


ARTIFICIAL INTELLIGENCE



RICHARD E. NEAPOLITAN
XIA JIANG



CRC Press
Taylor & Francis Group

A CHAPMAN & HALL BOOK

CONTEMPORARY ARTIFICIAL INTELLIGENCE

RICHARD E.



XIA JIANG



CRC Press

Taylor & Francis Group

Boca Raton London New York

CRC Press is an imprint of the
Taylor & Francis Group, an Informa business

A CHAPMAN & HALL BOOK

CRC Press
Taylor & Francis Group
6000 Broken Sound Parkway NW, Suite 300
Boca Raton, FL 33487-2742

© 2013 by Taylor & Francis Group, LLC
CRC Press is an imprint of Taylor & Francis Group, an Informa business

No claim to original U.S. Government works

Printed in the United States of America on acid-free paper
Version Date: 20120612

International Standard Book Number: 978-1-4398-4469-4 (Hardback)

This book contains information obtained from authentic and highly regarded sources. Reasonable efforts have been made to publish reliable data and information, but the author and publisher cannot assume responsibility for the validity of all materials or the consequences of their use. The authors and publishers have attempted to trace the copyright holders of all material reproduced in this publication and apologize to copyright holders if permission to publish in this form has not been obtained. If any copyright material has not been acknowledged please write and let us know so we may rectify in any future reprint.

Except as permitted under U.S. Copyright Law, no part of this book may be reprinted, reproduced, transmitted, or utilized in any form by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying, microfilming, and recording, or in any information storage or retrieval system, without written permission from the publishers.

For permission to photocopy or use material electronically from this work, please access www.copyright.com (<http://www.copyright.com/>) or contact the Copyright Clearance Center, Inc. (CCC), 222 Rosewood Drive, Danvers, MA 01923, 978-750-8400. CCC is a not-for-profit organization that provides licenses and registration for a variety of users. For organizations that have been granted a photocopy license by the CCC, a separate system of payment has been arranged.

Trademark Notice: Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

Library of Congress Cataloging-in-Publication Data

Neapolitan, Richard E.

Contemporary artificial intelligence / Richard E. Neapolitan, Xia Jiang.
p. cm.

Includes bibliographical references and index.

ISBN 978-1-4398-4469-4 (hardback)

1. Artificial intelligence. I. Jiang, Xia, 1967- II. Title.

Q335.N425 2012

006.3--dc23

2012017897

Visit the Taylor & Francis Web site at
<http://www.taylorandfrancis.com>

and the CRC Press Web site at
<http://www.crcpress.com>

CONTEMPORARY ARTIFICIAL INTELLIGENCE

Preface

Over the years my view of an artificial intelligence (AI) course has changed significantly. I used to view it as a course that should discuss our efforts to develop an artificial entity that can learn and make decisions in a complex, changing environment, affect that environment, and communicate its knowledge and choices to humans; that is, an entity that can think. I would therefore, cover the weak AI methods that failed to scale up. However, as strong methods that solved challenging problems in limited domains became more predominant, my course increasingly concerned these methods. I would cover backward chaining, forward chaining, planning, inference in Bayesian networks, normative decision analysis, evolutionary computation, decision tree learning, Bayesian network learning, supervised and unsupervised learning, and reinforcement learning. I would show useful applications of these methods. These techniques have come to be as important to a computer science student's repertoire as techniques such as divide-and-conquer, greedy methods, branch-and-bound, etc. Yet a student would not see them unless the student took an AI course. So my AI course evolved into a course that undergraduate students would take either concurrently or following an analysis of algorithms course, and would cover what I viewed as important problem-solving strategies that have emerged from the field of AI. I feel such a course should be a standard component of every computer science curriculum just like data structures and analysis of algorithms.

No text satisfied my needs for the course I taught for two reasons:

1. AI is a vast field that has included the development of many and varied techniques in the past 50 years. Current texts tried to include most of what has been going on rather than simply providing useful methods and algorithms.
2. No current text was accessible to students at a mainstream university like Northeastern Illinois University. I had this same problem with my analysis of algorithms course, and was the reason I wrote *Foundations of Algorithms*.

