
Humble Book Bundle Data Science Presented By Oreilly

The Science of Getting from Where You Are to Where You Want to Be
Feature Engineering for Machine Learning
Analytical Skills for AI and Data Science
Learning ITIL Made Simple with Real-life Examples
Learning TensorFlow
Technology, Metaphor, and the Search for Meaning
Big data, machine learning, and more, using Python tools
Introducing Data Science
Hands-On Machine Learning for Algorithmic Trading
The Book of R
The Self-Service Data Roadmap
Hadoop: The Definitive Guide
Build a Career in Data Science
Design and implement investment strategies based on smart algorithms that learn from data using Python
From Gods and Goddesses to Monsters and Mortals, Your Guide to Ancient Mythology
Business Intelligence For Dummies
Malware Data Science
Expert techniques to train advanced neural networks using TensorFlow and Keras
Practical Google Analytics and Google Tag Manager for Developers
Machine Learning and Security
Building Skills for an AI-Driven Enterprise
Attack Detection and Attribution
Humble Pi
Building Machine Learning Powered Applications
A First Course in Programming and Statistics
The Gentleman Bastard Series 3-Book Bundle
Hands-On Data Analysis with Pandas
Scrape, Clean, Explore & Transform Your Data
Deep Learning for Computer Vision
A Guide to Building Deep Learning Systems
How to Change
Working with Structured Data in Python
Math for Programmers
Syntax, Concepts, and Examples
Smarter Decisions, Better Results
A Beginner's Guide to Building Blockchain Solutions
Deep Learning
Machine Learning in Action

GOODMAN KIDD

The Science of Getting from Where You Are to Where You Want to Be Apress

Scientific progress depends on good research, and good research needs good statistics. But statistical analysis is tricky to get right, even for the best and brightest of us. You'd be surprised how many scientists are doing it wrong. *Statistics Done Wrong* is a pithy, essential guide to statistical blunders in modern science that will show you how to keep your research blunder-free. You'll examine embarrassing errors and omissions in recent research, learn about the misconceptions and scientific politics that allow these mistakes to happen, and begin your quest to reform the way you and your peers do statistics. You'll find advice on: -Asking the right question, designing the right experiment, choosing the right statistical analysis, and sticking to the plan -How to think about p values, significance, insignificance, confidence intervals, and regression -Choosing the right sample size and avoiding false positives -Reporting your analysis and publishing your data and source code -Procedures to follow, precautions to take, and analytical software that can help Scientists: Read this concise, powerful guide to help you produce statistically sound research. Statisticians: Give this book to everyone you know. The first step toward statistics done right is *Statistics Done Wrong*.

Feature Engineering for Machine Learning "O'Reilly Media, Inc."

Learn the skills necessary to design, build, and deploy applications powered by machine learning (ML). Through the course of this hands-on book, you'll build an example ML-driven application from initial idea to deployed product. Data scientists, software engineers, and product managers—including experienced practitioners and novices alike—will learn the tools, best practices, and challenges involved in building a real-world ML application step by step. Author Emmanuel Ameisen, an experienced data scientist who led an AI education program, demonstrates practical ML concepts using code snippets, illustrations, screenshots, and interviews with industry leaders. Part I teaches you how to plan an ML application and measure success. Part II explains how to build a working ML model. Part III demonstrates ways to improve the model until it fulfills your original vision. Part IV covers deployment and monitoring strategies. This book will help you: Define your product goal and set up a machine learning problem Build your first end-to-end pipeline quickly and acquire an initial dataset Train and evaluate your ML models and address performance bottlenecks Deploy and monitor your models in a production environment

