

Neil C. Rowe

ARTIFICIAL INTELLIGENCE THROUGH PROLOG

NEIL C. ROWE

U.S. Naval Postgraduate School



PRENTICE HALL, Englewood Cliffs, New Jersey 07632

Rowe, Neil C.

Artificial intelligence through Prolog / Neil C. Rowe.

p. cm.

Bibliography: p.

Includes index.

ISBN 0-13-048679-5

1. Prolog (Computer program language) 2. Artificial intelligence.

I. Title.

QA76.73.P76R69 1988

005.13'3-dc19

87-26219

CIP

Editorial/production supervision: **Claudia Citarella**

Cover design: Lundgren Graphics, Ltd.

Manufacturing buyer: S. Gordon Osbourne



© 1988 by Prentice-Hall, Inc.

A Division of Simon & Schuster

Englewood Cliffs, New Jersey 07632

All rights reserved. No part of this book may be reproduced, in any form or by any means, without permission in writing from the publisher.

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

ISBN 0-13-048679-5 025

Prentice-Hall International (UK) Limited, *London*

Prentice-Hall of Australia Pty. Limited, *Sydney*

Prentice-Hall Canada Inc., *Toronto*

Prentice-Hall Hispanoamericana, S.A., *Mexico*

Prentice-Hall of India Private Limited, *New Delhi*

Prentice-Hall of Japan, Inc., *Tokyo*

Simon & Schuster Asia Pte. Ltd., *Singapore*

Editora Prentice-Hall do Brasil, Ltda., *Rio de Janeiro*

PREFACE

Artificial intelligence is a hard subject to learn. I have written a book to make it easier. I explain difficult concepts in a simple, concrete way. I have organized the material in a new and (I feel) clearer way, a way in which the chapters are in a logical sequence and not just unrelated topics. I believe that with this book, readers can learn the key concepts of artificial intelligence faster and better than with other books. This book is intended for all first courses in artificial intelligence at the undergraduate or graduate level, requiring background of only a few computer science courses. It can also be used on one's own. No prior knowledge of the language Prolog is assumed.

Students often complain that while they understand the terminology of artificial intelligence, they don't have a gut feeling for what's going on or how you apply the concepts to a situation. One cause is the complexity of artificial intelligence. Another is the unnecessary baggage, like overly formal logical calculi, that some books and teachers saddle students with. But an equally important cause is the often poor connection made between abstract concepts and their use. So I considered it essential to integrate practical programming examples into this book, in the style of programming language and data structures books. (I stress *practical*, not missionaries and cannibals, definitions of "grandfather," or rules for identifying animals in zoos—at least rarely.) This book has about 500 chunks of code. Clear, concrete formalization of artificial-intelligence ideas by programs and program fragments is all the more critical today with commercialization and media discovery of the field, which has caused a good deal of throwing around of artificial-intelligence terms by people who don't understand them.

But artificial intelligence is a tool for complex problems, and its program examples can easily be forbiddingly complicated. Books attempting to explain artificial intelligence with examples from the programming language Lisp have repeatedly demonstrated this. But I have come to see that the fault lies more with Lisp than with artificial intelligence. Lisp has been the primary language of artificial intelligence for many years, but it is a low-level language, too low for most students. Designed in the early 1960s, Lisp reflects the then-primitive understanding of good programming, and requires the programmer to worry considerably about actual memory references (pointers). Furthermore, Lisp has a weird, hard-to-read syntax unlike that of any other programming language. To make matters worse, the widespread adoption of Common Lisp as a de facto standard has discouraged research on improved Lisps.

Fortunately there is an alternative: Prolog. Developed in Europe in the 1970s, the language Prolog has steadily gained enthusiastic converts, bolstered by its surprise choice as the initial language of the Japanese Fifth Generation Computer project. Prolog has three positive features that give it key advantages over Lisp. First, Prolog syntax and semantics are much closer to formal logic, the most common way of representing facts and reasoning methods used in the artificial-intelligence literature. Second, Prolog provides automatic backtracking, a feature making for considerably easier “search,” the most central of all artificial-intelligence techniques. Third, Prolog supports multidirectional (or multiuse) reasoning, in which arguments to a procedure can freely be designated inputs and outputs in different ways in different procedure calls, so that the same procedure definition can be used for many different kinds of reasoning. Besides this, new implementation techniques have given current versions of Prolog close in speed to Lisp implementations, so efficiency is no longer a reason to prefer Lisp.

