

Artificial Intelligence

A Guide to Intelligent Systems

MICHAEL NEGNEVITSKY



Artificial Intelligence

A Guide to Intelligent Systems

Michael Negre

江苏工业学院图书馆
藏书章

 Addison-Wesley

An imprint of Pearson Education

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

Pearson Education Limited

Edinburgh Gate

Harlow

Essex CM20 2JE

England

and Associated Companies around the World.

Visit us on the World Wide Web at:

www.pearsoneduc.com

First edition 2002

© Pearson Education Limited 2002

The right of Michael Negnevitsky to be identified as author of this Work has been asserted by him in accordance with the Copyright, Designs and Patents Act 1988.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without either the prior written permission of the publisher or a licence permitting restricted copying in the United Kingdom issued by the Copyright Licensing Agency Ltd, 90 Tottenham Court Road, London W1T 4LP.

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

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Pearson Education has made every attempt to supply trademark information about manufacturers and their products mentioned in this book. A list of the trademark designations and their owners appears on page x.

ISBN 0201-71159-1

British Library Cataloguing-in-Publication Data

A catalogue record for this book can be obtained from the British Library

Library of Congress Cataloging-in-Publication Data

Negnevitsky, Michael.

Artificial intelligence : a guide to intelligent systems / Michael Negnevitsky.

p. cm.

Includes bibliographical references and index.

ISBN 0-201-71159-1

1. Expert systems (Computer science) 2. Artificial intelligence. I. Title.

QA76.76.E95 N445 2001

006.3--dc21

2001022092

10 9 8 7 6 5 4

05 04 03 02 01

Typeset by 43 in 9/12pt Stone Serif

Printed and bound in Great Britain by

Biddles Ltd, Guildford and King's Lynn

Artificial Intelligence



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

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

To find out more about the complete range of our publishing please visit us on the World Wide Web at: www.pearsoneduc.com

For my son, Vlad

Preface

'The only way not to succeed is not to try.'

Edward Teller

Another book on artificial intelligence ... I've already seen so many of them. Why should I bother with this one? What makes this book different from the others?

Each year hundreds of books and doctoral theses extend our knowledge of computer, or artificial, intelligence. Expert systems, artificial neural networks, fuzzy systems and evolutionary computation are major technologies used in intelligent systems. Hundreds of tools support these technologies, and thousands of scientific papers continue to push their boundaries. The contents of any chapter in this book can be, and in fact is, the subject of dozens of monographs. However, I wanted to write a book that would explain the basics of intelligent systems, and perhaps even more importantly, eliminate the fear of artificial intelligence.

Most of the literature on artificial intelligence is expressed in the jargon of computer science, and crowded with complex matrix algebra and differential equations. This, of course, gives artificial intelligence an aura of respectability, and until recently kept non-computer scientists at bay. But the situation has changed!

The personal computer has become indispensable in our everyday life. We use it as a typewriter and a calculator, a calendar and a communication system, an interactive database and a decision-support system. And we want more. We want our computers to act intelligently! We see that intelligent systems are rapidly coming out of research laboratories, and we want to use them to our advantage.

What are the principles behind intelligent systems? How are they built? What are intelligent systems useful for? How do we choose the right tool for the job? These questions are answered in this book.

Unlike many books on computer intelligence, this one shows that most ideas behind intelligent systems are wonderfully simple and straightforward. The book is based on lectures given to students who have little knowledge of calculus. And readers do not need to learn a programming language! The material in this book has been extensively tested through several courses taught by the author for the

past decade. Typical questions and suggestions from my students influenced the way this book was written.

The book is an introduction to the field of computer intelligence. It covers rule-based expert systems, fuzzy expert systems, frame-based expert systems, artificial neural networks, evolutionary computation, hybrid intelligent systems and knowledge engineering.

In a university setting, this book provides an introductory course for undergraduate students in computer science, computer information systems, and engineering. In the course I teach at the University of Tasmania, my students develop small rule-based and frame-based expert systems, design a fuzzy system, explore artificial neural networks, and implement a simple problem as a genetic algorithm. They use expert system shells (Leonardo and Level5 Object), MATLAB Fuzzy Logic Toolbox and MATLAB Neural Network Toolbox. I chose these tools because they can easily demonstrate the theory being presented. However, the book is not tied to any specific tool; the examples given in the book are easy to implement with different tools.

This book is also suitable as a self-study guide for non-computer science professionals. For them, the book provides access to the state of the art in knowledge-based systems and computational intelligence. In fact, this book is aimed at a large professional audience: engineers and scientists, managers and businessmen, doctors and lawyers – everyone who faces challenging problems and cannot solve them by using traditional approaches, everyone who wants to understand the tremendous achievements in computer intelligence. The book will help to develop a practical understanding of what intelligent systems can and cannot do, discover which tools are most relevant for your task and, finally, how to use these tools.

