Article from

Predictive Analytics & Futurism

December 2017

Issue 16

Teach Yourself Deep Learning

By Michael Niemerg

Deep learning has made incredible progress in several machine learning domains in the past few years. From image classification to speech recognition, to self-driving cars, to even art, the improvements in accuracy keep coming and the range of applications keeps widening. Deep learning is nothing more than a rebranded and "scaled up" version of the same neural networks that have been around for decades, but with increases in the size of both the neural networks themselves and the size of data they are trained with, along with some technical refinements thrown into the mix. These improvements have drastically improved the usefulness of neural networks, turning what was an intellectually interesting but impractical algorithm into deep learning—the state of the art for machine learning and something that is becoming a part of our daily lives. I personally talk to a neural network everyday (via Alexa on my Amazon Echo) and use it to create some really cool filters for pictures (check out an app called Prisma). And while deep learning has recently been subject to some overblown hype—I'm not too worried myself about a hostile artificial intelligence (AI) takeover—it seems a pretty safe bet that deep learning will continue to improve in capability and usage in the upcoming years.

Given the recent rise in the popularity of deep learning, I thought it would be a good time to dive into the topic and really understand what deep learning is all about. However, as the popularity of the topic has grown, so has the number of books, online courses, tutorials and references. When I started researching all the various introductions and resources for neural networks, I got a bit overwhelmed.

Instead of letting that daunt me, I embraced the feeling and dove right in.

Starting with only a cursory idea of what deep learning is and some practical experience in other predictive modeling algorithms, I've subsequently spent a lot of time with various deep learning resources trying to learn the basics and getting to know the language of the field. To help anyone else interested in learning about neural networks, I decided to rate and review some of the sources I've used (listed in the approximate order of their technical depth). Don't be intimidated. You absolutely don't need to use this many resources to get an introduction to neural networks; one or two should suffice. However, the topic of deep learning as a whole is rather deep itself and has many advanced applications; I discovered that at the end of the day each one of these books taught me some new facet about neural networks that the others didn't. In that respect, the more viewpoints you can approach deep learning from, the better.

Make Your Own Neural Network

Author: Tariq Rashid
Rating: 3/5

This book is perfect for: An absolute beginner who doesn't want to go too deep or anyone whose college math is a fuzzy memory and wants a gentle launching point.

This book is a rather short, breezy and super-gentle introduction to neural networks. The only real prerequisites are high

school math. I enjoy the conversational tone that the book is written in. It beats the more "stodgy" tone that many textbooks employ. This stuff is fun, why do we have to talk about it in such a serious and boring manner?

That said, this book is such a gentle introduction that it can be a little painful if you know much of anything about predictive models or programming in Python.

Deep Learning With Python

Author: Jason Brownlee
Rating: 4/5

This book is perfect for: Someone who wants to dive right into code and achieve almost immediate results. Someone who wants a practical companion to a more theory-oriented work.

This book was written by Jason Brownlee, at Machine Learning Mastery, who has a small suite of books, email courses, and blog topics on machine learning. I love lots of Jason's output—he does a great job of distilling topics down to their most basic forms. One of the things about Jason's general approach is his strong emphasis on intuition, pragmatic programming, and avoidance of heavy theory or mathematics.

Jason's philosophy really shows in this book. The explanations get right to the heart of the matter, giving you concise, easily understandable examples. The challenge is that this book always stays at the very shallow end of the pool. At times, it never even leaves the kiddie pool.

This is not a book that will give you a thorough understanding of deep learning. It will not even give you all that much in the way of intuition. It will, however, give you enough to get by if what you want to do is create models. What it excels at more than any other resource I've explored is being a great cookbook on how to actually code useful neural networks in painless fashion.

Neural Networks for Applied Sciences and Engineering

Author: Sandhya Samarasinghe
Rating: 3.5/5

This book is perfect for: Someone who wants a slow, steady introduction to neural networks with plenty of examples and visualizations.

This is my favorite introductory text to get a feel for the basic architecture of neural networks. Mostly, I like it because it is so gentle and deliberate. It contains a liberal amount of visuals and many examples done from the bottom up with thorough explanations. The explanations are simple and the author Sandhya Samarasinghe does a great job of explaining the concepts in an unassuming manner while being very careful about describing her terminology.

Based on the book title, it would seem that there would be a tilt toward engineering applications, but in reality the book is rather agnostic about application and doesn't really deal with the applied sciences or engineering in any direct way. With respect to software, this book is also agnostic, not making mention of it at all.

Because of the gentle approach and verbosity, this book makes a better learning aid than reference. Also, it's relatively old, first published over 10 years ago, and does feel dated. With deep learning being such a dynamic field, this text is simply too old to capture many of the new advancements over the last decade. That said, for the beginner, the age of this text isn't truly a concern because the basics detailed here haven't really changed. In terms of more advanced topics, this book is a little light, covering only three "intermediate-level" neural network topics: Bayesian parameter estimation, self-organizing maps and nonlinear time-series analysis.

