Antonio Gulli, Sujit Pal

Deep Learning with Keras

Implement neural networks with Keras on Theano and TensorFlow

Packt>

Deep Learning with Keras

This book starts by introducing you to supervised learning algorithms such as simple linear regression, the classical multilayer perceptron and more sophisticated deep convolutional networks. You will also explore image processing with recognition of hand written digit images, classification of images into different categories, and advanced objects recognition with related image annotations. An example of identification of salient points for face detection is also provided. Next you will be introduced to recurrent networks, which are optimized for processing sequence data such as text, audio or time series. Following that, you will learn about unsupervised learning algorithms such as autoencoders and the very popular Generative Adversarial Networks (GANs). You will also explore non-traditional uses of neural networks as style transfer.

Finally, you will look at reinforcement learning and its application to AI game playing, another popular direction of research and application of neural networks.

Things you will learn:

- Optimize step-by-step functions on a large neural network using the backpropagation algorithm
- Fine-tune a neural network to improve the quality of results
- Use deep learning for image and audio processing
- Use Recursive Neural Tensor Networks (RNTNs) to outperform standard word embedding in special cases
- Identify problems for which Recurrent Neural Network (RNN) solutions are suitable
- Explore the process required to implement autoencoders
- Evolve a deep neural network using reinforcement learning

Packt>
www.packtpub.com

$ **49.99** US
£ **41.99** UK

Prices do not include local sales
Tax or VAT where applicable

ISBN 978-1-78712-842-2
54999
9 781787 128422

Deep Learning with Keras

Antonio Gulli, Sujit Pal

Deep Learning with Keras

Implement neural networks with Keras on Theano and TensorFlow

Antonio Gulli

Sujit Pal

BIRMINGHAM - MUMBAI

Deep Learning with Keras

First published: April 2017

Production reference: 1240417

Published by Packt Publishing Ltd.
Livery Place
35 Livery Street
Birmingham
B3 2PB, UK.
ISBN 978-1-78712-842-2

www.packtpub.com

Credits

Authors

Antonio Gulli

Sujit Pal

Reviewers

Mike Dahlin

Nick McClure

Corrado Zocollo

Commissioning Editor

Amey Varangaonkar

Acquisition Editor

Divya Poojari

Content Development Editor

Cheryl Dsa

Technical Editor

Dinesh Pawar

Copy Editor

Vikrant Phadkay

Project Coordinator

Nidhi Joshi

Proofreader

Safis Editing

Indexer

Francy Puthiry

Graphics

Tania Dutta

Production Coordinator

Arvindkumar Gupta

About the Authors

Antonio Gulli is a software executive and business leader with a passion for establishing and managing global technological talent, innovation, and execution. He is an expert in search engines, online services, machine learning, information retrieval, analytics, and cloud computing. So far, he has been lucky enough to gain professional experience in four different countries in Europe and managed people in six different countries in Europe and America. Antonio served as CEO, GM, CTO, VP, director, and site lead in multiple fields spanning from publishing (Elsevier) to consumer internet (Ask.com and Tiscali) and high-tech R&D (Microsoft and Google).

I would like to thank my coauthor, Sujit Pal, for being a such talented colleague, always willing to help with a humble spirit. I constantly appreciate his dedication to teamwork, which made this book a real thing.
I would like to thank Francois Chollet (and the many Keras contributors) for taking the time and effort to build an awesome deep learning toolkit that is easy to use without sacrificing too much power.
I would also like to thank our editors from Packt, Divya Poojari, Cheryl Dsa, and Dinesh Pawar, and our reviewers from Packt and Google, for their support and valuable suggestions. This book would not have been possible without you.
I would like to thank my manager, Brad, and my colleagues Mike and Corrado at Google for encouraging me to write this book, and for their constant help in reviewing the content.
I would like to thank Same Fusy, Herbaciarnia i Kawiarnia in Warsaw. I got the initial inspiration to write this book in front of a cup of tea chosen among hundreds of different offers. This place is magic and I strongly recommend visiting it if you are in search of a place to stimulate creativeness (`http://www.samefusy.pl/`).
Then I would like to thank HRBP at Google for supporting my wish to donate all of this book's royalties in favor of a minority/diversity scholarship.
I would like to thank my friends Eric, Laura, Francesco, Ettore, and Antonella for supporting me when I was in need. Long-term friendship is a real thing, and you are true friends to me.
I would like to thank my son Lorenzo for encouraging me to join Google, my son Leonardo for his constant passion to discover new things, and my daughter Aurora for making me smile every day of my life. Finally thanks to my father Elio and my mother Maria for their love.

Sujit Pal is a technology research director at Elsevier Labs, working on building intelligent systems around research content and metadata. His primary interests are information retrieval, ontologies, natural language processing, machine learning, and distributed processing. He is currently working on image classification and similarity using deep learning models. Prior to this, he worked in the consumer healthcare industry, where he helped build ontology-backed semantic search, contextual advertising, and EMR data processing platforms. He writes about technology on his blog at *Salmon Run*.

I would like to thank my coauthor, Antonio Gulli, for asking me to join him in writing this book. This was an incredible opportunity and a great learning experience for me. Besides, had he not done so, I quite literally wouldn't have been here today.
I would like to thank Ron Daniel, the director of Elsevier Labs, and Bradley P Allen, chief architect at Elsevier, for introducing me to deep learning and making me a believer in its capabilities.
I would also like to thank Francois Chollet (and the many Keras contributors) for taking the time and effort to build an awesome deep learning toolkit that is easy to use without sacrificing too much power.
Thanks to our editors from Packt, Divya Poojari, Cheryl Dsa, and Dinesh Pawar, and our reviewers from Packt and Google, for their support and valuable suggestions. This book would not have been possible without you.
I would like to thank my colleagues and managers over the years, especially the ones who took their chances with me and helped me make discontinuous changes in my career.
Finally, I would like to thank my family for putting up with me these past few months as I juggled work, this book, and family, in that order. I hope you will agree that it was all worth it.