Analytical Skills for AI and Data Science No Starch Press

Feature engineering is a crucial step in the machine-learning pipeline, yet this topic is rarely examined on its own. With this practical book, you'll learn techniques for extracting and transforming features—the numeric representations of raw data—into formats for machine-learning models. Each chapter guides you through a single data problem, such as how to represent text or image data. Together, these examples illustrate the main principles of feature engineering. Rather than simply teach these principles, authors Alice Zheng and Amanda Casari focus on practical application with exercises throughout the book. The closing chapter brings everything together by

tackling a real-world, structured dataset with several feature-engineering techniques. Python packages including numpy, Pandas, Scikit-learn, and Matplotlib are used in code examples. You'll examine: Feature engineering for numeric data: filtering, binning, scaling, log transforms, and power transforms Natural text techniques: bag-of-words, n-grams, and phrase detection Frequency-based filtering and feature scaling for eliminating uninformative features Encoding techniques of categorical variables, including feature hashing and bin-counting Model-based feature engineering with principal component analysis The concept of model stacking, using k-means as a featurization technique Image feature extraction with manual and deep-learning techniques

Learning ITIL Made Simple with Real-life Examples "O'Reilly Media, Inc."

As part of the best-selling Pocket Primer series, this book is designed to introduce the reader to the basic concepts of data analytics using Python 3. It is intended to be a fast-paced introduction to some basic features of data analytics and also covers statistics, data visualization, and data cleaning. The book includes numerous code samples using NumPy, Pandas, Matplotlib, Seaborn, and features an appendix on regular expressions. Companion files with source code and color figures are available. FEATURES: Includes a concise introduction to Python 3 Provides a thorough introduction to data and data cleaning Covers NumPy and Pandas Introduces statistical concepts and data visualization (Matplotlib/Seaborn) Features an appendix on regular expressions Includes companion files with source code and figures

Learning TensorFlow O'Reilly Media

Roughly inspired by the human brain, deep neural networks trained with large amounts of data can solve complex tasks with unprecedented accuracy. This practical book provides an end-to-end guide to TensorFlow, the leading open source software library that helps you build and train neural networks for computer vision, natural language processing (NLP), speech recognition, and general predictive analytics. Authors Tom Hope, Yehezkel Resheff, and Itay Lieder provide a hands-on approach to TensorFlow fundamentals for a broad technical audience—from data scientists and engineers to students and researchers. You'll begin by working through some basic examples in TensorFlow before diving deeper into topics such as neural network architectures, TensorBoard visualization, TensorFlow abstraction libraries, and multithreaded input pipelines. Once you finish this book, you'll know how to build and deploy production-ready deep learning systems in TensorFlow. Get up and running with TensorFlow, rapidly and painlessly Learn how to use TensorFlow to build deep learning models from the ground up Train popular deep learning models for computer vision and NLP Use extensive abstraction libraries to make development easier and faster Learn how to scale TensorFlow, and use clusters to distribute model training Deploy TensorFlow in a production setting

Technology, Metaphor, and the Search for Meaning Apress

Whether you're a marketer with development skills or a full-on web developer/analyst, *Practical Google Analytics and Google Tag Manager for Developers* shows you how to implement Google Analytics using Google Tag Manager to jumpstart your web analytics measurement. There's a reason that so many organizations use Google Analytics. Effective collection of data with Google Analytics

can reduce customer acquisition costs, provide priceless feedback on new product initiatives, and offer insights that will grow a customer or client base. So where does Google Tag Manager fit in? Google Tag Manager allows for unprecedented collaboration between marketing and technical teams, lightning fast updates to your site, and standardization of the most common tags for on-site tracking and marketing efforts. To achieve the rich data you're really after to better serve your users' needs, you'll need the tools Google Tag Manager provides for a best-in-class implementation of Google Analytics measurement on your site. Written by data evangelist and Google Analytics expert Jonathan Weber and the team at LunaMetrics, this book offers foundational knowledge, a collection of practical Google Tag Manager recipes, well-tested best practices, and troubleshooting tips to get your implementation in tip-top condition. It covers topics including:

- Google Analytics implementation via Google Tag Manager
- How to customize Google Analytics for your unique situation
- Using Google Tag Manager to track and analyze interactions across multiple devices and touch points
- How to extract data from Google Analytics and use Google BigQuery to analyze Big Data questions

What You'll Learn Implementation approaches for Google Analytics, including common pitfalls and troubleshooting strategies. How to use tools like Google Tag Manager and jQuery to jumpstart your Google Analytics implementation. How to track metrics beyond page views to other critical user interactions, such as clicks on outbound links or downloads, scrolling and page engagement, usage of AJAX forms, and much more. How to incorporate additional, customized data into Google Analytics to track individual users or enrich data about their behavior. Who This Book Is For Web developers, data analysts, and marketers with a basic familiarity with Google Analytics from an end-user perspective, as well as some knowledge of HTML and JavaScript.

