



*Data Scientist teaches you a step-by-step approach to solving real-world data-centric problems. By breaking down carefully crafted examples, you'll learn to combine analytic, programming, and business perspectives into a repeatable process for extracting real knowledge from data. As you read, you'll discover (or remember) valuable statistical techniques and explore powerful data science software. More importantly, you'll put this knowledge together using a structured process for data science. When you've finished, you'll have a strong foundation for a lifetime of data science learning and practice. What's Inside The data science process, step-by-step How to anticipate problems Dealing with uncertainty Best practices in software and scientific thinking About the Reader Readers need beginner programming skills and knowledge of basic statistics. About the Author Brian Godsey has worked in software, academia, finance, and defense and has launched several data-centric start-ups. Table of Contents PART 1 - PREPARING AND GATHERING DATA AND KNOWLEDGE Philosophies of data science Setting goals by asking good questions Data all around us: the virtual wilderness Data wrangling: from capture to domestication Data assessment: poking and prodding PART 2 - BUILDING A PRODUCT WITH SOFTWARE AND STATISTICS Developing a plan Statistics and modeling: concepts and foundations Software: statistics in action Supplementary software: bigger, faster, more efficient Plan execution: putting it all together PART 3 - FINISHING OFF THE PRODUCT AND WRAPPING UP Delivering a product After product delivery: problems and revisions Wrapping up: putting the project away*

*Unlock the groundbreaking advances of deep learning with this extensively revised edition of the bestselling original. Learn directly from the creator of Keras and master practical Python deep learning techniques that are easy to apply in the real world. In Deep Learning with Python, Second Edition you will learn: Deep learning from first principles Image classification & image segmentation Timeseries forecasting Text classification and machine translation Text generation, neural style transfer, and image generation Deep Learning with Python has taught thousands of readers how to put the full capabilities of deep learning into action. This extensively revised second edition introduces deep learning using Python and Keras, and is loaded with insights for both novice and experienced ML practitioners. You'll learn practical techniques that are easy to apply in the real world, and important theory for perfecting neural networks. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Recent innovations in deep learning unlock exciting new software capabilities like automated language translation, image recognition, and more. Deep learning is becoming essential knowledge for every software developer, and modern tools like Keras and TensorFlow put it within your reach, even if you have no background in mathematics or data science. About the book Deep Learning with Python, Second Edition introduces the field of deep learning using Python and the powerful Keras library. In this new edition, Keras creator François Chollet offers insights for both novice and experienced machine learning practitioners. As you move through this book, you'll build your understanding through intuitive explanations, crisp illustrations, and clear examples. You'll pick up the skills to start developing deep-learning applications. What's inside Deep learning from first principles Image classification and image segmentation Time series forecasting Text classification and machine translation Text generation, neural style transfer, and image generation About the reader For readers with intermediate Python skills. No previous experience with Keras, TensorFlow, or machine learning is required. About the author François Chollet is a software engineer at Google and creator of the Keras deep-learning library. Table of Contents 1 What is deep learning? 2 The mathematical building blocks of neural networks 3 Introduction to Keras and TensorFlow 4 Getting started with neural networks: Classification and regression 5 Fundamentals of machine learning 6 The universal workflow of machine learning 7 Working with Keras: A deep dive 8 Introduction to deep learning for computer vision 9 Advanced deep learning for computer vision 10 Deep learning for timeseries 11 Deep learning for text 12 Generative deep learning 13 Best practices for the real world 14 Conclusions*

*As enterprise applications become larger and more distributed, new architectural approaches like reactive designs, microservices, and event streams are required knowledge. Vert.x in Action teaches you to build highly-scalable reactive enterprise applications using the mature, rock-solid Vert.x framework. Vert.x in Action gets you up to speed in the basics of asynchronous programming as you learn to design and code reactive applications. Using the Vert.x asynchronous APIs, you'll build services including web stack, messaging, authentication, and access control. You'll also dive into deployment of container-native components with Docker, Kubernetes, and OpenShift. Along the way, you'll check your app's health and learn to test its resilience to external service failures. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications.*

