

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!

Personal computers have become indispensable in our everyday life. We use them 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 courses I teach, 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, XpertRule, Level5 Object 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.

The book consists of ten chapters.

In Chapter 1, we briefly discuss the history of artificial intelligence from the era of great ideas and great expectations in the sixties to the disillusionment and funding cutbacks in the early seventies; 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 eighties and nineties; from a simple binary model of neurons proposed in the forties to a dramatic resurgence of the field of artificial neural networks in the eighties; from the introduction of fuzzy set theory and its being ignored by the West in the sixties to numerous ‘fuzzy’ consumer products offered by the Japanese in the eighties and world-wide acceptance of ‘soft’ computing and computing with words in the nineties.

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 the frame-based expert systems. We consider the concept of a frame and discussed how to use frames for knowledge representation. We find that inheritance is an essential feature of the 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 address 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 tolls. 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 with the subject of artificial intelligence and soft computing and 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 second 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 no programming experience.

In terms of the coverage, in this edition we demonstrate several new applications of intelligent tools for solving specific problems. The changes are in the following chapters:

- In Chapter 2, we introduce a new demonstration rule-based expert system, MEDIA ADVISOR.
- In Chapter 9, we add a new case study on classification neural networks with competitive learning.
- In Chapter 9, we introduce a section: ‘Will genetic algorithms work for my problem?’. The section includes a case study with the travelling salesman problem.
- Also in Chapter 9, we add a new section ‘Will a hybrid intelligent system work for my problem?’. This section includes two case studies: the first covers a neuro-fuzzy decision-support system with a heterogeneous structure, and the second explores an Adaptive Neuro-Fuzzy Inference System (ANFIS) with a homogeneous structure.

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

*Michael Negnevitsky
Hobart, Tasmania, Australia
January 2004*