HZ Books
华章教育

经典原版书库

人工智能

智能系统指南

（澳）Michael Negnevitsky 著
塔斯马尼亚大学

（英文版·第3版）

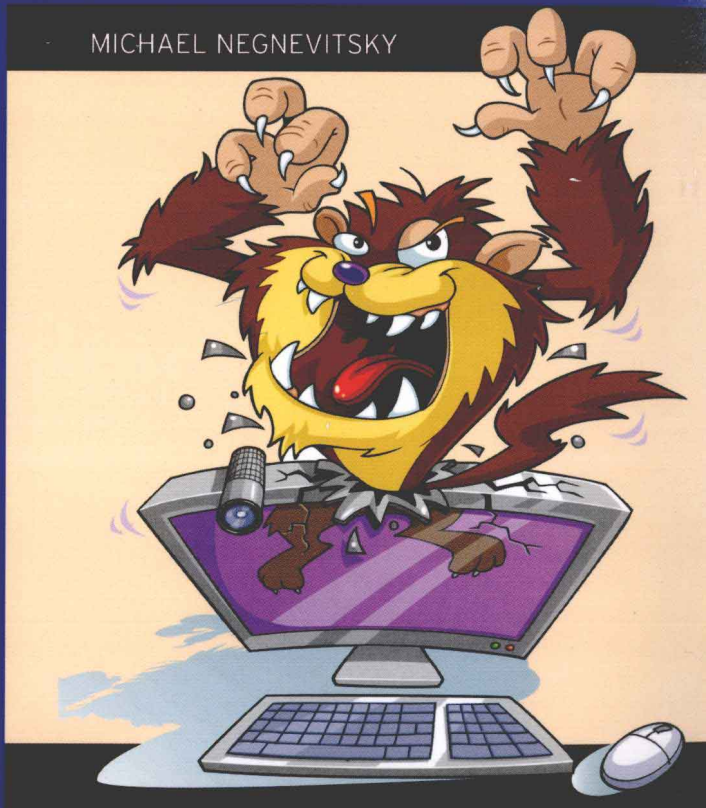

ARTIFICIAL INTELLIGENCE

A GUIDE TO INTELLIGENT SYSTEMS

机械工业出版社
ina Machine Press

经 典 原 版 书 库

人工智能

智能系统指南

（英文版·第3版）

Artificial Intelligence

A Guide to Intelligent Systems (Third Edition)

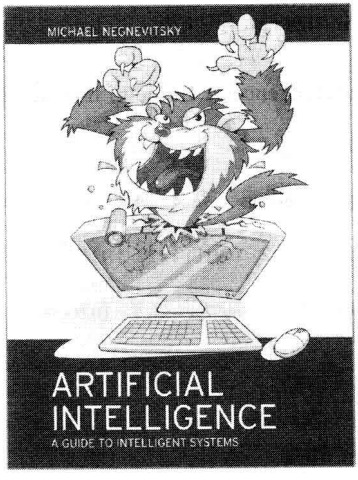

（澳）Michael Negnevitsky 著
塔斯马尼亚大学

机械工业出版社
China Machine Press

Michael Negnevitsky: Artificial Intelligence: A Guide to Intelligent Systems, Third Edition (ISBN 978-1-4082-2574-5).

本书版权登记号：图字：01-2011-4256

图书在版编目（CIP）数据

人工智能：智能系统指南（英文版·第 3 版）/（澳）尼格尼维斯基（Negnevitsky, M.）著. —北京：机械工业出版社，2011.9

（经典原版书库）

书名原文：Artificial Intelligence: A Guide to Intelligent Systems, Third Edition

ISBN 978-7-111-35822-0

I. 人… II. 尼… III. 人工智能－英文 IV. TP18

中国版本图书馆 CIP 数据核字（2011）第 181825 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

北京京师印务有限公司印刷

2011 年 9 月第 1 版第 1 次印刷

150mm × 214mm • 15.5 印张

标准书号：ISBN 978-7-111-35822-0

定价：49.00 元

凡购本书，如有缺页、倒页、脱页，由本社发行部调换

客服热线：(010) 88378991；88361066

购书热线：(010) 68326294；88379649；68995259

投稿热线：(010) 88379604

读者信箱：hzjsj@hzbook.com

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，还揭示了学术的源变，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和发展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到积极的推动作用，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公司较早意识到“出版要为教育服务”。自 1998 年开始，我们就将工作重点放在了遴选、移译国外优秀教材上。经过多年的不懈努力，我们与 Pearson，McGraw-Hill，Elsevier，MIT，John Wiley & Sons，Cengage 等世界著名出版公司建立了良好的合作关系，从他们现有的数百种教材中甄选出 Andrew S. Tanenbaum，Bjarne Stroustrup，Brain W. Kernighan，Dennis Ritchie，Jim Gray，Afred V. Aho，John E. Hopcroft，Jeffrey D. Ullman，Abraham Silberschatz，William Stallings，Donald E. Knuth，John L. Hennessy，Larry L. Peterson 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、研究及珍藏。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”已经出版了近两百个品种，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。其影印版“经典原版书库”作为姊妹篇也被越来越多实施双语教学的学校所采用。

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。随着计算机科学与技术专业学科建设的不断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将步入一个新的阶段，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。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，我们的联系方法如下：

华章网站：www.hzbook.com
电子邮件：hzjsj@hzbook.com
联系电话：(010) 88379604
联系地址：北京市西城区百万庄南街1号
邮政编码：100037

华章教育
华章科技图书出版中心

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 content 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

last 15 years. 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, knowledge engineering and data mining.