*Summary Express in Action is a carefully designed tutorial that teaches you how to build web applications using Node and Express. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Express.js is a web application framework for Node.js. Express organizes your server-side JavaScript into testable, maintainable modules. It provides a powerful set of features to efficiently manage routes, requests, and views along with beautiful boilerplate for your web applications. Express helps you concentrate on what your application does instead of managing time-consuming technical details. About the Book Express in Action teaches you how to build web applications using Node and Express. It starts by introducing Node's powerful traits and shows you how they map to the features of Express. You'll explore key development techniques, meet the rich ecosystem of companion tools and libraries, and get a glimpse into its inner workings. By the end of the book, you'll be able to use Express to build a Node app and know how to test it, hook it up to a database, and automate the dev process. What's Inside Simplify Node app setup with Express Testing Express applications Use Express for easy access to Node features Data storage with MongoDB Covers Express 4 and Express 5 alpha About the Reader To get the most out of this book, you'll need to know the basics of web application design and be proficient with JavaScript. About the Author Evan Hahn is an active member of the Node and Express community and contributes to many open source JavaScript projects. Table of Contents PART 1 INTRO What is Express? The basics of Node.js Foundations of Express PART 2 CORE Middleware Routing Building APIs Views and templates: Pug and EJS PART 3 EXPRESS IN CONTEXT Persisting your data with MongoDB Testing Express applications Security Deployment: assets and Heroku Best practices*

*Streaming Data*

*Writing Great Specifications*

*Learning Domain-Driven Design*

*The Programmer's Brain*

*DA Pam*

*Understanding the real-time pipeline*

*MongoDB in Action*

*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 ScienceIntroducing 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 Dave Clenden, Arno D. B. Weyman, and Mohamed Ali are the founders and managing partners of Optimately and Matlon, 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*

*Summary Relevant Search demystifies relevance work. Using Elasticsearch, it teaches you how to return engaging search results to your users, helping you understand and leverage the internals of Lucene-based search engines. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Users are accustomed to and expect instant, relevant search results. To achieve this, you must master the search engine. Yet for many developers, relevance ranking is mysterious or confusing. About the Book Relevant Search demystifies the subject and shows you that a search engine is a programmable relevance framework. You'll learn how to apply Elasticsearch or Solr to your business's unique ranking problems. The book demonstrates how to program relevance and how to incorporate secondary data sources, taxonomies, text analytics, and personalization. In practice, a relevance framework requires softer skills as well, such as collaborating with stakeholders to discover the right relevance requirements for your business. By the end, you'll be able to achieve a virtuous cycle of provable, measurable relevance improvements over a search product's lifetime. What's Inside Techniques for debugging relevance? Applying search engine features to real problems? Using the user interface to guide searchers? A systematic approach to relevance? A business culture focused on improving search About the Reader For developers trying to build smarter search with Elasticsearch or Solr. About the Authors Doug Turnbull is lead relevance consultant at OpenSource Connections, where he frequently speaks and blogs. John Beryman is a data engineer at Eventbrite, where he specializes in recommendations and search. Foreword author, Trey Grainger, is a director of engineering at CareerBuilder and author of Solr in Action. Table of Contents The search relevance problem Search under the hood Debugging your first relevance problem Taming tokens Basic multifield search Term-centric search Shaping the relevance function Providing relevance feedback Designing a relevance-focused search application The relevance-centered enterprise Semantic and personalized search*

