies, and the knowledge representations they employ. Symbol-based learning includes induction, concept learning, version-space search, and ID3. The role of inductive bias is considered, generalizations from patterns of data, as well as the effective use of knowledge to learn from a single example in explanation-based learning. Category learning, or conceptual clustering, is presented with unsupervised learning. Reinforcement learning, or the ability to integrate feedback from the environment into a policy for making new decisions concludes the chapter.

In Chapter 10 we present neural networks, often referred to as sub-symbolic or connectionist models of learning. In a neural net, information is implicit in the organization and weights on a set of connected processors, and learning involves a re-arrangement and modification of the overall weighting of nodes and structure of the system. We present a number of connectionist architectures, including perceptron learning, backpropagation, and counterpropagation. We demonstrate Kohonen, Grossberg, and Hebbian network models. We present associative learning and attractor models, including Hopfield networks.

Genetic algorithms and evolutionary approaches to learning are introduced in Chapter 11. On this viewpoint, learning is cast as an emerging and adaptive process. After several examples of problem solutions based on genetic algorithms, we introduce the application of genetic techniques to more general problem solvers. These include classifier systems and genetic programming. We then describe society-based learning with examples from artificial life, called *a-life*, research. We conclude the chapter with an example of emergent computation from research at the Santa Fe Institute. We compare and contrast the three

approaches we present to machine learning (symbol-based, connectionist, social and emergent) in Chapter 16.

Part V, Chapters 12 and 13, continues our presentation of important AI application areas. Theorem proving, often referred to as automated reasoning, is one of the oldest areas of AI research. In Chapter 12, we discuss the first programs in this area, including the Logic Theorist and the General Problem Solver. The primary focus of the chapter is binary resolution proof procedures, especially resolution refutations. More advanced inferencing with hyper-resolution and paramodulation is also presented. Finally, we describe the PROLOG interpreter as a Horn clause and resolution-based inferencing system, and see PROLOG computing to as an instance of the logic programming paradigm.

Chapter 13 presents natural language understanding. Our traditional approach to language understanding, exemplified by many of the semantic structures presented in Chapter 6, is complemented with the stochastic approach. These include Markov models, CART trees, mutual information clustering, and statistics-based parsing. The chapter concludes with examples applying these natural language techniques to database query systems and also to a text summarization system for use on the WWW.

Part VI presents LISP and PROLOG. Chapter 14 covers PROLOG, and Chapter 15, LISP. We demonstrate these languages as tools for AI problem solving by building the search and representation techniques of the earlier chapters, including breadth-, depth-, and best-first search algorithms. We implement these search techniques in a problem-independent fashion so that they may be extended to create shells for search in rule-based expert systems, to build semantic networks, natural language understanding systems, and learning applications.

Finally, Chapter 16 serves as an epilogue for the book. It addresses the issue of the possibility of a science of intelligent systems, and considers contemporary challenges to AI; it discusses AI's current limitations, and projects its exciting future.

## Using This Book

---

Artificial intelligence is a big field, and consequently, this is a big book. Although it would require more than a single semester to cover all of the material in the text, we have designed it so that a number of paths may be taken through the material. By selecting subsets of the material, we have used this text for single semester and full year (two semester) courses.

We assume that most students will have had introductory courses in discrete mathematics, including predicate calculus and introductory graph theory. If this is not true, the instructor should spend more time on these concepts in the sections at the beginning of the text (2.1, 3.1). We also assume that students have had courses in data structures including trees, graphs, and recursion-based search, using stacks, queues, and priority queues. If they have not, then spend more time on the beginning sections of Chapters 3, 4, and 5.

In a one semester course, we go quickly through the first two parts of the book. With this preparation, students are able to appreciate the material in Part III. We then consider the PROLOG and LISP in Part VI and require students to build many of the representation

and search techniques of the first sections. Alternatively, one of the languages, PROLOG, for example, can be introduced early in the course and be used to test out the data structures and search techniques as they are encountered. We feel the meta-interpreters presented in the language chapters are very helpful for building rule-based and other knowledge-intensive problem solvers. PROLOG is an excellent tool for building natural language understanding systems.

In a two-semester course, we are able to cover the application areas of Parts IV and V, especially the machine learning chapters, in appropriate detail. We also expect a much more detailed programming project from students. We think that it is very important in the second semester for students to revisit many of the primary sources in the AI literature. It is crucial for students to see both where we are, as well as how we got here, and to have an appreciation of the future promises of artificial intelligence. We use a collected set of readings for this purpose, *Computation and Intelligence* (Luger 1995).

