

---

# Undergraduate Topics in Computer Science

## **Series editor**

Ian Mackie

## **Advisory Board**

Samson Abramsky, University of Oxford, Oxford, UK

Karin Breitman, Pontifical Catholic University of Rio de Janeiro, Rio de Janeiro, Brazil

Chris Hankin, Imperial College London, London, UK

Dexter Kozen, Cornell University, Ithaca, USA

Andrew Pitts, University of Cambridge, Cambridge, UK

Hanne Riis Nielson, Technical University of Denmark, Kongens Lyngby, Denmark

Steven Skiena, Stony Brook University, Stony Brook, USA

Iain Stewart, University of Durham, Durham, UK

Undergraduate Topics in Computer Science (UTiCS) delivers high-quality instructional content for undergraduates studying in all areas of computing and information science. From core foundational and theoretical material to final-year topics and applications, UTiCS books take a fresh, concise, and modern approach and are ideal for self-study or for a one- or two-semester course. The texts are all authored by established experts in their fields, reviewed by an international advisory board, and contain numerous examples and problems. Many include fully worked solutions.

More information about this series at <http://www.springer.com/series/7592>

---

Wolfgang Ertel

# Introduction to Artificial Intelligence

Second Edition

Translated by Nathanael Black  
With illustrations by Florian Mast

 Springer

Wolfgang Ertel  
Hochschule Ravensburg-Weingarten  
Weingarten  
Germany

ISSN 1863-7310 ISSN 2197-1781 (electronic)  
Undergraduate Topics in Computer Science  
ISBN 978-3-319-58486-7 ISBN 978-3-319-58487-4 (eBook)  
DOI 10.1007/978-3-319-58487-4

Library of Congress Control Number: 2017943187

1st edition: © Springer-Verlag London Limited 2011  
2nd edition: © Springer International Publishing AG 2017

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

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

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Printed on acid-free paper

This Springer imprint is published by Springer Nature  
The registered company is Springer International Publishing AG  
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

---

## Preface to the Second Edition

After 60 years, Artificial Intelligence (AI) has now reached industry and the consciousness of the population. The impressive successes and new AI methods are now so relevant that they should be taught even in a basic course. In about 30 new pages, I report mainly on *deep learning*, a consistent further development of neural networks, which finally enables image processing systems to recognize almost any object in pixel images. Among other benefits, this lead to the first computer program that could beat one of the world's best Go players.

In the new section on Deep Learning, we must not leave out a short report about the fascinating new subarea of *creativity*. For the first time neural networks can creatively generate texts, music pieces, and even paintings in the style of the old masters. These achievements are based on many years of research on neural networks and machine learning. Practical AI has developed into an engineering discipline in which programs are developed in large industrial teams by experts from various specializations.

Self-driving cars, service robots, and smart homes—which are all applications of AI—will greatly change our lives. However, in addition to great rays of hope, there will be a dark side. Though we live in a time of rapid technological progress, we have long since exceeded the limits of growth. We must therefore think about sustainability when implementing each new invention. In Chap. 1, I would like to give you some food for thought about this topic.

Other new additions to the book include a section on performance evaluation of clustering algorithms and two practical examples explaining Bayes' theorem and its relevance in everyday life. Finally, in a section on search algorithms, we analyze the cycle check, explain route planning for car navigation systems, and briefly introduce Monte Carlo Tree Search.

All known errors have been corrected and updates have been made in many places.

I would like to sincerely thank the readers who have given me feedback and all those who contributed to this new edition through proofreading and suggestions.

I would especially like to thank Adrian Batzill for the route planning measurements and graphs, as well as Nate Black, Nicole Dathe, Markus Schneider, Robin Lehmann, Ankita Agrawal, Wenzel Massag, Lars Berge, Jonas Lang, and Richard Cubek.

Ravensburg  
March 2017

Wolfgang Ertel

---

## Preface to the First Edition

Artificial Intelligence (AI) has the definite goal of understanding intelligence and building intelligent systems. However, the methods and formalisms used on the way to this goal are not firmly set, which has resulted in AI consisting of a multitude of subdisciplines today. The difficulty in an introductory AI course lies in conveying as many branches as possible without losing too much depth and precision.


Russell and Norvig's book [RN10] is more or less the standard introduction into AI. However, since this book has 1,152 pages, and since it is too extensive and costly for most students, the requirements for writing this book were clear: it should be an accessible introduction to modern AI for self-study or as the foundation of a four-hour lecture, with at most 300 pages. The result is in front of you.

