


The book was found

Deep Learning (Adaptive Computation And Machine Learning Series)



Synopsis

"Written by three experts in the field, Deep Learning is the only comprehensive book on the subject." -- Elon Musk, cochair of OpenAI; cofounder and CEO of Tesla and SpaceX
Deep learning is a form of machine learning that enables computers to learn from experience and understand the world in terms of a hierarchy of concepts. Because the computer gathers knowledge from experience, there is no need for a human computer operator to formally specify all the knowledge that the computer needs. The hierarchy of concepts allows the computer to learn complicated concepts by building them out of simpler ones; a graph of these hierarchies would be many layers deep. This book introduces a broad range of topics in deep learning. The text offers mathematical and conceptual background, covering relevant concepts in linear algebra, probability theory and information theory, numerical computation, and machine learning. It describes deep learning techniques used by practitioners in industry, including deep feedforward networks, regularization, optimization algorithms, convolutional networks, sequence modeling, and practical methodology; and it surveys such applications as natural language processing, speech recognition, computer vision, online recommendation systems, bioinformatics, and videogames. Finally, the book offers research perspectives, covering such theoretical topics as linear factor models, autoencoders, representation learning, structured probabilistic models, Monte Carlo methods, the partition function, approximate inference, and deep generative models. Deep Learning can be used by undergraduate or graduate students planning careers in either industry or research, and by software engineers who want to begin using deep learning in their products or platforms. A website offers supplementary material for both readers and instructors.

Book Information

Series: Adaptive Computation and Machine Learning series

Hardcover: 800 pages

Publisher: The MIT Press (November 18, 2016)

Language: English

ISBN-10: 0262035618

ISBN-13: 978-0262035613

Product Dimensions: 7 x 1 x 9 inches

Shipping Weight: 2.9 pounds (View shipping rates and policies)

Average Customer Review: 4.7 out of 5 stars 93 customer reviews

Best Sellers Rank: #763 in Books (See Top 100 in Books) #2 in [Books > Computers &](#)

Technology > Computer Science > AI & Machine Learning > Intelligence & Semantics #4
in [AI & Books](#) > [Textbooks](#) > [Computer Science](#) #33 in [AI & Books](#) > [Education & Teaching](#) >
[Schools & Teaching](#)

Customer Reviews

Written by three experts in the field, Deep Learning is the only comprehensive book on the subject. It provides much-needed broad perspective and mathematical preliminaries for software engineers and students entering the field, and serves as a reference for authorities. (Elon Musk, cochair of OpenAI; cofounder and CEO of Tesla and SpaceX) This is the definitive textbook on deep learning. Written by major contributors to the field, it is clear, comprehensive, and authoritative. If you want to know where deep learning came from, what it is good for, and where it is going, read this book. (Geoffrey Hinton FRS, Emeritus Professor, University of Toronto; Distinguished Research Scientist, Google) Deep learning has taken the world of technology by storm since the beginning of the decade. There was a need for a textbook for students, practitioners, and instructors that includes basic concepts, practical aspects, and advanced research topics. This is the first comprehensive textbook on the subject, written by some of the most innovative and prolific researchers in the field. This will be a reference for years to come. (Yann LeCun, Director of AI Research, Facebook; Silver Professor of Computer Science, Data Science, and Neuroscience, New York University) [T]he AI bible... the text should be mandatory reading by all data scientists and machine learning practitioners to get a proper foothold in this rapidly growing area of next-gen technology. (Daniel D. Gutierrez insideBIGDATA)

Ian Goodfellow is Research Scientist at OpenAI. Yoshua Bengio is Professor of Computer Science at the [Université de Montréal](#). Aaron Courville is Assistant Professor of Computer Science at the [Université de Montréal](#).

This is, to invoke a technical reviewer cliché, a 'valuable' book. Read it and you will have a detailed and sophisticated practical understanding of the state of the art in neural networks technology. Interestingly, I also suspect it will remain current for a long time, because reading it I came to more and more of an impression that neural network technology (at least in the current iteration) is plateauing. Why? Because this book also makes very clear - is completely honest - that neural networks are a 'folk' technology (though they do not use those words): Neural networks work (in fact they work unbelievably well - at least, as Geoffrey Hinton himself has remarked, given

unbelievably powerful computers), but the underlying theory is very limited and there is no reason to think that it will become less limited, and the lack of a theory means that there is no convincing 'gradient', to use an appropriate metaphor, for future development. A constant theme here is that 'this works better than that' for practical reasons not for underlying theoretical reasons. Neural networks are engineering, they are not applied mathematics, and this is very much, and very effectively, an engineer's book.