*Summary MongoDB in Action, Second Edition is a completely revised and updated version. It introduces MongoDB 3.0 and the document-oriented database model. This perfectly paced book gives you both the big picture you'll need as a developer and enough low-level detail to satisfy system engineers. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology This document-oriented database was built for high availability, supports rich, dynamic schemas, and lets you easily distribute data across multiple servers. MongoDB 3.0 is flexible, scalable, and very fast, even with big data loads. About the Book MongoDB in Action, Second Edition is a completely revised and updated version. It introduces MongoDB 3.0 and the document-oriented database model. This perfectly paced book gives you both the big picture you'll need as a developer and enough low-level detail to satisfy system engineers. Lots of examples will help you develop confidence in the crucial area of data modeling. You'll also love the deep explanations of each feature, including replication, auto-sharding, and deployment. What's Inside Indexes, queries, and standard DB operations Aggregation and text searching Map-reduce for custom aggregations and reporting Deploying for scale and high availability Updated for Mongo 3.0 About the Reader Written for developers. No previous MongoDB or NoSQL experience is assumed. About the Authors After working at MongoDB, Kyle Banker is now at a startup. Peter Bakkum is a developer with MongoDB expertise. Shaun Verch has worked on the core server team at MongoDB. A Genentech engineer, Doug Garrett is one of the winners of the MongoDB Innovation Award for Analytics. A software architect, Tim Hawkins has led search engineering at Yahoo Europe. Technical Contributor: Wouter Thielen. Technical Editor: Mihails Teoukalos. Table of Contents PART 1 GETTING STARTED A database for the modern web MongoDB through the JavaScript shell Writing programs using MongoDB PART 2 APPLICATION DEVELOPMENT IN MONGODB Document-oriented data Constructing queries Aggregation Updates, atomic operations, and deletes PART 3 MONGODB MASTERY*

*Indexing and query optimization Text search Wicketizer and pluggable storage Replication Scaling your system with sharding Deployment and administration*

*Summary Sails.js in Action is a comprehensive guide to buildingenterprise-capable web applications using Node and Sails. Written by the creators of the Sails.js framework, this book carefully introduces each concept, technique, and tool with real-world examples and crystal clear explanations. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Sails makes professional web development a breeze. This instantly familiar MVC framework automatically handles the tedious application boilerplate, so you can concentrate on developing features and creating business value. You get powerful tools for rapid API development, task automation, an ORM, and easy integration with any web, mobile, or IoT frontend. And because you're using Node.js, it's JavaScript all the way down. About the Book Sails.js in Action is a comprehensive guide on how to build enterprise-capable web applications. Written by the creators of Sails.js, this book introduces each concept and technique with real-world examples and thorough explanations. As you read, you'll learn to build the backend of a typical web application while you explore real-time programming with WebSockets, security fundamentals, and best practices for building Sails/Node.js apps. What's Inside Creating the backend for a web, mobile, or IoT app Real-time programming with WebSockets User management, authentication, and password recovery Using Sails to autogenerate REST APIs Custom backend development and third-party API integrations About the Reader Readers should be comfortable with JavaScript and frontend web development. About the Author Mike McNeil is the creator of Sails.js. Irl Nathan is the producer of sailsCasts, a series focused on using Sails. Table of Contents Getting started First steps Using static assets Using the blueprint API Custom backend code*

*Using models Custom actions Server-rendered views Authentication and sessions Policies and access control Refactoring Embedded data and associations Ratings, followers, and search Realtime with WebSockets Deployment, testing, and security*

*Storm Applied*

*African American Films Through 1959*

*Grokking Simplicity*

*CoreOS in Action*

*Knative in Action*

*Summary Streaming Data introduces the concepts and requirements of streaming and real-time data systems. The book is an idea-rich tutorial that teaches you to think about how to efficiently interact with fast-flowing data. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology As humans, we're constantly filtering and deciphering the information streaming toward us. In the same way, streaming data applications can accomplish amazing tasks like reading live location data to recommend nearby services, tracking faults with machinery in real time, and sending digital receipts before your customers leave the shop. Recent advances in streaming data technology and techniques make it possible for any developer to build these applications if they have the right mindset. This book will let you join them. About the Book Streaming Data is an idea-rich tutorial that teaches you to think about efficiently interacting with fast-flowing data. Through relevant examples and illustrated use cases, you'll explore designs for applications that read, analyze, share, and store streaming data. Along the way, you'll discover the roles of key technologies like Spark, Storm, Kafka, Flink, RabbitMQ, and more. This book offers the perfect balance between big-picture thinking and implementation details. What's Inside The right way to collect real-time data Architecting a streaming pipeline Analyzing the data Which technologies to use and when About the Reader Written for developers familiar with relational database concepts. No experience with streaming or real-time applications required. About the Author Andrew Psaltis is a software engineer focused on massively scalable real-time analytics. Table of Contents PART 1 A NEW HOLISTIC APPROACH Introducing streaming data Getting data from clients: data ingestion Transposing the data from collection tier: decoupling the data pipeline Analyzing streaming data Algorithms for data analysis Storing the analyzed or collected data Making the data available Consumer device capabilities and limitations accessing the data PART 2 TAKING IT REAL WORLD Analyzing Meetup RSVPs in real time*