So I taught the course using my own Bayesian network texts and class notes. I finally decided to turn these notes into this textbook so professors at other mainstream universities could provide a similar course. This text is not meant to be an encyclopedia or history of AI, but rather is meant to be a text that

can be covered in one semester; and, in the amount of time one semester allows, provide the student with what I consider the most useful techniques that have emerged from AI. These techniques include the following:

1. Logic-based methods
2. Probability-based methods
3. Evolutionary computation and methods based on swarm intelligence
4. Language understanding

The text clearly reflects my own bias. I have not discussed neural networks, fuzzy logic, support vector machines, and many other endeavors in AI. I also do not include searching because most searching techniques appear in data structures and algorithms texts. Almost half the text is about probabilistic methods. This is perhaps partially due to the fact that I know these methods best because they are my own area of research, but also due to the fact that I view them as most important (that is why they are my area of research).

I have written the material in the order I teach it. So I recommend simply covering the chapters from first to last. If there is not time to cover the entire book, I recommend leaving out Sections 8.3 through 8.7 where I have explored more advanced topics in decision analysis, and Section 12.3 which concerns causal learning. Sections marked with a ★ contain material that is inherently more difficult than the other material in the text. However, they do cover important topics and should not be skipped if the students have the sophistication to grasp them.

I thank Dawn Holmes and Kevin Korb for reading the manuscript and providing useful comments. I also thank Prentice Hall for allowing me to include excerpts from my text *Learning Bayesian Networks*, and Morgan Kaufmann for allowing me to include excerpts from my texts *Probabilistic Methods for Financial and Marketing Informatics* and *Probabilistic Methods for Bioinformatics*. Finally, I thank Relu I. Jianu for designing an attractive and thought-provoking cover.

Richard E. Neapolitan
RE-Neapolitan@neu.edu

About the Authors



Richard E. Neapolitan is a professor of computer science at Northeastern Illinois University. His research interests include probability and statistics, expert systems, cognitive science, and applications of probabilistic modeling to fields such as medicine, biology, and finance. Dr. Neapolitan is a prolific author and has published in the most prestigious journals in the broad area of reasoning under uncertainty. He has previously written five books, including the seminal 1990 Bayesian network text *Probabilistic Reasoning in Expert Systems; Learning Bayesian Networks* (2004); *Foundations of Algorithms* (1996, 1998, 2003, 2010), which has been translated into three languages; *Probabilistic Methods for Financial and Marketing Informatics* (2007); and *Probabilistic Methods for Bioinformatics* (2009). His approach to textbook writing is innovative; his books have the reputation for making difficult concepts easy to understand while still remaining thought-provoking.



Xia Jiang is an assistant professor in the Department of Biomedical Informatics at the University of Pittsburgh. She has over 13 years of teaching and research experience in Bayesian network modeling, machine learning, and algorithm design. She is the co-author of the book *Probabilistic Methods for Financial and Marketing Informatics* (2007). Dr. Jiang is currently focusing on developing novel algorithms/systems that improve the computational efficiency of large data analysis, and network modeling of cancer genome data.

Contents

Preface	xi
About the Authors	xiii
1 Introduction to Artificial Intelligence	1
1.1 History of Artificial Intelligence	2
1.1.1 What Is Artificial Intelligence?	2
1.1.2 Emergence of AI	4
1.1.3 Cognitive Science and AI	4
1.1.4 Logical Approach to AI	5
1.1.5 Knowledge-Based Systems	6
1.1.6 Probabilistic Approach to AI	7
1.1.7 Evolutionary Computation and Swarm Intelligence	7
1.1.8 A Return to Creating HAL	8
1.2 Contemporary Artificial Intelligence	8
I Logical Intelligence	11
2 Propositional Logic	13
2.1 Basics of Propositional Logic	14
2.1.1 Syntax	14
2.1.2 Semantics	16
2.1.3 Tautologies and Logical Implication	20
2.1.4 Logical Arguments	22
2.1.5 Derivation Systems	25
2.2 Resolution	28
2.2.1 Normal Forms	28
2.2.2 Derivations Using Resolution	31
2.2.3 Resolution Algorithm	34
2.3 Artificial Intelligence Applications	35
2.3.1 Knowledge-Based Systems	35
2.3.2 Wumpus World	48
2.4 Discussion and Further Reading	56

