

artificial intelligence programming

eugene charniak
christopher k. riesbeck
drew v. mcdermott
james r. meehan

second edition



ARTIFICIAL INTELLIGENCE PROGRAMMING

SECOND EDITION

EUGENE CHARNIAK

Brown University

CHRISTOPHER K. RIESBECK

Yale University

DREW V. McDERMOTT

Yale University

JAMES R. MEEHAN

Cognitive Systems, Inc.



LAWRENCE ERLBAUM ASSOCIATES, PUBLISHERS

1987 Hillsdale, New Jersey

Hove and London

Copyright © 1987 by Lawrence Erlbaum Associates, Inc.

All rights reserved. No part of this book may be reproduced in any form, by photostat, microform, retrieval system, or any other means, without the prior written permission of the publisher.

Lawrence Erlbaum Associates, Inc., Publishers
365 Broadway
Hillsdale, New Jersey 07642

Library of Congress Cataloging in Publication Data

Artificial intelligence programming.

Rev. ed. of: Artificial intelligence programming /
Eugene Charniak, Christopher K. Riesbeck, Drew V.
McDermott. 1980.

Bibliography: p.

Includes index.

1. Artificial intelligence—Data processing.
2. Programming (Electronic computers) 3. LISP
(Computer program language) I. Charniak, Eugene.
Q336.A78 1987 006.3 87-8981
ISBN 0-89859-609-2

Printed in the United States of America

10 9 8 7 6 5 4 3 2

ARTIFICIAL INTELLIGENCE PROGRAMMING

SECOND EDITION

PREFACE TO THE SECOND EDITION

Since the first edition of this book appeared, some things in the AI programming world have changed a great deal, and some things are almost exactly the way they were five years ago. Perhaps the most significant development has been the appearance of COMMON LISP, documented in abundant detail by Guy Steele [102]. All of the LISP code in this new edition has been rewritten in COMMON LISP. COMMON LISP is a pleasant surprise, given the normal result of compromise solutions designed by committees. It manages to be a synthesis of many of the best ideas present in modern LISP dialects, rather than a fossilization of the worst. While there are other dialects that have a more coherent semantics, such as SCHEME [82] and T [81], COMMON LISP is more than adequate for our needs.

The choice of COMMON LISP has affected the contents of this book in several ways. First, many features that we spent some time developing and adding to our earlier dialect of LISP are already available in COMMON LISP. In some cases, we have therefore just described the COMMON LISP feature. In other cases, we have retained the developmental material in order to explain the underlying principles.

COMMON LISP is a “large” language, and we cover only part of it; some of the best-designed features allowed us to remove material from the first edition that dealt with the friendly but limited dialect we used then, UCI LISP. Gone are the sections on FEXPRs and LEXPRs; while we still discuss the issue of extending the language by adding new data types, COMMON LISP’s DEFSTRUCT is an example of a tool that we had to build from scratch in UCI LISP.

We considered both SCHEME and T for this edition. SCHEME has the right essential semantics for programming, such as lexical scoping and closures, and

T extends that to include the right primitives for object-oriented programming and language-extension via macros, as well as a host of well-designed support features. (The next edition (!) will probably be in T.) Happily, many of these ideas are also present in COMMON LISP, and given its greater visibility, we chose it instead. In particular, the availability of lexical closures has allowed us to re-implement a number of disparate ideas in a uniform manner.

One problem with COMMON LISP is that it is not yet widely available. Most implementations exist only on very large machines. While this situation will certainly change in the next few years, we have tried to ameliorate the problem in two ways. First, we have used only a subset of COMMON LISP in our code. Second, we have provided a glossary describing the subset we used.

In the LISP chapters in Part 1 and in the AI topics in Part 2, five years' additional experience has led us to provide completely new explanations, examples, and implementations. The original chapter on alternative control structures, which described the implementation of a variation of SCHEME, served two purposes: to introduce some of the power behind lexically scoped languages with procedures as first-class objects, and to give an example of how other languages with very different control structures could be implemented in LISP. Since COMMON LISP is lexically scoped, there is no longer any need to treat that topic separately. The topic of control structure in SCHEME is now described briefly in a new chapter on higher-order functions, continuations, and coroutines. Some of the flavor of implementing interpreters can be gleaned from a new chapter on production systems. (For more information on SCHEME and implementing a SCHEME interpreter, the interested reader is referred to Abelson and Sussman [2].)