*In the classic style of Manning's "In Action" series, NHibernate in Action shows.NET developers how to use the NHibernate Object/Relational Mapping tool.This book is a translation from Java to .NET, as well as an expansion. ofManning's bestselling Hibernate in Action. All traces of Java have been carefullyreplaced by their .NET equivalents. The book shows how to implementcomplex business objects, and later teaches advanced techniques like cachingand session management. Readers will discover how to implement persistence in a .NET application, and how to configure NHibernate to specify the mappininformation between business objects and database tables. Readers will also beintroduced to the internal architecture of NHibernate by progressively buildings complete sample application using Agile methodologies. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book.*

*AI as a Service is a practical handbook to building and implementing serverless AI applications, without bogging you down with a lot of theory. Instead, you'll find easy-to-digest instruction and two complete hands-on serverless AI builds in this must-have guide! Summary Companies everywhere are moving everyday business processes over to the cloud, and AI is increasingly being given the reins in these tasks. As this massive digital transformation continues, the combination of serverless computing and AI promises to become the de facto standard for business-to-consumer platform development—and developers who can design, develop, implement, and maintain these systems will be in high demand! AI as a Service is a practical handbook to building and implementing serverless AI applications, without bogging you down with a lot of theory. Instead, you'll find easy-to-digest instruction and two complete hands-on serverless AI builds in this must-have guide! Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Cloud-based AI services can automate a variety of labor intensive business tasks in areas such as customer service, data analysis, and financial reporting. The secret is taking advantage of pre-built tools like Amazon Rekognition for image analysis or AWS Comprehend for natural language processing. That way, there's no need to build expensive custom software. Artificial Intelligence (AI), a machine's ability to learn and make predictions based on patterns it identifies, is already being leveraged by businesses around the world in areas like targeted product recommendations, financial forecasting and resource planning, customer service chatbots, healthcare diagnostics, data security, and more. With the exciting combination of serverless computing and AI, software developers now have enormous power to improve their businesses' existing systems and rapidly deploy new AI-enabled platforms. And to get on this fast-moving train, you don't have to invest loads of time and effort in becoming a data scientist or AI expert, thanks to cloud planning, customer service chatbots, healthcare diagnostics, data security, and more. About the book AI as a Service is a fast-paced guide to harnessing the power of cloud-based solutions. You'll learn to build real-world apps—such as chatbots and text-to-speech services—by stitching together cloud components. Work your way from small projects to large data-intensive applications. What's inside - Apply cloud AI services to existing platforms - Design and build scalable data pipelines - Debug and troubleshoot AI services - Start fast with serverless templates About the reader For software developers familiar with cloud basics. About the author Peter Elger and Eoin Shanaghy are founders and CEO/CTO of fourTheorem, a software solutions company providing expertise on architecture, DevOps, and machine learning. Table of Contents PART 1 - FIRST STEPS 1 A tale of two technologies 2 Building a serverless image recognition system, part 1 3 Building a serverless image recognition system, part 2 PART 2 - TOOLS OF THE TRADE 4 Building and securing a web application the serverless way 5 Adding AI interfaces to a web application 6 How to be effective with AI as a Service 7 Applying AI to existing platforms PART 3 - BRINGING IT ALL TOGETHER 8 Gathering data at scale for real-world AI 9 Extracting value from large data sets with AI*