


The book was found

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)



Synopsis

The first 2 books in this series focused on word embeddings using 2 novel techniques: Word2Vec and GloVe. In this book, we return to a classic NLP problem: sentiment analysis. Classification performance on the sentiment analysis task had plateaued for many years, due to not being able to handle negation, which is essentially because existing models failed to account for the structure of language. The bag-of-words vectors for "I love this movie", "I don't love this movie", and "Don't you love this movie" are very similar. In this book, we return to the fundamentals of language - the parse tree - and structure our neural networks to mirror the tree. It makes sense that a neural network created to classify language would have the same structure as language. These neural networks are called "recursive neural networks" and I will show you how they work both mathematically and with a full implementation in Theano. A naive solution to recursive neural networks would be to use recursion to implement them. This is however, a very poor solution because both Theano and TensorFlow require you to compile a graph of the neural network. If every sentence is a different tree, then every sentence will require a different neural network graph, which would be very inefficient for both Theano or TensorFlow to compute. No one would blame you for attempting this solution first. In fact, I will demonstrate why it's bad by having you run code that implements it. Once you understand why recursion is not ideal for recursive neural networks, I will show you a "trick" that will help you implement them more efficiently. We will then run the recursive neural net on our sentiment analysis data and achieve state-of-the-art performance. Finally, we discuss a modification to the vanilla recursive neural network called the recursive neural tensor network or RNTN. It was invented by the guys at Stanford, who have created and published many NLP tools throughout the years that are now considered standard. I will show you how the model is structured mathematically and then I will show you how to implement it in Theano. You'll see that it's just a simple modification to our recursive neural network. Amazingly, all the technologies we discuss in this book can be downloaded and installed for FREE. That means all you need to invest after purchasing this book is your effort and your time. The only prerequisites are that you are comfortable with Python, Numpy, and Theano coding and you know the basics of deep learning. "Hold up... what's deep learning and all this other crazy stuff you're talking about?" If you are completely new to deep learning, you might want to check out my earlier books and courses on the subject, since they are required in order to understand this book: Deep Learning in Python <https://www.udemy.com/data-science-deep-learning-in-python> Deep Learning in Python Prerequisites <https://www.udemy.com/data-science-logistic-regression-in-python> Much like how IBM's Deep

Blue beat world champion chess player Garry Kasparov in 1996, Google's AlphaGo recently made headlines when it beat world champion Lee Sedol in March 2016. What was amazing about this win was that experts in the field didn't think it would happen for another 10 years. The search space of Go is much larger than that of chess, meaning that existing techniques for playing games with artificial intelligence were infeasible. Deep learning was the technique that enabled AlphaGo to correctly predict the outcome of its moves and defeat the world champion. Deep learning progress has accelerated in recent years due to more processing power (see: Tensor Processing Unit or TPU), larger datasets, and new algorithms like the ones discussed in this book. Book 1 in the series can be found at: <https://www.amazon.com/dp/B01KQ0ZN0A> Book 2 in the series can be found at: <https://www.amazon.com/dp/B01KRBOO4Y>

Book Information

File Size: 228 KB

Print Length: 69 pages

Simultaneous Device Usage: Unlimited

Publication Date: August 20, 2016

Sold by: Digital Services LLC

Language: English

ASIN: B01KS5AEXO

Text-to-Speech: Enabled

X-Ray: Not Enabled

Word Wise: Not Enabled

Lending: Not Enabled

Enhanced Typesetting: Enabled

Best Sellers Rank: #86,347 Paid in Kindle Store (See Top 100 Paid in Kindle Store) #6 in Books

> Computers & Technology > Computer Science > AI & Machine Learning > Natural Language

Processing #11 in Books > Computers & Technology > Computer Science > AI & Machine

Learning > Neural Networks #64 in Kindle Store > Kindle Short Reads > Two hours or more

(65-100 pages) > Computers & Technology

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

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: 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: 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)
Python: Python Programming Course: Learn the Crash Course to Learning the Basics of Python
(Python Programming, Python Programming Course, Python Beginners Course) 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 Python:
Python Programming For Beginners - The Comprehensive Guide To Python Programming:
Computer Programming, Computer Language, Computer Science (Machine Language) Deep
Learning for Business with R: A Very Gentle Introduction to Business Analytics Using Deep Neural
Networks Neural Smithing: Supervised Learning in Feedforward Artificial Neural Networks (MIT
Press) Python: Python Programming For Beginners - The Comprehensive Guide To Python
Programming: Computer Programming, Computer Language, Computer Science Learn Python in
One Day and Learn It Well: Python for Beginners with Hands-on Project. The only book you need to
start coding in Python immediately Beginning Python Programming: Learn Python Programming in 7
Days: Treading on Python, Book 1 Maya Python for Games and Film: A Complete Reference for
Maya Python and the Maya Python API Artificial Intelligence for Humans, Volume 3: Deep Learning
and Neural Networks Deep Learning Neural Networks: Design and Case Studies Speech and
Language Processing: An Introduction to Natural Language Processing, Computational Linguistics
and Speech Recognition

[Dmca](#)