Deep Learning With Python

Author: Francois Chollet
"Preview" rating: 4.5/5

This book is perfect for: Someone who wants a one-stop shop to both understand and code neural networks.

This book has yet to be released as of the time of this writing (it's due in October), but the publisher Manning allows an "early look" if you preorder, so I have read draft versions of all the chapters. Consider this less of a review and more of a preview, though I've read virtually all of what will consist of the finished product.

This book is authored by Francois Chollet, the author of the Keras package. Keras is a relatively new programming wrapper that sits atop TensorFlow, CNTK or Theano, and is quickly becoming the "go-to method" to program neural networks. This book shuns mathematical formulas, preferring to explain algorithms programmatically in Keras code. At times, I find this approach preferable and at other times a simple formula wouldn't have been such a bad thing.

Overall, this book is a great compromise between understanding neural networks and actually programming them, which none of the other books I've looked at quite manages. However, it does require some basic familiarity with Python (or a willingness to pick it up). It does a great job of giving you explanations with its code (instead of throwing you huge chunks of it 50 lines at a time) and also does a good job of interspersing images wisely into the text.

This may already be my favorite book on neural networks and it's not even finished, so it should only get better as it moves toward its final published form. I'm eagerly anticipating my copy of the finished product.

Fundamentals of Deep Learning: Designing Next Generation Machine Learning Algorithms

Author: Nikhil Buduma
Rating: 2/5

This book is perfect for: Someone who prefers TensorFlow and wants some functioning code to play with for some advanced applications.

This book was just released over the summer. Being so recent, it has some good explanations on recent developments in neural networks that are hard to get elsewhere and it even has code on how to create some of these architectures. For me, that was the highlight of this book.

However, I have several challenges with this book. Some of the flow of the presentation and ordering of topics I found a bit unintuitive. The book also contains large masses of TensorFlow code with little commentary. However, its worst transgression is its poor formatting. When viewed on a kindle, the code wraps all over the place, making it very challenging to follow what it is even supposed to be doing. There is a similar problem with the mathematical notation: it is too small. There were several formulas I simply couldn't read no matter how hard I squinted. While you can increase the font size on a Kindle, you can't increase the size of formulas. These issues made the book a struggle to follow at times.

I don't want to sound overly harsh. I quite enjoyed sections of this book, but it probably serves better as a supplement than as someone's first introduction to neural networks. Definitely avoid the kindle version.

Deep Learning

Authors: Ian Goodfellow, Yoshua Bengio and Aaron Courville
Rating: 4.5/5

This book is perfect for: Someone who is comfortable with more mathematical and technical depth and wants a broad and comprehensive exposure to different uses and architectures for neural networks. Someone on a budget.

This book was my launching pad into neural networks. In retrospect, there are gentler starting points. But while this book gets denser the further you progress, it doesn't exactly throw you into the deep end of the pool at the onset.

This book has three meaty sections. The first is a background section that serves as a good refresher on machine learning basics for those of us who haven't taken a probability course and are unfamiliar with machine learning. The second section focuses on the basics of neural networks. Finally, the third section talks about recent advancements in neural network topics and more advanced applications. Because of this structure, this book offers something for everyone, from absolute beginners to advanced practitioners and researchers.

Overall, this book is generally quite readable, with a few really dense sections being the exception. However, it is a tougher entry point for newbies than some of the other books I've looked at. At times, it can get a bit verbose and covers some more obscure topics. At other times it was even hard for a newbie like myself to grasp the key takeaways versus the more ancillary details from some of the advanced sections. A big downside of this book is that there are no real exercises or examples, either theoretical or programming. It doesn't provide any opportunity to grapple with and learn the material.

What it excels at more than any other book in this list is in the amount of advanced topics that it covers and the sheer volume of material it covers. It also does a better job of getting you closer to the actual math and mechanics of neural networks than the other books in this list.

Overall, this book contains a lot to recommend it. I plan on revisiting it sometime soon. It is also downloadable for free online, so the price is just right.

ONLINE RESOURCES

I have a preference for learning via reading and so haven't completed any deep learning massive open online courses (MOOCs). However, for those who prefer the lecture format, I do want to call out several online learning resources that I've heard very positive things about. Andrew Ng, a powerhouse in the artificial intelligence community, has a new "deeplearning.ai" project that has a specialization on *Coursera* consisting of five mini-courses. Geoffrey Hinton, whose advances in the field are primary reasons behind the recent renaissance of neural networks, is also on *Coursera* with "Neural Networks for Machine Learning." Additionally, there is a course titled "Deep Learning A-Z: Hands-On Artificial Neural Networks" on *Udemy* that looks really promising. There are many more out there. My advice is to do your research. With so many options, there's bound to be something out there that fits your learning style and goals. ■

Michael Niemerg, FSA, MAAA, is an actuary at Milliman in Chicago. He can be reached at *michael.niemerg@milliman.com.*