But Prolog also, I believe, makes teaching artificial intelligence easier. This book is a demonstration. This book is an organic whole, not a random collection of chapters on random topics. My chapters form a steady, logical progression, from knowledge representation to inferences on the representation, to rule-based systems codifying classes of inferences, to search as an abstraction of rule-based systems, to extensions of the methodology, and finally to evaluation of systems. Topics hard to understand like search, the cut predicate, relaxation, and resolution are introduced late and only with careful preparation. In each chapter, details of Prolog are integrated with major concepts of artificial intelligence. For instance, Chapter 2 discusses the kinds of facts about the world that one can put into computers as well as the syntax of Prolog’s way; Chapter 3 discusses automatic backtracking as well as Prolog querying; Chapter 4 discusses inference and inheritance as well as the definition of procedures in Prolog; Chapter 5 discusses multidirectional reasoning as well as the syntax of Prolog arithmetic; and so on. This constant tying of theory to practice makes artificial intelligence a lot more concrete. Learning is better motivated since one doesn’t need to master a lot of mumbo-jumbo to get to the good stuff. I can’t take much of the credit myself: the very nature of Prolog, and particularly the advantages of the last paragraph, make it easy.

Despite my integrated approach to the material, I think I have covered nearly

all the topics in ACM and IEEE guidelines for a first course in artificial intelligence. Basic concepts mentioned in those guidelines appear toward the beginning of chapters, and applications mentioned in the guidelines appear toward the ends. Beyond the guidelines however, I have had to make tough decisions about what to leave out—a coherent book is better than an incoherent book that covers everything. Since this is a first course, I concentrate on the hard core of artificial intelligence. So I don't discuss much how humans think (that's psychology), or how human language works (that's linguistics), or how sensor interpretation and low-level visual processing are done (that's pattern recognition), or whether computers will ever really think (that's philosophy). I have also cut corners on hard noncentral topics like computer learning and the full formal development of predicate calculus. On the other hand, I emphasize more than other books do the central computer science concepts of procedure calls, variable binding, list processing, tree traversal, analysis of processing efficiency, compilation, caching, and recursion. This is a computer science textbook.

A disadvantage of my integrated approach is that chapters can't so easily be skipped. To partially compensate, I mark some sections within chapters (usually sections toward the end) with asterisks to indicate that they are optional to the main flow of the book. In addition, all of Chapters 7, 10, and 14 can be omitted, and perhaps Chapters 12 and 13 too. (Chapters 7, 10, 13, and 14 provide a good basis for a second course in artificial intelligence, and I have used them that way myself.) Besides this, I cater to the different needs of different readers in the exercises. Exercises are essential to learning the material in a textbook. Unfortunately, there is little consensus about what kind of exercises to give for courses in artificial intelligence. So I have provided a wide variety: short-answer questions for checking basic understanding of material, programming exercises for people who like to program, "play computer" exercises that have the reader simulate techniques described, application questions that have the reader apply methods to new areas (my favorite kind of exercise because it tests real understanding of the material), essay questions, fallacies to analyze, complexity analysis questions, and a few extended projects suitable for teams of students. There are also some miscellaneous questions drawing on the entire book, at the end of Chapter 15. Answers to about one third of the exercises are provided in Appendix G, to offer readers immediate feedback on their understanding, something especially important to those tackling this book on their own.

To make learning the difficult material of this book even easier, I provide other learning aids. I apportion the book into short labeled sections, to make it easier for readers to chunk the material into mind-sized bites. I provide reinforcement of key concepts with some novel graphical and tabular displays. I provide "glass box" computer programs (that is, the opposite of "black box") for readers to study. I mark key terms in italics where they are defined in the text, and then group the most important of these terms into keyword lists at the end of every chapter. I give appendices summarizing the important background material needed for this book, concepts in logic, recursion, and data structures. In other appendices, I summarize the Prolog dialect of the book, make a few comments on Micro-Prolog, and provide a short bibliography (most of the artificial intelligence literature is now either too hard

or too easy for readers of this book). The major programs of the book are available on tape or diskette from the publisher for a small fee. Also, I have prepared an instructor's manual.

It's not necessary to have a Prolog interpreter or compiler available to use this book, but it does make learning easier. This book uses a limited subset of the most common dialect of Prolog, the "standard Prolog" of *Programming in Prolog* by Clocksin and Mellish (second edition, Springer-Verlag, 1984). But most exercises do not require programming.

I've tried to doublecheck all examples, programs, and exercises, but some errors may have escaped me. If you find any, please write me in care of the publisher, or send computer mail to rowe@nps-cs.arpa.