Big data, machine learning, and more, using Python tools "O'Reilly Media, Inc."

Learn how to model and train advanced neural networks to implement a variety of Computer Vision tasks Key Features Train different kinds of deep learning model from scratch to solve specific problems in Computer Vision Combine the power of Python, Keras, and TensorFlow to build deep learning models for object detection, image classification, similarity learning, image captioning, and more Includes tips on optimizing and improving the performance of your models under various constraints Book Description Deep learning has shown its power in several application areas of Artificial Intelligence, especially in Computer Vision. Computer Vision is the science of understanding and manipulating images, and finds enormous applications in the areas of robotics, automation, and so on. This book will also show you, with practical examples, how to develop Computer Vision applications by leveraging the power of deep learning. In this book, you will learn different techniques related to object classification, object detection, image segmentation, captioning, image generation, face analysis, and more. You will also explore their applications using popular Python libraries such as TensorFlow and Keras. This book will help you master state-of-the-art, deep learning algorithms and their implementation. What you will learn Set up an environment for deep learning with Python, TensorFlow, and Keras Define and train a model for image and video classification Use features from a pre-trained Convolutional Neural Network model for image retrieval Understand and implement object detection using the real-world Pedestrian Detection scenario Learn about various problems in image captioning and how to overcome them by training images and text together Implement similarity matching and train a model for face recognition

Understand the concept of generative models and use them for image generation Deploy your deep learning models and optimize them for high performance Who this book is for This book is targeted at data scientists and Computer Vision practitioners who wish to apply the concepts of Deep Learning to overcome any problem related to Computer Vision. A basic knowledge of programming in Python—and some understanding of machine learning concepts—is required to get the best out of this book.

Introducing Data Science Packt Publishing Ltd

Summary Introducing Data Science teaches you how to accomplish the fundamental tasks that occupy data scientists. Using the Python language and common Python libraries, you'll experience firsthand the challenges of dealing with data at scale and gain a solid foundation in data science. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Many companies need developers with data science skills to work on projects ranging from social media marketing to machine learning. Discovering what you need to learn to begin a career as a data scientist can seem bewildering. This book is designed to help you get started. About the Book Introducing Data Science Introducing Data Science explains vital data science concepts and teaches you how to accomplish the fundamental tasks that occupy data scientists. You'll explore data visualization, graph databases, the use of NoSQL, and the data science process. You'll use the Python language and common Python libraries as you experience firsthand the challenges of dealing with data at scale. Discover how Python allows you to gain insights from data sets so big that they need to be stored on multiple machines, or from data moving so quickly that no single machine can handle it. This book gives you hands-on experience with the most popular Python data science libraries, Scikit-learn and StatsModels. After reading this book, you'll have the solid foundation you need to start a career in data science. What's Inside Handling large data Introduction to machine learning Using Python to work with data Writing data science algorithms About the Reader This book assumes you're comfortable reading code in Python or a similar language, such as C, Ruby, or JavaScript. No prior experience with data science is required. About the Authors Davy Cielen, Arno D. B. Meysman, and Mohamed Ali are the founders and managing partners of Optimately and Maiton, where they focus on developing data science projects and solutions in various sectors. Table of Contents Data science in a big data world The data science process Machine learning Handling large data on a single computer First steps in big data Join the NoSQL movement The rise of graph databases Text mining and text analytics Data visualization to the end user