In the space of 300 pages, a field as extensive as AI cannot be fully covered. To avoid turning the book into a table of contents, I have attempted to go into some depth and to introduce concrete algorithms and applications in each of the following branches: agents, logic, search, reasoning with uncertainty, machine learning, and neural networks.

The fields of image processing, fuzzy logic, and natural language processing are not covered in detail. The field of image processing, which is important for all of computer science, is a stand-alone discipline with very good textbooks, such as [GW08]. Natural language processing has a similar status. In recognizing and generating text and spoken language, methods from logic, probabilistic reasoning, and neural networks are applied. In this sense this field is part of AI. On the other hand, computer linguistics is its own extensive branch of computer science and has much in common with formal languages. In this book we will point to such appropriate systems in several places, but not give a systematic introduction. For a first introduction in this field, we refer to Chaps. 22 and 23 in [RN10]. Fuzzy logic, or fuzzy set theory, has developed into a branch of control theory due to its primary application in automation technology and is covered in the corresponding books and lectures. Therefore we will forego an introduction here.

The dependencies between chapters of the book are coarsely sketched in the graph shown below. To keep it simple, Chap. 1, with the fundamental introduction for all further chapters, is left out. As an example, the thicker arrow from 2 to 3 means that propositional logic is a prerequisite for understanding predicate logic.

The thin arrow from 9 to 10 means that neural networks are helpful for understanding reinforcement learning, but not absolutely necessary. Thin backward arrows should make clear that later chapters can give more depth of understanding to topics which have already been learned.



This book is applicable to students of computer science and other technical natural sciences and, for the most part, requires high school level knowledge of mathematics. In several places, knowledge from linear algebra and multidimensional analysis is needed. For a deeper understanding of the contents, actively working on the exercises is indispensable. This means that the solutions should only be consulted after intensive work with each problem, and only to check one's solutions, true to Leonardo da Vinci's motto "Study without devotion damages the brain". Somewhat more difficult problems are marked with \*, and especially difficult ones with \*\*. Problems which require programming or special computer science knowledge are labeled with ➡.

On the book's web site at <http://www.hs-weingarten.de/~ertel/aibook> digital materials for the exercises such as training data for learning algorithms, a page with references to AI programs mentioned in the book, a list of links to the covered topics, a clickable list of the bibliography, an errata list, and presentation slides for lecturers can be found. I ask the reader to please send suggestions, criticisms, and tips about errors directly to [ertel@hs-weingarten.de](mailto:ertel@hs-weingarten.de).

This book is an updated translation of my German book "Grundkurs Künstliche Intelligenz" published by Vieweg Verlag. My special thanks go to the translator Nathan Black who in an excellent trans-Atlantic cooperation between Germany and California via SVN, Skype and Email produced this text. I am grateful to Franz Kurfes, who introduced me to Nathan; to Matthew Wight for proofreading the translated book and to Simon Rees from Springer Verlag for his patience.

I would like to thank my wife Evelyn for her support and patience during this time consuming project. Special thanks go to Wolfgang Bibel and Chris Lobenschuss, who carefully corrected the German manuscript. Their suggestions and discussions lead to many improvements and additions. For reading the corrections and other valuable services, I would like to thank Richard Cubek, Celal Döven, Joachim Feßler, Nico Hochgeschwender, Paul Kirner, Wilfried Meister, Norbert Perk, Peter Radtke, Markus Schneider, Manfred Schramm, Uli Stärk, Michel Tokic, Arne Usadel and all interested students. My thanks also go out to Florian Mast for the priceless cartoons and very effective collaboration.

I hope that during your studies this book will help you share my fascination with Artificial Intelligence.



---

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	What Is Artificial Intelligence?	1
1.1.1	Brain Science and Problem Solving	3
1.1.2	The Turing Test and Chatterbots	5
1.2	The History of AI	5
1.2.1	The First Beginnings	7
1.2.2	Logic Solves (Almost) All Problems	8
1.2.3	The New Connectionism	9
1.2.4	Reasoning Under Uncertainty	9
1.2.5	Distributed, Autonomous and Learning Agents	10
1.2.6	AI Grows Up	11
1.2.7	The AI Revolution	11
1.3	AI and Society	11
1.3.1	Does AI Destroy Jobs?	11
1.3.2	AI and Transportation	14
1.3.3	Service Robotics	15
1.4	Agents	17
1.5	Knowledge-Based Systems	19
1.6	Exercises	20
<b>2</b>	<b>Propositional Logic</b>	<b>23</b>
2.1	Syntax	23
2.2	Semantics	24
2.3	Proof Systems	26
2.4	Resolution	30
2.5	Horn Clauses	33
2.6	Computability and Complexity	36
2.7	Applications and Limitations	37
2.8	Exercises	37