ACKNOWLEDGMENTS

Many people contributed ideas to this book. Michael Genesereth first suggested to me the teaching of introductory artificial intelligence in a way based on logic. David H. Warren gradually eroded my skepticism about Prolog. Harold Abelson and Seymour Papert have steered my teaching style toward student activity rather than passivity.

Judy Quesenberry spent many long hours helping me with the typing and correction of this book, and deserves a great deal of thanks, even if she ate an awful lot of my cookies. Robert Richbourg has been helpful in many different ways, in suggesting corrections and improvements and in testing out some of the programs, despite his having to jump through all the hoops Ph.D. students must jump through. Richard Hamming provided valuable advice on book production. Other people who provided valuable comments include Chris Carlson, Daniel Chester, Ernest Davis, Eileen Entin, Robert Grant, Mike Goyden, Simon Hart, Greg Hoppenstand, Kirk Jennings, Grace Mason, Bruce MacLennan, Norman McNeal, Bob McGhee, James Milojkovic, Doug Owen, Jim Peak, Olen Porter, Brian Rodeck, Jean Sando, Derek Sleeman, Amnon Shefi, and Steve Weingart. Mycke Moore made the creative suggestion that I put a lot of sex into this book to boost sales.

Besides those named, I am grateful to all my students over the years at the Massachusetts Institute of Technology, Stanford University, and the Naval Postgraduate School for providing valuable feedback. They deserve a good deal of credit for the quality of this book—but sorry, people, I’m poor and unwilling to share royalties.

TO THE READER

Artificial intelligence draws on many different areas of computer science. It is hard to recommend prerequisites because what you need to know is bits and pieces scattered over many different courses. At least two quarters or semesters of computer programming in a higher-level language like Pascal is strongly recommended, since we will introduce here a programming language several degrees more difficult, Prolog. If you can get programming experience in Prolog, Lisp, or Logo, that's even better. It also helps to have a course in formal logic, though we won't use much of the fancy stuff they usually cover in those courses; see Appendix A for what you do need to know. Artificial intelligence uses sophisticated data structures, so a data structures course helps; see Appendix C for a summary. Finally, you should be familiar with recursion, because Prolog is well suited to this way of writing programs. Recursion is a difficult concept to understand at first, but once you get used to it you will find it easy and natural; Appendix B provides some hints.

Solving problems is the best way to learn artificial intelligence. So there are lots of exercises in this book, at the ends of chapters. Please take these exercises seriously; many of them are hard, but you can really learn from them, much more than by just passively reading the text. Artificial intelligence is difficult to learn, and feedback really helps, especially if you're working on your own. (But don't plan to do all the exercises: there are too many.) Exercises have code letters to indicate their special features:

- R means a particularly good problem recommended for all readers;
- A means a question that has an answer in Appendix G;
- H means a particularly hard problem;
- P means a problem requiring actual programming in Prolog;
- E means an essay question;
- G means a good group project.

In addition to exercises, each chapter has a list of key terms you should know. Think of this list, at the end of the text for each chapter, as a set of “review questions.”

The symbol * on a section of a chapter means optional reading. These sections are either significantly harder than the rest of the text or significantly far from the core material.

CONTENTS

PREFACE *xiii*

ACKNOWLEDGMENTS *xvii*

TO THE READER *xix*

1 INTRODUCTION *1*

- 1.1 What artificial intelligence is about *1*
- 1.2 Understanding artificial intelligence *12*
- 1.3 Preview *3*

2 REPRESENTING FACTS *4*

- 2.1 Predicates and predicate expressions *4*
- 2.2 Predicates indicating types *6*
- 2.3 About types *8*
- 2.4 Good naming *8*
- 2.5 Property predicates *9*
- 2.6 Predicates for relationships *10*
- 2.7 Semantic networks *12*
- 2.8 Getting facts from English descriptions *13*
- 2.9 Predicates with three or more arguments *14*

- 2.10 Probabilities 15
- 2.11 How many facts do we need? 15

3 VARIABLES AND QUERIES 19

- 3.1 Querying the facts 19
- 3.2 Queries with one variable 21
- 3.3 Multidirectional queries 21
- 3.4 Matching alternatives 22
- 3.5 Multicondition queries 23
- 3.6 Negative predicate expressions 25
- 3.7 Some query examples 26
- 3.8 Loading a database 27
- 3.9 Backtracking 27
- 3.10 A harder backtracking example:
superbosses 30
- 3.11 Backtracking with “not”s 32
- 3.12 The generate-and-test scheme 34
- 3.13 Backtracking with “or”s* 34
- 3.14 Implementation of backtracking 35
- 3.15 About long examples 36

