

Extracted from:

Programming Machine Learning

From Zero to Deep Learning

This PDF file contains pages extracted from *Programming Machine Learning*, published by the Pragmatic Bookshelf. For more information or to purchase a paperback or PDF copy, please visit <http://www.pragprog.com>.

Note: This extract contains some colored text (particularly in code listing). This is available only in online versions of the books. The printed versions are black and white. Pagination might vary between the online and printed versions; the content is otherwise identical.

Copyright © 2019 The Pragmatic Programmers, LLC.

All rights reserved.

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

The Pragmatic Bookshelf

Raleigh, North Carolina

The
Pragmatic
Programmers

Programming Machine Learning

From Zero to
Deep Learning



Paolo Perrotta
edited by Meghan Blanchette

Programming Machine Learning

From Zero to Deep Learning

Paolo Perrotta

The Pragmatic Bookshelf

Raleigh, North Carolina



Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and The Pragmatic Programmers, LLC was aware of a trademark claim, the designations have been printed in initial capital letters or in all capitals. The Pragmatic Starter Kit, The Pragmatic Programmer, Pragmatic Programming, Pragmatic Bookshelf, PragProg and the linking *g* device are trademarks of The Pragmatic Programmers, LLC.

Every precaution was taken in the preparation of this book. However, the publisher assumes no responsibility for errors or omissions, or for damages that may result from the use of information (including program listings) contained herein.

Our Pragmatic books, screencasts, and audio books can help you and your team create better software and have more fun. Visit us at <https://pragprog.com>.

For sales, volume licensing, and support, please contact support@pragprog.com.

For international rights, please contact rights@pragprog.com.

Copyright © 2019 The Pragmatic Programmers, LLC.

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

ISBN-13: 978-1-68050-660-0

Book version: B1.0—March 20, 2019

How the Heck Is That Possible?

Machine learning can seem like magic. How can a computer recognize the objects in an image? How can a car drive itself?

Those feats are baffling—not just to the layman, but to many software developers like you and me. Even after writing code for many years, I had no idea how machine learning could possibly work. While I tinkered with the latest web framework, someone out there was writing amazing software that looked like science fiction—and I couldn't even comprehend it.

I wanted in on the action. I wanted to be able to build those things myself.

I knew how to write software, so I assumed that I would grok machine learning quickly. I mean, how hard could it be? I put on a confident smile and started studying. Then I kept smiling confidently as I slammed my muzzle into a long sequence of brick walls.

To us developers, machine learning feels... *foreign*. The field is teeming with math jargon, researchy conventions and, frankly, bad code. Instead of tutorials, people point you at lectures and research papers. For many of us, machine learning is as intimidating as it's intriguing.

This is the book I missed when I got started with machine learning: an introduction for developers, written in our own language. After reading it, you'll be comfortable with the fundamentals, and able to write machine learning programs. No, you probably won't be able to build your own self-driving car just yet—but at least you will know how the heck that's possible.

Come in.

About This Book

This is a book for developers who want to learn machine learning from scratch.

Machine learning is a broad field, and no book can cover it all. We'll focus on the three facets of machine learning that are most important today: *supervised*

learning, *neural networks*, and *deep learning*. No worries if you don't know what those terms mean—we'll define them as we go through the book.

For now, just know this:

- *Supervised learning* is a specific type of machine learning. Machine learning comes in a few different flavors, but supervised learning is the most popular one these days. Part I of this book, *From Zero to Image Recognition*, is a hands-on supervised learning tutorial. Within a couple of chapters, we'll write a minimal learning program. Then we'll evolve the program step by step, until it's powerful enough to recognize handwritten characters. We'll craft this program all by ourselves, without using libraries. You'll understand each single line of code.
- There are many algorithms that you can use to implement a supervised learning system. Part II of this book, *Neural Networks*, is dedicated to the most popular of those algorithms. We'll grow the program from Part I into a full-fledged neural network. We'll have to overcome a few challenges along the way, but the payoff will be worth it: the final neural network will be way more powerful than the starting program. Once again, we'll write the code ourselves, line by line. Its inner workings will be open for you to play with.
- Neural networks got a big boost in recent years, when researchers came up with breakthrough techniques to design and use them. This souped-up technology is vastly more powerful than the simple neural networks of old—so much so, that it got its own name: *deep learning*. That's also the title of Part III of this book. In it, we'll rewrite our neural network using a modern machine learning library. The resulting code will be our starting point to understand what deep learning is about. Finally, as we wrap up the book, we'll take a look at a few advanced deep learning techniques, paving the way for your future explorations.

Before We Begin

This book cannot turn you into a machine learning pro overnight, but it can give you an intuitive, practical understanding of how machine learning works. I want to open the hood of this discipline, show you the gears, and demystify the magic. Once you grasp the fundamental principles of machine learning, you'll find it much easier to dig deeper, incorporate these techniques in your daily job, and maybe even embark on a career as a data scientist.

You don't need to be a senior developer to read this book. However, you should be comfortable writing short programs. If you know Python, then you just lucked out: that's the language that I will use throughout, so you'll feel right at home. Even if you don't know Python yet, no worries. It's a friendly language, and the code in this book will never get too complicated. Read [Appendix 1, Just Enough Python, on page ?](#) to get up to speed, and be ready to Google for more information if you get stuck.

Machine learning involves a lot of mathematics. I won't dumb down the math, but I'll make it as intuitive as I can. You will need some high school math: I'm going to assume that you can read a Cartesian chart, that you know what the "axes" and their "origin" are, and that you can make sense of a function plot. Other than that, you don't need much math knowledge. In fact, you might be able to read through even if you consider yourself terrible at math... But be ready to be proven wrong about that!

Math Deep Dives

We all love intuitive math, but sometimes you might strive for a more formal explanation. If you ever feel lost while parsing a formula, or if you like mathematics and want to dig deeper, then look for "Math Deep Dive" boxes like this one. There are only a few of them—three or four in the entire book. They'll point you at relevant math screencasts on the excellent Kahn Academy.^a No matter what your current level of math is, this site has got you covered.

Just to be clear, these additional lessons are optional. You don't need them to read this book—only if you wish to really wrap your mind about the mathematics of machine learning.

a. www.khanacademy.org

On the other hand, if you have a solid background in linear algebra and calculus, then you might find some of the math obvious. In that case, feel free to breeze over the explanations that you don't need.

Machine learning has a rich and specific vocabulary. You're likely to stumble upon new words, or new meanings for old words. Take it easy and don't feel like you have to remember everything. I will remind you of many of those words' meaning the next time we encounter them. Whenever a term gives you that obnoxious "I cannot quite remember what this means" feeling, you can look it up in *the (as yet) unwritten Appendix 2, The Words of Machine Learning*,

One word about the datasets that we'll use in our examples: many of them are collections of images. Stay assured that machine learning can do much more besides image recognition: it can analyze text, generate music, or even hold conversations. However, image recognition makes for very intuitive examples, so it will be our go-to application throughout the book.

That's enough public service announcements. Let's dive into Part I.