Effective Machine Learning Teams
by David Tan, Ada Leung, and David Colls

Printed in the United States of America.

Published by O'Reilly Media, Inc., 1005 Gravenstein Highway North, Sebastopol, CA 95472.

O'Reilly books may be purchased for educational, business, or sales promotional use. Online editions are also available for most titles (*https://oreilly.com*). For more information, contact our corporate/institutional sales department: 800-998-9938 or *corporate@oreilly.com*.

Acquisitions Editor: Nicole Butterfield
Development Editor: Melissa Potter
Production Editor: Gregory Hyman
Copyeditor: Nicole Taché
Proofreader: M & R Consultants Corporation

Indexer: Judith McConville
Interior Designer: David Futato
Cover Designer: Karen Montgomery
Illustrator: Kate Dullea

February 2024: First Edition

Revision History for the First Edition
2024-02-29: First Release

See *https://oreilly.com/catalog/errata.csp?isbn=9781098144630* for release details.

The views expressed in this work are those of the authors, and do not represent the publisher's views. While the publisher and the authors have used good faith efforts to ensure that the information and instructions contained in this work are accurate, the publisher and the authors disclaim all responsibility for errors or omissions, including without limitation responsibility for damages resulting from the use

978-1-098-14463-0

[LSI]

Table of Contents

Part II. Engineering

Preface

It

the capabilities necessary for data scientists and ML engineers to design and deliver ML solutions in the real world.

In the past, we've presented the principles, practices, and hands-on exercises in this book to data scientists, ML engineers, PhD students, software engineers, quality analysts, and product managers, and we've consistently received positive you to try out the practices in your teams, and we hope you'll experience firsthand the value that they bring in real-world projects.

Approach this book with a continuous improvement mindset, not a perfectionist mindset. There is no perfect project where everything works perfectly without challenges. There will always be complexity and challenges (and we know a healthy amount of challenge is essential for growth), but the practices in this book will help you minimize *accidental* complexity so that you can focus on the *essential* complexity of your ML solutions and on delivering value responsibly.

How This Book Is Organized

Chapter 1, "Challenges and Better Paths in Delivering ML Solutions", is a distillation of the entire book. We explore high-level and low-level reasons for why and how ML projects fail. We then lay out a more reliable path for delivering value in ML solutions by adopting Lean delivery practices across five key disciplines: product, delivery, machine learning, engineering, and data.

In the remaining chapters, we describe practices of effective ML teams and ML practitioners. In Part I, "Product and Delivery", we elaborate on practices in other subsystems that are necessary for delivering ML solutions, such as product thinking

Part II: Engineering

Chapters 3 and 4: Effective dependency management

Here, we describe principles and practices—along with a hands-on example that you can code along with—for creating consistent, reproducible, secure, and production-like runtime environments for running your code. When we hit the ground running and start delivering solutions, you'll see how the practices in this chapter will enable you and your teammates to "check out and go" and create consistent environments effortlessly, instead of getting trapped in dependency hell.

Chapters 5 and 6: Automated testing for ML systems

These chapters provide you with a rubric for testing components of your ML solution—be they software tests, model tests, or data tests. We demonstrate how automated tests help us shorten our feedback cycles and reduce the tedious effort of manual testing, or worse, fixing production defects that slipped through the cracks of manual testing. We describe the limits of the software testing paradigm on ML models, and how ML fitness functions and behavioral tests can help us scale the automated testing of ML models. We also cover techniques for comprehensively testing large language models (LLMs) and LLM applications.

Chapter 7, "Supercharging Your Code Editor with Simple Techniques"

We'll show you how to configure your code editor (PyCharm or VS Code) to help you code more effectively. After we've configured our IDE in a few steps, we'll go through a series of keyboard shortcuts that can help you to automate refactoring, automatically detect and fix issues, and navigate your codebase without getting lost in the weeds, among other things.

Chapter 8, "Refactoring and Technical Debt Management"

In this chapter, we draw from the wisdom of software design to help us design readable, testable, maintainable, and evolvable code. In the spirit of "learning by doing," you'll see how we can take a problematic, messy, and brittle notebook and apply refactoring techniques to iteratively improve our codebase to a modular, tested, and readable state. You'll also learn techniques that can help you and your team make technical debt visible and take actions to keep it at a healthy level.

This element indicates a warning or caution.

Using Code Examples

Supplemental material (code examples, exercises, etc.) is available for download at:

- *https://github.com/davified/loan-default-prediction*
- *https://github.com/davified/ide-productivity*
- *https://github.com/davified/refactoring-exercise*

If you have a technical question or a problem using the code examples, please send email to *support@oreilly.com*.

This book is here to help you get your job done. In general, if example code is offered with this book, you may use it in your programs and documentation. You do not need to contact us for permission unless you're reproducing a significant portion of the code. For example, writing a program that uses several chunks of code from this book does not require permission. Selling or distributing examples from O'Reilly books does require permission. Answering a question by citing this book and quoting example code does not require permission. Incorporating a significant amount of example code from this book into your product's documentation does require permission.

We appreciate, but generally do not require, attribution. An attribution usually includes the title, author, publisher, and ISBN. For example: "*Effective Machine Learning Teams* platform gives you on-demand access to live training courses, in-depth learning paths, interactive coding environments, and a vast collection of text and video from O'Reilly and 200+ other publishers. For more information, visit *https://oreilly.com*.

How to Contact Us

Please address comments and questions concerning this book to the publisher:

O'Reilly Media, Inc.
1005 Gravenstein Highway North
Sebastopol, CA 95472
800-889-8969 (in the United States or Canada)
707-827-7019 (international or local)
707-829-0104 (fax)
support@oreilly.com
https://www.oreilly.com/about/contact.html

We have a web page for this book, where we list errata, examples, and any additional information. You can access this page at *https://oreil.ly/effective-ml-teams*.

For news and information about our books and courses, visit *https://oreilly.com*.

Find us on LinkedIn: *https://linkedin.com/company/oreilly-media*

Follow us on Twitter: *https://twitter.com/oreillymedia*

Watch us on YouTube: *https://youtube.com/oreillymedia*

Acknowledgments

When we started writing this book, we set out to share a collection of point practices that have helped 0.5

From David Tan

Thank you Nhung for being so patient and supportive through the late nights that I spent on this book. I would not have finished this book without your support.

From David Colls

I'd like to thank my family for putting up with a husband and father writing, reviewing, and researching content for a few months straight, on weekends, during movie nights, and on the basketball sidelines.

I'd like to thank the many Thoughtworks colleagues around the world who have gone before us in writing books and creating transformative perspectives on technology, stoking our determination to do the same, and showing us what good looks like. Closer to home, I'd like to thank all the Australian Thoughtworkers I've worked alongside over more than a decade for broadening my perspectives and enriching me professionally and as a human being.

I'd especially like to thank all those members of the Thoughtworks Australia Data & AI practice who I've had the privilege to work with as we built something new together—there's a little bit of each of you in this book. I'd also like to thank our clients for the trust they place in us to develop new approaches to their biggest opportunities and challenges.

Finally, I'd like to thank my coauthors David and Ada for their expertise and insight, feedback on my ideas, and structured approach to distilling and sharing our knowledge in this book. It has been a pleasure working with you.