The book consists of nine chapters.

In Chapter 1, we briefly discuss the history of artificial intelligence from the era of great ideas and great expectations in the 1960s to the disillusionment and funding cutbacks in the early 1970s; from the development of the first expert systems such as DENDRAL, MYCIN and PROSPECTOR in the seventies to the maturity of expert system technology and its massive applications in different areas in the 1980s and 1990s; from a simple binary model of neurons proposed in the 1940s to a dramatic resurgence of the field of artificial neural networks in the 1980s; from the introduction of fuzzy set theory and its being ignored by the West in the 1960s to numerous ‘fuzzy’ consumer products offered by the Japanese in the 1980s and world-wide acceptance of ‘soft’ computing and computing with words in the 1990s.

In Chapter 2, we present an overview of rule-based expert systems. We briefly discuss what knowledge is, and how experts express their knowledge in the form of production rules. We identify the main players in the expert system development team and show the structure of a rule-based system. We discuss fundamental characteristics of expert systems and note that expert systems can make mistakes. Then we review the forward and backward chaining inference techniques and debate conflict resolution strategies. Finally, the advantages and disadvantages of rule-based expert systems are examined.

In Chapter 3, we present two uncertainty management techniques used in expert systems: Bayesian reasoning and certainty factors. We identify the main sources of uncertain knowledge and briefly review probability theory. We consider the Bayesian method of accumulating evidence and develop a simple expert system based on the Bayesian approach. Then we examine the certainty factors theory (a popular alternative to Bayesian reasoning) and develop an expert system based on evidential reasoning. Finally, we compare Bayesian reasoning and certainty factors, and determine appropriate areas for their applications.

In Chapter 4, we introduce fuzzy logic and discuss the philosophical ideas behind it. We present the concept of fuzzy sets, consider how to represent a fuzzy set in a computer, and examine operations of fuzzy sets. We also define linguistic variables and hedges. Then we present fuzzy rules and explain the main differences between classical and fuzzy rules. We explore two fuzzy inference techniques – Mamdani and Sugeno – and suggest appropriate areas for their application. Finally, we introduce the main steps in developing a fuzzy expert system, and illustrate the theory through the actual process of building and tuning a fuzzy system.

In Chapter 5, we present an overview of frame-based expert systems. We consider the concept of a frame and discuss how to use frames for knowledge representation. We find that inheritance is an essential feature of frame based systems. We examine the application of methods, demons and rules. Finally, we consider the development of a frame-based expert system through an example.

In Chapter 6, we introduce artificial neural networks and discuss the basic ideas behind machine learning. We present the concept of a perceptron as a simple computing element and consider the perceptron learning rule. We explore multilayer neural networks and discuss how to improve the computational efficiency of the back-propagation learning algorithm. Then we introduce recurrent neural networks, consider the Hopfield network training algorithm and bidirectional associative memory (BAM). Finally, we present self-organising neural networks and explore Hebbian and competitive learning.

In Chapter 7, we present an overview of evolutionary computation. We consider genetic algorithms, evolution strategies and genetic programming. We introduce the main steps in developing a genetic algorithm, discuss why genetic algorithms work, and illustrate the theory through actual applications of genetic algorithms. Then we present a basic concept of evolutionary strategies and determine the differences between evolutionary strategies and genetic algorithms. Finally, we consider genetic programming and its application to real problems.

In Chapter 8, we consider hybrid intelligent systems as a combination of different intelligent technologies. First we introduce a new breed of expert systems, called neural expert systems, which combine neural networks and rule-based expert systems. Then we consider a neuro-fuzzy system that is functionally equivalent to the Mamdani fuzzy inference model, and an adaptive neuro-fuzzy inference system (ANFIS), equivalent to the Sugeno fuzzy inference

model. Finally, we discuss evolutionary neural networks and fuzzy evolutionary systems.

In Chapter 9, we consider knowledge engineering and data mining. First we discuss what kind of problems can be addressed with intelligent systems and introduce six main phases of the knowledge engineering process. Then we study typical applications of intelligent systems, including diagnosis, classification, decision support, pattern recognition and prediction. Finally, we examine an application of decision trees in data mining.

The book also has an appendix and a glossary. The appendix provides a list of commercially available AI tools. The glossary contains definitions of over 250 terms used in expert systems, fuzzy logic, neural networks, evolutionary computation, knowledge engineering and data mining.

I hope that the reader will share my excitement on the subject of artificial intelligence and soft computing and will find this book useful.