Hands-On Machine Learning for Algorithmic Trading John Wiley & Sons

With the resurgence of neural networks in the 2010s, deep learning has become essential for machine learning practitioners and even many software engineers. This book provides a comprehensive introduction for data scientists and software engineers with machine learning experience. You'll start with deep learning basics and move quickly to the details of important advanced architectures, implementing everything from scratch along the way. Author Seth Weidman shows you how neural networks work using a first principles approach. You'll learn how to apply multilayer neural networks, convolutional neural networks, and recurrent neural networks from the ground up. With a thorough understanding of how neural networks work mathematically,

computationally, and conceptually, you'll be set up for success on all future deep learning projects. This book provides: Extremely clear and thorough mental models—accompanied by working code examples and mathematical explanations—for understanding neural networks Methods for implementing multilayer neural networks from scratch, using an easy-to-understand object-oriented framework Working implementations and clear-cut explanations of convolutional and recurrent neural networks Implementation of these neural network concepts using the popular PyTorch framework

The Book of R Random House

A guide for data managers and analyzers shares guidelines for identifying patterns, predicting future outcomes, and presenting findings to others; drawing on current research in cognitive science and learning theory while covering such additional topics as assessing data quality, handling ambiguous information, and organizing data within market groups. Original.

The Self-Service Data Roadmap "O'Reilly Media, Inc."

You're intelligent, right? So you've already figured out that Business Intelligence can be pretty valuable in making the right decisions about your business. But you've heard at least a dozen definitions of what it is, and heard of at least that many BI tools. Where do you start? Business Intelligence For Dummies makes BI understandable! It takes you step by step through the technologies and the alphabet soup, so you can choose the right technology and implement a successful BI environment. You'll see how the applications and technologies work together to access, analyze, and present data that you can use to make better decisions about your products, customers, competitors, and more. You'll find out how to: Understand the principles and practical elements of BI Determine what your business needs Compare different approaches to BI Build a solid BI architecture and roadmap Design, develop, and deploy your BI plan Relate BI to data warehousing, ERP, CRM, and e-commerce Analyze emerging trends and developing BI tools to see what else may be useful Whether you're the business owner or the person charged with developing and implementing a BI strategy, checking out Business Intelligence For Dummies is a good business decision.

Hadoop: The Definitive Guide Apress

Ready to unlock the power of your data? With this comprehensive guide, you'll learn how to build and maintain reliable, scalable, distributed systems with Apache Hadoop. This book is ideal for programmers looking to analyze datasets of any size, and for administrators who want to set up and run Hadoop clusters. You'll find illuminating case studies that demonstrate how Hadoop is used to solve specific problems. This third edition covers recent changes to Hadoop, including material on the new MapReduce API, as well as MapReduce 2 and its more flexible execution model (YARN). Store large datasets with the Hadoop Distributed File System (HDFS) Run distributed computations with MapReduce Use Hadoop's data and I/O building blocks for compression, data integrity, serialization (including Avro), and persistence Discover common pitfalls and advanced features for writing real-world MapReduce programs Design, build, and administer a dedicated Hadoop cluster—or run Hadoop in the cloud Load data from relational databases into HDFS, using Sqoop Perform large-scale data processing with the Pig query language Analyze datasets with Hive, Hadoop's data warehousing system Take advantage of HBase for structured and semi-structured

data, and ZooKeeper for building distributed systems

Build a Career in Data Science Simon and Schuster

While several market-leading companies have successfully transformed their business models by following data- and AI-driven paths, the vast majority have yet to reap the benefits. How can your business and analytics units gain a competitive advantage by capturing the full potential of this predictive revolution? This practical guide presents a battle-tested end-to-end method to help you translate business decisions into tractable prescriptive solutions using data and AI as fundamental inputs. Author Daniel Vaughan shows data scientists, analytics practitioners, and others interested in using AI to transform their businesses not only how to ask the right questions but also how to generate value using modern AI technologies and decision-making principles. You'll explore several use cases common to many enterprises, complete with examples you can apply when working to solve your own issues. Break business decisions into stages that can be tackled using different skills from the analytical toolbox Identify and embrace uncertainty in decision making and protect against common human biases Customize optimal decisions to different customers using predictive and prescriptive methods and technologies Ask business questions that create high value through AI- and data-driven technologies