About the Reviewer

Nick McClure is currently a senior data scientist at PayScale Inc. in Seattle, Washington, USA. Prior to that, he worked at Zillow and Caesars Entertainment. He got his degrees in applied mathematics from the University of Montana and the College of Saint Benedict and Saint John's University. Nick has also authored *TensorFlow Machine Learning Cookbook* by Packt Publishing.

He has a passion for learning and advocating for analytics, machine learning, and artificial intelligence. Nick occasionally puts his thoughts and musing on his blog, `fromdata.org`, or through his Twitter account at `@nfmcclure`.

www.PacktPub.com

For support files and downloads related to your book, please visit `www.PacktPub.com`.

Did you know that Packt offers eBook versions of every book published, with PDF and ePub files available? You can upgrade to the eBook version at `www.PacktPub.com` and as a print book customer, you are entitled to a discount on the eBook copy. Get in touch with us at `service@packtpub.com` for more details.

At `www.PacktPub.com`, you can also read a collection of free technical articles, sign up for a range of free newsletters and receive exclusive discounts and offers on Packt books and eBooks.

`https://www.packtpub.com/mapt`

Get the most in-demand software skills with Mapt. Mapt gives you full access to all Packt books and video courses, as well as industry-leading tools to help you plan your personal development and advance your career.

Why subscribe?

- Fully searchable across every book published by Packt
- Copy and paste, print, and bookmark content
- On demand and accessible via a web browser

Customer Feedback

Thanks for purchasing this Packt book. At Packt, quality is at the heart of our editorial process. To help us improve, please leave us an honest review on this book's Amazon page at `https://www.amazon.com/dp/1787128423`.

If you'd like to join our team of regular reviewers, you can e-mail us at `customerreviews@packtpub.com`. We award our regular reviewers with free eBooks and videos in exchange for their valuable feedback. Help us be relentless in improving our products!

Table of Contents

Preface

Hands-on deep learning with Keras is a concise yet thorough introduction to modern neural networks, artificial intelligence, and deep learning technologies designed especially for software engineers and data scientists.

Mission

The book presents more than 20 working deep neural networks coded in Python using Keras, a modular neural network library that runs on top of either Google's TensorFlow or Lisa Lab's Theano backends.

The reader is introduced step by step to supervised learning algorithms such as simple linear regression, classical multilayer perceptron, and more sophisticated deep convolutional networks and generative adversarial networks. In addition, the book covers unsupervised learning algorithms such as autoencoders and generative networks. Recurrent networks and **long short-term memory** (**LSTM**) networks are also explained in detail. The book goes on to cover the Keras functional API and how to customize Keras in case the reader's use case is not covered by Keras's extensive functionality. It also looks at larger, more complex systems composed of the building blocks covered previously. The book concludes with an introduction to deep reinforcement learning and how it can be used to build game playing AIs.

Practical applications include code for the classification of news articles into predefined categories, syntactic analysis of texts, sentiment analysis, synthetic generation of texts, and parts of speech annotation. Image processing is also explored, with recognition of handwritten digit images, classification of images into different categories, and advanced object recognition with related image annotations. An example of identification of salient points for face detection will be also provided. Sound analysis comprises recognition of discrete speeches from multiple speakers. Reinforcement learning is used to build a deep Q-learning network capable of playing games autonomously.

Experiments are the essence of the book. Each net is augmented by multiple variants that progressively improve the learning performance by changing the input parameters, the shape of the network, loss functions, and algorithms used for optimizations. Several comparisons between training on CPUs and GPUs are also provided.

How deep learning is different from machine learning and artificial intelligence

Artificial intelligence (**AI**) is a very large research field, where machines show *cognitive* capabilities such as learning behaviours, proactive interaction with the environment, inference and deduction, computer vision, speech recognition, problem solving, knowledge representation, perception, and many others (for more information, refer to this article: *Artificial Intelligence: A Modern Approach*, by S. Russell and P. Norvig, Prentice Hall, 2003). More colloquially, AI denotes any activity where machines mimic *intelligent* behaviors typically shown by humans. Artificial intelligence takes inspiration from elements of computer science, mathematics, and statistics.

Machine learning (**ML**) is a subbranch of AI that focuses on teaching computers how to learn without the need to be programmed for specific tasks (for more information refer to *Pattern Recognition and Machine Learning*, by C. M. Bishop, Springer, 2006). In fact, the key idea behind ML is that it is possible to create algorithms that learn from and make predictions on data. There are three different broad categories of ML. In supervised learning, the machine is presented with input data and desired output, and the goal is to learn from those training examples in such a way that meaningful predictions can be made for fresh unseen data. In unsupervised learning, the machine is presented with input data only and the machine has to find some meaningful structure by itself with no external supervision. In reinforcement learning, the machine acts as an agent interacting with the environment and learning what are the behaviours that generate rewards.

Deep learning (**DL**) is a particular subset of ML methodologies using **artificial neural networks** (**ANN**) slightly inspired by the structure of neurons located in the human brain (for more information, refer to the article *Learning Deep Architectures for AI*, by Y. Bengio, Found. Trends, vol. 2, 2009). Informally, the word *deep* refers to the presence of many layers in the artificial neural network, but this meaning has changed over time. While 4 years ago, 10 layers were already sufficient to consider a network as *deep*, today it is more common to consider a network as *deep* when it has hundreds of layers.