The website can be accessed at: <http://www.booksites.net/negnevitsky>

*Michael Negnevitsky
Hobart, Tasmania, Australia*

Acknowledgements

I am deeply indebted to many individuals who, directly or indirectly, are responsible for this book coming into being. I am most grateful to Dr Vitaly Faybisovich for his constructive criticism of my research on soft computing, and most of all for his friendship and support in all my endeavours for the last twenty years.

I am also very grateful to numerous reviewers of my book for their comments and helpful suggestions; and to the Pearson Education editors, particularly Keith Mansfield, who led me through the process of publishing this book.

I also thank my undergraduate and postgraduate students from the University of Tasmania – especially my former Ph.D. students Tan Loc Le, Quang Ha and Steven Carter, whose desire for new knowledge has been both a challenge and an inspiration to me.

I am indebted to Professor Stephen Grossberg from Boston University, Professor Frank Palis from the Otto-von-Guericke-Universität Magdeburg, Germany and Professor Hiroshi Sasaki from Hiroshima University, Japan for giving me the opportunity to test the book's material on their students.

Last, but by no means least, I am truly grateful to Dr Vivienne Mawson and Margaret Eldridge for proof-reading the draft text.

Trademark notice

The following are trademarks or registered trademarks of their respective companies:

KnowledgeSEEKER is a trademark of Angoss Software Corporation; Outlook and Windows are trademarks of Microsoft Corporation; MATLAB is a trademark of The MathWorks, Inc; Unix is a trademark of the Open Group.

See Appendix for AI tools and their respective vendors.

Contents

Preface	xi
Acknowledgements	xv
1 Introduction to knowledge-based intelligent systems	1
1.1 Intelligent machines, or what machines can do	1
1.2 The history of artificial intelligence, or from the ‘Dark Ages’ to knowledge-based systems	4
1.3 Summary	17
Questions for review	21
References	22
2 Rule-based expert systems	25
2.1 Introduction, or what is knowledge?	25
2.2 Rules as a knowledge representation technique	26
2.3 The main players in the expert system development team	28
2.4 Structure of a rule-based expert system	30
2.5 Fundamental characteristics of an expert system	33
2.6 Forward chaining and backward chaining inference techniques	35
2.7 THERMOSTAT: a demonstration rule-based expert system	41
2.8 Conflict resolution	46
2.9 Advantages and disadvantages of rule-based expert systems	49
2.10 Summary	51
Questions for review	53
References	53
3 Uncertainty management in rule-based expert systems	55
3.1 Introduction, or what is uncertainty?	55
3.2 Basic probability theory	57
3.3 Bayesian reasoning	61
3.4 FORECAST: Bayesian accumulation of evidence	65

3.5	Bias of the Bayesian method	72
3.6	Certainty factors theory and evidential reasoning	74
3.7	FORECAST: an application of certainty factors	80
3.8	Comparison of Bayesian reasoning and certainty factors	82
3.9	Summary	83
	Questions for review	85
	References	85
4	Fuzzy expert systems	87
4.1	Introduction, or what is fuzzy thinking?	87
4.2	Fuzzy sets	89
4.3	Linguistic variables and hedges	94
4.4	Operations of fuzzy sets	97
4.5	Fuzzy rules	103
4.6	Fuzzy inference	106
4.7	Building a fuzzy expert system	114
4.8	Summary	125
	Questions for review	126
	References	127
	Bibliography	127
5	Frame-based expert systems	129
5.1	Introduction, or what is a frame?	129
5.2	Frames as a knowledge representation technique	131
5.3	Inheritance in frame-based systems	136
5.4	Methods and demons	140
5.5	<i>Interaction of frames and rules</i>	144
5.6	Buy Smart: a frame-based expert system	147
5.7	Summary	159
	Questions for review	161
	References	161
	Bibliography	162
6	Artificial neural networks	163
6.1	Introduction, or how the brain works	163
6.2	The neuron as a simple computing element	166
6.3	The perceptron	168
6.4	Multilayer neural networks	173
6.5	<i>Accelerated learning in multilayer neural networks</i>	183
6.6	The Hopfield network	186
6.7	Bidirectional associative memory	194
6.8	Self-organising neural networks	198
6.9	Summary	210
	Questions for review	213
	References	214

