

人工智能

复杂问题求解的结构和策略


新版

(英文版 · 第5版)

GEORGE F LUGER

ARTIFICIAL INTELLIGENCE

Structures and Strategies for Complex Problem Solving



Fifth
Edition



(美) George F. Luger 著

经典原版书库

人工智能

复杂问题求解的结构和策略

(英文版·第5版)

Artificial Intelligence
Structures and Strategies for Complex Problem Solving
(Fifth Edition)

(美) George F. Luger 著



机械工业出版社
China Machine Press

George F. Luger: Artificial Intelligence: Structures and Strategies for Complex Problem Solving, Fifth Edition (ISBN 0-321-26318-9).

Copyright © 2005, 1997 by Pearson Education Limited.

This edition of Artificial Intelligence: Structures and Strategies for Complex Problem Solving, Fifth Edition is published by arrangement with Pearson Education Limited. Licensed for sale in the mainland territory of the People's Republic of China only, excluding Hong Kong, Macau, and Taiwan.

本书英文影印版由英国Pearson Education (培生教育出版集团) 授权出版。未经出版者书面许可, 不得以任何方式复制或抄袭本书内容。

此影印版只限在中国大陆地区销售 (不包括香港、澳门、台湾地区)。

版权所有, 侵权必究。

本书法律顾问 北京市展达律师事务所

本书版权登记号: 图字: 01-2005-4391

图书在版编目 (CIP) 数据

人工智能: 复杂问题求解的结构和策略 (英文版·第5版) / (美) 卢格尔 (Luger, G. F.) 著. -北京: 机械工业出版社, 2005. 9

(经典原版书库)

书名原文: Artificial Intelligence: Structures and Strategies for Complex Problem Solving, Fifth Edition

ISBN 7-111-17268-X

I. 人… II. 卢… III. 人工智能-英文 IV. TP18

中国版本图书馆CIP数据核字 (2005) 第100798号

机械工业出版社 (北京市西城区百万庄大街22号 邮政编码 100037)

责任编辑: 迟振春

北京瑞德印刷有限公司印刷·新华书店北京发行所发行

2005年9月第1版第1次印刷

787mm × 1092mm 1/16 · 58印张

印数: 0 001-3 000册

定 价: 88.00元

凡购本书, 如有倒页、脱页、缺页, 由本社发行部调换
本社购书热线: (010) 68326294

出版者的话

文艺复兴以降，源远流长的科学精神和逐步形成的学术规范，使西方国家在自然科学的各个领域取得了垄断性的优势；也正是这样的传统，使美国在信息技术发展的六十多年间名家辈出、独领风骚。在商业化的进程中，美国的产业界与教育界越来越紧密地结合，计算机学科中的许多泰山北斗同时身处科研和教学的最前线，由此而产生的经典科学著作，不仅擘划了研究的范畴，还揭橥了学术的源变，既遵循学术规范，又自有学者个性，其价值并不会因年月的流逝而减退。

近年，在全球信息化大潮的推动下，我国的计算机产业发展迅猛，对专业人才的需求日益迫切。这对计算机教育界和出版界都既是机遇，也是挑战；而专业教材的建设在教育战略上显得举足轻重。在我国信息技术发展时间较短、从业人员较少的现状下，美国等发达国家在其计算机科学发展的几十年间积淀的经典教材仍有许多值得借鉴之处。因此，引进一批国外优秀计算机教材将对我国计算机教育事业的发展起积极的推动作用，也是与世界接轨、建设真正的世界一流大学的必由之路。

机械工业出版社华章图文信息有限公司较早意识到“出版要为教育服务”。自1998年开始，华章公司就将工作重点放在了遴选、移译国外优秀教材上。经过几年的不懈努力，我们与Prentice Hall, Addison-Wesley, McGraw-Hill, Morgan Kaufmann等世界著名出版公司建立了良好的合作关系，从它们现有的数百种教材中甄选出Tanenbaum, Stroustrup, Kernighan, Jim Gray等大师名家的一批经典作品，以“计算机科学丛书”为总称出版，供读者学习、研究及收藏。大理石纹理的封面，也正体现了这套丛书的品位和格调。

