

Giancarlo Zaccone, Md. Rezaul Karim,
Ahmed Menshawy

Deep Learning with TensorFlow

Take your machine learning knowledge to the next level
with the power of TensorFlow 1.x



Packt >

Deep Learning with TensorFlow

Deep learning is the step that comes after machine learning, and has more advanced implementations. Machine learning is not just for academics anymore, but is becoming a mainstream practice through wide adoption, and deep learning has taken the front seat. As a data scientist, if you want to explore data abstraction layers, this book will be your guide. This book shows how this can be exploited in the real world with complex raw data using TensorFlow 1.x.

Throughout the book, you'll learn how to implement deep learning algorithms for machine learning systems and integrate them into your product offerings, including search, image recognition, and language processing. Additionally, you'll learn how to analyze and improve the performance of deep learning models. This can be done by comparing algorithms against benchmarks, along with machine intelligence, to learn from the information and determine ideal behaviors within a specific context.

After finishing the book, you will be familiar with machine learning techniques, in particular the use of TensorFlow for deep learning, and will be ready to apply your knowledge to research or commercial projects.

Things you will learn:

- Learn about machine learning landscapes along with the historical development and progress of deep learning
- Learn about deep machine intelligence and GPU computing with the latest TensorFlow 1.x
- Access public datasets and utilize them using TensorFlow to load, process, and transform data
- Use TensorFlow on real-world datasets, including images, text, and more
- Learn how to evaluate the performance of your deep learning models
- Using deep learning for scalable object detection and mobile computing
- Train machines quickly to learn from data by exploring reinforcement learning techniques
- Explore active areas of deep learning research and applications

Packt
www.packtpub.com

\$ 49.99 US
£ 41.99 UK

Prices do not include local sales
Tax or VAT where applicable



Deep Learning with TensorFlow

Ciancarlo Zacccone, Md. Rezaul Karim, Ahmed Menshawy



Deep Learning with TensorFlow

Take your machine learning knowledge to the next level with the power of TensorFlow 1.x

Giancarlo Zaccone
Md. Rezaul Karim
Ahmed Menshawy

Packt >

BIRMINGHAM - MUMBAI

Deep Learning with TensorFlow

Copyright © 2017 Packt Publishing

All rights reserved. No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, without the prior written permission of the publisher, except in the case of brief quotations embedded in critical articles or reviews.

Every effort has been made in the preparation of this book to ensure the accuracy of the information presented. However, the information contained in this book is sold without warranty, either express or implied. Neither the authors, nor Packt Publishing, and its dealers and distributors will be held liable for any damages caused or alleged to be caused directly or indirectly by this book.

Packt Publishing has endeavored to provide trademark information about all of the companies and products mentioned in this book by the appropriate use of capitals. However, Packt Publishing cannot guarantee the accuracy of this information.

First published: April 2017

Production reference: 1200417

Published by Packt Publishing Ltd.
Livery Place
35 Livery Street
Birmingham
B3 2PB, UK.

ISBN 978-1-78646-978-6

www.packtpub.com

Credits

Authors

Giancarlo Zaccone
Md. Rezaul Karim
Ahmed Menshawy

Copy Editor

Safis Editing

Reviewers

Swapnil Ashok Jadhav
Chetan Khatri

Project Coordinator

Shweta H Birwatkar

Commissioning Editor

Veena Pagare

Proofreader

Safis Editing

Acquisition Editor

Vinay Agrekar

Indexer

Aishwarya Gangawane

Content Development Editor

Amrita Norohna

Graphics

Tania Dutta

Technical Editor

Deepti Tuscano

Production Coordinator

Nilesh Mohite

About the Authors

Giancarlo Zaccone has more than ten years of experience in managing research projects both in scientific and industrial areas. He worked as researcher at the C.N.R, the National Research Council, where he was involved in projects relating to parallel computing and scientific visualization.

Currently, he is a system and software engineer at a consulting company developing and maintaining software systems for space and defense applications.

He is author of the following Packt volumes: *Python Parallel Programming Cookbook* and *Getting Started with TensorFlow*.

You can follow him at <https://it.linkedin.com/in/giancarlozaccone>.

Md. Rezaul Karim has more than 8 years of experience in the area of research and development with a solid knowledge of algorithms and data structures, focusing C/C++, Java, Scala, R, and Python and big data technologies such as Spark, Kafka, DC/OS, Docker, Mesos, Hadoop, and MapReduce. His research interests include machine learning, deep learning, Semantic Web, big data, and bioinformatics. He is the author of the book titled *Large-Scale Machine Learning with Spark*, Packt Publishing.

He is a Software Engineer and Researcher currently working at the Insight Center for Data Analytics, Ireland. He is also a Ph.D. candidate at the National University of Ireland, Galway. He also holds a BS and an MS degree in Computer Engineering. Before joining the Insight Centre for Data Analytics, he had been working as a Lead Software Engineer with Samsung Electronics, where he worked with the distributed Samsung R&D centers across the world, including Korea, India, Vietnam, Turkey, and Bangladesh. Before that, he worked as a Research Assistant in the Database Lab at Kyung Hee University, Korea. He also worked as an R&D Engineer with BMTech21 Worldwide, Korea. Even before that, he worked as a Software Engineer with i2SoftTechnology, Dhaka, Bangladesh.