7	Evolutionary computation	217
7.1	Introduction, or can evolution be intelligent?	217
7.2	Simulation of natural evolution	217
7.3	Genetic algorithms	220
7.4	Why genetic algorithms work	230
7.5	Case study: maintenance scheduling with genetic algorithms	233
7.6	Evolution strategies	240
7.7	Genetic programming	243
7.8	Summary	252
	Questions for review	253
	References	254
8	Hybrid intelligent systems	257
8.1	Introduction, or how to combine German mechanics with Italian love	257
8.2	Neural expert systems	259
8.3	Neuro-fuzzy systems	266
8.4	ANFIS: Adaptive Neuro-Fuzzy Inference System	275
8.5	Evolutionary neural networks	283
8.6	Fuzzy evolutionary systems	288
8.7	Summary	294
	Questions for review	295
	References	296
9	Knowledge engineering and data mining	299
9.1	Introduction, or what is knowledge engineering?	299
9.2	Will an expert system work for my problem?	306
9.3	Will a fuzzy expert system work for my problem?	315
9.4	Will a neural network work for my problem?	321
9.5	Data mining and knowledge discovery	330
9.6	Summary	341
	Questions for review	342
	References	343
	Glossary	345
	Appendix	371
	Index	387

Introduction to knowledge-based intelligent systems

In which we consider what it means to be intelligent and whether machines could be such a thing.

1.1 Intelligent machines, or what machines can do

Philosophers have been trying for over two thousand years to understand and resolve two big questions of the universe: how does a human mind work, and can non-humans have minds? However, these questions are still unanswered.

Some philosophers have picked up the computational approach originated by computer scientists and accepted the idea that machines can do everything that humans can do. Others have openly opposed this idea, claiming that such highly sophisticated behaviour as love, creative discovery and moral choice will always be beyond the scope of any machine.

The nature of philosophy allows for disagreements to remain unresolved. In fact, engineers and scientists have already built machines that we can call 'intelligent'. So what does the word 'intelligence' mean? Let us look at a dictionary definition.

- 1 Someone's **intelligence** is their ability to understand and learn things.
- 2 **Intelligence** is the ability to think and understand instead of doing things by instinct or automatically.

(Essential English Dictionary, Collins, London, 1990)

Thus, according to the first definition, intelligence is the quality possessed by humans. But the second definition suggests a completely different approach and gives some flexibility; it does not specify whether it is **someone** or **something** that has the ability to think and understand. Now we should discover what thinking means. Let us consult our dictionary again.

Thinking is the activity of using your brain to consider a problem or to create an idea.

(Essential English Dictionary, Collins, London, 1990)

So, in order to think, **someone** or **something** has to have a brain, or in other words, an organ that enables **someone** or **something** to learn and understand things, to solve problems and to make decisions. So we can define intelligence as 'the ability to learn and understand, to solve problems and to make decisions'.

The very question that asks whether computers can be intelligent, or whether machines can think, came to us from the 'dark ages' of artificial intelligence (from the late 1940s). The goal of **artificial intelligence** (AI) as a science is to make machines do things that would require intelligence if done by humans (Boden, 1977). Therefore, the answer to the question 'Can machines think?' was vitally important to the discipline. However, the answer is not a simple 'Yes' or 'No', but rather a vague or **fuzzy** one. Your everyday experience and common sense would have told you that. Some people are smarter in some ways than others. Sometimes we make very intelligent decisions but sometimes we also make very silly mistakes. Some of us deal with complex mathematical and engineering problems but are moronic in philosophy and history. Some people are good at making money, while others are better at spending it. As humans, we all have the ability to learn and understand, to solve problems and to make decisions; however, our abilities are not equal and lie in different areas. Therefore, we should expect that if machines can think, some of them might be smarter than others in some ways.

One of the earliest and most significant papers on machine intelligence, 'Computing machinery and intelligence', was written by the British mathematician Alan Turing over fifty years ago (Turing, 1950). However, it has stood up well to the test of time, and Turing's approach remains universal.

Alan Turing began his scientific career in the early 1930s by rediscovering the Central Limit Theorem. In 1937 he wrote a paper on computable numbers, in which he proposed the concept of a universal machine. Later, during the Second World War, he was a key player in deciphering Enigma, the German military encoding machine. After the war, Turing designed the 'Automatic Computing Engine'. He also wrote the first program capable of playing a complete chess game; it was later implemented on the Manchester University computer. Turing's theoretical concept of the universal computer and his practical experience in building code-breaking systems equipped him to approach the key fundamental question of artificial intelligence. He asked: Is there thought without experience? Is there mind without communication? Is there language without living? Is there intelligence without life? All these questions, as you can see, are just variations on the fundamental question of artificial intelligence, Can machines think?

Turing did not provide definitions of machines and thinking, he just avoided semantic arguments by inventing a game, the **Turing imitation game**. Instead of asking, 'Can machines think?', Turing said we should ask, 'Can machines pass a behaviour test for intelligence?' He predicted that by the year 2000, a computer could be programmed to have a conversation with a human interrogator for five minutes and would have a 30 per cent chance of deceiving the interrogator that it was a human. Turing defined the intelligent behaviour of a computer as the ability to achieve the human-level performance in cognitive tasks. In other