<b>3</b>	<b>First-order Predicate Logic</b>	39
3.1	Syntax	40
3.2	Semantics	41
3.2.1	Equality	45
3.3	Quantifiers and Normal Forms	45
3.4	Proof Calculi	49
3.5	Resolution	51
3.5.1	Resolution Strategies	55
3.5.2	Equality	55
3.6	Automated Theorem Provers	56
3.7	Mathematical Examples	57
3.8	Applications	60
3.9	Summary	63
3.10	Exercises	63
<b>4</b>	<b>Limitations of Logic</b>	65
4.1	The Search Space Problem	65
4.2	Decidability and Incompleteness	67
4.3	The Flying Penguin	69
4.4	Modeling Uncertainty	71
4.5	Exercises	73
<b>5</b>	<b>Logic Programming with PROLOG</b>	75
5.1	PROLOG Systems and Implementations	76
5.2	Simple Examples	76
5.3	Execution Control and Procedural Elements	79
5.4	Lists	81
5.5	Self-modifying Programs	82
5.6	A Planning Example	83
5.7	Constraint Logic Programming	85
5.8	Summary	87
5.9	Exercises	88
<b>6</b>	<b>Search, Games and Problem Solving</b>	91
6.1	Introduction	91
6.2	Uninformed Search	97
6.2.1	Breadth-First Search	97
6.2.2	Depth-First Search	99
6.2.3	Iterative Deepening	100
6.2.4	Comparison	102
6.2.5	Cycle Check	102
6.3	Heuristic Search	103
6.3.1	Greedy Search	106
6.3.2	A <sup>★</sup> -Search	107
6.3.3	Route Planning with the A <sup>★</sup> Search Algorithm	109

6.3.4	IDA <sup>★</sup> -Search . . . . .	111
6.3.5	Empirical Comparison of the Search Algorithms . . . . .	111
6.3.6	Summary . . . . .	113
6.4	Games with Opponents . . . . .	114
6.4.1	Minimax Search . . . . .	114
6.4.2	Alpha-Beta-Pruning . . . . .	115
6.4.3	Non-deterministic Games . . . . .	117
6.5	Heuristic Evaluation Functions . . . . .	118
6.5.1	Learning of Heuristics . . . . .	118
6.6	State of the Art . . . . .	119
6.6.1	Chess . . . . .	120
6.6.2	Go . . . . .	121
6.7	Exercises . . . . .	122
<b>7</b>	<b>Reasoning with Uncertainty . . . . .</b>	<b>125</b>
7.1	Computing with Probabilities . . . . .	127
7.1.1	Conditional Probability . . . . .	130
7.2	The Principle of Maximum Entropy . . . . .	136
7.2.1	An Inference Rule for Probabilities . . . . .	136
7.2.2	Maximum Entropy Without Explicit Constraints . . . . .	141
7.2.3	Conditional Probability Versus Material Implication . . . . .	142
7.2.4	MaxEnt-Systems . . . . .	143
7.2.5	The Tweety Example . . . . .	144
7.3	LEXMED, an Expert System for Diagnosing Appendicitis . . . . .	145
7.3.1	Appendicitis Diagnosis with Formal Methods . . . . .	145
7.3.2	Hybrid Probabilistic Knowledge Base . . . . .	146
7.3.3	Application of LEXMED . . . . .	149
7.3.4	Function of LEXMED . . . . .	150
7.3.5	Risk Management Using the Cost Matrix . . . . .	153
7.3.6	Performance . . . . .	155
7.3.7	Application Areas and Experiences . . . . .	157
7.4	Reasoning with Bayesian Networks . . . . .	158
7.4.1	Independent Variables . . . . .	158
7.4.2	Graphical Representation of Knowledge as a Bayesian Network . . . . .	160
7.4.3	Conditional Independence . . . . .	160
7.4.4	Practical Application . . . . .	162
7.4.5	Software for Bayesian Networks . . . . .	163
7.4.6	Development of Bayesian Networks . . . . .	165
7.4.7	Semantics of Bayesian Networks . . . . .	168
7.5	Summary . . . . .	170
7.6	Exercises . . . . .	171

