

Deep Learning with PyTorch 1.x *Second Edition*

Implement deep learning techniques and neural network architecture variants using Python

Laura Mitchell
Sri. Yogesh K.
Vishnu Subramanian

Packt>

BIRMINGHAM - MUMBAI

Deep Learning with PyTorch 1.x

Second Edition

Copyright © 2019 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 or its dealers and distributors, will be held liable for any damages caused or alleged to have been 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.

Commissioning Editor: Sunith Shetty
Acquisition Editor: Devika Battike
Content Development Editor: Athikho Sapuni Rishana
Senior Editor: Sofi Rogers
Technical Editor: Joseph Sunil
Copy Editor: Safis Editing
Project Coordinator: Aishwarya Mohan
Proofreader: Safis Editing
Indexer: Pratik Shirodkar
Production Designer: Jyoti Chauhan

First published: February 2018
Second edition: November 2019

Production reference: 1291119

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

ISBN 978-1-83855-300-5

www.packt.com



Packt.com

Subscribe to our online digital library for full access to over 7,000 books and videos, as well as industry leading tools to help you plan your personal development and advance your career. For more information, please visit our website.

Why subscribe?

- Spend less time learning and more time coding with practical eBooks and Videos from over 4,000 industry professionals
- Improve your learning with Skill Plans built especially for you
- Get a free eBook or video every month
- Fully searchable for easy access to vital information
- Copy and paste, print, and bookmark content

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.packt.com and as a print book customer, you are entitled to a discount on the eBook copy. Get in touch with us at customer-care@packtpub.com for more details.

At www.packt.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.

About the authors

Laura Mitchell graduated with a degree in mathematics from the University of Edinburgh. With 15 years of experience in the tech and data science space, Laura is the lead data scientist at MagicLab whose brands have connected the lives of over 500 million people through dating, social and business. Laura has hands-on experience in the delivery of projects surrounding natural language processing, image classification and recommender systems, from initial conception to production. She has a passion for learning new technologies and keeping herself up to date with industry trends.

Sri. Yogesh K. is an experienced data scientist with a history of working in higher education. He is skilled in Python, Apache Spark, deep learning, Hadoop, and machine learning. He is a strong engineering professional with a Certificate of Engineering Excellence from the International School of Engineering (INSOFE) and is focused on big data analytics. Sri has trained over 500 working professionals in data science and deep learning from companies including Flipkart, Honeywell, GE, and Rakuten. Additionally, he has worked on various projects that involved deep learning and PyTorch.

Vishnu Subramanian has experience in leading, architecting, and implementing several big data analytical projects using artificial intelligence, machine learning, and deep learning. He specializes in machine learning, deep learning, distributed machine learning, and visualization. He has experience in retail, finance, and travel domains. Also, he is good at understanding and coordinating between businesses, AI, and engineering teams.

About the reviewers

Mingfei Ma is a senior deep learning software engineer from Intel Asia-Pacific Research & Development Ltd and he has plenty of experience in high-performance computation. Mingfei contributed extensively to the CPU performance optimization of PyTorch and its predecessor, Torch. He also has expertise in computer graphics, heterogeneous computing, microarchitecture detection, high-performance computation libraries, and more.

Ajit Pratap Kundan is at the forefront of innovative technologies in the world of IT. He's worked with HPE, VMware, Novell, Redington, and PCS to help their customers in transforming their data centers through software-defined services. Ajit is an innovative pre-sales tech enthusiast with over 19 years of experience in technologies such as Lotus, SUSE Linux, Platespin, and all VMware solutions. Ajit is a valued author on cloud technologies and has authored two books, *VMware Cross-Cloud Architecture* and *Intelligent Automation with VMware*, published by Packt.

Packt is searching for authors like you

If you're interested in becoming an author for Packt, please visit authors.packtpub.com and apply today. We have worked with thousands of developers and tech professionals, just like you, to help them share their insight with the global tech community. You can make a general application, apply for a specific hot topic that we are recruiting an author for, or submit your own idea.

