




The book was found

Artificial Intelligence: A Modern Approach (3rd Edition)



Synopsis

Artificial Intelligence: A Modern Approach, 3e offers the most comprehensive, up-to-date introduction to the theory and practice of artificial intelligence. Number one in its field, this textbook is ideal for one or two-semester, undergraduate or graduate-level courses in Artificial Intelligence. Dr. Peter Norvig, contributing Artificial Intelligence author and Professor Sebastian Thrun, a Pearson author are offering a free online course at Stanford University on artificial intelligence. According to an article in The New York Times, the course on artificial intelligence is “one of three being offered experimentally by the Stanford computer science department to extend technology knowledge and skills beyond this elite campus to the entire world.” One of the other two courses, an introduction to database software, is being taught by Pearson author Dr. Jennifer Widom. Artificial Intelligence: A Modern Approach, 3e is available to purchase as an eText for your Kindle, NOOK, and the iPhone/iPad. To learn more about the course on artificial intelligence, visit <http://www.ai-class.com>. To read the full New York Times article, click here.

Book Information

Hardcover: 1152 pages

Publisher: Pearson; 3 edition (December 11, 2009)

Language: English

ISBN-10: 0136042597

ISBN-13: 978-0136042594

Product Dimensions: 8.1 x 1.7 x 10.1 inches

Shipping Weight: 4.8 pounds (View shipping rates and policies)

Average Customer Review: 4.0 out of 5 stars 183 customer reviews

Best Sellers Rank: #5,185 in Books (See Top 100 in Books) #5 in Books > Textbooks >

Computer Science > Artificial Intelligence #7 in Books > Computers & Technology > Computer Science > AI & Machine Learning > Intelligence & Semantics

Customer Reviews

The long-anticipated revision of this #1 selling book offers the most comprehensive, state of the art introduction to the theory and practice of artificial intelligence for modern applications. Intelligent Agents. Solving Problems by Searching. Informed Search Methods. Game Playing. Agents that Reason Logically. First-order Logic. Building a Knowledge Base. Inference in First-Order Logic. Logical Reasoning Systems. Practical Planning. Planning and Acting. Uncertainty. Probabilistic

Reasoning Systems. Making Simple Decisions. Making Complex Decisions. Learning from Observations. Learning with Neural Networks. Reinforcement Learning. Knowledge in Learning. Agents that Communicate. Practical Communication in English. Perception. Robotics. For computer professionals, linguists, and cognitive scientists interested in artificial intelligence.

Stuart Russell was born in 1962 in Portsmouth, England. He received his B.A. with first-class honours in physics from Oxford University in 1982, and his Ph.D. in computer science from Stanford in 1986. He then joined the faculty of the University of California at Berkeley, where he is a professor of computer science, director of the Center for Intelligent Systems, and holder of the Smith–Petersen Zadeh Chair in Engineering. In 1990, he received the Presidential Young Investigator Award of the National Science Foundation, and in 1995 he was cowinner of the Computers and Thought Award. He was a 1996 Miller Professor of the University of California and was appointed to a Chancellor’s Professorship in 2000. In 1998, he gave the Forsythe Memorial Lectures at Stanford University. He is a Fellow and former Executive Council member of the American Association for Artificial Intelligence. He has published over 100 papers on a wide range of topics in artificial intelligence. His other books include *The Use of Knowledge in Analogy and Induction* and (with Eric Wefald) *Do the Right Thing: Studies in Limited Rationality*. Peter Norvig is currently Director of Research at Google, Inc., and was the director responsible for the core Web search algorithms from 2002 to 2005. He is a Fellow of the American Association for Artificial Intelligence and the Association for Computing Machinery. Previously, he was head of the Computational Sciences Division at NASA Ames Research Center, where he oversaw NASA’s research and development in artificial intelligence and robotics, and chief scientist at Junglee, where he helped develop one of the first Internet information extraction services. He received a B.S. in applied mathematics from Brown University and a Ph.D. in computer science from the University of California at Berkeley. He received the Distinguished Alumni and Engineering Innovation awards from Berkeley and the Exceptional Achievement Medal from NASA. He has been a professor at the University of Southern California and a research faculty member at Berkeley. His other books are *Paradigms of AI Programming: Case Studies in Common Lisp* and *Verbmobil: A Translation System for Face-to-Face Dialog and Intelligent Help Systems for UNIX*.

Since this version of the book is for the Indian market, I was a bit worried about the potential differences from the U.S. 3rd edition. But after doing a first-hand comparison, I found a LOT more similarities than differences. The biggest difference is the absence of the last two chapters. The U.S.

edition includes: Chapter 26, "Philosophical Foundations", which covers arguments over consciousness in machines and the possibility of robot uprisings; and Chapter 27, "AI: The Present and Future", which *briefly* describes some things AI researchers need to work on before we can build a "general-purpose intelligent agent" (a.k.a. one single AI that will be good enough at a lot of different tasks). These two chapters are interesting, but I wouldn't call them core material, so I'm not surprised they got left out. Other than that, the differences are astonishingly minor. The chapters are unnumbered, and some of them swapped places for no reason, but all the content from chapters 1 through 25 is here, plus the two appendixes. The ONLY edits to the text are removals of cross-chapter references. The U.S. version will say something like, "When we discussed whatzits in Chapter 4, we mentioned that they come in two flavors, X and Y", while the same line in this book will say, "Whatzits come in two flavors, X and Y". Again, those are the ONLY edits. All the other sentences are the same. The equations are the same. The diagrams are the same. The exercises at the end of each chapter are the same. In short, if you can do without those last two chapters, buy this version and save your money.

