




Ebook Directory
the best source of ebook

The book was found

Fundamentals Of Deep Learning: Designing Next-Generation Machine Intelligence Algorithms



Synopsis

With the reinvigoration of neural networks in the 2000s, deep learning has become an extremely active area of research, one that's paving the way for modern machine learning. In this practical book, author Nikhil Buduma provides examples and clear explanations to guide you through major concepts of this complicated field. Companies such as Google, Microsoft, and Facebook are actively growing in-house deep-learning teams. For the rest of us, however, deep learning is still a pretty complex and difficult subject to grasp. If you're familiar with Python, and have a background in calculus, along with a basic understanding of machine learning, this book will get you started. Examine the foundations of machine learning and neural networks
Learn how to train feed-forward neural networks
Use TensorFlow to implement your first neural network
Manage problems that arise as you begin to make networks deeper
Build neural networks that analyze complex images
Perform effective dimensionality reduction using autoencoders
Dive deep into sequence analysis to examine language
Understand the fundamentals of reinforcement learning

Book Information

Paperback: 298 pages

Publisher: O'Reilly Media; 1 edition (June 29, 2017)

Language: English

ISBN-10: 1491925612

ISBN-13: 978-1491925614

Product Dimensions: 6.9 x 0.5 x 9.1 inches

Shipping Weight: 1.2 pounds (View shipping rates and policies)

Average Customer Review: 4.3 out of 5 stars 22 customer reviews

Best Sellers Rank: #12,844 in Books (See Top 100 in Books) #8 in Books > Computers & Technology > Databases & Big Data > Data Modeling & Design #8 in Books > Textbooks > Computer Science > Artificial Intelligence #9 in Books > Computers & Technology > Databases & Big Data > Data Mining

Customer Reviews

From the Preface With the reinvigoration of neural networks in the 2000s, deep learning has become an extremely active area of research that is paving the way for modern machine learning. This book uses exposition and examples to help you understand major concepts in this complicated field. Large companies such as Google, Microsoft, and Facebook have taken notice and are actively growing in-house deep learning teams. For the rest of us, deep learning is still a pretty complex and

difficult subject to grasp. Research papers are filled to the brim with jargon, and scattered online tutorials do little to help build a strong intuition for why and how deep learning practitioners approach problems. Our goal is to bridge this gap. **Prerequisites and Objectives** This book is aimed at an audience with a basic operating understanding of calculus, matrices, and Python programming. Approaching this material without this background is possible, but likely to be more challenging. Background in linear algebra may also be helpful in navigating certain sections of mathematical exposition. By the end of the book, we hope that our readers will be left with an intuition for how to approach problems using deep learning, the historical context for modern deep learning approaches, and a familiarity with implementing deep learning algorithms using the TensorFlow open source library.

How to Simulate the Mind

Nikhil Buduma is a computer science student at MIT with deep interests in machine learning and the biomedical sciences. He is a two time gold medalist at the International Biology Olympiad, a student researcher, and a "hacker." He was selected as a finalist in the 2012 International BioGENEius Challenge for his research on the pertussis vaccine, and served as the lab manager of the Veregge Lab at San Jose State University at the age of 16. At age 19, he had a first author publication on using protist models for high throughput drug screening using flow cytometry. Nikhil also has a passion for education, regularly writing technical posts on his blog, teaching machine learning tutorials at hackathons, and recently, received the Young Innovator Award from the Gordon and Betty Moore Foundation for re-invisioning the traditional chemistry set using augmented reality.

As an executive managing teams deploying deep learning in various ways, I need to stay up to date on effective implementation strategies. This book has been an invaluable resource for me in our recent machine learning projects. Having read this book, I gained situational awareness, and respect for nuance in deploying a deep learning application. Now, I am better equipped to understand the challenges faced by my team and more effective in guiding them toward a solution.

Buduma is a talented teacher. The book is highly intuitive and easy to read. **Con (singular):** as tensorflow is updated regularly, a lot of the code in this book is already deprecated. If you are a hands-on person, figuring out rewrites will only serve to better your understanding of the code.

I was looking forward to this for some time, hoping it would be a clean practical description of how to implement a basic deep network. This is more of an introductory tutorial on the basics that uses the TensorFlow library for illustrations. Though this isn't what I was looking for, I assume the objective was to produce a good such tutorial. But it's written in a wordy manner, spends many pages reviewing basic machine learning, non-deep networks, and misc topics like reinforcement learning, defers details of algorithms such as AdaGrad and many other things to the Tensorflow implementations, and could use more/better examples and justifications of why things are done a certain way. On the other hand, the presentation is clear, and you may get pointers to some new ideas, particularly to try with a library such as Tensorflow, though need to research further on your own. This is much better than other "simple" deep learning books that I have seen, but I can't think of what this brings to someone who got thru Bengio's book.