Design and implement investment strategies based on smart algorithms that learn from data using Python Simon and Schuster

A strikingly original exploration of what it might mean to be authentically human in the age of artificial intelligence, from the author of the critically-acclaimed Interior States. "Meghan O'Gieblyn is a brilliant and humble philosopher, and her book is an explosively thought-provoking, candidly personal ride I wished never to end ... This book is such an original synthesis of ideas and disclosures. It introduces what will soon be called the O'Gieblyn genre of essay writing." —Heidi Julavits, author of *The Folded Clock* For most of human history the world was a magical and enchanted place ruled by forces beyond our understanding. The rise of science and Descartes's division of mind from world made materialism our ruling paradigm, in the process asking whether our own consciousness—i.e., souls—might be illusions. Now the inexorable rise of technology, with artificial intelligences that surpass our comprehension and control, and the spread of digital metaphors for self-understanding, the core questions of existence—identity, knowledge, the very nature and purpose of life itself—urgently require rethinking. Meghan O'Gieblyn tackles this challenge with philosophical rigor, intellectual reach, essayistic verve, refreshing originality, and an ironic sense of contradiction. She draws deeply and sometimes humorously from her own personal experience as a formerly religious believer still haunted by questions of faith, and she serves as the best possible guide to navigating the territory we are all entering.

From Gods and Goddesses to Monsters and Mortals, Your Guide to Ancient Mythology "O'Reilly Media, Inc."

Hailed by George R. R. Martin as "a fresh, original, and engrossing tale by a bright new voice in the fantasy genre," *The Lies of Locke Lamora* introduced one of the most compelling protagonists in years and established Scott Lynch as a household name. Years later, the Gentleman Bastard series continues to shine with Lynch's signature blend of swashbuckling adventure, sparkling wit, and seemingly impossible heists. It's all here in this rollicking eBook bundle, which collects the first three

volumes of Lynch's one-of-a-kind series: *THE LIES OF LOCKE LAMORA RED SEAS UNDER RED SKIES THE REPUBLIC OF THIEVES* An orphan's life is harsh—and often short—in the mysterious island city of Camorr. But young Locke Lamora dodges death and slavery, becoming a thief under the tutelage of a gifted con artist. As leader of the band of light-fingered brothers known as the Gentleman Bastards, Locke is soon infamous, fooling even the underworld's most feared ruler. But in the shadows lurks someone still more ambitious and deadly. Faced with a bloody coup that threatens to destroy everyone and everything that holds meaning in his mercenary life, Locke vows to beat the enemy at his own brutal game—or die trying. Praise for the Gentleman Bastard series *The Lies of Locke Lamora* "Remarkable . . . Scott Lynch's first novel, *The Lies of Locke Lamora*, exports the suspense and wit of a cleverly constructed crime caper into an exotic realm of fantasy, and the result is engagingly entertaining."—*The Times* (London) "Right now, in the full flush of a second reading, I think *The Lies of Locke Lamora* is probably in my top ten favorite books ever. Maybe my top five. If you haven't read it, you should. If you have read it, you should probably read it again."—Patrick Rothfuss, New York Times bestselling author of *The Name of the Wind* "A unique fantasy milieu peopled by absorbing, colorful characters . . . Locke's wit and audacity endear him to victims and bystanders alike."—*The Seattle Times* *Red Seas Under Red Skies* "Lynch hasn't merely imagined a far-off world, he's created it, put it all down on paper—the smells, the sounds, the people, the feel of the place. The novel is a virtuoso performance, and sf/fantasy fans will gobble it up."—Booklist (starred review) "It'll only be a matter of time before Scott Lynch is mentioned in the same breath as George R. R. Martin and Steven Erikson."—Fantasy Book Critic "The kind of witty romp that reminds you exactly how much fun heroic fantasy is supposed to be."—SFX *The Republic of Thieves* "Fast-paced, fun, and impossible to put down . . . Locke and company remain among the most engaging protagonists in fantasy."—Publishers Weekly (starred review) "The Republic of Thieves has all the colorful action, witty repartee, and devious scheming that fans of the series have come to expect."—Wired "A fantasy world unique among its peers . . . If you're looking for a great new fantasy series this is one you won't want to miss. . . . In a word: AWESOME!"—SF Revu