I would like to thank my parents (Mr. Razzaque and Mrs. Monoara) for their continuous encouragement and motivation throughout my life. I would also like to thank my wife (Saroar) and my kid (Shadman) for their never-ending support, which keeps me going. I would like to give special thanks to Ahmed Menshawy and Giancarlo Zaccone for authoring this book. Without their contributions, the writing would have been impossible. Overall, I would like to dedicate this book to my elder brother Md. Mamtaz Uddin (Manager, International Business, Biopharma Ltd., Bangladesh) for his endless contributions to my life.

Further, I would like to thank the acquisition, content development and technical editors of Packt Publishing (and others who were involved in this book title) for their sincere cooperation and coordination. Additionally, without the work of numerous researchers and deep learning practitioners who shared their expertise in publications, lectures, and source code, this book might not exist at all! Finally, I appreciate the efforts of the TensorFlow community and all those who have contributed to APIs, whose work ultimately brought the deep learning to the masses.

Ahmed Menshawy is a Research Engineer at the Trinity College Dublin, Ireland. He has more than 5 years of working experience in the area of Machine Learning and Natural Language Processing (NLP). He holds an MSc in Advanced Computer Science. He started his Career as a Teaching Assistant at the Department of Computer Science, Helwan University, Cairo, Egypt. He taught several advanced ML and NLP courses such as Machine Learning, Image Processing, Linear Algebra, Probability and Statistics, Data structures, Essential Mathematics for Computer Science. Next, he joined as a research scientist at the Industrial research and development lab at IST Networks, based in Egypt. He was involved in implementing the state-of-the-art system for Arabic Text to Speech. Consequently, he was the main machine learning specialist in that company. Later on, he joined the Insight Centre for Data Analytics, the National University of Ireland at Galway as a Research Assistant working on building a Predictive Analytics Platform. Finally, he joined ADAPT Centre, Trinity College Dublin as a Research Engineer. His main role in ADAPT is to build prototypes and applications using ML and NLP techniques based on the research that is done within ADAPT.

I would like to thank my parents, my Wife Sara and daughter Asma for their support and patience during the book. Also I would like to sincerely thank Md. Rezaul Karim and Giancarlo Zaccone for authoring this book.

Further, I would like to thank the acquisition, content development and technical editors of Packt Publishing (and others who were involved in this book title) for their sincere cooperation and coordination. Additionally, without the work of numerous researchers and deep learning practitioners who shared their expertise in publications, lectures, and source code, this book might not exist at all! Finally, I appreciate the efforts of the TensorFlow community and all those who have contributed to APIs, whose work ultimately brought the machine learning to the masses.

About the Reviewers

Swapnil Ashok Jadhav is a Machine Learning and NLP enthusiast. He enjoys learning new Machine Learning and Deep Learning technologies and solving interesting data science problems and has around 3 years of working experience in these fields.

He is currently working at Haptik Infotech Pvt. Ltd. as a Machine Learning Scientist. Swapnil holds Masters degree in Information Security from NIT Warangal and Bachelors degree from VJTI Mumbai.

You can follow him at <https://www.linkedin.com/in/swapnil-jadhav-9448872a>.

Chetan Khatri is a data science researcher with having total of five years of experience in research and development. He works as a lead technology at Accionlabs India. Prior to that he worked with Nazara Games, where he lead data science practice as a principal big data engineer for Gaming and Telecom Business. He has worked with a leading data companies and a Big 4 companies, where he managed the data science practice platform and one of the Big 4 company's resources team.

He completed his master's degree in computer science and minor data science at KSKV Kachchh.

University, and was awarded as "Gold Medalist" by the Govenor of Gujarat for his "University 1st Rank" achievements.

He contributes to society in various ways, including giving talks to sophomore students at universities and giving talks on the various fields of data science, machine learning, AI, IoT in academia and at various conferences. He has excellent correlative knowledge of both academic research and industry best practices. Hence, He always come forward to remove gap between Industry and Academia where he has good number of achievements. He was core co-author of various courses such as data science, IoT, machine learning/AI, distributed databases at PG/UG cariculla at university of Kachchh. Hence, university of Kachchh become first government university in Gujarat to introduce Python as a first programming language in Cariculla and India's first government university to introduce data science, AI, IoT courses in Cariculla entire success story presented by Chetan at Pycon India 2016 conference. He is one of the founding members of PyKutch— A Python Community.

Currently, he is working on intelligent IoT devices with deep learning , reinforcement learning and distributed computing with various modern architectures. He is committer at Apache HBase and Spark HBase connector.