This is an excellent book however I cannot recommend purchasing the Kindle version of this text. It is atrocious. There are subject headings inserted after the subject is spoken about and, quite often, many heading stacked up on top of each other taking up almost an entire page with useless titles that are in the wrong order anyway. There are no page numbers, which is unacceptable for a text that is used by many college AI programs across the country. There are tons of hyphenation errors. The delineations between figure notes and the text are almost imperceptible so it is difficult to tell what text goes where. In general it is difficult to read and navigate due to this horrible Kindle conversion.

This review is organized into two parts. The first part gives a general overview of the textbook and my thoughts about the content, while the second part gives a summary of the problems I've encountered while reading the Kindle version. This book was made by multiple authors. Fortunately, the book's definitions and summaries are generally coherent. It has a through coverage on the history of AI, the definitions (those that are controversial have been noted as such), and provides pseudocode of many algorithms. Similar to other textbooks on this topic, it does not provide run-throughs of those algorithms. The organizations of the book can be improved. I have some problems with the way they organized the material on the second part of the book (II: Problem Solving). Chapter 3 is named "Solving Problems by Searching", but I didn't realize all they did was

on classical searches in Chapter 3 until I read Chapter 4. They could have explained why uniform-cost search is uninformed search even though it has a cost function better than they did in the textbook. A better naming scheme for those chapters would be "Ch3: Global Search ", "Ch4: Local Search/..." (I would suggest breaking Chapter 4 and put it into different places.)".I bought the Kindle version, and it had many formatting problems. Most of them are minor (such as missing a "space" between two words, images appearing in different locations, image sizes too small, etc.) However, I had caught one specific error in the second part of the book (somewhere between chapter 3-5) where the fact is exactly the opposite than the one written (forgot to add the negation?) I should have marked and noted it so I can write down which error it was in the review but I was too lazy...But it's a good purchase overall!

So I'm in a class that requires this text book. So far I have to say they provide good examples and explain the complex parts pretty well. They also add some humor in here on example which I love. This book is a easy read compared to most text, and the HW question inside the book are fun to answer. This book makes me feel like I'm learning a lot.Ill update this again once the class is over and I'm farther in the book.

This book was quite excellent, especially in presenting concepts to people who had no previous knowledge or experience with artificial intelligence. It provides an excellent overview of artificial intelligence concepts, though they gloss over the statistical parts of it (I guess they assumed you should already have knowledge of these).The reason why this book got 4 instead of 5 stars was for two reasons: (1) the statistical knowledge that one needed to know in order to understand a concept was not adequately covered in the book; and (2) it is pretty much impossible to skip around the book (you're forced to learn chapter by chapter, otherwise you'll get lost in the later ones, so in terms of learning methods, it's not really flexible).If you want to learn about artificial neural networks and reinforcement learning, the Mitchell text is pretty decent as well. Wikipedia also helps.

I purchased this product for my Artificial Intelligence Class. This is a great introduction book to AI algorithms. For a graduating Computer Science, I felt that it presented many examples and algorithms in a simple easy to understand way with great real world applications.My only complaint with this book is that a lot of the pseudocode is hard to follow and implements a lot of algorithms without any explanation. For a more experienced programmer this may not be a problem, but as a student I found it hard to follow at times. Overall great book!

[Download to continue reading...](#)

Readings in Medical Artificial Intelligence. The First Decade (Addison-Wesley Series in Artificial Intelligence) Emotional Intelligence: Why You're Smarter But They Are More Successful (Emotional intelligence leadership, Emotional Quotient, emotional intelligence depression, emotional intelligence workbook) Artificial Intelligence: A Modern Approach (3rd Edition) Artificial Intelligence: A Modern Approach (2nd Edition) Artificial Intelligence: A Modern Approach Emotional Intelligence: 3 Manuscripts - Emotional Intelligence Definitive Guide, Mastery, Complete Step by Step Guide (Social Engineering, Leadership, ... (Emotional Intelligence Series Book 4) Artificial Intelligence: Structures and Strategies for Complex Problem Solving (6th Edition) Prolog Programming for Artificial Intelligence (4th Edition) (International Computer Science Series) Life 3.0: Being Human in the Age of Artificial Intelligence Artificial Intelligence for Games Robots and Artificial Intelligence (Technology Behind) The Fourth Transformation: How Augmented Reality & Artificial Intelligence Will Change Everything Humans Need Not Apply: A Guide to Wealth and Work in the Age of Artificial Intelligence Artificial Intelligence: What Everyone Needs to Know Economy Monitor Guide to Smart Contracts: Blockchain Examples (Artificial Intelligence, Law and Finance) The Most Human Human: What Artificial Intelligence Teaches Us About Being Alive Applications of Artificial Intelligence for Decision-Making: Multi-Strategy Reasoning Under Uncertainty Essentials of Game Theory: A Concise, Multidisciplinary Introduction (Synthesis Lectures on Artificial Intelligence and Machine Learning) Forbidden Gates: How Genetics, Robotics, Artificial Intelligence, Synthetic Biology, Nanotechnology, & Human Enhancement Herald The Dawn Of Techno-Dimensional Spiritual Warfare Artificial Intelligence in Label-free Microscopy: Biological Cell Classification by Time Stretch

[Contact Us](#)

[DMCA](#)

[Privacy](#)

[FAQ & Help](#)