The chapter on production systems illustrates several aspects of AI programming. A production system is much easier to implement than a deductive retriever, and correspondingly more limited in power. Production systems, however, have proved to be very useful in the development of expert systems, which are the basis for most of the commercial AI work at the moment.

As before, the intended audience of *Artificial Intelligence Programming* remains the advanced undergraduate or early graduate student in AI. Although the material involved requires only modest knowledge of programming, the student who has had the most experience in creating AI programs already will understand best the benefits of the techniques described. In addition, in response to the changing nature of AI in industry, we have changed the text to be a bit more like a cookbook. It has turned out that many people learning AI programming are doing so on their own, either at home or at work. Many of them prefer to begin with working pieces of code that they can then extend, rather than inadequate pieces of code that are then corrected in the text and exercises. This latter technique works in classrooms, but not in a self-study situation. Therefore, we have eliminated almost all figures with deliberately incorrect code, and added an appendix with the answers to nearly all the exercises in the book. We have also removed the chapters on the sample course project.

ACKNOWLEDGMENTS

We would like to thank the teachers, students, and other programmers who suggested changes (and supplied corrections!) to the first edition, as well as reviewers who gave us criticisms prior to this revision. Thanks to Paul Hudak and Bill Ferguson for comments on Chapter 7, and to Robert Farrell for comments on Chapter 12. We are indebted to Bob Strong at Cognitive Systems Inc. for reviewing the new edition, recommending improvements, and helping us test the code. Answers to some of the exercises were written by Gregory Parkinson and Rika Yoshii when they were students at UCI. The Yale Artificial Intelligence Project and Cognitive Systems Inc. provided computer support for the preparation of the new manuscript. Finally, Chris Riesbeck would like to dedicate this book to Maxine.

Christopher K. Riesbeck
Jim Meehan

PREFACE TO THE FIRST EDITION

Artificial Intelligence (henceforth AI) is still a field where disagreement is more common than solid theory, and interesting ideas more common than polished programs. Yet there is slowly coming into being a small core of accepted (though not universally accepted) theory and practice. This book is an attempt to gather together the “practice” aspect of this “core” AI.

The “practice” of AI is, of course, the writing of programs. AI problems are usually ill-defined and the theories proposed are often too complex and complicated to be verified by intuitive or formal arguments. Sometimes the only way to understand and evaluate a theory is to see what comes next. To find this out, and to check for obvious inconsistencies and contradictions, we write programs that are intended to reflect our theories. If these programs work, our theories are not proved, of course, but at least we gain some understanding of how they behave. When the programs don’t work (or we find ourselves unable to program the theories at all), then we learn what we have yet to define or re-define.

With this emphasis on programming, it becomes important that an AI researcher have a wide library of programming tools available. This is particularly true because of the “level” problem. That is, your theory describes what to do at a fairly high level, but you need to tell the machine what to do at a low level. So a theory of, say, coherency in conversation, will in all probability say nothing about pattern matching or efficient data retrieval. It is not that these topics are not worthy of their own theory. How people manage to retrieve knowledge efficiently under a wide variety of circumstances is a fascinating

question. But if you are worried about conversation, it is simply not your department.

The intent of this book is to give you a wide variety of commonly used tools for programming Artificial Intelligence theories: discrimination nets, agendas, deduction, data dependencies, backtracking, etc. By having these tools, we hope you will find that your programs better reflect your intentions.

Almost all of the ideas that are described here are in common use, particularly at the larger Artificial Intelligence centers. But very few of them have ever been written down in one place. There are a number of books that introduce you to LISP (although none of them are completely satisfactory), and there are a number of books on theories and algorithms in Artificial Intelligence. Until now, however, there have been no books that fill in the middle ground and present the methods that all the old-timers know for getting from theory to practice. That is what this book is all about.

