Manpreet Singh Ghotra, Rajdeep Dua

Neural Network Programming with TensorFlow

Unleash the power of TensorFlow to train efficient neural networks

Packt>

Neural Network Programming with TensorFlow

If you're aware of the buzz surrounding terms such as machine learning, artificial intelligence, or deep learning, you might know what neural networks are. Ever wondered how they help in solving complex computational problems efficiently, or how to train efficient neural networks? This book will teach you just that.

You will start by getting a quick overview of the popular TensorFlow library and how it is used to train different neural networks. You will develop a thorough understanding of the fundamentals and basic math for neural networks, and why TensorFlow is a popular choice. Next, you will master optimization techniques and algorithms for neural networks using TensorFlow. Then, you will learn to implement some more complex types of neural network such as CNNs, RNNs, and deep belief networks. You will also get to train generative models and learn the application of autoencoders.

By the end of this book, you will have a fair understanding of how you can leverage the power of TensorFlow to train neural networks of varying complexities without any hassle. While you are learning about various neural network implementations, you will learn the underlying mathematics and linear algebra, and how they map to the appropriate TensorFlow constructs.

Things you will learn:

- Learn linear algebra and the mathematics behind neural networks
- Dive deep into neural networks from basic concepts, to advanced ones such as CNNs, RNNs, deep belief networks, and deep feedforward networks
- Explore optimization techniques for solving problems such as local minima, global minima, and saddle points
- Learn through real-world examples such as sentiment analysis
- Train different types of generative models and explore autoencoders
- Explore TensorFlow as an example of a deep learning implementation

ISBN 978-1-78839-039-2

Packt>
www.packtpub.com

Neural Network Programming with TensorFlow

Manpreet Singh Ghotra, Rajdeep Dua

Neural Network Programming with TensorFlow

Unleash the power of TensorFlow to train efficient neural networks

Manpreet Singh Ghotra

Rajdeep Dua

BIRMINGHAM - MUMBAI

Neural Network Programming 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: November 2017

Production reference: 1081117

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

ISBN 978-1-78839-039-2

www.packtpub.com

此为试读,需要完整PDF请访问: www.ertongbook.com

Credits

Authors
Manpreet Singh Ghotra
Rajdeep Dua

Reviewer
Giancarlo Zaccone

Commissioning Editor
Amey Varangaonkar

Acquisition Editor
Varsha Shetty

Content Development Editor
Tejas Limkar

Technical Editor
Sagar Sawant

Copy Editor
Tasneem Fatehi

Project Coordinator
Manthan Patel

Proofreader
Safis Editing

Indexer
Rekha Nair

Graphics
Tania Dutta

Production Coordinator
Aparna Bhagat

About the Authors

Manpreet Singh Ghotra has more than 15 years of experience in software development for both enterprise and big data software. He is currently working on developing a machine learning platform/API's using open source libraries and frameworks like TensorFlow, Keras, Apache Spark and PredictionIO at Salesforce. He has worked on various machine learning systems like sentiment analysis, spam detection, image moderation and anomaly detection. He was part of the machine learning group at one of the largest online retailers in the world, working on transit time calculations using R and Apache Mahout. With a master's and a postgraduate degree in machine learning, he has contributed to and worked for the machine learning community.

His GitHub profile is `https://github.com/badlogicmanpreet` and you can find him on LinkedIn at `https://in.linkedin.com/in/msghotra`.

This book is dedicated to my parents, Amrik Singh and Nirmal Kaur. I would like to thank my family for bearing with me while I spent time writing this book.

Rajdeep Dua has over 18 years of experience in the cloud, big data, and machine learning space. He has worked in the advocacy team for Google's big data tools, BigQuery. He has worked on the Greenplum big data platform at VMware in the developer evangelist team. He also worked closely with a team on porting Spark to run on VMware's public and private cloud as a feature set. He has taught Spark and big data at some of the most prestigious tech schools in India--IIIT Hyderabad, ISB, IIIT Delhi, and College of Engineering, Pune.

Currently, he leads the developer relations team at Salesforce India. He has published numerous blogs and hands-on labs for Einstein, a machine learning library from Salesforce and Apache Prediction IO.

He has published several big data and Spark tutorials at cluddatalab. He has also presented BigQuery and Google App Engine at the W3C conference in Hyderabad. He has led the developer relations teams at Google, VMware, and Microsoft, and he has spoken at hundreds of other conferences on the cloud.

His contributions to the open source community are for Docker, Kubernetes, Android, OpenStack, Prediction IO, and cloudfoundry projects.

This book is dedicated to my parents. I would like to thank my wife, Manjula, and sons, Navtej and Kairav, for bearing with me while I spent time writing this book.

About the Reviewer

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

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

Giancarlo holds a master's degree in physics from the Federico II of Naples and a second-level postgraduate master's degree in scientific computing from La Sapienza of Rome.

He is the author of the following Packt books: *Python Parallel Programming Cookbook, Getting Started with TensorFlow,* and *Deep Learning with TensorFlow.*

You can contact him at `https://it.linkedin.com/in/giancarlozaccone`.

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.in/dp/1788390393`. If you'd like to join our team of regular reviewers, you can email 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!