Very good intro to the ideas behind deep learning systems. I'm a beginner in this field, I'm still only part of the way through the text but I think I'll finish it and learn a lot. One issue right now, that could be easily solved, is that using the accompanying source code at github can be frustrating. The book itself is not clear about exactly which version of Python and TensorFlow is required to run the examples. The downloaded code I tried so far uses Python 2.x, and a much earlier version of TensorFlow than I'm using (1.2), it seems to use pre-1.0 tensorflow but I could be wrong.. This is not a huge problem if you have some patience and time, but it gets in the way of progressing through the book. I recommend this as a good, short intro to deep learning.

The Fundamentals of Deep Learning is a fantastic guide to deep learning for anyone looking for a solid understanding of an emerging field moving at a breakneck pace. It begins with a clear introduction to the primary building blocks (perceptron, basic calculus, gradient descent, etc.) and ends with cutting-edge techniques in visual recognition, natural language processing, and artificial intelligence (reinforcement learning). It conveys these complex and incredibly useful ideas through a steady progression of topics that is surprisingly digestible -- along with exercises for the reader to actually implement what they learn. I'd highly recommend this book for anyone looking for a comprehensive yet concise overview of modern deep learning.

If you're looking to get into deep learning, this is the book for you. It balances and in depth education with a breadth of examples, helping readers really understand the fundamentals. I don't have significant AI/ML experience, but this book still felt like it was at a level I could handle. You

won't regret buying it!

Fundamentals of Deep Learning: Designing Next-Generation Machine Intelligence Algorithms I have taken several Coursera/Udemy/Udacity courses, study CS & Deep Learning at Princeton, and have worked in Data Science at Uber, and can confidently say this is one of the best resources out there to delve into deep learning. Buduma does an excellent job in this book of balancing intuition and mathematics - starting off with great use cases, common-sense explanations of algorithms, thorough explanation of the mathematics behind them, and then explaining and clearly delineating implementation and examples in Python/Tensorflow. Beyond explaining concepts through these multiple levels, he also includes great visuals on basically every other page to bolster his already clear prose. Rarely do you see a resource that takes you from use case to detailed code implementation without skipping a step. I think many of the criticisms of this book are due to people's very specific expectations and complaints about Tensorflow. On the first point, this is a great intro and overview of deep learning, but you need some basic familiarity with linear algebra; on the other end of the spectrum, this isn't a research paper that's going to be absurdly intricate. On the second point, Tensorflow changes often, and it's not about the code - it's about the concepts, which are very well-explained! If you have basic programming experience, changing Tensorflow is not an issue. I highly recommend! It is a great overview that already gives a reader a lot of power to use deep learning, and opens the door to further study.

The book is actually quite good as an introduction to deep learning and author is very knowledgeable. I agree with other reviewers about the code on github falling behind on pre-1.0 Tensorflow. I have tried to change the code by making cosmetic changes to API calls but this approach stops to work in chapter 5 when there is no error and code is not running (i.e. both CPU and GPU are sitting idle). Unless I am willing to spend a lot of time to debug sample code, I cannot get maximum benefits of the book except as a reference book. Now I am switching to another book with TensorFlow 1.0 support.

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

Fundamentals of Deep Learning: Designing Next-Generation Machine Intelligence Algorithms
Algorithms for Reinforcement Learning (Synthesis Lectures on Artificial Intelligence and Machine Learning)
Machine Learning: For Beginners: Definitive Guide for Neural Networks, Algorithms, Random Forests and Decision Trees Made Simple (Machine Learning, Book 1)
Emotional Intelligence: Why You're Smarter But They Are More Successful (Emotional intelligence

leadership, Emotional Quotient, emotional intelligence depression, emotional intelligence workbook)

MATLAB Deep Learning: With Machine Learning, Neural Networks and Artificial Intelligence

Introduction to Deep Learning Using R: A Step-by-Step Guide to Learning and Implementing Deep Learning Models Using R

Fundamentals of Machine Learning for Predictive Data Analytics: Algorithms, Worked Examples, and Case Studies (MIT Press)

Machine Learning: Fundamental Algorithms for Supervised and Unsupervised Learning With Real-World Applications

Bundle of Algorithms in C++, Parts 1-5: Fundamentals, Data Structures, Sorting, Searching, and Graph Algorithms (3rd Edition) (Pts. 1-5)

Emotional Intelligence: 3 Manuscripts - Emotional Intelligence Definitive Guide, Mastery, Complete Step by Step Guide (Social Engineering, Leadership, ... (Emotional Intelligence Series Book 4)

Machine Learning: An Algorithmic Perspective, Second Edition (Chapman & Hall/Crc Machine Learning & Pattern Recognition)

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)

Practical Machine Learning with H2O: Powerful, Scalable Techniques for Deep Learning and AI

Deep Thinking: Where Machine Intelligence Ends and Human Creativity Begins

The Next IQ: The Next Level of Intelligence for 21st Century Leaders

The Fifty-Year Mission: The Next 25 Years: From the Next Generation to J. J. Abrams: The Complete, Uncensored, and Unauthorized Oral History of Star Trek

Numerical Algorithms: Methods for Computer Vision, Machine Learning, and Graphics

Machine Learning for Hackers: Case Studies and Algorithms to Get You Started

[Contact Us](#)

[DMCA](#)

[Privacy](#)

[FAQ & Help](#)