“计算机科学丛书”的出版工作得到了国内外学者的鼎力襄助，国内的专家不仅提供了中肯的选题指导，还不辞劳苦地担任了翻译和审校的工作；而原书的作者也相当关注其作品在中国的传播，有的还专程为其书的中译本作序。迄今，“计算机科学丛书”已经出版了近百个品种，这些书籍在读者中树立了良好的口碑，并被许多高校采用为正式教材和参考书籍，为进一步推广与发展打下了坚实的基础。

随着学科建设的初步完善和教材改革的逐渐深化，教育界对国外计算机教材的需求和应用都步入一个新的阶段。为此，华章公司将加大引进教材的力度，在“华章教育”的总规划之下出版三个系列的计算机教材：除“计算机科学丛书”之外，对影印版的教材，则单独开辟出“经典原版书库”；同时，引进全美通行的教学辅导书“Schaum's Outlines”系列组成“全美经典学习指导系列”。为了保证这三套丛书的权威性，同时也为了更好地为学校和老师服务，华章公司聘请了中国科学院、北京大学、清华大学、国防科技大学、复旦大学、上海交通大学、南京大学、浙江大学、中国科技大学、哈尔滨工业大学、西安交通大学、中国人民大学、北京航空航天大学、北京邮电大学、中山大学、解放军理工大学、郑州大学、湖北工学院、中国国

家信息安全测评认证中心等国内重点大学和科研机构在计算机的各个领域的著名学者组成“专家指导委员会”，为我们提供选题意见和出版监督。

这三套丛书是响应教育部提出的使用外版教材的号召，为国内高校的计算机及相关专业的教学度身订造的。其中许多教材均已为M. I. T., Stanford, U.C. Berkeley, C. M. U. 等世界名牌大学所采用。不仅涵盖了程序设计、数据结构、操作系统、计算机体系结构、数据库、编译原理、软件工程、图形学、通信与网络、离散数学等国内大学计算机专业普遍开设的核心课程，而且各具特色——有的出自语言设计者之手、有的历经三十年而不衰、有的已被全世界的几百所高校采用。在这些圆熟通博的名师大作的指引之下，读者必将在计算机科学的宫殿中由登堂而入室。

权威的作者、经典的教材、一流的译者、严格的审校、精细的编辑，这些因素使我们的图书有了质量的保证，但我们的目标是尽善尽美，而反馈的意见正是我们达到这一终极目标的重要帮助。教材的出版只是我们的后续服务的起点。华章公司欢迎老师和读者对我们的工作提出建议或给予指正，我们的联系方式如下：

电子邮件: hzjsj@hzbook.com

联系电话: (010) 68995264

联系地址: 北京市西城区百万庄南街1号

邮政编码: 100037

专家指导委员会

(按姓氏笔画顺序)

尤晋元
石教英
张立昂
邵维忠
周克定
郑国梁
高传善
裘宗燕

王 珊
吕 建
李伟琴
陆丽娜
周傲英
施伯乐
梅 宏
戴 葵

冯博琴
孙玉芳
李师贤
陆鑫达
孟小峰
钟玉琢
程 旭

史忠植
吴世忠
李建中
陈向群
岳丽华
唐世渭
程时端

史美林
吴时霖
杨冬青
周伯生
范 明
袁崇义
谢希仁

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

Si quid est in me ingenii, judices . . .

Cicero, Pro Archia Poeta

GFL

PREFACE

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

—ARISTOTLE, *Ethics*

Welcome to the Fifth Edition!