4 DEFINITIONS AND INFERENCES 43

- 4.1 Rules for definitions 43
- 4.2 Rule and fact order 46
- 4.3 Rules as programs 47
- 4.4 Rules in natural language 47
- 4.5 Rules without right sides 48
- 4.6 Postponed binding 49
- 4.7 Backtracking with rules 49
- 4.8 Transitivity inferences 52
- 4.9 Inheritance inferences 54
- 4.10 Some implementation problems
for transitivity and inheritance 58
- 4.11 A longer example: some traffic laws 60
- 4.12 Running the traffic lights program 64
- 4.13 Declarative programming 65

5 ARITHMETIC AND LISTS IN PROLOG 74

- 5.1 Arithmetic comparisons 74
- 5.2 Arithmetic assignment 75
- 5.3 Reversing the “Is” 76
- 5.4 Lists in Prolog 78
- 5.5 Defining some list-processing predicates 80
- 5.6 List-creating predicates 83

- 5.7 Combining list predicates 87
- 5.8 Redundancy in definitions 88
- 5.9 An example: dejargonizing
bureaucratese* 89

6 CONTROL STRUCTURES FOR RULE-BASED SYSTEMS 99

- 6.1 Backward-chaining control structures 100
- 6.2 Forward chaining 102
- 6.3 A forward chaining example 103
- 6.4 Hybrid control structures 105
- 6.5 Order variants 108
- 6.6 Partitioned control structures 109
- 6.7 Meta-rules 109
- 6.8 Decision lattices 110
- 6.9 Concurrency in control structures 112
- 6.10 And-or-not lattices 113
- 6.11 Randomness in control structures 115
- 6.12 Grammars for interpreting languages* 115

7 IMPLEMENTATION OF RULE-BASED SYSTEMS 125

- 7.1 Implementing backward chaining 125
- 7.2 Implementing virtual facts
and caching 126
- 7.3 Input coding 127
- 7.4 Output coding 128
- 7.5 Intermediate predicates 130
- 7.6 An example program 131
- 7.7 Running the example program 133
- 7.8 Partitioned rule-based systems 134
- 7.9 Implementing the rule-cycle hybrid 135
- 7.10 Implementing pure forward chaining* 137
- 7.11 Forward chaining with “not”’s* 140
- 7.12 General iteration with “forall”
and “doall”* 141
- 7.13 Input and output of forward chaining* 143
- 7.14 Rule form conversions* 146
- 7.15 Indexing of predicate expressions* 148
- 7.16 Implementing meta-rules* 149
- 7.17 Implementing concurrency* 151
- 7.18 Decision lattices: a compilation
of a rule-based system* 151
- 7.19 Summary of the code described
in the chapter* 155

8 REPRESENTING UNCERTAINTY IN RULE-BASED SYSTEMS 164

- 8.1 Probabilities in rules 164
- 8.2 Some rules with probabilities 166
- 8.3 Combining evidence assuming
statistical independence 167
- 8.4 Prolog implementation of independence-assumption
“and-combination” 170
- 8.5 Prolog implementation of independence-assumption
“or-combination” 171
- 8.6 The conservative approach 173
- 8.7 The liberal approach and others 175
- 8.8 Negation and probabilities 177
- 8.9 An example: fixing televisions 177
- 8.10 Graphical representation of probabilities
in rule-based systems 180
- 8.11 Getting probabilities from statistics 180
- 8.12 Probabilities derived from others 182
- 8.13 Subjective probabilities 183
- 8.14 Maximum-entropy probabilities* 184
- 8.15 Consistency* 185

9 SEARCH 191

- 9.1 Changing words 191
- 9.2 States 192
- 9.3 Three examples 193
- 9.4 Operators 193
- 9.5 Search as graph traversal 194
- 9.6 The simplest search strategies:
depth-first and breadth-first 197
- 9.7 Heuristics 199
- 9.8 Evaluation functions 200
- 9.9 Cost functions 202
- 9.10 Optimal-path search 203
- 9.11 A route-finding example 204
- 9.12 Special cases of search 205
- 9.13 How hard is a search problem? 207
- 9.14 Backward chaining
versus forward chaining* 208
- 9.15 Using probabilities in search* 211
- 9.16 Another example: visual
edge-finding as search* 211