3	First-Order Logic	61
3.1	Basics of First-Order Logic	62
3.1.1	Syntax	62
3.1.2	Semantics	65
3.1.3	Validity and Logical Implication	69
3.1.4	Derivation Systems	71
3.1.5	Modus Ponens for First-Order Logic	75
3.2	Artificial Intelligence Applications	79
3.2.1	Wumpus World Revisited	79
3.2.2	Planning	80
3.3	Discussion and Further Reading	84
4	Certain Knowledge Representation	89
4.1	Taxonomic Knowledge	90
4.1.1	Semantic Nets	90
4.1.2	Model of Human Organization of Knowledge	92
4.2	Frames	92
4.2.1	Frame Data Structure	92
4.2.2	Planning a Trip Using Frames	93
4.3	Nonmonotonic Logic	96
4.3.1	Circumscription	96
4.3.2	Default Logic	97
4.3.3	Difficulties	99
4.4	Discussion and Further Reading	99
II	Probabilistic Intelligence	101
5	Probability	103
5.1	Probability Basics	106
5.1.1	Probability Spaces	106
5.1.2	Conditional Probability and Independence	109
5.1.3	Bayes' Theorem	112
5.2	Random Variables	113
5.2.1	Probability Distributions of Random Variables	113
5.2.2	Independence of Random Variables	118
5.3	Meaning of Probability	122
5.3.1	Relative Frequency Approach to Probability	122
5.3.2	Subjective Approach to Probability	125
5.4	Random Variables in Applications	127
5.5	Probability in the Wumpus World	131
6	Uncertain Knowledge Representation	137
6.1	Intuitive Introduction to Bayesian Networks	139
6.2	Properties of Bayesian Networks	142
6.2.1	Definition of a Bayesian Network	142
6.2.2	Representation of a Bayesian Network	145

6.3	Causal Networks as Bayesian Networks	148
6.3.1	Causality	148
6.3.2	Causality and the Markov Condition	149
6.3.3	Markov Condition without Causality	154
6.4	Inference in Bayesian Networks	154
6.4.1	Examples of Inference	155
6.4.2	Inference Algorithms and Packages	157
6.4.3	Inference Using Netica	159
6.5	Networks with Continuous Variables	161
6.5.1	Gaussian Bayesian Networks	161
6.5.2	Hybrid Networks	163
6.6	Obtaining the Probabilities	165
6.6.1	Difficulty Inherent in Multiple Parents	166
6.6.2	Basic Noisy OR-Gate Model	167
6.6.3	Leaky Noisy OR-Gate Model	169
6.6.4	Further Models	171
6.7	Large-Scale Application: Promedas	171
7	Advanced Properties of Bayesian Network	177
7.1	Entailed Conditional Independencies	178
7.1.1	Examples of Entailed Conditional Independencies	178
7.1.2	d-Separation	182
7.2	Faithfulness	185
7.2.1	Unfaithful Probability Distributions	186
7.2.2	Faithfulness Condition	188
7.3	Markov Equivalence	189
7.4	Markov Blankets and Boundaries	192
8	Decision Analysis	199
8.1	Decision Trees	200
8.1.1	Simple Examples	200
8.1.2	Solving More Complex Decision Trees	204
8.2	Influence Diagrams	217
8.2.1	Representing with Influence Diagrams	217
8.2.2	Solving Influence Diagrams	225
8.2.3	Techniques for Solving Influence Diagrams	225
8.2.4	Solving Influence Diagrams Using Netica	229
8.3	Modeling Risk Preferences	234
8.3.1	Exponential Utility Function	234
8.3.2	Assessing r	236
8.4	Analyzing Risk Directly	237
8.4.1	Using the Variance to Measure Risk	237
8.4.2	Risk Profiles	239
8.4.3	Dominance	240
8.5	Good Decision versus Good Outcome	245
8.6	Sensitivity Analysis	245
8.7	Value of Information	248