I was very pleased to be asked to produce a fifth edition of our artificial intelligence book. It is a compliment to the earlier editions, started almost twenty years ago, that our approach to AI has been so highly valued. It is also exciting that, as new development in the field emerges, 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. We created a new introductory chapter on stochastic methods for the fifth edition. We feel that the stochastic technology is having an increasingly larger impact on AI, especially in areas such as diagnostic reasoning, natural language analysis, and learning. To support these technologies we expand the presentation of Bayes' theorem, Bayesian networks and related graphical models, probabilistic finite state machines, dynamic programming with the Viterbi algorithm, and Markov modeling. 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. The changes in the fifth edition reflect the evolving research questions and are evidence of the continued vitality of the field of artificial intelligence.

As the scope of our AI project grew, we have been 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 remain excited at the writing opportunity we are afforded: scientists are rarely encouraged to look up from their own, narrow research interests and chart the larger trajectories of their chosen field. Our readers have asked us to do just that. We are grateful to them for this opportunity. We are also encouraged that our earlier editions have been used in AI communities worldwide and translated into a number of languages including German, Polish, Portuguese, Russian, and two dialects of Chinese!

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 fifth 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 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 that 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 address 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. At 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. Finally, as philosophers we are charged to critique the epistemological validity of the AI enterprise. For this task we discuss the rationalist project, the empiricists dilemma, and propose a constructivist rapprochement. In this fifth edition, we have touched on all these levels of the developing AI endeavour.

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. We also feel that an overview of the historical traditions and the skills acquired through the study of representation and search are critical components for AI education. Furthermore, these 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 17.

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 The Fifth Edition

The most general change for the fifth edition was the extension of the material related to the stochastic approaches to AI. To accomplish we added a completely new Chapter 5 that introduces the stochastic methodology. From the basic foundations of set theory and counting we develop the notions of probabilities, random variables, and independence. We

present and use Bayes' theorem first with one symptom and one disease and then in its full general form. We examine the hypotheses that underlie the use of Bayes and then present the *argmax* and *naive Bayes* approaches. We present examples of stochastic reasoning, including several from work in the analysis of language phenomena. We also introduce the idea of conditional independence that leads to our presentation of Bayesian belief networks (BBNs) and d-separation in Chapter 9.

We supplemented the presentation of materials throughout the book with more on these stochastic methods. This included introducing *probabilistic finite state machines* and *probabilistic acceptors*, and the presentation of algorithms for *dynamic programming*, especially using stochastic measures (sometimes referred to as the *Viterbi algorithm*). We extended the material in Chapter 9 (reasoning under conditions of uncertainty) to include Bayesian belief nets, hidden Markov, and other graphical models. We added a stochastic English language parser (based on the work of Mark Steedman at the University of Edinburgh) to Chapter 15, the PROLOG material in the book.

We also extended material in several sections of the book to recognize the continuing importance of agent-based problem solving and embodiment in 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 7 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 the situated and social aspects of created intelligence. Chapter 17 contains a critique of each of these representational approaches.

Chapter 14 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 for the WWW.

Finally, in a revised Chapter 17, 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 (Part I) introduces artificial intelligence. We begin with a brief history of attempts to understand mind and intelligence in philosophy, psychology, and other related areas. 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 applications of 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, 5, and 6 (Part II) introduce the research tools for AI problem solving. These include, in Chapter 2, the predicate calculus presented both as a mathematical system as well as a representation language to describe the essential features

of a problem. Search, and the algorithms and data structures used to implement search, are introduced in Chapter 3, to organize the exploration of problem situations. In Chapter 4, we discuss the essential role of heuristics in focusing and constraining search-based problem solving. In Chapter 5, we introduce the stochastic methodology, important technology for reasoning in situations of uncertainty. In Chapter 6, we present a number of software architectures, including the blackboard and production system, for implementing these search algorithms.

Chapters 7, 8, and 9 make up Part III: representations for AI, knowledge-intensive problem solving, and reasoning in changing and ambiguous situations. In Chapter 7 we present the evolving story of AI representational schemes. We begin with a discussion of association-based 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. Expanding on this formalism in Chapter 14, we show how conceptual graphs can be used to implement a natural language database front end. We conclude Chapter 7 with more modern approaches to representation, including Copycat and agent-oriented architectures.