I would like to thank Prof. Devji Chhanga, Head of the Computer Science, University of Kachchh, for routing me to the correct path and for his valuable guidance in the field of data science research.

I would also like to thanks Prof. Shweta Gorania for being the first to introduce genetic algorithm and neural networks.

Last but not least, I would like to thank my beloved family for their support.

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/1786469786>.

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	1
<hr/>	
Chapter 1: Getting Started with Deep Learning	7
<hr/>	
Introducing machine learning	8
Supervised learning	8
Unsupervised learning	9
Reinforcement learning	9
What is deep learning?	10
How the human brain works	10
Deep learning history	11
Problems addressed	12
Neural networks	12
The biological neuron	12
An artificial neuron	13
How does an artificial neural network learn?	15
The backpropagation algorithm	15
Weights optimization	16
Stochastic gradient descent	17
Neural network architectures	19
Multilayer perceptron	19
DNNs architectures	20
Convolutional Neural Networks	20
Restricted Boltzmann Machines	21
Autoencoders	22
Recurrent Neural Networks	23
Deep learning framework comparisons	24
Summary	28
<hr/>	
Chapter 2: First Look at TensorFlow	29
<hr/>	
General overview	30
What's new with TensorFlow 1.x?	30
How does it change the way people use it?	31
Installing and getting started with TensorFlow	32
Installing TensorFlow on Linux	32
Which TensorFlow to install on your platform?	33
Requirements for running TensorFlow with GPU from NVIDIA	33

Step 1: Install NVIDIA CUDA	34
Step 2: Installing NVIDIA cuDNN v5.1+	34
Step 3: GPU card with CUDA compute capability 3.0+	36
Step 4: Installing the libcupti-dev library	36
Step 5: Installing Python (or Python3)	36
Step 6: Installing and upgrading PIP (or PIP3)	37
Step 7: Installing TensorFlow	38
How to install TensorFlow	38
Installing TensorFlow with native pip	38
Installing with virtualenv	39
Installing TensorFlow on Windows	41
Installation from source	41
Install on Windows	43
Test your TensorFlow installation	43
Computational graphs	43
Why a computational graph?	44
Neural networks as computational graphs	45
The programming model	47
Data model	49
Rank	50
Shape	50
Data types	51
Variables	54
Fetches	55
Feeds	55
TensorBoard	56
How does TensorBoard work?	56
Implementing a single input neuron	57
Source code for the single input neuron	64
Migrating to TensorFlow 1.x	64
How to upgrade using the script	65
Limitations	69
Upgrading code manually	69
Variables	69
Summary functions	69
Simplified mathematical variants	70
Miscellaneous changes	71
Summary	71
Chapter 3: Using TensorFlow on a Feed-Forward Neural Network	73

Introducing feed-forward neural networks	74
Feed-forward and backpropagation	75
Weights and biases	76
Transfer functions	76
Classification of handwritten digits	78
Exploring the MNIST dataset	79
Softmax classifier	81
Visualization	87
How to save and restore a TensorFlow model	89
Saving a model	89
Restoring a model	89
Softmax source code	92
Softmax loader source code	93
Implementing a five-layer neural network	94
Visualization	97
Five-layer neural network source code	99
ReLU classifier	100
Visualization	102
Source code for the ReLU classifier	103
Dropout optimization	105
Visualization	108
Source code for dropout optimization	109
Summary	112
Chapter 4: TensorFlow on a Convolutional Neural Network	113
<hr/>	
Introducing CNNs	114
CNN architecture	116
A model for CNNs - LeNet	118
Building your first CNN	119
Source code for a handwritten classifier	128
Emotion recognition with CNNs	130
Source code for emotion classifier	141
Testing the model on your own image	145
Source code	148
Summary	150
Chapter 5: Optimizing TensorFlow Autoencoders	151
<hr/>	
Introducing autoencoders	152
Implementing an autoencoder	153
Source code for the autoencoder	159

Improving autoencoder robustness	160
Building a denoising autoencoder	161
Source code for the denoising autoencoder	168
Convolutional autoencoders	170
Encoder	171
Decoder	171
Source code for convolutional autoencoder	181
Summary	184
Chapter 6: Recurrent Neural Networks	185
<hr/>	
RNNs basic concepts	185
RNNs at work	187
Unfolding an RNN	187
The vanishing gradient problem	188
LSTM networks	189
An image classifier with RNNs	190
Source code for RNN image classifier	196
Bidirectional RNNs	198
Source code for the bidirectional RNN	203
Text prediction	205
Dataset	206
Perplexity	206
PTB model	206
Running the example	207
Summary	209
Chapter 7: GPU Computing	211
<hr/>	
GPGPU computing	212
GPGPU history	212
The CUDA architecture	213
GPU programming model	214
TensorFlow GPU set up	215
Update TensorFlow	218
TensorFlow GPU management	218
Programming example	219
Source code for GPU computation	220
GPU memory management	222
Assigning a single GPU on a multi-GPU system	222
Source code for GPU with soft placement	224
Using multiple GPUs	225