Good book. I came in with only a cursory understanding of machine learning and now have successfully designed several to solve somewhat challenging problems.

Very well written and very accessible. The authors review linear algebra and some statistics in detail but be forewarned; you need to have seen all of those things before. It's only a review. The thing I like the most is that the descriptions of all concepts and math are written in an intuitive language instead of terse cryptic statements that are usually in these advanced technical texts.

I was pretty excited when I heard earlier this year that Yoshua Bengio, Ian Goodfellow, and Aaron Courville were writing this textbook. Just got my copy in the mail today. Read a few chapters and skimmed through some of the book. This is a very comprehensive, well-written, and easy-to-understand textbook on the theoretical foundations, current research, and applications of deep learning. I've read a lot of research papers (DeepMind, Google Brain, Facebook, NYU, Stanford, etc.), blogs (Nervana Systems, Indico, Colah, Otoro's Blog, etc.), lecture notes (Stanford cs231n, cs224d, cs229), and tutorials (Quoc Le's tutorial, TensorFlow, etc.), and have watched a lot of videos (Hugo Larochelle's tutorials, Stanford cs229, TedTalks, lectures by Yann LeCun

An excellent book! A comprehensive yet efficient approach to Deep Learning.

concise. gets complicated math-wise fast - being realistic here, this isn't an easy subject. This is the best book on the subject.

This book is written by people who have been influential in the field and as you read you understand how well they provide a great insight of the concepts.

Clearly explains underlying derivations

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

Machine Learning: A Probabilistic Perspective (Adaptive Computation and Machine Learning series) Introduction to Machine Learning (Adaptive Computation and Machine Learning series) Deep Learning (Adaptive Computation and Machine Learning series) Reinforcement Learning with Python: An Introduction (Adaptive Computation and Machine Learning series) Semi-Supervised Learning (Adaptive Computation and Machine Learning series) Reinforcement Learning: An Introduction (Adaptive Computation and Machine Learning) Probabilistic Graphical Models: Principles and Techniques (Adaptive Computation and Machine Learning series) Introduction to Deep Learning Using R: A Step-by-Step Guide to Learning and Implementing Deep Learning Models Using R IntAR, Interventions Adaptive Reuse, Volume 03; Adaptive Reuse in Emerging Economies Human Computation (Synthesis Lectures on Artificial Intelligence and Machine Learning) DEEP HEALING SLEEP CD: Deep Relaxation, Guided Imagery Meditation and Affirmations Proven to Help Induce Deep, Restful Sleep BREAD MACHINE COOKBOOK: 120 Most Delicious Bread Machine Recipes (bread, bread bible, bread makers, breakfast, bread machine cookbook, bread baking, bread making, healthy, healthy recipes) Diving Deep: A Beginners Guide to Deep Sea Diving: (Scuba, Snorkelling, Diving, Scuba Diver, Deep Sea Diving, Swimming, Scuba Diving) Machine Learning: For Beginners: Definitive Guide for Neural Networks, Algorithms, Random Forests and Decision Trees Made Simple (Machine Learning, Book 1) Practical Machine Learning with H2O: Powerful, Scalable Techniques for Deep Learning and AI MATLAB Deep Learning: With Machine Learning, Neural Networks and Artificial Intelligence Machine Learning: An Algorithmic Perspective, Second Edition (Chapman & Hall/Crc Machine Learning & Pattern Recognition) Fundamentals of Deep Learning: Designing Next-Generation Machine Intelligence Algorithms Bread Machine Recipes: Hot and Fresh Bread Machine Recipes Anyone Can Make in the Comfort of Their Own Home (The Essential Kitchen Series Book 82) Number Talks: Whole Number Computation, Grades K-5: A Multimedia Professional Learning Resource

[Contact Us](#)

[DMCA](#)

[Privacy](#)

[FAQ & Help](#)