

Online Library Sample Chapter 3 Manning Publications File Type Pdf Free Copy

Introduction to Information Retrieval Introducing Data Science Data Science Bookcamp Deep Learning with Python Machine Learning in Action Feature Engineering Bookcamp Deep Learning and the Game of Go Build a Career in Data Science Math for Programmers Spark in Action Grokking Machine Learning Machine Learning in Action Deep Learning for Search Practices of the Python Pro Grokking Deep Learning Speech & Language Processing Deep Learning with R Machine Learning Bookcamp Graph-Powered Machine Learning Data Analysis with Python and PySpark Real-World Machine Learning Grokking Simplicity The Quick Python Book Ensemble Methods for Machine Learning MLOps Engineering at Scale Graph Databases in Action Machine Learning with TensorFlow, Second Edition Machine Learning with TensorFlow Deep Learning with PyTorch Advanced Algorithms and Data Structures Machine Learning Systems Data Pipelines with Apache Airflow Human-in-the-Loop Machine Learning The Quick Python Book Hadoop in Practice Grokking Deep Reinforcement Learning Deep Reinforcement Learning in Action Grokking Artificial Intelligence Algorithms HTTP/2 in Action Natural Language Processing 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 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 Summary Natural Language Processing in Action is your guide to creating machines that understand human language using the power of Python with its ecosystem of packages dedicated to NLP and AI. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Recent advances in deep learning empower applications to understand text and speech with extreme accuracy. The result? Chatbots that can imitate real people, meaningful resume-to-job matches, superb predictive search, and automatically generated document summaries—all at a low cost. New techniques, along with accessible tools like Keras and TensorFlow, make professional-quality NLP easier than ever before. About the Book Natural Language Processing in Action is your guide to building machines that can read and interpret human language. In it, you'll use readily available Python packages to capture the meaning in text and react accordingly. The book expands traditional NLP approaches to include neural networks, modern deep learning algorithms, and generative techniques as you tackle real-world problems like extracting dates and names, composing text, and answering free-form questions. What's inside Some sentences in this book were written by NLP! Can you guess which ones? Working with Keras, TensorFlow, gensim, and scikit-learn Rule-based and data-based NLP Scalable pipelines About the Reader This book requires a basic understanding of deep learning and intermediate Python skills. About the Author Hobson Lane, Cole Howard, and Hannes Max Hapke are experienced NLP engineers who use these techniques in production. Table of Contents PART 1 - WORDY MACHINES Packets of thought (NLP overview) Build your vocabulary (word tokenization) Math with words (TF-IDF vectors) Finding meaning in word counts (semantic analysis) PART 2 - DEEPER LEARNING (NEURAL NETWORKS) Baby steps with neural networks (perceptrons and backpropagation) Reasoning with word vectors (Word2vec) Getting words in order with convolutional neural networks (CNNs) Loopy (recurrent) neural networks (RNNs) Improving retention with long short-term memory networks Sequence-to-sequence models and attention PART 3 - GETTING REAL (REAL-WORLD NLP CHALLENGES) Information extraction (named entity extraction and question answering) Getting chatty (dialog engines) Scaling up (optimization, parallelization, and batch processing) Ensemble machine learning combines the power of multiple machine learning approaches, working together to deliver models that are highly performant and highly accurate. Inside Ensemble Methods for Machine Learning you will find: Methods for classification, regression, and recommendations Sophisticated off-the-shelf ensemble implementations Random forests, boosting, and gradient boosting Feature engineering and ensemble diversity Interpretability and explainability for ensemble methods Ensemble machine learning trains a diverse group of machine learning models to work together, aggregating their output to deliver richer results than a single model. Now in Ensemble Methods for Machine Learning you'll discover core ensemble methods that have proven records in both data science competitions and real-world applications. Hands-on case studies show you how each algorithm works in production. By the time you're done, you'll know the benefits, limitations, and practical methods of applying ensemble machine learning to real-world data, and be ready to build more explainable ML systems. About the Technology Automatically compare, contrast, and blend the output from multiple models to squeeze the best results from your data. Ensemble machine learning applies a “wisdom of crowds” method that dodges the inaccuracies and limitations of a single model. By basing responses on multiple perspectives, this innovative approach can deliver robust predictions even without massive datasets. About the Book Ensemble Methods for Machine Learning teaches you practical techniques for applying multiple ML approaches simultaneously. Each chapter contains a unique case study that demonstrates a fully functional ensemble method, with examples including medical diagnosis, sentiment analysis, handwriting classification, and more. There's no complex math