<b>8</b>	<b>Machine Learning and Data Mining</b>	175
8.1	Data Analysis	180
8.2	The Perceptron, a Linear Classifier	183
8.2.1	The Learning Rule	185
8.2.2	Optimization and Outlook	188
8.3	The Nearest Neighbor Method	189
8.3.1	Two Classes, Many Classes, Approximation	193
8.3.2	Distance Is Relevant	194
8.3.3	Computation Times	195
8.3.4	Summary and Outlook	196
8.3.5	Case-Based Reasoning	197
8.4	Decision Tree Learning	198
8.4.1	A Simple Example	199
8.4.2	Entropy as a Metric for Information Content	200
8.4.3	Information Gain	203
8.4.4	Application of C4.5	205
8.4.5	Learning of Appendicitis Diagnosis	207
8.4.6	Continuous Attributes	210
8.4.7	Pruning—Cutting the Tree	211
8.4.8	Missing Values	212
8.4.9	Summary	213
8.5	Cross-Validation and Overfitting	213
8.6	Learning of Bayesian Networks	215
8.6.1	Learning the Network Structure	215
8.7	The Naive Bayes Classifier	218
8.7.1	Text Classification with Naive Bayes	220
8.8	One-Class Learning	222
8.8.1	Nearest Neighbor Data Description	223
8.9	Clustering	224
8.9.1	Distance Metrics	225
8.9.2	k-Means and the EM Algorithm	226
8.9.3	Hierarchical Clustering	228
8.9.4	How is the Number of Clusters Determined?	230
8.10	Data Mining in Practice	233
8.10.1	The Data Mining Tool KNIME	233
8.11	Summary	236
8.12	Exercises	238
<b>9</b>	<b>Neural Networks</b>	245
9.1	From Biology to Simulation	246
9.1.1	The Mathematical Model	247
9.2	Hopfield Networks	250

9.2.1	Application to a Pattern Recognition Example . . . . .	251
9.2.2	Analysis. . . . .	252
9.2.3	Summary and Outlook. . . . .	255
9.3	Neural Associative Memory . . . . .	256
9.3.1	Correlation Matrix Memory . . . . .	257
9.3.2	The Binary Hebb Rule. . . . .	259
9.3.3	A Spelling Correction Program . . . . .	261
9.4	Linear Networks with Minimal Errors . . . . .	263
9.4.1	Least Squares Method . . . . .	264
9.4.2	Application to the Appendicitis Data. . . . .	265
9.4.3	The Delta Rule . . . . .	266
9.4.4	Comparison to the Perceptron . . . . .	268
9.5	The Backpropagation Algorithm. . . . .	269
9.5.1	NETtalk: A Network Learns to Speak. . . . .	272
9.5.2	Learning of Heuristics for Theorem Provers . . . . .	273
9.5.3	Problems and Improvements . . . . .	274
9.6	Support Vector Machines . . . . .	275
9.7	Deep Learning . . . . .	277
9.7.1	Nature as Example. . . . .	278
9.7.2	Stacked Denoising Autoencoder . . . . .	279
9.7.3	Other Methods. . . . .	280
9.7.4	Systems and Implementations . . . . .	281
9.7.5	Applications of Deep Learning . . . . .	281
9.8	Creativity . . . . .	282
9.9	Applications of Neural Networks . . . . .	284
9.10	Summary and Outlook . . . . .	285
9.11	Exercises. . . . .	286
<b>10</b>	<b>Reinforcement Learning . . . . .</b>	<b>289</b>
10.1	Introduction . . . . .	289
10.2	The Task. . . . .	291
10.3	Uninformed Combinatorial Search . . . . .	293
10.4	Value Iteration and Dynamic Programming . . . . .	295
10.5	A Learning Walking Robot and Its Simulation. . . . .	298
10.6	Q-Learning . . . . .	300
10.6.1	Q-Learning in a Nondeterministic Environment. . . . .	303
10.7	Exploration and Exploitation. . . . .	304
10.8	Approximation, Generalization and Convergence . . . . .	305
10.9	Applications . . . . .	306
10.10	AlphaGo, the Breakthrough in Go . . . . .	306
10.11	Curse of Dimensionality . . . . .	309
10.12	Summary and Outlook . . . . .	310
10.13	Exercises. . . . .	310

---

<b>11</b>	<b>Solutions for the Exercises</b>	313
11.1	Introduction	313
11.2	Propositional Logic	314
11.3	First-Order Predicate Logic	316
11.4	Limitations of Logic	317
11.5	PROLOG	317
11.6	Search, Games and Problem Solving	319
11.7	Reasoning with Uncertainty	322
11.8	Machine Learning and Data Mining	329
11.9	Neural Networks	335
11.10	Reinforcement Learning	337
	<b>References</b>	339
	<b>Index</b>	351