


The book was found

Deep Learning In Python: Master Data Science And Machine Learning With Modern Neural Networks Written In Python, Theano, And TensorFlow (Machine Learning In Python)



Synopsis

Deep Learning Deep learning is making waves. At the time of this writing (March 2016), Google's AlphaGo program just beat 9-dan professional Go player Lee Sedol at the game of Go, a Chinese board game. Experts in the field of Artificial Intelligence thought we were 10 years away from achieving a victory against a top professional Go player, but progress seems to have accelerated! While deep learning is a complex subject, it is not any more difficult to learn than any other machine learning algorithm. I wrote this book to introduce you to the basics of neural networks. You will get along fine with undergraduate-level math and programming skill. All the materials in this book can be downloaded and installed for free. We will use the Python programming language, along with the numerical computing library Numpy. I will also show you in the later chapters how to build a deep network using Theano and TensorFlow, which are libraries built specifically for deep learning and can accelerate computation by taking advantage of the GPU. Unlike other machine learning algorithms, deep learning is particularly powerful because it automatically learns features. That means you don't need to spend your time trying to come up with and test "kernels" or "interaction effects" - something only statisticians love to do. Instead, we will let the neural network learn these things for us. Each layer of the neural network learns a different abstraction than the previous layers. For example, in image classification, the first layer might learn different strokes, and in the next layer put the strokes together to learn shapes, and in the next layer put the shapes together to form facial features, and in the next layer have a high level representation of faces. On top of all this, deep learning is known for winning its fair share of Kaggle contests. These are machine learning contests that are open to anyone in the world who are allowed to use any machine learning technique they want. Deep learning is that powerful. Do you want a gentle introduction to this "dark art", with practical code examples that you can try right away and apply to your own data? Then this book is for you. Who is this book NOT for? Deep Learning and Neural Networks are usually taught at the upper-year undergraduate level. That should give you some idea of the type of knowledge you need to understand this kind of material. You absolutely need exposure to calculus to understand deep learning, no matter how simple the instructor makes things. Linear algebra would help. I will assume familiarity with Python (although it is an easy language to pick up). You will need to have some concept of machine learning. If you know about algorithms like logistic regression already, this book is perfect for you. If not, you might want to check out my "prerequisites" book, at: <http://amzn.com/B01D7GDRQ2> On the other hand, this book is more like a casual primer than a dry textbook. If you are looking for material on more advanced topics, like LSTMs, convolutional neural

networks, or reinforcement learning, I have online courses that teach this material, for example: <https://www.udemy.com/deep-learning-convolutional-neural-networks-theano-tensorflow> New libraries like TensorFlow are being updated constantly. This is not an encyclopedia for these libraries (as such a thing would be impossible to keep up to date). In the one (1!!!) month since the book was first published, no less than THREE new wrapper libraries for TensorFlow have been released to make coding deep networks easier. To try and incorporate every little update would not only be impossible, but would continually cause parts of the book to be obsolete. Nobody wants that. This book, rather, includes fundamentals. Understanding these building blocks will make tackling these new libraries and features a piece of cake - that is my goal.

Book Information

File Size: 258 KB

Print Length: 58 pages

Simultaneous Device Usage: Unlimited

Publication Date: March 11, 2016

Sold by: Digital Services LLC

Language: English

ASIN: B01CVJ19E8

Text-to-Speech: Enabled

X-Ray: Not Enabled

Word Wise: Not Enabled

Lending: Not Enabled

Enhanced Typesetting: Not Enabled

Best Sellers Rank: #105,510 Paid in Kindle Store (See Top 100 Paid in Kindle Store) #14

in Books > Computers & Technology > Computer Science > AI & Machine Learning > Neural Networks #64 in Kindle Store > Kindle Short Reads > 90 minutes (44-64 pages) > Computers & Technology #105 in Books > Computers & Technology > Computer Science > AI & Machine Learning > Intelligence & Semantics

Customer Reviews

Awful, awful book. Terrible writing, no insight, and incredibly brief -- the section on Tensorflow is four pages and says, of a 10 line example: "While these functions probably all seem unfamiliar and foreign, with enough consultation of the TensorFlow documentation, you will acclimate yourself to them." Yeah. Even the simplest examples are handed off to the documentation. Why even write this

book except to benefit from the mistaken 1-clicks of programmers who don't know any better? Throw a dart at any of the search results for "Theano tutorial" or "Tensorflow tutorial" on Google and you will get something much, much better than this.