The major problem in writing a book such as this is that of selection. In some cases it is easy. It seems unlikely that anyone would seriously contest our inclusion of discrimination nets, pattern matching, or agendas. These techniques have been used by many researchers in the field and in a variety of problem areas—from natural language comprehension to problem solving to medical diagnosis. However once one moves beyond this handful of topics, or even starts getting specific about the type of pattern matching, or agenda, then consensus is not so easy. So to some degree the selections made in this book are personal ones. Of course, to say they are personal is not to say they cannot be defended on scientific grounds, but rather that the defense would take the form of an extended debate on the nature of AI and where it is going. For example, data dependencies receive a chapter to themselves here in spite of the fact that they are fairly new on the scene and hence relatively untested, at least compared to something like unification pattern matching. Naturally we try to show the usefulness of these ideas, but only to show how the ideas are motivated, not to defend particular approaches against competitors. Such a defense would be well worth having, but it would be out of place in a text such as this.

Selection also implies that some things are omitted, and there are at least two notable omissions from these chapters. One of these is inadvertent. The techniques discussed here all come from what might be thought of as “abstract” AI. That is, if we think of AI programs on a spectrum from “concrete” programs which must deal with the real world in terms of sound and light input (or sound and muscle output) to “abstract” programs which only deal with abstractions, the techniques described here fall most naturally towards the abstract end. This book does not have the space and the authors do not have the expertise to do justice to the concrete end of things.

A second omission is quite deliberate. We have made no attempt to survey, much less teach, the many AI languages (CONNIVER, QA-4, KRL, etc.). This stems from our conviction that at present there is no commonly agreed-upon set

of functions above the level of list processing which everyone would agree is useful in a wide variety of AI settings. Experience has shown that each major project has found it necessary to build up its own tools, starting, typically, from LISP. We do not see this situation as likely to change in the foreseeable future. Hence rather than covering the basics of the various languages we have tried instead to explain the techniques which typically lie behind these languages.

The book is divided into two parts. Since almost all serious programming in Artificial Intelligence is done in the language LISP, Part 1 tells you how to improve your general abilities as a LISP programmer. [The first chapter covers] most of the basic LISP concepts needed for the rest of the book. We intend the introductory material to cover all the concepts of LISP needed later, but if you have never programmed in LISP before, we recommend that you spend some time writing simple LISP programs until you get a feel for the language.

The [second through seventh] chapters are concerned with the many features found (or implementable) in LISP that make the language an attractive one to use. Many of the ideas that pass under the rubric of “structured programming” will be found here. Although LISP is almost as old as FORTRAN, it is surprisingly amenable to things like top-down programming and data types.

Part 2 contains more advanced and complex techniques. Since this book is intended not just to be a description of ideas, but also to give you a chance to learn the craft of Artificial Intelligence, we present actual LISP implementations of all the ideas discussed, along with exercises which modify and extend the code. These exercises are intended to make you familiar, in a practical hands-on way, with the techniques involved. We hope that the exercises will inspire you to experiment and learn on your own.

This book is intended mainly for use as a textbook for an AI course in which programming is emphasized. This could be either an advanced or a fast elementary course. The book might also be used as an auxiliary text for a systems course; for this purpose, the chapters on macros, structured programming, and alternative control structures would be most useful.

ACKNOWLEDGMENTS

The first edition of this book was the outgrowth of a graduate course given by the authors in the spring of 1978 at Yale University for students in the Artificial Intelligence Project of the Yale Computer Science Department. We would like to acknowledge Roger Schank who originated the idea of the course. We'd also like to thank Dave Barstow and Walter Stutzman for their detailed comments on the complete manuscript, Laury Miller and Glenn Edelson for their comments on the introductory chapters, and Jon Doyle for his comments on the chapter on data dependencies.

The Yale Artificial Intelligence Project is funded by the Advanced Research Projects Agency of the Department of Defense and the Office of Naval Research.

During the writing of [the first edition of] this book, Eugene Charniak and Christopher Riesbeck were supported by the Advanced Research Projects Agency monitored under the Office of Naval Research under contract N00014-75-C-1111.

