

DEEP LEARNING

with Python

François Chollet



Deep Learning with Python

François Chollet

Machine learning has made remarkable progress in recent years. We went from near-unusable speech and image recognition, to near-human accuracy. We went from machines that couldn't beat a serious Go player, to defeating a world champion. Behind this progress is deep learning—a combination of engineering advances, best practices, and theory that enables a wealth of previously impossible smart applications.

Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. You'll explore challenging concepts and practice with applications in computer vision, natural-language processing, and generative models. By the time you finish, you'll have the knowledge and hands-on skills to apply deep learning in your own projects.


What's Inside

- Deep learning from first principles
- Setting up your own deep-learning environment
- Image-classification models
- Deep learning for text and sequences
- Neural style transfer, text generation, and image generation

Readers need intermediate Python skills. No previous experience with Keras, TensorFlow, or machine learning is required.

François Chollet is an AI researcher on the Google Brain Team and author of the Keras deep-learning library.

To download their free eBook in PDF, ePub, and Kindle formats, owners of this book should visit
www.manning.com/books/deep-learning-with-python



“The clearest explanation of deep learning I have come across ... it was a joy to read.”

—Richard Tobias, Cephasonics

“An excellent hands-on introductory title, with great depth and breadth.”

—David Blumenthal-Barby
Babbel

“Bridges the gap between the hype and a functioning deep-learning system.”

—Peter Rabinovitch, Akamai

“The best resource for becoming a master of Keras and deep learning.”

—Claudio Rodriguez
Cox Media Group

ISBN-13: 978-1-61729-443-3
 ISBN-10: 1-61729-443-8



9 781617 294433

Chollet



Deep Learning with Python



MANNING

Deep Learning with Python

FRANÇOIS CHOLLET



MANNING
SHELTER ISLAND

For online information and ordering of this and other Manning books, please visit www.manning.com. The publisher offers discounts on this book when ordered in quantity. For more information, please contact


Special Sales Department
Manning Publications Co.
20 Baldwin Road
PO Box 761
Shelter Island, NY 11964
Email: orders@manning.com

©2018 by Manning Publications Co. All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by means electronic, mechanical, photocopying, or otherwise, without prior written permission of the publisher.

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in the book, and Manning Publications was aware of a trademark claim, the designations have been printed in initial caps or all caps.

© Recognizing the importance of preserving what has been written, it is Manning's policy to have the books we publish printed on acid-free paper, and we exert our best efforts to that end. Recognizing also our responsibility to conserve the resources of our planet, Manning books are printed on paper that is at least 15 percent recycled and processed without the use of elemental chlorine.

 Manning Publications Co.
20 Baldwin Road
PO Box 761
Shelter Island, NY 11964

Development editor: Toni Arritola
Technical development editor: Jerry Gaines
Review editor: Aleksandar Dragosavljević
Project editor: Tiffany Taylor
Copyeditor: Tiffany Taylor
Proofreader: Katie Tennant
Technical proofreaders: Alex Ott and Richard Tobias
Typesetter: Dottie Marsico
Cover designer: Marija Tudor

ISBN 9781617294433

Printed in the United States of America

2 3 4 5 6 7 8 9 10 – EBM – 22 21 20 19 18 17



Get the eBooks FREE!

(PDF, ePub, and Kindle all included)

We believe that once you buy a book from us, you should be able to read it in any format we have available. To get electronic versions of this book at no additional cost to you, purchase and then register this book at the Manning website following the instructions inside this insert.

That's it!

Thanks from Manning!





Deep Learning with Python

preface

If you've picked up this book, you're probably aware of the extraordinary progress that deep learning has represented for the field of artificial intelligence in the recent past. In a mere five years, we've gone from near-unusable image recognition and speech transcription, to superhuman performance on these tasks.

The consequences of this sudden progress extend to almost every industry. But in order to begin deploying deep-learning technology to every problem that it could solve, we need to make it accessible to as many people as possible, including non-experts—people who aren't researchers or graduate students. For deep learning to reach its full potential, we need to radically democratize it.

When I released the first version of the Keras deep-learning framework in March 2015, the democratization of AI wasn't what I had in mind. I had been doing research in machine learning for several years, and had built Keras to help me with my own experiments. But throughout 2015 and 2016, tens of thousands of new people entered the field of deep learning; many of them picked up Keras because it was—and still is—the easiest framework to get started with. As I watched scores of newcomers use Keras in unexpected, powerful ways, I came to care deeply about the accessibility and democratization of AI. I realized that the further we spread these technologies, the more useful and valuable they become. Accessibility quickly became an explicit goal in the development of Keras, and over a few short years, the Keras developer community has made fantastic achievements on this front. We've put deep learning into the hands of tens of thousands of people, who in turn are using it to solve important problems we didn't even know existed until recently.

