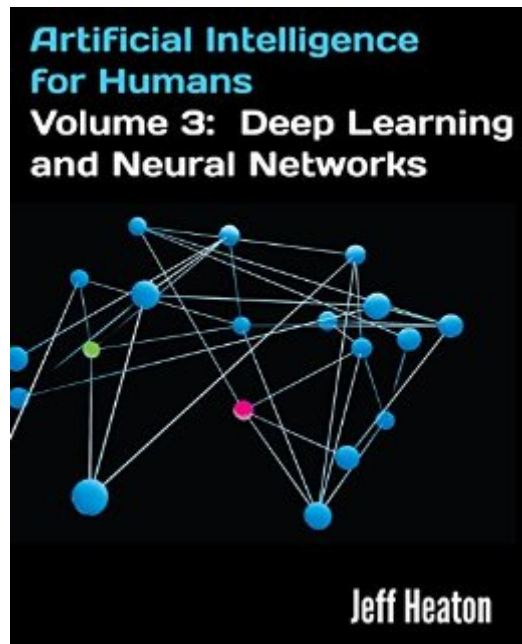# Artificial Intelligence For Humans, Volume 3: Deep Learning And Neural Networks

Synopsis

Neural networks have been a mainstay of artificial intelligence since its earliest days. Now, exciting new technologies such as deep learning and convolution are taking neural networks in bold new directions. In this book, we will demonstrate the neural networks in a variety of real-world tasks such as image recognition and data science. We examine current neural network technologies, including ReLU activation, stochastic gradient descent, cross-entropy, regularization, dropout, and visualization.

Book Information

File Size: 13653 KB

Print Length: 345 pages

Simultaneous Device Usage: Unlimited

Publisher: Heaton Research, Inc.; 1 edition (November 17, 2015)

Publication Date: November 17, 2015

Sold by:Â Digital Services LLC

Language: English

ASIN: B0184WRDEQ

Text-to-Speech: Enabled

X-Ray: Not Enabled

Word Wise: Enabled

Lending: Not Enabled

Enhanced Typesetting: Not Enabled

Best Sellers Rank: #40,856 Paid in Kindle Store (See Top 100 Paid in Kindle Store)   #5 inÂ Books > Computers & Technology > Computer Science > AI & Machine Learning > Computer Vision & Pattern Recognition   #6 inÂ Books > Computers & Technology > Computer Science > AI & Machine Learning > Neural Networks   #223 inÂ Kindle Store > Kindle eBooks > Computers & Technology

Customer Reviews

The book describes few concepts. A beginner would expect to see a bigger picture rather than a seemingly random set of methods or models. A professional would look for detailed explanation of concepts.The book provides neither. Descriptions are shallow, formulas contain errors, pseudo code is often wrong. Some description would better be replaced with concise and clear formulas, some formulas require explanations. The book needs a proof editing. However, the bigger problem

is how the author presents the material.I'd recommend a much better book to start with: "Neural Networks and Deep Learning" by Michael Nielsen.

I ran through the book quickly (kindle) and it seems good although I expected much more mathematics involved (Integrals, derivatives, linear algebra). I hope that a new version or a new book will go more in depth with the actual problems...

The book is more like a quick compilation of a college student's note. Concepts are presented in a relatively isolated manner; connections between concepts are, for the large part, missing. Furthermore, if the materials presented are rather shallow like in this book, readers will expect to see a strong emphasis on, or hands on exercises of, practical applications. But this book doesn't seem to help much in that regard either, despite what the book claims. The book does give introduction to a bunch of models, which can be useful for a beginner. But at least this edition I wouldn't suggest any one to buy because of poor editing. For example, the pseudo code fragments listed in the book are often wrong, or different than what is described by the text. References to formulas, etc need some proof reading, too.

The book gives a brief description of the large number of types and techniques for applying neural networks, but never gives the reader any examples of their implementation and how effective they are at solving any of the classic problems. Instead the reader is fed trivia like the definitions of sgn(x), mean, standard deviation, partial derivative or the amazing fact that the second derivative is the derivative of the first derivative. There is nothing like a tutorial anywhere in the book.

Very basic high school level introduction to deep learning. The math although elementary is also poor rendered (no LaTeX here) and sometimes wrong (starting with equation 1.1). On the bright side it mentions a wide range of topics in the field (to be studied elsewhere)

The author explained the algorithms/concepts poorly. When the author presents an algorithm, he simply puts some equations there, without giving any clue of how they are used and why they are needed. The readers have to figure out them by reading online tutorials.The book lacks of the intuition and the mathematical fundamentals.

Definitly worth its price for the content and the GitHub examples. The content is pretty minimalistic

in of itself and is insufficient to have a full understand of deep learning, but there are some really good pointers and simplified explanations to help you get there. I enjoyed the fact that each concept was presented seperatly which helped me get a better undestanding on how neural networks works and what is there underlying philosophie.

DO NOT buy this book. Its content quality is just above Scientific America article. Useless for professionals.

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) Artificial Intelligence for Humans, Volume 3: Deep Learning and Neural Networks Neural Smithing: Supervised Learning in Feedforward Artificial Neural Networks (MIT Press) Deep Learning: Recurrent Neural Networks in Python: LSTM, GRU, and more RNN machine learning architectures in Python and Theano (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) Convolutional Neural Networks in Python: Master Data Science and Machine Learning with Modern Deep Learning in Python, Theano, and TensorFlow (Machine Learning in Python) Deep Learning in Python: Master Data Science and Machine Learning with Modern Neural Networks written in Python, Theano, and TensorFlow (Machine Learning in Python) Deep Learning for Business with R: A Very Gentle Introduction to Business Analytics Using Deep Neural Networks Deep Learning Step by Step with Python: A Very Gentle Introduction to Deep Neural Networks for Practical Data Science Java: Artificial Intelligence; Made Easy, w/ Java Programming; Learn to Create your * Problem Solving * Algorithms! TODAY! w/ Machine Learning & Data Structures (Artificial Intelligence Series) Javascript Artificial Intelligence: Made Easy, w/ Essential Programming; Create your * Problem Solving * Algorithms! TODAY! w/ Machine Learning & Data Structures (Artificial Intelligence Series) Deep Learning Neural Networks: Design and Case Studies Elements of Artificial Neural Networks (Complex Adaptive Systems) Neural Network Training Using Genetic Algorithms (Series in Machine Perception and Artificial Intelligence) Humans Need Not Apply: A Guide to Wealth and Work in the Age of Artificial Intelligence Fusion of Neural Networks, Fuzzy Systems and Genetic Algorithms: Industrial Applications (International Series on Computational Intelligence) Gene Expression Programming: Mathematical Modeling by an Artificial Intelligence (Studies in Computational Intelligence) 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) Principles of Neural Science, Fifth Edition (Principles of Neural Science (Kandel)) Deep Learning: Natural Language Processing in Python with GLoVe: From Word2Vec to GLoVe in Python and Theano (Deep Learning and Natural Language Processing)