Eugene Charniak
Christopher K. Riesbeck
Drew V. McDermott



LISP PROGRAMMING

Contents

Preface xi

PART I: LISP PROGRAMMING

1. LISP REVIEW	1
1.1 Data Structures	1
1.2 Program Structures	2
1.3 Primitive Operations on S-expressions	8
1.4 Tree Structures	10
1.5 Lists	12
1.6 Mapping Functions	14
1.7 LAMBDA Expressions	15
1.8 Atoms	17
1.9 Inside Lisp	18
1.10 Equality	23
1.11 Local versus free variables	25
1.12 Lisp style	27
1.13 Keywords	29
1.14 More lambda-list keywords	32
2. MACROS AND READ-MACROS	33
2.1 Read-Macros	33
2.2 The Backquote Read-Macro	34
2.3 Terminating read-macros	39

2.4	Macros	39	
2.5	Generating new symbols	45	
2.6	BIND: A Macro For Special Variables	47	
3.	DATA STRUCTURES AND CONTROL STRUCTURES IN LISP		48
3.1	The Need for Data Types	48	
3.2	The Conservative Approach to Type Definition	50	
3.3	The Liberal Approach	51	
3.4	The Radical Approach	53	
3.5	Control Structures	55	
3.6	Basic control structure	56	
3.7	Local functions: FLET and LABELS	56	
3.8	Interrupting the normal flow	58	
3.9	Redirecting the flow of control: GO	59	
3.10	ITERATE	60	
3.11	DO	63	
3.12	Iteration and lists: the mapping functions	64	
3.13	Defining the FOR-macro	69	
3.14	CATCH and THROW	74	
3.15	UNWIND-PROTECT	74	
3.16	BIND	75	
3.17	Conclusion	82	
4.	INPUT/OUTPUT IN LISP		83
4.1	Streams	84	
4.2	READ-CHAR, READ-LINE, and READ	85	
4.3	PRIN1, PRINC, and TERPRI	88	
4.4	PRINT, PPRINT, and FORMAT	88	
4.5	The MSG Macro	90	
4.6	Separating I/O from Your Functions	93	
5.	COMPILING YOUR PROGRAM AND YOUR PROGRAM'S PROGRAM		97
5.1	What Is Compilation?	97	
5.2	Implications for AI Programs	99	
5.3	Example: Regular Expressions	100	
5.4	What the Lisp Compiler Does	101	
5.5	Compiler Declarations	105	
5.6	Macros in Compiled Code	107	
5.7	Variables in Compiled Code	107	

5.8	Lexical Scoping Versus EVAL and SET	108
5.9	Ignored Variables	110
6.	DATA-DRIVEN PROGRAMMING AND OTHER PROGRAMMING TECHNIQUES	111
6.1	Data-driven Programming	111
6.2	Association Lists, Property Lists, and Hash Tables	112
6.3	Reimplementing MSG	116
6.4	Data-Driven Programming as an Organizational Device	119
6.5	Set Operations on Lists	121
6.6	Headed Lists and Queues	123
7.	HIGHER-ORDER FUNCTIONS, CONTINUATIONS, AND COROUTINES	127
7.1	Passing Procedures In and Out	127
7.2	Continuations and Tail-Recursion	129
7.3	Continuations and Multiple Values	132
7.4	Coroutines	136
7.5	Continuations and Control Flow: CALL-WITH-CURRENT-CONTINUATION	142
7.6	Problems with Continuation Passing	147
PART II: AI PROGRAMMING TECHNIQUES		
8.	SIMPLE DISCRIMINATION NETS	149
8.1	The General Discrimination Net	149
8.2	Database Discrimination Nets—Lists of Atoms	153
8.3	Database Discrimination Nets—General S-expressions	157
8.4	Implementing Discrimination Trees with Continuations	161
9.	AGENDA CONTROL STRUCTURES	166
9.1	Introduction	166
9.2	Best-First Tree Search	166
9.3	Coroutines and Agendas	171
9.4	Design Alternatives for Agendas	176
9.5	Generated Lists	177
10.	DEDUCTIVE INFORMATION RETRIEVAL	182
10.1	Introduction	182

