



Artificial Intelligence: A Modern Approach

Stuart Russell, Peter Norvig

Download now

Read Online 

Artificial Intelligence: A Modern Approach

Stuart Russell , Peter Norvig

Artificial Intelligence: A Modern Approach Stuart Russell , Peter Norvig

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(TM), NOOK(TM), and the iPhone(R)/iPad(R).

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.

Artificial Intelligence: A Modern Approach Details

Date : Published November 1st 2009 by Vital Source (for Pearson) VST E+p (first published December 13th 1994)

ISBN : 9780136042594

Author : Stuart Russell , Peter Norvig

Format : Paperback 1132 pages

Genre : Science, Computer Science, Artificial Intelligence, Programming, Nonfiction, Reference, Technical

 [Download Artificial Intelligence: A Modern Approach ...pdf](#)

 [Read Online Artificial Intelligence: A Modern Approach ...pdf](#)

Download and Read Free Online Artificial Intelligence: A Modern Approach Stuart Russell , Peter Norvig

From Reader Review Artificial Intelligence: A Modern Approach for online ebook

Drew says

A comprehensive course in modern AI topics. While the book is dense with information, the authors provide clear explanations that will be easily picked up by the careful reader. An excellent companion to an undergraduate course in artificial intelligence.

Erik says

OK so I did not read this cover to cover, but I did look closely at much of what you might call the foundational chapters, just to see 1. is there such a thing as AI, or are we just hoping there will be and 2. what can I learn as a philosopher from AI, whether it exists or not. Goal 2 was much more important as I teach a logic of induction class and of course one major pillar of AI would be developing machines that can perform judgments under uncertainty and apply rational heuristics as well as humans do (which is not very well at all by the way). I found out that I already knew most of this, from studies of Bayesian reasoning (which is very tricky by the way and should not be blindly implemented like this without a clear view of the limitations), and the study of acyclic causal graphs (which is standard academy reading for philosophers). These graphs also admit of howlers and counterexamples as anyone knows. I am more interested in the idea of developing "stupid machines" that function more like neural networks and less like probability maximizers. The human brain is fundamentally (in my view anyway) a stupid-machine, full of crazy workarounds and faulty logic. The correct solution or path is virtually never the one evolution comes up with, it just grinds it out with massive armies of neurons and interconnections and lots of trial and error, but nothing one would call a computation, as in Turing machines. Elegant algorithms for computer vision have, I believe, nothing to do with the way the brain constructs the visual image. One philosopher's take.

James Ravenscroft says

This is **THE** book to read on anything to do with modern artificial intelligence. I regard this as my personal bible and would recommend it to anyone who is involved in technical artificial intelligence.

Luis says

A bit boring. Lacks good solved exercises. Very short on detail in some areas such as Neural Networks. Very "theoretical".

However it does provide a good theoretical introduction to many subjects. I liked the chapters on search.

Paul says

5 stars because there is, quite simply, no substitute.

Artificial Intelligence is, in the context of the infant science of computing, a very old and very broad subdiscipline, the "Turing test" having arisen, not only at the same time, but from the same person as many of the foundations of computing itself. Those of us students of a certain age will recall terms like "symbolic" vs. "connectionist" vs. "probabilistic," as well as "scruffies" and "neats." Key figures, events, and schools of thought span multiple institutions on multiple continents. In short, a major challenge facing anyone wishing to survey Artificial Intelligence is simply coming up with a unifying theme.

The major accomplishment, in my opinion, of AIMA, then, is that: Russell and Norvig take the hodge-podge of AI research, manage to fit it sensibly into a narrative structure centered on the notion of different kinds of "agents" (not to be confused with that portion of AI research that explicitly refers to its constructs as "agents!") and, having dug the pond and filled it with water, skip a stone across the surface. It's up to the reader whether to follow the arcs of the stone from major subject to major subject, forego depth, or whether to pick a particular contact point and concentrate on the eddies propagating from it. For the latter purpose, the extensive bibliography is indispensable.

With all of this said, I have to acknowledge that Russell and Norvig are not entirely impartial AI practitioners. Norvig, in particular, is well-known by now as a staunch Bayesian probabilist who, as Director of Search Quality or Machine Learning or whatever Google has decided to call it today, has made Google the Bayesian powerhouse that it is. (Less known is Norvig's previous stint at high-tech startup Junglee, which was acquired by Amazon. So to some extent Peter Norvig powers both Google and Amazon.) So one can probably claim, not without justification, that AIMA emphasizes Bayesian probability over other approaches.

Finally, as good as AIMA is, it is still a *survey*. Even with respect to Bayesian probability, the treatment is introductory, as I discovered with some shock upon reading *Probability Theory: The Logic of Science*. That's OK, though: it's the best introduction I've ever seen.

So read it once for the survey, keep it on your shelf for the bibliography, and refer back to it whenever you find yourself thinking "hey, didn't I read about that somewhere before?"

Arjun says

A fantastic textbook that's not only a great introduction to AI but also serves as a survey course in technical writing. I only read about 75% of it but definitely plan on revisiting it. Re-reading some earlier chapters taught me how much I missed on a first read (or forgot).