Business Intelligence For Dummies Simon and Schuster

In *Math for Programmers* you'll explore important mathematical concepts through hands-on coding. Filled with graphics and more than 300 exercises and mini-projects, this book unlocks the door to interesting—and lucrative!—careers in some of today's hottest fields. As you tackle the basics of linear algebra, calculus, and machine learning, you'll master the key Python libraries used to turn them into real-world software applications. Summary To score a job in data science, machine learning, computer graphics, and cryptography, you need to bring strong math skills to the party. *Math for Programmers* teaches the math you need for these hot careers, concentrating on what you need to know as a developer. Filled with lots of helpful graphics and more than 200 exercises and mini-projects, this book unlocks the door to interesting—and lucrative!—careers in some of today's hottest programming fields. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Skip the mathematical jargon: This one-of-a-kind book uses Python to teach the math you need to build games, simulations, 3D graphics, and machine learning algorithms. Discover how algebra and calculus come alive when you see them in code! About the book In *Math for Programmers* you'll explore important mathematical concepts

through hands-on coding. Filled with graphics and more than 300 exercises and mini-projects, this book unlocks the door to interesting—and lucrative!—careers in some of today's hottest fields. As you tackle the basics of linear algebra, calculus, and machine learning, you'll master the key Python libraries used to turn them into real-world software applications. What's inside Vector geometry for computer graphics Matrices and linear transformations Core concepts from calculus Simulation and optimization Image and audio processing Machine learning algorithms for regression and classification About the reader For programmers with basic skills in algebra. About the author Paul Orland is a programmer, software entrepreneur, and math enthusiast. He is co-founder of Tachyus, a start-up building predictive analytics software for the energy industry. You can find him online at www.paulor.land. Table of Contents 1 Learning math with code PART I - VECTORS AND GRAPHICS 2 Drawing with 2D vectors 3 Ascending to the 3D world 4 Transforming vectors and graphics 5 Computing transformations with matrices 6 Generalizing to higher dimensions 7 Solving systems of linear equations PART 2 - CALCULUS AND PHYSICAL SIMULATION 8 Understanding rates of change 9 Simulating moving objects 10 Working with symbolic expressions 11 Simulating force fields 12 Optimizing a physical system 13 Analyzing sound waves with a Fourier series PART 3 - MACHINE LEARNING APPLICATIONS 14 Fitting functions to data 15 Classifying data with logistic regression 16 Training neural networks

Malware Data Science No Starch Press

With detailed notes, tables, and examples, this handy reference will help you navigate the basics of structured machine learning. Author Matt Harrison delivers a valuable guide that you can use for additional support during training and as a convenient resource when you dive into your next machine learning project. Ideal for programmers, data scientists, and AI engineers, this book includes an overview of the machine learning process and walks you through classification with structured data. You'll also learn methods for clustering, predicting a continuous value (regression), and reducing dimensionality, among other topics. This pocket reference includes sections that cover: Classification, using the Titanic dataset Cleaning data and dealing with missing data Exploratory data analysis Common preprocessing steps using sample data Selecting features useful to the model Model selection Metrics and classification evaluation Regression examples using k-nearest neighbor, decision trees, boosting, and more Metrics for regression evaluation Clustering Dimensionality reduction Scikit-learn pipelines *Expert techniques to train advanced neural networks using TensorFlow and Keras* "O'Reilly Media, Inc."