8.7.1	Expected Value of Perfect Information	249
8.7.2	Expected Value of Imperfect Information	252
8.8	Discussion and Further Reading	253
8.8.1	Academics	254
8.8.2	Business and Finance	254
8.8.3	Capital Equipment	254
8.8.4	Computer Games	254
8.8.5	Computer Vision	254
8.8.6	Computer Software	255
8.8.7	Medicine	255
8.8.8	Natural Language Processing	256
8.8.9	Planning	256
8.8.10	Psychology	256
8.8.11	Reliability Analysis	256
8.8.12	Scheduling	256
8.8.13	Speech Recognition	256
8.8.14	Vehicle Control and Malfunction Diagnosis	257
 III Emergent Intelligence		 265
9	Evolutionary Computation	267
9.1	Genetics Review	268
9.2	Genetic Algorithms	270
9.2.1	Algorithm	271
9.2.2	Illustrative Example	271
9.2.3	Traveling Salesperson Problem	274
9.3	Genetic Programming	284
9.3.1	Illustrative Example	285
9.3.2	Artificial Ant	288
9.3.3	Application to Financial Trading	291
9.4	Discussion and Further Reading	294
10	Swarm Intelligence	297
10.1	Ant System	298
10.1.1	Real Ant Colonies	298
10.1.2	Artificial Ants for Solving the TSP	298
10.2	Flocks	302
10.3	Discussion and Further Reading	306
 IV Learning		 309
11	Learning Deterministic Models	311
11.1	Supervised Learning	312
11.2	Regression	312
11.2.1	Simple Linear Regression	312

11.2.2	Multiple Linear Regression	315
11.2.3	Overfitting and Cross Validation	317
11.3	Learning a Decision Tree	319
11.3.1	Information Theory	321
11.3.2	Information Gain and the ID3 Algorithm	323
11.3.3	Overfitting	326
12	Learning Probabilistic Model Parameters	331
12.1	Learning a Single Parameter	332
12.1.1	Binomial Random Variables	332
12.1.2	Multinomial Random Variables	335
12.2	Learning Parameters in a Bayesian Network	336
12.2.1	Procedure for Learning Parameters	337
12.2.2	Equivalent Sample Size	339
12.3	Learning Parameters with Missing Data★	341
13	Learning Probabilistic Model Structure	351
13.1	Structure Learning Problem	352
13.2	Score-Based Structure Learning	353
13.2.1	Bayesian Score	353
13.2.2	BIC Score	360
13.2.3	Consistent Scoring Criteria	362
13.2.4	How Many DAGs Must We Score?	362
13.2.5	Using the Learned Network to Do Inference	363
13.2.6	Learning Structure with Missing Data★	364
13.2.7	Approximate Structure Learning	373
13.2.8	Model Averaging	377
13.2.9	Approximate Model Averaging★	380
13.3	Constraint-Based Structure Learning	383
13.3.1	Learning a DAG Faithful to P	383
13.3.2	Learning a DAG in which P Is Embedded Faithfully	389
13.4	Application: MENTOR	390
13.4.1	Developing the Network	390
13.4.2	Validating MENTOR	394
13.5	Software Packages for Learning	394
13.6	Causal Learning	394
13.6.1	Causal Faithfulness Assumption	395
13.6.2	Causal Embedded Faithfulness Assumption	397
13.6.3	Application: College Student Retention Rate	401
13.7	Class Probability Trees	404
13.7.1	Theory of Class Probability Trees	405
13.7.2	Application to Targeted Advertising	407
13.8	Discussion and Further Reading	410
13.8.1	Biology	410
13.8.2	Business and Finance	411
13.8.3	Causal Learning	411
13.8.4	Data Mining	411