Chapter 8 presents the rule-based expert system along with case-based and model-based reasoning, including examples from the NASA space program. These approaches to problem solving are presented as a natural evolution of the material in Part II: 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 9 presents models for reasoning with uncertainty as well as the use of unreliable information. We introduce Bayesian models, belief networks, Dempster-Shafer, causal models, and the Stanford certainty algebra for reasoning in uncertain situations. Chapter 9 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 10 through 12, is an extensive presentation of issues in machine learning. In Chapter 10 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 11 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 models. We present associative learning and attractor models, including Hopfield networks.

Genetic algorithms and evolutionary approaches to learning are introduced in Chapter 12. 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, contrast, and critique the three approaches we present to machine learning (symbol-based, connectionist, social and emergent) in Chapter 17.

Part V, Chapters 13 and 14, presents automated reasoning and natural language understanding. Theorem proving, often referred to as automated reasoning, is one of the oldest areas of AI research. In Chapter 13, 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 as an instance of the logic programming paradigm.

Chapter 14 presents natural language understanding. Our traditional approach to language understanding, exemplified by many of the semantic structures presented in Chapter 7, 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 develops in LISP and PROLOG many of the algorithms presented in the earlier chapters. Chapter 15 covers PROLOG and Chapter 16 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 17 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 large book. Although it would require more than a single semester to cover all of the material offered, we have designed our book 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, set theory, counting, and graph theory. If this is not true, the instructor should spend more time on these concepts in the optional sections at the beginning of the introductory chapters (2.1, 3.1, and 5.1). 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 6.

In a one quarter or 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 and LISP are both excellent tools for building natural language understanding and learning systems.

In a two-semester or three-quarter 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, such as *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 our pseudo-code 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 fifth edition has an important attached web site maintained by Addison-Wesley Pearson. This site, built by two UNM graduate students, Alejandro CdeBaca and Cheng Liu, includes supplementary ideas for most chapters, some sample problems with their solutions, and many ideas for student projects. Besides the LISP and PROLOG materials

in Chapters 15 and 16, we have included many AI algorithms in Java and C++ on the web site. Students are welcome to use these materials and supplement them with their own comments, code, and critiques. The web url for this book is www.booksites.net/luger.

The PROLOG and LISP code presented in the book is available to readers on the book web page just described as well as via the internet at www.cs.unm.edu/~luger/. Follow the pointers to the fifth edition.

Addison-Wesley and Pearson Education: www.aw-bc.com/cs/ and www.booksites.net/luger support this book with an Instructor's Guide, which has many of the book's exercises worked out, several practice tests with solutions, and ideas for teaching the material. They also have a full set of PowerPoint presentation materials for use by instructors. See your local A-W Pearson representative for access to these materials.

My e-mail address is luger@cs.unm.edu, and I very much enjoy hearing from my readers.

Acknowledgements

Although I am the sole author of the fifth edition, this book has always been the product of my efforts as Professor of Computer Science, Psychology, and Linguistics at the University of New Mexico together with the contributions of my fellow faculty, my professional colleagues, graduate students, and friends, especially the members of the UNM artificial intelligence community. The fifth edition is also the product of the many readers that have e-mailed me comments, corrections, and suggestions. The book will continue this way, reflecting this "community" effort; consequently, I will continue using the prepositions *we*, *our*, and *us* when presenting material.

I thank Bill Stubblefield, the co-author for the first three editions, for more than fifteen years of contributions, but even more importantly, for his friendship over the past twenty-five years. I also thank the many reviewers that have helped develop these five editions. These include Dennis Bahler, Leonardo Bottaci, Skona Brittain, Philip Chan, Peter Collingwood, Mehdi Dastani, John Donald, Sarah Douglas, Christophe Giraud-Carrier, Andrew Kosoresow, Terran Lane, Chris Malcolm, Ray Mooney, Marek Perkowski, Barak Pearlmutter, Dan Pless, Bruce Porter, Julian Richardson, Jude Shavlik, John Sheppard, Carl Stern, Leon van der Torre, Marco Valtorta, and Bob Veroff. We also appreciate the numerous suggestions and comments sent directly by e-mail from people using the book. Finally, Chris Malcolm, Brendan McGonnigle, and Akasha Tang, critiqued Chapter 17.