Malware Data Science explains how to identify, analyze, and classify large-scale malware using machine learning and data visualization. Security has become a "big data" problem. The growth rate of malware has accelerated to tens of millions of new files per year while our networks generate an ever-larger flood of security-relevant data each day. In order to defend against these advanced attacks, you'll need to know how to think like a data scientist. In *Malware Data Science*, security data scientist Joshua Saxe introduces machine learning, statistics, social network analysis, and data visualization, and shows you how to apply these methods to malware detection and analysis. You'll learn how to: - Analyze malware using static analysis - Observe malware behavior using dynamic analysis - Identify adversary groups through shared code analysis - Catch 0-day vulnerabilities by

building your own machine learning detector - Measure malware detector accuracy - Identify malware campaigns, trends, and relationships through data visualization Whether you're a malware analyst looking to add skills to your existing arsenal, or a data scientist interested in attack detection and threat intelligence, *Malware Data Science* will help you stay ahead of the curve. [Practical Google Analytics and Google Tag Manager for Developers](#) Mercury Learning and Information

Wall Street Journal bestseller "A welcome revelation." --The Financial Times Award-winning Wharton Professor and Choiceology podcast host Katy Milkman has devoted her career to the study of behavior change. In this ground-breaking book, Milkman reveals a proven path that can take you from where you are to where you want to be, with a foreword from psychologist Angela Duckworth, the best-selling author of *Grit*. Change comes most readily when you understand what's standing between you and success and tailor your solution to that roadblock. If you want to work out more but find exercise difficult and boring, downloading a goal-setting app probably won't help. But what if, instead, you transformed your workouts so they became a source of pleasure instead of a chore? Turning an uphill battle into a downhill one is the key to success. Drawing on Milkman's original research and the work of her world-renowned scientific collaborators, *How to Change* shares strategic methods for identifying and overcoming common barriers to change, such as impulsivity, procrastination, and forgetfulness. Through case studies and engaging stories, you'll learn:

- Why timing can be everything when it comes to making a change
- How to turn temptation and inertia into assets
- That giving advice, even if it's about something you're struggling with, can help you achieve more

Whether you're a manager, coach, or teacher aiming to help others change for the better or are struggling to kick-start change yourself, *How to Change* offers an invaluable, science-based blueprint for achieving your goals, once and for all.

Machine Learning and Security Simon and Schuster

Summary You are going to need more than technical knowledge to succeed as a data scientist. *Build a Career in Data Science* teaches you what school leaves out, from how to land your first job to the lifecycle of a data science project, and even how to become a manager. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology What are the keys to a data scientist's long-term success? Blending your technical know-how with the right "soft skills" turns out to be a central ingredient of a rewarding career. About the book *Build a Career in Data Science* is your guide to landing your first data science job and developing into a valued senior employee. By following clear and simple instructions, you'll learn to craft an amazing resume and ace your interviews. In this demanding, rapidly changing field, it can be challenging to keep projects on track, adapt to company needs, and manage tricky stakeholders. You'll love the insights on how to handle expectations, deal with failures, and plan your career path in the stories from seasoned data scientists included in the book. What's inside Creating a portfolio of data science projects Assessing and negotiating an offer Leaving gracefully and moving up the ladder Interviews with professional data scientists About the reader For readers who want to begin or advance a data science career. About the author Emily Robinson is a data scientist at Warby Parker. Jacqueline Nolis is a data science consultant and mentor. Table of Contents: PART 1 - GETTING STARTED WITH DATA SCIENCE 1. What is data science? 2. Data science companies 3. Getting the skills 4. Building a portfolio PART 2 - FINDING YOUR DATA SCIENCE JOB 5. The search: Identifying the right job for you 6. The application: Résumés and cover letters 7. The interview: What to expect and how to handle it 8. The offer: Knowing what to accept PART 3 - SETTLING INTO DATA SCIENCE 9. The first months on the job 10. Making an effective analysis 11. Deploying a model into production 12. Working with stakeholders PART 4 - GROWING IN YOUR DATA SCIENCE ROLE 13. When your data science project fails 14. Joining the data science community 15. Leaving your job gracefully 16. Moving up the ladder