Logic Programming with Prolog

Max Bramer

Logic Programming with Prolog

Springer

Max Bramer, BSc, PhD, CEng, FBCS, FIEE, FRSA, ILTM
Department of Computer Science and Software Engineering
University of Portsmouth
United Kingdom

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

Library of Congress Control Number: 2005923526

Apart from any fair dealing for the purposes of research or private study, or criticism or review, as permitted under the Copyright, Designs and Patents Act 1988, this publication may be reproduced, stored or transmitted, in any form or by any means, with the prior permission in writing of the publishers, or in the case of reprographic reproduction in accordance with the terms of licenses issued by the Copyright Licensing Agency. Enquiries concerning reproduction outside those terms should be sent to the publishers.

ISBN-10: 1-85233-938-1
ISBN-13: 978-1852-33938-8
Springer Science+Business Media
springeronline.com

© Max Bramer 2005

The use of registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant laws and regulations and therefore free for general use.

The publisher makes no representation, express or implied, with regard to the accuracy of the information contained in this book and cannot accept any legal responsibility or liability for any errors or omissions that may be made.

Printed in the United States of America
34-543210 Printed on acid-free paper SPIN 11344445

Introduction

Logic Programming is the name given to a distinctive style of programming, very different from that of conventional programming languages such as C++ and Java. Fans of Logic Programming would say that 'different' means clearer, simpler and generally better!

Although there are other Logic Programming languages, by far the most widely used is **Prolog**. The name stands for Programming in Logic. This book teaches the techniques of Logic Programming through the Prolog language. Prolog is based on research by computer scientists in Europe in the 1960s and 1970s, notably at the Universities of Marseilles, London and Edinburgh. The first implementation was at the University of Marseilles in the early 1970s. Further development at the University of Edinburgh led to a *de facto* standard version, now known as Edinburgh Prolog. Prolog has been widely used for developing complex applications, especially in the field of Artificial Intelligence. Although it is a general-purpose language, its main strengths are for symbolic rather than for numerical computation.

The developers of the language were researchers working on automating mathematical theorem proving. This field is often known as *computational logic*. But if you are not a Computer Scientist, a logician or a mathematician do not let this deter you! This book is aimed at the 99.9% of the population who are none of these. Those who are, already have a number of excellent textbooks from which to choose.

The idea that the methods developed by computational logicians could be used as the basis for a powerful general purpose programming language was revolutionary 30 years ago. Unfortunately most other programming languages have not yet caught up.

The most striking feature of Prolog for the newcomer is how much simpler the programs look than in other languages. Many language designers started out with good intentions but could not resist, or perhaps could not avoid, making their creations over elaborate. In some languages even writing the customary test program to print out the words *Hello World!* to the user's screen is hard work. All the user has to do in Prolog is to enter **write('Hello World!')**.

Traditional programming languages have one feature in common. They all contain a series of instructions that are to be performed ('executed') one after another. This style of programming is called *procedural*. It corresponds closely to the way computers are generally built. This is a tempting approach that has been used since the 1950s but is ultimately flawed. The way users write programs should depend as little as possible on how the underlying machine was built and as much as possible on what the user is trying to do. In just the same way, the facilities I use when I drive my car should depend as little as possible on how the engine was designed or the carburettor works. I want all that level of detail hidden from me, although in practice I understand that it may not be completely possible to ignore it.

Prolog programs are often described as *declarative*, although they unavoidably also have a procedural element. Programs are based on the techniques developed by logicians to form valid conclusions from available evidence. There are only two components to any program: facts and rules. The Prolog system reads in the program and simply stores it. The user then types in a series of questions (strictly known as *queries*), which the system answers using the facts and rules available to it. This is a simple example, a series of queries and answers about animals. The program consists of just seven lines (blank lines are ignored).

```
dog(fido).

dog(rover).

dog(henry).

cat(felix).

cat(michael).

cat(jane).

animal(X):-dog(X).
```

This program is not too hard to decipher. The first three lines are *facts*, with the obvious interpretation that fido, rover and henry are all dogs. The next three facts say that felix, michael and jane are all cats.

The final line is a *rule* saying that anything (let us call it X) is an animal if it is a dog. Cat lovers may feel that cats can also claim to be called animals, but the program is silent about this.

Having loaded the program, the user is then faced with the two character symbol **?-** which is called the *system prompt*. To check whether fido is a dog all that is necessary is to type the query **dog(fido)** followed by a full stop and press the 'return' key, which indicates to the system that a response is needed. This gives the complete dialogue:

?- dog(fido).
yes

The user can enter a series of queries at the prompt, asking for further information.