AIMA doesn't presume a ton of background beyond some programming experience, exposure to mathematical notation, and a basic understanding of computational complexity/algorithmic efficiency.

The first 10 chapters or so are the best and the second half of the book can be a bit of a trudge as it devolves into mathematical masturbation. A lot of the chapters are better served by other resources – I highly recommend the CS188 lectures from UC Berkeley for supplementation. Unfortunately, some chapters are straight up bad (the chapter on *Philosophical Foundations* comes to mind), but these tend to be few and far between.

Despite that, there is no more comprehensive book on AI. Read this, re-read this, and treat it with care – you

will reap the rewards for a long time to come.

Carl says

For a textbook, this is amazingly accessible and interesting. if you have any interest in the topic, this is the book to read. It's \$100 or more, but it's very popular for AI classes, so any good college library should have a copy.

Hasnaa says

?? ??????? ??????? ?? ?????? ??? ?????? ??? ?????? ??? ?????? ?????? :D
????? ??? ?????? ??? ?????? ?? ?? ?????? ?? ??? ?????? ?? ?????? ??? ?????? ???
?? ??????? ????

?????? ??????? ??? ??????? ??? ?????? ??? ?????? ??? ?????? ??? ?????? ??? ????????

Koen Crolla says

Holy balls this book has a lot of pages. I also don't know why these things always have to have separate ``international'' editions.

It starts off strongly for a few hundred pages, but then for no reason at all devotes several chapters to high school-level probability and statistics, before devolving into essentially pointless mathematical show-boating for another few hundred pages. Then it finishes off with an interesting but not really relevant and highly unrigorous (not to mention typo-ridden) overview of Google's various products (mostly PageRank and Google Translate).

There's a few more chapters after that, but I think it's best to pretend they don't exist. Chapter 26 (*Philosophical Foundations*), in particular, was a fucking embarrassment, giving more unnecessary to idiots like John Searle and Ray Kurzweil, and wasting paper on absurd hand-wringing over off-the-wall science-fiction scenarios. AI is too legitimate and interesting a field to justify that sort of crap in a university textbook.

In spite of all that, though, it's still a very good book, and a good overview of the field. I particularly liked that each chapter had an extensive section with historical and biographical notes at the end. If nothing else, it at least demonstrates that if the AI winter was ever a real thing (at least in terms of research activity and progress), it's far behind us now.

Manny says

This monumental work, which completely dominates the AI textbook market, has been compared with

classics like Watson's *Molecular Biology of the Cell*, and eminently succeeds in its goal of providing a clear, single-volume summary of the whole field of Artificial Intelligence. As pointed out on the book's home page, it is used in over 1200 universities in over 100 countries, and is the 25th most cited publication on Citeseer and the 2nd most cited publication of this century. The occasional suggestion you may hear that it "has passed its sell-by" or "gives a decent picture of Good Old-Fashioned AI" can unhesitatingly be written off as envious carping from academics who wish they'd got something even a tenth as impressive on their CVs.

What was that? Ah, yes, as a matter of fact it does cite one of my papers. How did you guess?

Nakosy says

It was written more like a text book for undergrads with extensive coverage of many topics. However, I was looking for more in-depth information on knowledge representation. But, it was too superficial for my need. May be, in 3rd edition it encompassed the latest ideas in this area.

Wooi Hen Yap says

Wants a book that explains broad and deep AI yet in laymen term (nearly)? This is IT. Of all the AI books I have read, this one is arguably the most accessible to undergrads (CS, EE background) It assumes only minimal mathematical formalities and pretty much the maths things are self-contained. The authors did a great job of keeping the contents up-to-date with the latest happenings in AI, while keeping the readers sane. Overall, thumbs up!

Nick Black says

Heh, I opened this up to find the ISBN and found dried blood all over the pages, suggesting I read this during my cocaine-intensive period back in 1999-2000. That's fitting, since cocaine and the study of artificial intelligence seem to enjoy several similarities -- incredible expense as a barrier to entry, exciting short-term effects (see: euphoria, A* search) but letdowns upon prolonged use (see: addiction, combinatorial explosions), and they've both ruined plenty of fine careers in computer science. We used this book for CS4600, but I only got halfway through that semester and remember little of it (see: careers in computer science, *aforementioned negative effects of cocaine on*). I went back and read most of this in 2003, and found solid coverage of most everything useful I'm aware of from AI.

Mohammed says

A.N. Mignan says

The Bible on computational decision-making. I use this term as this book is not just about the AI/machine learning we consistently hear about, it's much more. This textbook tends to perfection, with no stone left unturned. Looking forward to the next edition, which, at the accelerating rate of innovation, looks overdue (the following sentence surely feels outdated: "Current Go programs play at the master level on a reduced 9 × 9 board, but are still at advanced amateur level on a full board"). There are 2 aspects I particularly enjoyed, (1) the historical sections at the end of each chapter; the introduction also gave a fascinating history of AI and its relationship to other fields (neurology, logics, cybernetics...). (2) I also liked all the gaming aspects, such as the Wumpus World which I didn't know before. I truly wish I had discovered that book when it was first published in 1995, sigh