13.8.5	Medicine	411
13.8.6	Weather Forecasting	412
14	More Learning	417
14.1	Unsupervised Learning	417
14.1.1	Clustering	418
14.1.2	Automated Discovery	419
14.2	Reinforcement Learning	420
14.2.1	Multi-Armed Bandit Algorithms	420
14.2.2	Dynamic Networks★	423
14.3	Discussion and Further Reading	436
V	Language Understanding	439
15	Natural Language Understanding	441
15.1	Parsing	442
15.1.1	Recursive Parser	445
15.1.2	Ambiguity	447
15.1.3	Dynamic Programming Parser	449
15.1.4	Probabilistic Parser	454
15.1.5	Obtaining Probabilities for a PCFG	456
15.1.6	Lexicalized PCFG	456
15.2	Semantic Interpretation	458
15.3	Concept/Knowledge Interpretation	460
15.4	Information Extraction	461
15.4.1	Applications of Information Extraction	462
15.4.2	Architecture for an Information Extraction System	463
15.5	Discussion and Further Reading	464
	Bibliography	467
	Index	493

Chapter 1

Introduction to Artificial Intelligence



In 1990, I (Richard Neapolitan) was relatively new to the field of **Artificial Intelligence (AI)**. At the *6th Conference on Uncertainty in Artificial Intelligence*, which was held at MIT, I met Eugene Charniak, who at the time was a well-known researcher in AI. During a conversation while strolling along the campus, I said, “I heard that the Japanese are applying fuzzy logic to AI.” Gene responded “I don’t believe the Japanese have been any more successful at AI than us.” This comment substantiated that which I had already begun to realize, namely that AI seemed to be a dismal failure.

Dr. Charniak’s comment was correct in 1990 and it is still correct today. If we consider AI to be the development of an artificial entity such as the Termi-

nator in the movie by the same name or HAL in the classic sci-fi movie *Space Odyssey*, then we have not developed anything close. The Terminator and HAL are artificial entities that can learn and make decisions in a complex, changing environment, affect that environment, and communicate their knowledge and choices to humans. We have no such entities.

So why does the field of AI persist, and why was this book written? In their efforts to develop artificial intelligence, researchers looked at the behavior/reasoning of intelligent entities such as humans, and developed algorithms based on that behavior. These algorithms have been used to solve many interesting problems, including the development of systems that behave intelligently in limited domains. Such systems includes ones that can perform medical diagnosis, diagnose problems with software, make financial decisions, navigate a difficult terrain, monitor the possible failure of a space shuttle, recognize speech, understand text, plan a trip, track a target, learn an individual's preferences from the preferences of similar individuals, learn the causal relationships among genes, and learn which genes affect a phenotype. This book concerns these algorithms and applications. Before discussing the content of this book further, we provide a brief history of AI.

1.1 History of Artificial Intelligence

We start by discussing early efforts to define artificial intelligence.

1.1.1 What Is Artificial Intelligence?

Abandoning the philosophical question of what it means for an artificial entity to think or have intelligence, Alan Turing [1950] developed an empirical test of artificial intelligence, which is more appropriate to the computer scientist endeavoring to implement artificial intelligence on a computer. The **Turing test** is an operational test; that is, it provides a concrete way to determine whether the entity is intelligent. The test involves a human interrogator who is in one room, another human being in a second room, and an artificial entity in a third room. The interrogator is allowed to communicate with both the other human and the artificial entity only with a textual device such as a terminal. The interrogator is asked to distinguish the other human from the artificial entity based on answers to questions posed by the interrogator. If the interrogator cannot do this, the Turing test is passed and we say that the artificial entity is intelligent.