?-dog(jane).	[Is jane a dog? No - a cat]
no	
?- animal(fido).	[Is fido an animal?]
yes	[yes - because it is a dog and any dog is an animal]
?- dog(X).	[Is it possible to find anything, let us call it X, that is a dog?]
X = fido ;	
X = rover ;	
X = henry	[All 3 possible answers are provided]
?-animal(felix).	[felix is a cat and so does not qualify as an animal, as far as the program is concerned]
no	

Although straightforward, this example shows the two components of any Prolog program, *rules* and *facts*, and also the use of *queries* that make Prolog search through its facts and rules to work out the answer. Determining that fido is an animal involves a very simple form of logical reasoning:

GIVEN THAT
any X is an animal if it is a dog

AND
fido is a dog

DEDUCE
fido must be an animal

This type of reasoning is fundamental to theorem proving in Mathematics and to writing programs in Prolog.

Even very simple queries such as:

?-dog(fido).

can be looked at as asking the Prolog system to prove something, in this case that fido is a dog. In the simplest cases it can do so simply by finding a fact such as dog(fido) that it has been given. The system answers yes to indicate that this simple 'theorem' has been proved.

You have now seen all three elements needed for logic programming in Prolog: facts, rules and queries. There are no others. Everything else is built from them.

A word of warning is necessary at this stage. It is easy for the newcomer to get started with Prolog, but do not be fooled by these examples into thinking that Prolog is only capable of handling simple (Mickey Mouse?) problems. By putting

these very basic building blocks together Prolog provides a very powerful facility for building programs to solve complex problems, especially ones involving reasoning, but all Prolog programs are simple in form and are soundly based on the mathematical idea of proving results from the facts and rules available.

Prolog has been used for a wide variety of applications. Many of these are in Mathematics and Logic but many are not. Some examples of the second type of application are

- programs for processing a 'natural language' text, to answer questions about its meaning, translate it to another language etc.
- advisory systems for legal applications
- applications for training
- maintaining databases for the Human Genome project
- a personnel system for a multi-national computer company
- automatic story generation
- analysing and measuring 'social networks'
- a software engineering platform for supporting the development of complex software systems
- automatically generating legally correct licences and other documents in multiple languages
- an electronic support system for doctors.

Prolog is also being used as the basis for a standard 'knowledge representation language' for the Semantic Web – the next generation of internet technology.

Prolog is one of the principal languages used by researchers in Artificial Intelligence, with many applications developed in that field, especially in the form of *Expert Systems* – programs that 'reason out' the solution to complex problems using rules.

Many textbooks on Prolog assume that you are an experienced programmer with a strong background in Mathematics, Logic or Artificial Intelligence (preferably all three). This book makes no such assumptions. It starts from scratch and aims to take you to a point where you can write quite powerful programs in the language, often with considerably fewer lines of program 'code' than would be needed in other languages.

You do not need to be an experienced programmer to learn Prolog. Some initial familiarity with basic computing concepts such as program, variable, constant and function would make it easier to achieve, but paradoxically too much experience of writing programs in other languages may make the task harder – it may be necessary to unlearn bad habits of thinking learnt elsewhere.

Some Technical Details

Experienced programmers will search this book in vain for such standard language features as variable declarations, subroutines, methods, for loops, while loops or assignment statements. (If you don't know what these terms mean, don't worry – you will not be needing them.)

On the other hand experienced readers may like to know that Prolog has a straightforward uniform syntax, programs that are equivalent to a database of facts and rules, a built-in theorem prover with automatic backtracking, list processing, recursion and facilities for modifying programs (or databases) at run-time. (You probably will not know what most of these mean either – but you will be using all of them by the end of this book.)

Prolog lends itself to a style of programming making particular use of two powerful techniques: recursion and list processing. In many cases algorithms that would require substantial amounts of coding in other languages can be implemented in a few lines in Prolog.

There are many versions of Prolog available for PC, Macintosh and Unix systems, including versions for Microsoft Windows, to link Prolog to an Oracle relational database and for use with 'object-oriented' program design. These range from commercial systems with many features to public domain and 'freeware' versions. Some of these are listed (in alphabetical order) below, together with web addresses at which more information can be found.