The book you're holding is another step on the way to making deep learning available to as many people as possible. Keras had always needed a companion course to

simultaneously cover fundamentals of deep learning, Keras usage patterns, and deep-learning best practices. This book is my best effort to produce such a course. I wrote it with a focus on making the concepts behind deep learning, and their implementation, as approachable as possible. Doing so didn't require me to dumb down anything—I strongly believe that there are no difficult ideas in deep learning. I hope you'll find this book valuable and that it will enable you to begin building intelligent applications and solve the problems that matter to you.

acknowledgments

I'd like to thank the Keras community for making this book possible. Keras has grown to have hundreds of open source contributors and more than 200,000 users. Your contributions and feedback have turned Keras into what it is today.

I'd also like to thank Google for backing the Keras project. It has been fantastic to see Keras adopted as TensorFlow's high-level API. A smooth integration between Keras and TensorFlow greatly benefits both TensorFlow users and Keras users and makes deep learning accessible to most.

I want to thank the people at Manning who made this book possible: publisher Marjan Bace and everyone on the editorial and production teams, including Christina Taylor, Janet Vail, Tiffany Taylor, Katie Tennant, Dottie Marsico, and many others who worked behind the scenes.

Many thanks go to the technical peer reviewers led by Aleksandar Dragosavljević—Diego Acuña Rozas, Geoff Barto, David Blumenthal-Barby, Abel Brown, Clark Dorman, Clark Gaylord, Thomas Heiman, Wilson Mar, Sumit Pal, Vladimir Pasman, Gustavo Patino, Peter Rabinovitch, Alvin Raj, Claudio Rodriguez, Srdjan Santic, Richard Tobias, Martin Verzilli, William E. Wheeler, and Daniel Williams—and the forum contributors. Their contributions included catching technical mistakes, errors in terminology, and typos, and making topic suggestions. Each pass through the review process and each piece of feedback implemented through the forum topics shaped and molded the manuscript.

On the technical side, special thanks go to Jerry Gaines, who served as the book's technical editor; and Alex Ott and Richard Tobias, who served as the book's technical proofreaders. They're the best technical editors I could have hoped for.

Finally, I'd like to express my gratitude to my wife Maria for being extremely supportive throughout the development of Keras and the writing of this book.

about this book

This book was written for anyone who wishes to explore deep learning from scratch or broaden their understanding of deep learning. Whether you're a practicing machine-learning engineer, a software developer, or a college student, you'll find value in these pages.

This book offers a practical, hands-on exploration of deep learning. It avoids mathematical notation, preferring instead to explain quantitative concepts via code snippets and to build practical intuition about the core ideas of machine learning and deep learning.

You'll learn from more than 30 code examples that include detailed commentary, practical recommendations, and simple high-level explanations of everything you need to know to start using deep learning to solve concrete problems.

The code examples use the Python deep-learning framework Keras, with TensorFlow as a backend engine. Keras, one of the most popular and fastest-growing deep-learning frameworks, is widely recommended as the best tool to get started with deep learning.

After reading this book, you'll have a solid understand of what deep learning is, when it's applicable, and what its limitations are. You'll be familiar with the standard workflow for approaching and solving machine-learning problems, and you'll know how to address commonly encountered issues. You'll be able to use Keras to tackle real-world problems ranging from computer vision to natural-language processing: image classification, timeseries forecasting, sentiment analysis, image and text generation, and more.

Who should read this book

This book is written for people with Python programming experience who want to get started with machine learning and deep learning. But this book can also be valuable to many different types of readers:

- If you're a data scientist familiar with machine learning, this book will provide you with a solid, practical introduction to deep learning, the fastest-growing and most significant subfield of machine learning.
- If you're a deep-learning expert looking to get started with the Keras framework, you'll find this book to be the best Keras crash course available.
- If you're a graduate student studying deep learning in a formal setting, you'll find this book to be a practical complement to your education, helping you build intuition around the behavior of deep neural networks and familiarizing you with key best practices.

Even technically minded people who don't code regularly will find this book useful as an introduction to both basic and advanced deep-learning concepts.