The algorithms of our book are described using a Pascal-like pseudo-code. This notation uses the control structures of Pascal along with English descriptions of the tests and operations. We have added two useful constructs to the Pascal control structures. The first is a modified **case** statement that, rather than comparing the value of a variable with constant case labels, as in standard Pascal, lets each item be labeled with an arbitrary boolean test. The **case** evaluates these tests in order until one of them is true and then performs the associated action; all other actions are ignored. Those familiar with LISP will note that this has the same semantics as the LISP **cond** statement.

The other addition to the language is a **return** statement which takes one argument and can appear anywhere within a procedure or function. When the **return** is encountered, it causes the program to immediately exit the function, returning its argument as a result. Other than these modifications we used Pascal structure, with a reliance on the English descriptions, to make the algorithms clear.

## Supplemental Material Available via the Internet

---

The code, in PROLOG and LISP, presented in the book is available to readers via the internet. There is also a full instructor's guide available for faculty using this book. These files are available through the web address: [www.booksites.net/luger](http://www.booksites.net/luger), or indirectly through my personal homepage: [www.cs.unm.edu/~luger/](http://www.cs.unm.edu/~luger/).

WWW sites for supporting course materials and software include those at Addison-Wesley and Pearson Education: [www.aw.com/cs/](http://www.aw.com/cs/) and [www.pearsoneduc.com/computing](http://www.pearsoneduc.com/computing). My e-mail address is [luger@cs.unm.edu](mailto:luger@cs.unm.edu), and I enjoy hearing from readers.

## Acknowledgements

---

First, we would like to thank Bill Stubblefield, the co-author for the first three editions, for more than a decade of contributions. We also thank the many reviewers that have helped develop these four editions. These include Dennis Bahler, Skona Brittain, Philip Chan,

Peter Collingwood, John Donald, Sarah Douglas, Christophe Giraud-Carrier, Andrew Kosoresow, Chris Malcolm, Ray Mooney, Bruce Porter, Jude Shavlik, Carl Stern, Marco Valtorta, and Bob Veroff. We also appreciate the numerous suggestions and comments sent directly to us by e-mail from people using the book.

For the fourth edition, we thank our AI PhD students at UNM for helping reorganize the topic presentation. The rework of Part III with the evolution of AI representational schemes first, the creating of a separate section for machine learning, and moving languages to the back were at their suggestion. We thank Dan Pless for his major role in developing material in Chapter 8 on abductive inference, Carl Stern for his help in developing Chapter 10 on connectionist learning, and Jared Saia for helping with the stochastic models of Chapter 13. Our “outside” reviewers for the fourth edition included Leon van der Torre and Mehdi Dastani from Vrije University in the Netherlands, Leonardo Bottaci, Hull University, and Julian Richardson, Heriot-Watt University, UK. Our American reviewers included Marek Perkowski of Portland State and John Sheppard of Johns Hopkins Universities. Barak Pearmutter, UNM, reviewed the chapters on machine learning. Finally, Joseph Lewis, Chris Malcolm, Brendan McGonnigle, and Akasha Tang, critiqued Chapter 16.

We thank Academic Press for permission to reprint much of the material of Chapter 10; this appeared in the book *Cognitive Science: The Science of Intelligent Systems* (Luger 1994). Finally, we thank more than a decade of students who have used this text and software at UNM for their help in expanding our horizons, as well as in removing typos and bugs from the book.

We thank our many friends at Addison-Wesley for their support and encouragement in completing our writing task, especially Alan Apt in helping us with the first edition, Lisa Moller and Mary Tudor for their help on the second, Victoria Henderson, Louise Wilson, and Karen Mosman for their assistance on the third, and Keith Mansfield, Karen Sutherland, and Anita Atkinson for support on this fourth edition. We thank Linda Cicarella of the University of New Mexico for her help in preparing figures for publication.

We thank Thomas Barrow, internationally recognized artist and University of New Mexico Professor of Art, who created the seven photograms for this book.

In a number of places, we have used figures or quotes from the work of other authors. We thank the authors and publishers for their permission to use this material. These contributions are listed in the Publisher’s Acknowledgements.

Artificial intelligence is an exciting and rewarding discipline; may you enjoy your study as you come to appreciate its power and challenges.