In a university setting, this book provides an introductory course for undergraduate students in computer science, computer information systems, and engineering. In the courses I teach, my students develop small rule-based and frame-based expert systems, design fuzzy systems, explore artificial neural networks, solve simple optimisation problems using genetic algorithms and develop hybrid neuro-fuzzy systems. They use expert system shells (XpertRule, Exsys Corvid and Visual Rule Studio), 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.

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
February 2001

Preface to the third edition

The main objective of the book remains the same as in the first edition – to provide the reader with practical understanding of the field of computer intelligence. It is intended as an introductory text suitable for a one-semester course, and assumes the students have only limited knowledge of calculus and little or no programming experience.

In terms of the coverage, this edition introduces a new chapter on data mining and demonstrates several new applications of intelligent tools for solving complex real-world problems. The major changes are as follows:

- In the new chapter, 'Data mining and knowledge discovery', we introduce data mining as an integral part of knowledge discovery in large databases. We consider the main techniques and tools for turning data into knowledge, including statistical methods, data visualisation tools, Structured Query Language, decision trees and market basket analysis. We also present several case studies on data mining applications.
- In Chapter 9, we add a new case study on clustering with a self-organising neural network.

Finally, we have expanded the book's references and bibliographies, and updated the list of AI tools and vendors in the appendix.

Michael Negnevitsky
Hobart, Tasmania, Australia
September 2010

Overview of the book

The book consists of 10 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 1970s to the maturity of expert system technology and its massive application 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 evolution strategies and determine the differences between evolution 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. 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 examine typical applications of expert systems, fuzzy systems, neural networks and genetic algorithms. We demonstrate how to build intelligent systems for solving diagnosis, selection, prediction, classification, clustering and optimisation problems. Finally, we discuss applications of hybrid neuro-fuzzy systems for decision support and time-series prediction.

In Chapter 10, we present an overview of data mining and consider the main techniques for turning data into knowledge. First, we broadly define data mining, and explain the process of data mining and knowledge discovery in large databases. We introduce statistical methods, including principal component analysis, and discuss their limitations. We then examine an application of Structured Query Language in relational databases, and introduce data warehouse and multidimensional data analysis. Finally, we consider the most popular tools of data mining – decision trees and market basket analysis.

The book also has a glossary and an appendix. The glossary contains definitions for over 300 terms used in expert systems, fuzzy logic, neural networks, evolutionary computation, knowledge engineering and data mining. The appendix provides a list of commercially available AI tools.

The book's website can be accessed at: http://www.booksites.net/negnevitsky

Acknowledgements

I am deeply indebted to many people 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 30 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, Owen Knight, Liz Johnson and Rufus Curnow, who led me through the process of publishing and editing 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, Steven Carter and Mark Lim, whose desire for new knowledge was both a challenge and an inspiration to me.

I am indebted to Professor Stephen Grossberg from Boston University, USA, Professor Frank Palis from the Otto-von-Guericke-Universität Magdeburg, Germany, Professor Hiroshi Sasaki from Hiroshima University, Japan, Professor Walter Wolf from the Rochester Institute of Technology, USA, and Professor Kaoru Hirota from Tokyo Institute of Technology, Japan, for giving me the opportunity to test the book's material on their students.

I am also truly grateful to Dr Vivienne Mawson and Margaret Eldridge for proof-reading the draft text.

Although the first edition of this book appeared a few years ago, I cannot possibly thank all the people who have already used it and sent me their comments. However, I must acknowledge at least those who made especially helpful suggestions: Martin Beck (University of Plymouth, UK), Mike Brooks (University of Adelaide, Australia), Genard Catalano (Columbia College, USA), Warren du Plessis (University of Pretoria, South Africa), Salah Amin Elewa (American University, Egypt), Michael Fang (Zhejiang University, China), John Fronckowiak (Medaille College, USA), Patrick B. Gibson (University of Windsor, Canada), Lev Goldfarb (University of New Brunswick, Canada), Susan Haller (University of Wisconsin, USA), Evor Hines (University of Warwick, UK), Philip Hingston (Edith Cowan University, Australia), Sam Hui (Stanford University, USA), Yong-Hyuk Kim (Kwangwoon University, Korea), David Lee (University of Hertfordshire, UK), Andrew Nunekpeku (University of Ghana), Vasile Palade (Oxford University, UK), Leon Reznik (Rochester Institute of Technology, USA), Simon Shiu (Hong Kong Polytechnic University), Boris Stilman (University of

Colorado, USA), Thomas Uthmann (Johannes Gutenberg-Universität Mainz, Germany), Anne Venables (Victoria University, Australia), Brigitte Verdonk (University of Antwerp, Belgium), Ken Vollmar (Southwest Missouri State University, USA), Kok Wai Wong (Nanyang Technological University, Singapore) and Georgios N. Yannakakis (IT-University of Copenhagen, Denmark).

Contents

Introduction to knowledge-based intelligent systems

1

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 2000 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, 2008)

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 some**one** or some**thing** 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, 2008)