此为试读,需要完整PDF请访问: www.ertongbook.com

Table of Contents

Preface 1

Chapter 1: Maths for Neural Networks 7

Understanding linear algebra 7
Environment setup 8
Setting up the Python environment in Pycharm 8
Linear algebra structures 9
Scalars, vectors, and matrices 9
Tensors 10
Operations 10
Vectors 10
Matrices 11
Matrix multiplication 12
Trace operator 13
Matrix transpose 13
Matrix diagonals 14
Identity matrix 15
Inverse matrix 15
Solving linear equations 16
Singular value decomposition 18
Eigenvalue decomposition 21
Principal Component Analysis 21
Calculus 23
Gradient 23
Hessian 31
Determinant 32
Optimization 33
Optimizers 33
Summary 36

Chapter 2: Deep Feedforward Networks 37

Defining feedforward networks 37
Understanding backpropagation 38
Implementing feedforward networks with TensorFlow 39
Analyzing the Iris dataset 42
Code execution 45
Implementing feedforward networks with images 48

Analyzing the effect of activation functions on the feedforward networks accuracy 58
Summary 63
Chapter 3: Optimization for Neural Networks 65
What is optimization? 66
Types of optimizers 66
Gradient descent 67
Different variants of gradient descent 69
Algorithms to optimize gradient descent 69
Which optimizer to choose 72
Optimization with an example 72
Summary 75
Chapter 4: Convolutional Neural Networks 77
An overview and the intuition of CNN 78
Single Conv Layer Computation 78
CNN in TensorFlow 83
Image loading in TensorFlow 83
Convolution operations 84
Convolution on an image 85
Strides 87
Pooling 89
Max pool 89
Example code 90
Average pool 91
Image classification with convolutional networks 92
Defining a tensor for input images and the first convolution layer 98
Input tensor 98
First convolution layer 99
Second convolution layer 101
Third convolution layer 102
Flatten the layer 103
Fully connected layers 103
Defining cost and optimizer 104
Optimizer 105
First epoch 105
Plotting filters and their effects on an image 108
Summary 119
Chapter 5: Recurrent Neural Networks 121
Introduction to RNNs 121
RNN implementation 123

Computational graph 127
RNN implementation with TensorFlow 128
Computational graph 132
Introduction to long short term memory networks 133
Life cycle of LSTM 135
LSTM implementation 137
Computational graph 140
Sentiment analysis 142
Word embeddings 142
Sentiment analysis with an RNN 149
Computational graph 154
Summary 156
Chapter 6: Generative Models 157
Generative models 157
Discriminative versus generative models 158
Types of generative models 159
Autoencoders 159
GAN 160
Sequence models 162
GANs 162
GAN with an example 163
Types of GANs 173
Vanilla GAN 173
Conditional GAN 174
Info GAN 174
Wasserstein GAN 174
Coupled GAN 175
Summary 175
Chapter 7: Deep Belief Networking 177
Understanding deep belief networks 178
DBN implementation 178
Class initialization 178
RBM class 182
Pretraining the DBN 184
Model training 185
Predicting the label 186
Finding the accuracy of the model 186
DBN implementation for the MNIST dataset 187
Loading the dataset 187
Input parameters for a DBN with 256-Neuron RBM layers 188
Output for a DBN with 256-neuron RBN layers 189

Effect of the number of neurons in an RBM layer in a DBN 190
An RBM layer with 512 neurons 190
An RBM layer with 128 neurons 190
Comparing the accuracy metrics 191
DBNs with two RBM layers 192
Classifying the NotMNIST dataset with a DBN 194
Summary 197
Chapter 8: Autoencoders 199
Autoencoder algorithms 200
Under-complete autoencoders 201
Dataset 202
Basic autoencoders 204
Autoencoder initialization 204
AutoEncoder class 205
Basic autoencoders with MNIST data 207
Basic autoencoder plot of weights 209
Basic autoencoder recreated images plot 210
Basic autoencoder full code listing 212
Basic autoencoder summary 214
Additive Gaussian Noise autoencoder 215
Autoencoder class 215
Additive Gaussian Autoencoder with the MNIST dataset 217
Training the model 218
Plotting the weights 218
Plotting the reconstructed images 220
Additive Gaussian autoencoder full code listing 222
Comparing basic encoder costs with the Additive Gaussian Noise autoencoder 224
Additive Gaussian Noise autoencoder summary 224
Sparse autoencoder 225
KL divergence 225
KL divergence in TensorFlow 226
Cost of a sparse autoencoder based on KL Divergence 227
Complete code listing of the sparse autoencoder 227
Sparse autoencoder on MNIST data 229
Comparing the Sparse encoder with the Additive Gaussian Noise encoder 231
Summary 232
Chapter 9: Research in Neural Networks 233

Avoiding overfitting in neural networks 233
Problem statement 233
Solution 234
Results 235
Large-scale video processing with neural networks 236
Resolution improvements 237
Feature histogram baselines 237
Quantitative results 238
Named entity recognition using a twisted neural network 238
Example of a named entity recognition 239
Defining Twinet 239
Results 240
Bidirectional RNNs 241
BRNN on TIMIT dataset 242
Summary 243
Appendix: Getting started with TensorFlow 245
Environment setup 245
TensorFlow comparison with Numpy 247
Computational graph 247
Graph 248
Session objects 248
Variables 249
Scope 250
Data input 251
Placeholders and feed dictionaries 252
Auto differentiation 252
TensorBoard 254
Index 255

此为试读,需要完整PDF请访问: www.ertongbook.com