10.2	Database-defined Predicates	183	
10.3	Connectives, Variables, and Inference Rules	183	
10.4	Existential Quantification	187	
10.5	Coming to Terms	188	
10.6	Issues	189	
10.7	Deductive Retrieval	190	
10.8	A Unification Algorithm	195	
10.9	A Deductive Retriever	202	
10.10	Forward Chaining	213	
10.11	Programming with Logic	215	
10.12	Extending the Retriever	218	
10.13	Deductive Retrieval versus Theorem Proving	223	
10.14	Pattern Matching	224	
10.15	The Pros and Cons of Deductive Information Retrieval	226	
11.	DISCRIMINATION NETS WITH VARIABLES		230
11.1	Plan Retrieval	231	
11.2	Fetching Facts	235	
11.3	Variations on the Discrimination Net Theme	240	
12.	PRODUCTION SYSTEMS		248
12.1	Representation of Rules in XPS	249	
12.2	Representation of Working Memory Elements in XPS	250	
12.3	Establishing the Conflict Set	251	
12.4	Conflict Resolution in XPS	251	
12.5	The Implementation of XPS	254	
12.6	An Example	266	
12.7	Improvements to XPS	271	
12.8	XPS and OPS5	272	
12.9	Deductive Retrieval and Production Systems	275	
13.	SLOT AND FILLER DATA BASES		276
13.1	Expanding Property List Facilities	276	
13.2	An Introduction to XRL	277	
13.3	IF-ADDED Methods	287	
13.4	Pattern Matching in XRL	288	
13.5	Indexing Forms	294	
13.6	Retrieving Forms	296	
13.7	Extending XRL	298	

14. CHRONOLOGICAL BACKTRACKING	304
14.1 Introduction	304
14.2 A Basic Transition Network Grammar without Backup	305
14.3 Representing an ATN Network in Lisp	308
14.4 An ATN without Backtracking	312
14.5 Backtracking In ATNs—A State-Saving Approach	320
14.6 Backtracking with a Transition-Saving Approach	327
14.7 Defining the ATN with Continuations	328
14.8 Using Problem-Solving Techniques	332
15. DATA DEPENDENCIES AND REASON MAINTENANCE SYSTEMS	337
15.1 The Need for Reason Maintenance	337
15.2 Data-Dependency Clauses	338
15.3 Data-Dependency Network Graphs	339
15.4 Labeling Data-Dependency Networks	339
15.5 Propagating Labels from New Clauses	341
15.6 Incompleteness of the RMS Algorithm	343
15.7 Retracting Clauses and Delabeling Nodes	345
15.8 The Implementation of an RMS	347
15.9 The Eight Queens Problem	355
15.10 Dependency-directed Backtracking	367
15.11 Finding Nogoods	372
15.12 Subsuming Clauses	378
15.13 Dealing with Interruptions	379
15.14 RMS to Lisp Communications	385
Appendix 1: A Glossary of Common Lisp Functions	389
Answers to Selected Exercises	441
Bibliography	495
Author Index	505
Index of Defined LISP Items	507
Subject Index	515

1

Lisp REVIEW

LISP has jokingly been called “the most intelligent way to misuse a computer.” I think that description is a great compliment because it transmits the full flavor of liberation: it has assisted a number of our most gifted fellow humans in thinking previously impossible thoughts.

— Edsger Dijkstra

LISP was the world’s first elegant language, in the sense that it provided a parsimonious base with rich possibilities for extension. LISP has been applied mainly to problems of symbolic manipulation and artificial intelligence, partly because manipulating symbols is so easy in LISP, and partly because AI programmers tend to be lazy and undisciplined, like pilots who refuse to file a flight plan before taking off, and LISP’s interactive structure allows them to get away with this.

1.1 Data Structures

LISP data structures are called “S-expressions.” The S stands for “symbolic.” In this text, the terms “S-expression” and “expression” are used interchangeably. An S-expression is

1. a *number*, e.g., 15, written as an optional plus or minus sign, followed by one or more digits.
2. a *symbol*, e.g., FOO, written as a letter followed by zero or more letters or digits.