George Luger  
1 July 2001  
Albuquerque

---

# PUBLISHER'S ACKNOWLEDGEMENTS

---

We are grateful to the following for permission to reproduce copyright material:

Figures 3.22, 4.17, 4.18, 4.19 and 5.6 from *Problem Solving Methods in Artificial Intelligence*, copyright © 1971 by McGraw-Hill, reproduced with permission of The McGraw-Hill Companies (Nilsson, N. J. 1971); Figures 4.21 and 4.22 adapted from *Principles in Artificial Intelligence*, copyright © 1980 by Morgan Kaufmann, reproduced by permission of the publisher (Nilsson, N. J. 1980); Figure 6.1 from *Expert Systems: Artificial Intelligence in Business*, copyright © 1985 by John Wiley and Sons, Inc., used by permission from John Wiley & Sons, Inc. (Harmon, P. and King, D. 1985); Figures 6.6, 6.9 and 6.10 from Inference and the computer understanding of natural language in *Artificial Intelligence* 5(4): 373–412, copyright 1974, reprinted with permission from Elsevier Science (Schank, R. C. and Reiger, C. J. 1974); Figure 6.8 adapted from *Computer Models of Thought and Language* by Roger C. Schank and Kenneth Mark Colby © 1973 by W. H. Freeman and Company, used with permission (Schank, R. C. and Colby, K. M.); Figure 6.13 from *Psychology of Computer Vision* by P. Winston, copyright © 1975 by The McGraw-Hill Companies, Inc. (Minsky, M. 1975a); Figures 6.27 and 6.28 adapted from *Analogy-Making as Perception: A Computer Model*, The MIT Press (Mitchell, M. 1993); Figures 7.15 and 7.16 from A model-based approach to reactive self-configuring systems in *Proceedings of the AAAI-96*, pp. 971–978, pub The MIT Press, copyright © 1996 American Association for Artificial Intelligence, reprinted with permission from the American Association for Artificial Intelligence (Williams, B. C. and Nayak, P. P. 1996); Figures 7.23, 7.24, 7.25 and 7.26 from Immobile robots: AI in the new millennium in *AI Magazine* 17(3), pp. 17–34, copyright © 1996 American Association for Artificial Intelligence, reprinted with permission from the American Association for Artificial Intelligence (Williams, B. C. and Nayak, P. P. 1996a); Figure 8.2 from An improved algorithm for non-monotonic dependency net update in *Technical Report LITH-MAT-R-82-23*, reprinted by permission of the author (Goodwin, J. 1982); Figure 8.16 from Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference, copyright © 1988 by Morgan Kaufmann, reproduced by permission of the publisher (Pearl, J. 1988); Figure 9.12 adapted from Learning by experimentation: Acquiring and refining problem solving heuristics in *Machine Learning: An Artificial Intelligence Approach, Volume 1* edited by R. S. Michalski, J. G. Carbonell and T. M. Mitchell, copyright © 1983 by Morgan Kaufmann, a division of

Academic Press, reproduced by permission of the publisher. This material may not be reproduced in any form or by any means without the prior written permission of the publisher (Mitchell, T. M., Ugtoff, P. E. and Banarji, R. 1983); Figure 9.21 adapted from Models of incremental concept formation in *Artificial Intelligence* 40 (1–3): 11–62, copyright 1989, reprinted with permission from Elsevier Science (Gennari, J. H., Langley, P. and Fisher, D. 1989); Figure 12.2 from GPS: a program that stimulates human thought, in *Computers and Thought* by E. A. Feigenbaum and J. Feldman (eds), copyright © 1963 by McGraw-Hill, reproduced with permission of The McGraw-Hill companies (Newell, A. and Simon, H. 1963b).

Academic Press for Chapter 10, adapted from *Cognitive Science: The Science of Intelligent Systems* (Luger, G. F. 1994); American Society for Public Administration for an abridged extract from Decision-making and administrative organization in *Public Administration Review* 4, Winter (Simon, H. A. 1944); Carlin Music Corp. for an excerpt from *Searchin'* by Jerry Leiber and Mike Stoller © 1957 Jerry Leiber Music and Mike Stoller Music – All Rights Reserved – world-wide authorization for lyric reproduction granted by kind permission of Carlin Music Corp., London NW1 8BD; Faber & Faber Ltd for excerpts from 'Little Gidding' and 'Burnt Norton' in *Four Quartets* by T. S. Eliot, copyright © 1943 by T. S. Eliot, renewed 1971 by Esme Valerie Eliot, and excerpts from 'Choruses from the Rock' and 'The Love Song of J. Alfred Prufrock' from *Collected Poems 1909–1962* by T. S. Eliot, copyright © 1963, 1964 by T. S. Eliot; Henry Holt and Company LLC for an excerpt from 'The Road Not Taken' by Robert Frost from *The Poetry of Robert Frost* edited by Edward Connery Lathem. Copyright 1916, © 1969 by Henry Holt and Co., © 1944 by Robert Frost. Reprinted by permission of Henry Holt and Company, LLC.

While every effort has been made to trace the owners of copyright material, in a few cases this has proved impossible and we take this opportunity to offer our apologies to any copyright holders whose rights we may have unwittingly infringed.