or theory—you'll learn in a visuals-first manner, with ample code for easy experimentation! What's Inside Bagging, boosting, and gradient boosting Methods for classification, regression, and retrieval Interpretability and explainability for ensemble methods Feature engineering and ensemble diversity About the Reader For Python programmers with machine learning experience. About the Author Gautam Kunapuli has over 15 years of experience in academia and the machine learning industry. Table of Contents PART 1 - THE BASICS OF ENSEMBLES 1 Ensemble methods: Hype or hallelujah? PART 2 - ESSENTIAL ENSEMBLE METHODS 2 Homogeneous parallel ensembles: Bagging and random forests 3 Heterogeneous parallel ensembles: Combining strong learners 4 Sequential ensembles: Adaptive boosting 5 Sequential ensembles: Gradient boosting 6 Sequential ensembles: Newton boosting PART 3 - ENSEMBLES IN THE WILD: ADAPTING ENSEMBLE METHODS TO YOUR DATA 7 Learning with continuous and count labels 8 Learning with categorical features 9 Explaining your ensembles Summary Professional developers know the many benefits of writing application code that's clean, well-organized, and easy to maintain. By learning and following established patterns and best practices, you can take your code and your career to a new level. With Practices of the Python Pro, you'll learn to design professional-level, clean, easily maintainable software at scale using the incredibly popular programming language, Python. You'll find easy-to-grok examples that use pseudocode and Python to introduce software development best practices, along with dozens of instantly useful techniques that will help you code like a pro. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Professional-quality code does more than just run without bugs. It's clean, readable, and easy to maintain. To step up from a capable Python coder to a professional developer, you need to learn industry standards for coding style, application design, and development process. That's where this book is indispensable. About the book Practices of the Python Pro teaches you to design and write professional-quality software that's understandable, maintainable, and extensible. Dane Hillard is a Python pro who has helped many dozens of developers make this step, and he knows what it takes. With helpful examples and exercises, he teaches you when, why, and how to modularize your code, how to improve quality by reducing complexity, and much more. Embrace these core principles, and your code will become easier for you and others to read, maintain, and reuse. What's inside Organizing large Python projects Achieving the right levels of abstraction Writing clean, reusable code Inheritance and composition Considerations for testing and performance About the reader For readers familiar with the basics of Python, or another OO language. About the author Dane Hillard has spent the majority of his development career using Python to build web applications. Table of Contents: PART 1 WHY IT ALL MATTERS 1 | The bigger picture PART 2 FOUNDATIONS OF DESIGN 2 | Separation of concerns 3 | Abstraction and encapsulation 4 | Designing for high performance 5 | Testing your software PART 3 NAILING DOWN LARGE SYSTEMS 6 | Separation of concerns in practice 7 | Extensibility and flexibility 8 | The rules (and exceptions) of inheritance 9 | Keeping things lightweight 10 | Achieving loose coupling PART 4 WHAT'S NEXT? 11 | Onward and upward Summary Real-World Machine Learning is a practical guide designed to teach working developers the art of ML project execution. Without overdosing you on academic theory and complex mathematics, it introduces the day-to-day practice of machine learning, preparing you to successfully build and deploy powerful ML systems. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Machine learning systems help you find valuable insights and patterns in data, which you'd never recognize with traditional methods. In the real world, ML

techniques give you a way to identify trends, forecast behavior, and make fact-based recommendations. It's a hot and growing field, and up-to-speed ML developers are in demand. About the Book Real-World Machine Learning will teach you the concepts and techniques you need to be a successful machine learning practitioner without overdosing you on abstract theory and complex mathematics. By working through immediately relevant examples in Python, you'll build skills in data acquisition and modeling, classification, and regression. You'll also explore the most important tasks like model validation, optimization, scalability, and real-time streaming. When you're done, you'll be ready to successfully build, deploy, and maintain your own powerful ML systems. What's Inside Predicting future behavior Performance evaluation and optimization Analyzing sentiment and making recommendations About the Reader No prior machine learning experience assumed. Readers should know Python. About the Authors Henrik Brink, Joseph Richards and Mark Fetherolf are experienced data scientists engaged in the daily practice of machine learning. Table of Contents PART 1: THE MACHINE-LEARNING WORKFLOW What is machine learning? Real-world data Modeling and prediction Model evaluation and optimization Basic feature engineering PART 2: PRACTICAL APPLICATION Example: NYC taxi data Advanced feature engineering Advanced NLP example: movie review sentiment Scaling machine-learning workflows Example: digital display advertising "From start to finish, the best book to