10 IMPLEMENTING SEARCH 223

- 10.1 Defining a simple search problem 223
- 10.2 Defining a search problem
with fact-list states 226
- 10.3 Implementing depth-first search 229
- 10.4 A depth-first example 230
- 10.5 Implementing breadth-first search 232
- 10.6 Collecting all items that satisfy
a predicate expression 236
- 10.7 The cut predicate 238
- 10.8 Iteration with the cut predicate* 240
- 10.9 Implementing best-first search* 241
- 10.10 Implementing A* search* 244
- 10.11 Implementing search with heuristics* 249
- 10.12 Compilation of search* 250

11 ABSTRACTION IN SEARCH 263

- 11.1 Means-ends analysis 263
- 11.2 A simple example 264
- 11.3 Partial state description 268
- 11.4 Implementation of means-ends analysis 268
- 11.5 A harder example: flashlight repair 271
- 11.6 Running the flashlight program 274
- 11.7 Means-ends versus other search methods 282
- 11.8 Modeling real-world uncertainty* 283
- 11.9 Procedural nets* 283

12 ABSTRACTION OF FACTS 285

- 12.1 Partitioning facts 285
- 12.2 Frames and slots 286
- 12.3 Slots qualifying other slots 287
- 12.4 Frames with components 288
- 12.5 Frames as forms: memos 288
- 12.6 Slot inheritance 290
- 12.7 Part-kind inheritance 291
- 12.8 Extensions versus intensions 291
- 12.9 Procedural attachment 292
- 12.10 Frames in Prolog 292
- 12.11 Example of a frame lattice 293
- 12.12 Expectations from slots 297
- 12.13 Frames for natural language
understanding* 298

- 12.14 Multiple inheritance* 299
- 12.15 A multiple inheritance example:
custom operating systems* 299

13 PROBLEMS WITH MANY CONSTRAINTS 307

- 13.1 Two examples 307
- 13.2 Rearranging long queries
without local variables 311
- 13.3 Some mathematics 312
- 13.4 Rearranging queries
with local variables 313
- 13.5 Rearranging queries based
on dependencies 314
- 13.6 Summary of guidelines for optimal query
arrangements 315
- 13.7 Rearrangement and improvement
of the photo interpretation query 316
- 13.8 Dependency-based backtracking 319
- 13.9 Reasoning about possibilities 321
- 13.10 Using relaxation for the photo
interpretation example 322
- 13.11 Quantifying the effect* 324
- 13.12 Formalization of pure relaxation 325
- 13.13 Another relaxation example:
cryptarithmic 326
- 13.14 Implementation of pure relaxation* 330
- 13.15 Running a cryptarithmic relaxation* 334
- 13.16 Implementing double relaxation* 336

14 A MORE GENERAL LOGIC PROGRAMMING 349

- 14.1 Logical limitations of Prolog 349
- 14.2 The logical (declarative) meaning
of Prolog rules and facts 350
- 14.3 Extending Prolog rules 352
- 14.4 More about clause form 353
- 14.5 Resolution 354
- 14.6 Resolution with variables 356
- 14.7 Three important applications
of resolution 357
- 14.8 Resolution search strategies 357
- 14.9 Implementing resolution
without variables* 359

15	TESTING AND DEBUGGING OF ARTIFICIAL-INTELLIGENCE PROGRAMS	364
15.1	The gold standard	364
15.2	Cases	365
15.3	Focusing on bugs	366
15.4	Exploiting pairs of similar cases	367
15.5	Composite results	368
15.6	Numbers in comparisons	369
15.7	Preventive measures	370
15.8	Supporting intuitive debugging by explanations	370
15.9	Evaluating cooperativeness	371
15.10	On problems unsuitable for artificial intelligence	372
	Miscellaneous exercises covering the entire book	375
A	BASICS OF LOGIC	381
B	BASICS OF RECURSION	385
C	BASICS OF DATA STRUCTURES	389
D	SUMMARY OF THE PROLOG DIALECT USED IN THIS BOOK	393
D.1	Managing facts and rules	393
D.2	The format of facts, rules, and queries	395
D.3	Program layout	395
D.4	Lists	396
D.5	Numbers	396
D.6	Output and input	396
D.7	Strings	397
D.8	Treating rules and facts as data	397
D.9	Miscellaneous predicates	397
D.10	Definable predicates	397
D.11	Debugging	398
E	USING THIS BOOK WITH MICRO-PROLOG	399