- Amzi! Prolog
  http://www.amzi.com/products/prolog_products.htm
- B-Prolog
  http://www.probp.com/
- Ciao Prolog
  http://clip.dia.fi.upm.es/Software/Ciao/
- GNU Prolog
  http://gnu-prolog.inria.fr/
- Logic Programming Associates Prolog (versions for Windows, DOS and Macintosh)
  http://www.lpa.co.uk
- Open Prolog (for Apple Macintosh)
  http://www.cs.tcd.ie/open-prolog/
- PD Prolog (a public domain version for MS-DOS only)
  http://www-2.cs.cmu.edu/afs/cs/project/
  ai-repository/ai/lang/prolog/impl/prolog/pdprolog/0.html
- SICStus Prolog
  http://www.sics.se/isl/sicstuswww/site/index.html

- SWI Prolog
  http://www.swi-prolog.org/
- Turbo Prolog (an old version that only runs in MS-DOS)
  http://www.fraber.de/university/prolog/tprolog.html
- Visual Prolog
  http://www.visual-prolog.com/
- W-Prolog (a Prolog-like language that runs in a web browser)
  http://goanna.cs.rmit.edu.au/~winikoff/wp/
- YAP Prolog
  http://www.ncc.up.pt/~vsc/Yap/

The programs in this book are all written using the standard 'Edinburgh syntax' and should run unchanged in virtually any version of Prolog you encounter (unfortunately, finding occasional subtle differences between implementations is one of the occupational hazards of learning any programming language). Features such as graphical interfaces, links to external databases etc. have deliberately not been included, as these generally vary from one implementation to another. All the examples given have been tested using Win-Prolog version 4.110 from the British company Logic Programming Associates (www.lpa.co.uk).

Each chapter has a number of self-assessment exercises to enable you to check your progress. A full glossary of the technical terms used completes the book.

Contents

1. Getting Started **1**

1.1 Starting Prolog 1
1.2 Prolog Programs 3
1.3 Data Objects in Prolog: Prolog Terms 8
Practical Exercise 1 12

2. Clauses and Predicates **13**

2.1 Clauses 13
2.2 Predicates 15
2.3 Loading Clauses 19
2.4 Variables 21
Practical Exercise 2 27

3. Satisfying Goals **29**

Introduction 29
3.1 Unification 31
3.1.1 Unifying Call Terms 32
3.2 Evaluating Goals 35
3.3 Backtracking 39
3.4 Satisfying Goals: A Summary 50
3.5 Removing Common Variables 52
3.6 A Note on Declarative Programming 52
Practical Exercise 3 55

4. Operators and Arithmetic **57**

4.1 Operators 57
4.2 Arithmetic 59
4.3 Equality Operators 63
4.4 Logical Operators 66
Practical Exercise 4 68

5. Input and Output **69**

Introduction 69
5.1 Outputting Terms 69
5.2 Inputting Terms 71
5.3 Input and Output Using Characters 72
5.4 Outputting Characters 73
5.5 Inputting Characters 73
5.6 Using Characters: Examples 74
5.7 Input and Output Using Files 77
5.8 File Output: Changing the Current Output Stream 77
5.9 File Input: Changing the Current Input Stream 78
5.9.1 Reading from Files: End of File 78
5.9.2 Reading from Files: End of Record 79
5.10 Using Files: Examples 79
Practical Exercise 5 82

6. Loops **85**

Introduction 85
6.1 Looping a Fixed Number of Times 85
6.2 Looping Until a Condition Is Satisfied 89
6.2.1 Recursion 89
6.2.2 Using the 'repeat' Predicate 91
6.3 Backtracking with Failure 94
6.3.1 Searching the Prolog Database 94
6.3.2 Finding Multiple Solutions 96
Practical Exercise 6 97

7. Preventing Backtracking **99**

Introduction 99
7.1 The Cut Predicate 99
7.2 Cut with Failure 106
Practical Exercise 7 108

8. Changing the Prolog Database **109**

8.1 Changing the Database: Adding and Deleting Clauses 109
8.2 Adding Clauses 110
8.3 Deleting Clauses 111
8.4 Changing the Database: Example 112
8.5 Maintaining a Database of Facts 114
Practical Exercise 8 118

9. List Processing **119**

9.1 Representing Data as Lists 119
9.2 Notation for Lists 120
9.3 Decomposing a List 122
9.4 Built-in Predicate: member 125
9.5 Built-in Predicate: length 126
9.6 Built-in Predicate: reverse 126
9.7 Built-in Predicate: append 127
9.8 List Processing: Examples 129
9.9 Using findall/3 to Create a List 133
Practical Exercise 9 136

10. String Processing **137**

10.1 Converting Strings of Characters To and From Lists 137
10.2 Joining Two Strings 139
10.3 Trimming a String 139
10.4 Inputting a String of Characters 141
10.5 Searching a String 142
10.6 Dividing a String into Its Component Parts 144
Practical Exercise 10 147

11. More Advanced Features **149**

Introduction 149
11.1 Extending Prolog: Arithmetic 149
11.2 Extending Prolog: Operations on Strings 155
11.3 Extending Prolog: Sets 157
11.4 Processing Terms 160
Practical Exercise 11 166

Appendix 1. Built-in Predicates **167**

Appendix 2. Built-in Operators **179**

Appendix 3. Specimen Solutions to Practical Exercises **185**

Appendix 4. Glossary **209**

Index **221**