From our UNM colleagues, we thank Dan Pless for his major role in developing material in Chapters 5 and 9, Joseph Lewis for his efforts on Chapters 9 and 17, Carl Stern for his help in developing Chapter 10 on connectionist learning, Bob Veroff for his critique of the automated reasoning material in Chapter 13, and Jared Saia and Monique Morin for helping with the stochastic approaches to natural language understanding of Chapter 14. Alejandro CdeBaca and Cheng Liu checked the bibliographic references and did the indexing.

We thank Academic Press for permission to reprint much of the material of Chapter 11; 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 Benjamin-Cummings, Addison-Wesley-Longman, and Pearson Education for their support and encouragement in completing the writing task of our five editions, 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, Keith Mansfield, Karen Sutherland, and Anita Atkinson for support on the fourth, and Keith Mansfield, Owen Knight, Mary Lince, and Bridget Allen for their help on this fifth edition. Katherine Haratunian of Addison Wesley USA has had a huge role in seeing that Professors received their Instructor Guides and PowerPoint presentation materials. These are maintained by Addison-Wesley Pearson and available only through your local sales representative. We appreciate Linda Cicarella's work at the University of New Mexico helping prepare figures for publication.

Individual debts for permissions to include copyright materials in all five editions are listed in the *Publisher's Acknowledgments* section which follows this preface.

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

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 2004

Albuquerque

PUBLISHER'S ACKNOWLEDGEMENTS

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

Figures 4.8, 4.9, 9.21 and 9.22, Tables 5.2 and 5.3 adapted from Figure 5.6, p. 157, Figure 5.18, p.178, Figure 5.20, p.180 and data on p.167, from *Speech and Language Processing: an introduction to natural language processing, computational linguistics, and speech recognition*, Prentice Hall, (Pearson Education, Inc.), (Jurafsky, D., and Martin, J.H., 2000); Figure 5.3 adapted from Figure 5.12, p. 170, from *Speech and Language Processing: an introduction to natural language processing, computational linguistics, and speech recognition*, Prentice Hall, (Pearson Education, Inc.), (Jurafsky, D., and Martin, J.H., 2000), which was itself adapted from a figure from *Artificial Intelligence: A Modern Approach*, 1st Edition, Prentice Hall, (Pearson Education, Inc.), (Russell, S.J., and Norvig, P., 1995); Figure 7.1 from figure from *Expert Systems: Artificial Intelligence in Business*, by Harmon, P. and King, D., John Wiley & Sons, Inc., Copyright © 1985 Paul Harmon and David King. This material is used by permission of John Wiley & Sons, Inc.; Figures 7.6, 7.9, and 7.10 from figures from 'Inference and the computer understanding of natural language', in *Artificial Intelligence*, Vol. 5, No. 4, 1974, pp. 373–412, Copyright © 1974, reprinted with permission from Elsevier Science, (Schank, R.C., and Reiger, C.J., 1974); Figures 7.27 and 7.28 from figures from *Analogy-Making as Perception: A Computer Model*, The MIT Press, (Mitchell, M., 1993); Figure 9.2 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 10.21 adapted from figure from 'Models of incremental concept formation', in *Artificial Intelligence*, Vol. 40, Nos. 1–3, 1989, pp. 11–62, Copyright © 1989, reprinted with permission from Elsevier Science, (Gennari, J.H., Langley, P., and Fisher, D., 1989); Figure 11.18 from part of Figure 6.2, p. 102, from *Introduction to Support Vector Machines: and other kernel-based learning methods*, Cambridge University Press, (Cristianini, N., and Shawe-Taylor, J., 2000).

Academic Press for Chapter 11, 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*, Vol. 4, Winter 1944, (Simon, H.A., 1944).

In some instances we have been unable to trace the owners of copyright material, and we would appreciate any information that would enable us to do so.