In order to use Keras, you'll need reasonable Python proficiency. Additionally, familiarity with the Numpy library will be helpful, although it isn't required. You don't need previous experience with machine learning or deep learning: this book covers from scratch all the necessary basics. You don't need an advanced mathematics background, either—high school-level mathematics should suffice in order to follow along.

Roadmap

This book is structured in two parts. If you have no prior experience with machine learning, I strongly recommend that you complete part 1 before approaching part 2. We'll start with simple examples, and as the book goes on, we'll get increasingly close to state-of-the-art techniques.

Part 1 is a high-level introduction to deep learning, providing context and definitions, and explaining all the notions required to get started with machine learning and neural networks:

- Chapter 1 presents essential context and background knowledge around AI, machine learning, and deep learning.
- Chapter 2 introduces fundamental concepts necessary in order to approach deep learning: tensors, tensor operations, gradient descent, and backpropagation. This chapter also features the book's first example of a working neural network.
- Chapter 3 includes everything you need to get started with neural networks: an introduction to Keras, our deep-learning framework of choice; a guide for setting up your workstation; and three foundational code examples with detailed explanations. By the end of this chapter, you'll be able to train simple neural

networks to handle classification and regression tasks, and you'll have a solid idea of what's happening in the background as you train them.

- Chapter 4 explores the canonical machine-learning workflow. You'll also learn about common pitfalls and their solutions.

Part 2 takes an in-depth dive into practical applications of deep learning in computer vision and natural-language processing. Many of the examples introduced in this part can be used as templates to solve problems you'll encounter in the real-world practice of deep learning:

- Chapter 5 examines a range of practical computer-vision examples, with a focus on image classification.
- Chapter 6 gives you practice with techniques for processing sequence data, such as text and timeseries.
- Chapter 7 introduces advanced techniques for building state-of-the-art deep-learning models.
- Chapter 8 explains generative models: deep-learning models capable of creating images and text, with sometimes surprisingly artistic results.
- Chapter 9 is dedicated to consolidating what you've learned throughout the book, as well as opening perspectives on the limitations of deep learning and exploring its probable future.

Software/hardware requirements

All of this book's code examples use the Keras deep-learning framework (<https://keras.io>), which is open source and free to download. You'll need access to a UNIX machine; it's possible to use Windows, too, but I don't recommend it. Appendix A walks you through the complete setup.

I also recommend that you have a recent NVIDIA GPU on your machine, such as a TITAN X. This isn't required, but it will make your experience better by allowing you to run the code examples several times faster. See section 3.3 for more information about setting up a deep-learning workstation.

If you don't have access to a local workstation with a recent NVIDIA GPU, you can use a cloud environment, instead. In particular, you can use Google Cloud instances (such as an n1-standard-8 instance with an NVIDIA Tesla K80 add-on) or Amazon Web Services (AWS) GPU instances (such as a p2.xlarge instance). Appendix B presents in detail one possible cloud workflow that runs an AWS instance via Jupyter notebooks, accessible in your browser.

Source code

All code examples in this book are available for download as Jupyter notebooks from the book's website, www.manning.com/books/deep-learning-with-python, and on GitHub at <https://github.com/fchollet/deep-learning-with-python-notebooks>.

Book forum

Purchase of *Deep Learning with Python* includes free access to a private web forum run by Manning Publications where you can make comments about the book, ask technical questions, and receive help from the author and from other users. To access the forum, go to <https://forums.manning.com/forums/deep-learning-with-python>. You can also learn more about Manning's forums and the rules of conduct at <https://forums.manning.com/forums/about>.

Manning's commitment to our readers is to provide a venue where a meaningful dialogue between individual readers and between readers and the author can take place. It isn't a commitment to any specific amount of participation on the part of the author, whose contribution to the forum remains voluntary (and unpaid). We suggest you try asking him some challenging questions lest his interest stray! The forum and the archives of previous discussions will be accessible from the publisher's website as long as the book is in print.

about the author



François Chollet works on deep learning at Google in Mountain View, CA. He is the creator of the Keras deep-learning library, as well as a contributor to the TensorFlow machine-learning framework. He also does deep-learning research, with a focus on computer vision and the application of machine learning to formal reasoning. His papers have been published at major conferences in the field, including the Conference on Computer Vision and Pattern Recognition (CVPR), the Conference and Workshop on Neural Information Processing Systems (NIPS), the International Conference on Learning Representations (ICLR), and others.