Table of Contents

Preface	1
<hr/>	
Section 1: Building Blocks of Deep Learning with PyTorch	
1.x	
<hr/>	
Chapter 1: Getting Started with Deep Learning Using PyTorch	9
Exploring artificial intelligence	9
The history of AI	10
Machine learning in the real world	10
So, why DL?	11
Applications of deep learning	12
Automatic translation of text from images	12
Object detection in self-driving cars	13
Deep learning frameworks	14
Why PyTorch?	14
What's new in PyTorch v1.x?	15
CPU versus GPU	16
What is CUDA?	16
Which GPUs should we use?	16
What should you do if you don't have a GPU?	17
Setting up PyTorch v1.x	17
Installing PyTorch	18
Summary	18
Chapter 2: Building Blocks of Neural Networks	19
What is a neural network?	20
Understanding the structure of neural networks	21
Building a neural network in PyTorch	22
PyTorch sequential neural network	22
Building a PyTorch neural network using nn.Module	26
Understanding PyTorch Tensors	29
Understanding Tensor shapes and reshaping Tensors	32
Understanding tensor operations	34
Understanding Tensor types in PyTorch	36
Importing our dataset as a PyTorch Tensor	37
Training neural networks in PyTorch	38
Summary	41
Section 2: Going Advanced with Deep Learning	
<hr/>	

Chapter 3: Diving Deep into Neural Networks	45
Diving into the building blocks of neural networks	45
Layers – the fundamental blocks of neural networks	47
Non-linear activations	49
Sigmoid	49
Tanh	50
ReLU	51
Leaky ReLU	52
PyTorch non-linear activations	52
The PyTorch way of building deep learning algorithms	53
Model architecture for different machine learning problems	54
Loss functions	54
Optimizing network architecture	56
Image classification using deep learning	58
Loading data into PyTorch tensors	61
Loading PyTorch tensors as batches	63
Building the network architecture	63
Training the model	65
Summary	67
Chapter 4: Deep Learning for Computer Vision	69
Introduction to neural networks	70
MNIST – getting data	72
Building a CNN model from scratch	73
Conv2d	75
Pooling	78
Nonlinear activation – ReLU	79
View	80
Linear layer	81
Training the model	81
Classifying dogs and cats – CNN from scratch	84
Classifying dogs and cats using transfer learning	86
Creating and exploring a VGG16 model	88
Freezing the layers	90
Fine-tuning VGG16	90
Training the VGG16 model	90
Calculating pre-convoluted features	93
Understanding what a CNN model learns	96
Visualizing outputs from intermediate layers	96
Visualizing the weights of the CNN layer	99
Summary	100
Chapter 5: Natural Language Processing with Sequence Data	101
Working with text data	102
Tokenization	103

Converting text into characters	104
Converting text into words	104
N-gram representation	105
Vectorization	106
One-hot encoding	106
Word embedding	108
Training word embedding by building a sentiment classifier	109
Downloading IMDb data and performing text tokenization	110
Tokenizing with torchtext.data	111
Tokenizing with torchtext.datasets	111
Building vocabulary	112
Generating batches of vectors	114
Creating a network model with embedding	115
Training the model	116
Using pretrained word embeddings	118
Downloading the embeddings	118
Loading the embeddings in the model	119
Freezing the embedding layer weights	120
Recursive neural networks	121
Understanding how RNN works with an example	122
Solving text classification problem using LSTM	125
Long-term dependency	125
LSTM networks	125
Preparing the data	128
Creating batches	128
Creating the network	129
Training the model	130
Convolutional network on sequence data	131
Understanding one-dimensional convolution for sequence data	132
Creating the network	132
Training the model	133
Language modeling	134
Pretrained models	134
Embeddings from language models	135
Bidirectional Encoder Representations from Transformers	136
Generative Pretrained Transformer 2	137
PyTorch implementations	137
GPT-2 playground	138
Summary	140
Section 3: Understanding Modern Architectures in Deep Learning	
Chapter 6: Implementing Autoencoders	143
Applications of autoencoders	144
Bottleneck and loss functions	144

Coded example – standard autoencoder	145
Convolutional autoencoders	148
Coded example – convolutional autoencoder	148
Denoising autoencoders	150
Variational autoencoders	151
Training VAEs	153
Coded example – VAE	155
Restricted Boltzmann machines	157
Training RBMs	160
Theoretical example – RBM recommender system	161
Coded example – RBM recommender system	162
DBN architecture	166
Fine-tuning	168
Summary	169
Further reading	169
Chapter 7: Working with Generative Adversarial Networks	171
Neural style transfer	172
Loading the data	174
Creating the VGG model	176
Content loss	176
Style loss	177
Extracting the losses	179
Creating a loss function for each layer	182
Creating the optimizer	183
Training the model	183
Introducing GANs	185
DCGAN	187
Defining the generator network	187
Transposed convolutions	187
Batch normalization	188
Generator	191
Defining the discriminator network	192
Defining loss and optimizer	193
Training the discriminator	194
Training the discriminator with real images	194
Training the discriminator with fake images	194
Training the generator network	195
Training the complete network	195
Inspecting the generated images	197
Summary	199
Chapter 8: Transfer Learning with Modern Network Architectures	201
Modern network architectures	202
ResNet	202
Creating PyTorch datasets	206
Creating loaders for training and validation	207