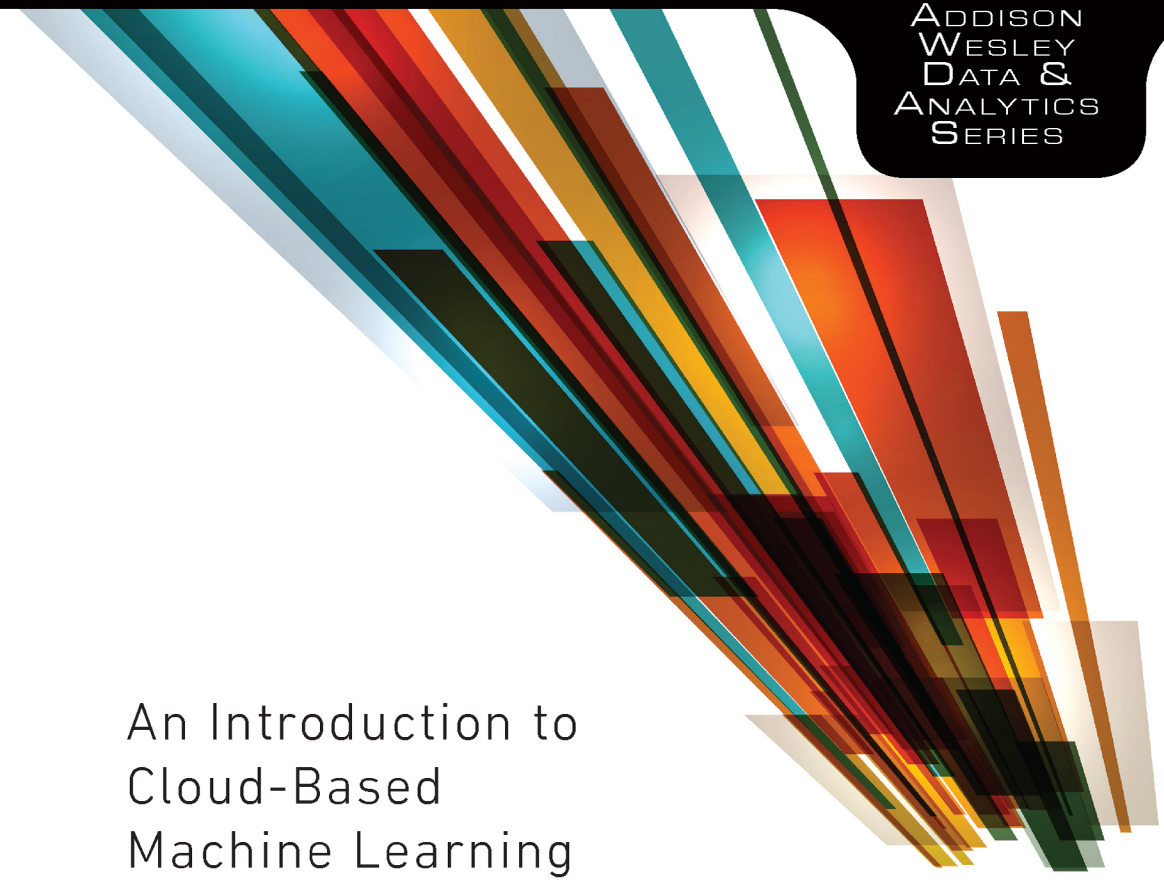

Addison Wesley Data & Analytics Series

An Introduction to Cloud-Based Machine Learning

Pragmatic AI

NOAH GIFT

Praise for *Pragmatic AI*

"[This is] a sweeping guide that bridges the gap between the promise of AI and solutions to the gritty problems that must be solved to deploy real-world projects. Clear and usable, *Pragmatic* AI covers much more than just Python and AI algorithms."

—Christopher Brousseau, founder and CEO of Surface Owl, the Enterprise AI platform

"A fantastic addition for any technology enthusiast! There is so much you can say about this book! Noah Gift really made this a practical guide for anyone involved with machine learning in the industry. Not only does it explain how one can apply machine learning on large data sets, it provides a valuable perspective on technology feedback loops. This book will benefit many data science and development teams so they can create and maintain their applications efficiently right from the beginning."

—Nivas Durairaj, technical account manager, AWS (Certified Professional Architect AWS)

"A great read if you want insights into building production-quality ML pipelines and tools that truly help your data engineering, data science, or data DevOps team. Even experienced developers often find themselves spinning their wheels on low-productivity tasks. Oftentimes, software books and university classes don't explain the steps needed to go into production. Noah has a gift for finding pragmatic approaches to software deployments that can really accelerate the development and delivery process. He has a focus and passion for enabling rapid software solutions that is very unique.

"The key to building production-quality ML pipelines is automation. Tasks and steps, which engineers may do manually during the research or prototype phase, must be automated and scaled in order to create a production system. This book is full of practical and fun exercises that will help any Python developer automate and extend their pipelines into the cloud.

"I'm currently working with big data, ML pipelines, Python, AWS, Google cloud, and Azure at Roofstock.com, an online real estate company. Our analytics database is approaching 500 million rows! Within this book I found many practical tips and exercises that will immediately improve my own productivity. Recommended!"

—Michael Vierling, lead engineer, Roofstock

Pragmatic AI: An Introduction to Cloud-Based Machine Learning

Table of Contents

Pearson

Table of Contents

Table of Contents

Table of Contents

Table of Contents

Pearson