Note that the Turing test avoids physical interaction between the interrogator and the artificial entity; the assumption is that physical interaction is not necessary for intelligence. For example, HAL in the movie *Space Odyssey* is simply an entity with which the crew communicates, and HAL would pass the Turing test. If the interrogator is provided with visual information about the artificial entity so that the interrogator can test the entity's ability to perceive and navigate in the world, we call the test the **total Turing test**. The Terminator in the movie of the same name would pass this test.



Figure 1.1 The Chinese room experiment.

Searle [1980] took exception to the Turing test with his **Chinese room** thought experiment. The experiment proceeds as follows. Suppose that we have successfully developed a computer program that appears to understand Chinese. That is, the program takes sentences written with Chinese characters as input, processes the characters, and outputs sentences written using Chinese characters. See Figure 1.1. If it is able to convince a Chinese interrogator that it is a human, then the Turing test would be passed.

Searle asks “Does the program literally *understand* Chinese, or is it only simulating the ability to understand Chinese?” To address this question, Searle proposes that he could sit in a closed room holding a book with an English version of the program, and adequate paper and pencils to carry out the instructions of the program by hand. The Chinese interrogator could then provide Chinese sentences through a slot in the door, and Searle could process them using the program’s instructions and send Chinese sentences back through the same slot. Searle says that he has performed the exact same task as the computer that passed the Turing test. That is, each is following a program that simulates intelligent behavior. However, Searle notes that he does not speak Chinese. Therefore, because he does not understand Chinese, the reasonable conclusion is that the computer does not understand Chinese either. Searle argues that if the computer is not understanding the conversation, then it is not thinking, and therefore it does not have an intelligent mind.

Searle formulated the philosophical position known as **strong AI**, which is as follows:

The appropriately programmed computer really is a mind, in the sense that computers given the right programs can be literally said to understand and have other cognitive states.

Based on his Chinese room experiment, Searle concludes that strong AI is not possible. He states that “I can have any formal program you like, but I still understand nothing.” Searle’s paper resulted in a great deal of controversy and discussion for some time to come (See, for example, [Harnad, 2001]).

The position that computers could appear and behave intelligently, but not necessarily understand, is called **weak AI**. The essence of the matter is whether a computer could actually have a mind (strong AI) or could only simulate a mind (weak AI). This distinction is of greater concern to the philosopher who is discussing the notion of consciousness [Chalmers, 1996]. Perhaps facetiously, a philosopher could even argue that emergentism might take place in the Chinese room experiment, and a mind might arise from Searle performing all his manipulations. Practically speaking, none of this is of concern to the computer scientist. If the program for all purposes behaves as if it is intelligent, computer scientists have achieved their goal.

1.1.2 Emergence of AI

Initial efforts at AI involved modeling the neurons in the brain. An artificial neuron is treated as a binary variable that is switched to either *on* or *off*. This notion was first proposed in [McCulloch and Pitts, 1943], and was furthered by Donald Hebb [1949] when he developed **Hebbian learning** for **neural networks**. In 1951, Marvin Minsky and Dean Edmonds built SNARC, the first neural network computer.

Following this accomplishment and Turing’s development of the Turing test, researchers became increasingly interested in the study of neural networks and intelligent systems, resulting in John McCarthy¹ organizing a 2-month workshop involving interested researchers at Dartmouth University in 1956. He coined the term *Artificial Intelligence* at that workshop. Attendees included Minsky, Claude Shannon (the developer of information theory), and many others. AI emerged as a new discipline whose goal was to create computer systems that could learn, react, and make decisions in a complex, changing environment.

1.1.3 Cognitive Science and AI

Cognitive science is the discipline that studies the mind and its processes. It concerns how information is represented and processed by the mind. It is an interdisciplinary field spanning philosophy, psychology, artificial intelligence, neuroscience, linguistics, and anthropology, and emerged as its own discipline somewhat concurrently with AI. Cognitive science involves empirical studies of the mind, whereas AI concerns the development of an artificial mind. However, owing to their related endeavors, each field is able to borrow from the other.

¹ John McCarthy developed the LISP programming language for AI applications and is considered by many to be the father of AI.