It talks about almost nothing, only some very basic ideas about neural network. It has no relevant with deep learning. The advisement about the author's online pay course is everywhere.

It's like a summary, introduction section to another book, not a preview even. This book, includes popular keywords from the current ML scene, briefly summarizes the concepts not in a understandable manner because there is no depth to the subjects. Includes simple toy examples that you can find anywhere on web just by stroking your keyboard at once. Better off.

Good for beginners or some that just needed topics or terms to further study. It is an introductory book for new learner in NN

Pros: concise
Cons: feel like a long internet article that is paid per view

Good book for a beginner. And that's what matters. Being able to break into this dark science. Making it accessible. No need for thick math formulas etc., just intuitive understanding

Every programming language has its pros and cons, but I decided to dig deeper into Python simply because of its rising popularity and greater use throughout every IT platform. But this guide has merely allowed me to introduce myself with the pure basics of the language and I'm still not sure whether I am ready for something more.

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

Deep Learning in Python: Master Data Science and Machine Learning with Modern Neural Networks written in Python, Theano, and TensorFlow (Machine Learning in Python) Convolutional Neural Networks in Python: Master Data Science and Machine Learning with Modern Deep Learning in Python, Theano, and TensorFlow (Machine Learning in Python) Unsupervised Deep Learning in Python: Master Data Science and Machine Learning with Modern Neural Networks written in Python and Theano (Machine Learning in Python) Deep Learning: Natural Language Processing in Python with Recursive Neural Networks: Recursive Neural (Tensor) Networks in Theano (Deep Learning and Natural Language Processing Book 3) Deep Learning: Recurrent

Neural Networks in Python: LSTM, GRU, and more RNN machine learning architectures in Python and Theano (Machine Learning in Python) Deep Learning in Python Prerequisites: Master Data Science and Machine Learning with Linear Regression and Logistic Regression in Python (Machine Learning in Python) Deep Learning Step by Step with Python: A Very Gentle Introduction to Deep Neural Networks for Practical Data Science Deep Learning: Natural Language Processing in Python with Word2Vec: Word2Vec and Word Embeddings in Python and Theano (Deep Learning and Natural Language Processing Book 1) Deep Learning: Natural Language Processing in Python with GLoVe: From Word2Vec to GLoVe in Python and Theano (Deep Learning and Natural Language Processing) Deep Learning for Business with R: A Very Gentle Introduction to Business Analytics Using Deep Neural Networks Python: Python Programming Course: Learn the Crash Course to Learning the Basics of Python (Python Programming, Python Programming Course, Python Beginners Course) Neural Smithing: Supervised Learning in Feedforward Artificial Neural Networks (MIT Press) Unsupervised Machine Learning in Python: Master Data Science and Machine Learning with Cluster Analysis, Gaussian Mixture Models, and Principal Components Analysis A collection of Advanced Data Science and Machine Learning Interview Questions Solved in Python and Spark (II): Hands-on Big Data and Machine ... Programming Interview Questions) (Volume 7) Data Analytics: What Every Business Must Know About Big Data And Data Science (Data Analytics for Business, Predictive Analysis, Big Data) Artificial Intelligence for Humans, Volume 3: Deep Learning and Neural Networks Deep Learning Neural Networks: Design and Case Studies Principles of Neural Science, Fifth Edition (Principles of Neural Science (Kandel)) Data Analytics: Practical Data Analysis and Statistical Guide to Transform and Evolve Any Business. Leveraging the Power of Data Analytics, Data ... (Hacking Freedom and Data Driven) (Volume 2) Hello World en TensorFlow: Para iniciarse en la programación del Deep Learning (Spanish Edition)

[Dmca](#)