help you learn AI algorithms and recall why and how you use them." - Linda Ristevski, York Region District School Board "This book takes an impossibly broad area of computer science and communicates what working developers need to understand in a clear and thorough way." - David Jacobs, Product Advance Local Key Features Master the core algorithms of deep learning and AI Build an intuitive understanding of AI problems and solutions Written in simple language, with lots of illustrations and hands-on examples Creative coding exercises, including building a maze puzzle game and exploring drone optimization About The Book "Artificial intelligence" requires teaching a computer how to approach different types of problems in a systematic way. The core of AI is the algorithms that the system uses to do things like identifying objects in an image, interpreting the meaning of text, or looking for patterns in data to spot fraud and other anomalies. Mastering the core algorithms for search, image recognition, and other common tasks is essential to building good AI applications Grokking Artificial Intelligence Algorithms uses illustrations, exercises, and jargon-free explanations to teach fundamental AI concepts. You'll explore coding challenges like detecting bank fraud, creating artistic masterpieces, and setting a self-driving car in motion. All you need is the algebra you remember from high school math class and beginning programming skills. What You Will Learn Use cases for different AI algorithms Intelligent search for decision making Biologically inspired algorithms Machine learning and neural networks Reinforcement learning to build a better robot This Book Is Written For For software developers with high school-level math skills. About the Author Rishal Hurbans is a technologist, startup and AI group founder, and international speaker. Table of Contents 1 Intuition of artificial intelligence 2 Search fundamentals 3 Intelligent search 4 Evolutionary algorithms 5 Advanced evolutionary approaches 6 Swarm intelligence: Ants 7 Swarm intelligence: Particles 8 Machine learning 9 Artificial neural networks 10 Reinforcement learning with Q-learning Dodge costly and time-consuming infrastructure tasks, and rapidly bring your machine learning models to production with MLOps and pre-built serverless tools! In MLOps Engineering at Scale you will learn: Extracting, transforming, and loading datasets Querying datasets with SQL Understanding automatic differentiation in PyTorch Deploying model training pipelines as a service endpoint Monitoring and managing your pipeline's life cycle Measuring performance improvements MLOps Engineering at Scale shows you how to put machine learning into production efficiently by using pre-built services from AWS and other cloud vendors. You'll learn how to rapidly create flexible and scalable machine learning systems without laboring over time-consuming operational tasks or taking on the costly overhead of physical hardware. Following a real-world use case for calculating taxi fares, you will engineer an MLOps pipeline for a PyTorch model using AWS server-less capabilities. About the technology A production-ready machine learning system includes efficient data pipelines, integrated monitoring, and means to scale up and down based on demand. Using cloud-based services to implement ML infrastructure reduces development time and lowers hosting costs. Serverless MLOps eliminates the need to build and maintain custom infrastructure, so you can concentrate on your data, models, and algorithms. About the book MLOps Engineering at Scale teaches you how to implement efficient machine learning systems using pre-built services from AWS and other cloud vendors. This easy-to-follow book guides you step-by-step as you set up your serverless ML infrastructure, even if you've never used a cloud platform before. You'll also explore tools like PyTorch Lightning, Optuna, and MLFlow that make it easy to build pipelines and scale your deep learning models in production. What's inside Reduce or eliminate ML infrastructure management Learn state-of-the-art MLOps tools like PyTorch Lightning and MLFlow Deploy training pipelines as a service endpoint Monitor and manage your pipeline's life cycle Measure performance improvements About the reader Readers need to know Python, SQL, and the basics of machine learning. No cloud experience required. About the author Carl Osipov implemented his first neural net in 2000 and has worked on deep learning and machine learning at Google and IBM. Table of Contents PART 1 - MASTERING THE DATA SET 1 Introduction to serverless machine learning 2 Getting started with the data set 3 Exploring and preparing the data set 4 More exploratory data analysis and data preparation PART 2 - PYTORCH FOR SERVERLESS MACHINE LEARNING 5 Introducing PyTorch: Tensor basics 6 Core PyTorch: Autograd, optimizers, and utilities 7 Serverless machine learning at scale 8 Scaling out with distributed training PART 3 - SERVERLESS MACHINE LEARNING PIPELINE 9 Feature selection 10 Adopting PyTorch Lightning 11 Hyperparameter optimization 12 Machine learning pipeline Summary Machine Learning Systems: Designs that scale is an example-rich guide that teaches you how to implement reactive design solutions in your machine learning systems to make them as reliable as a well-built web app. Foreword by Sean Owen, Director of Data Science, Cloudera Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology If you're building machine learning models to be used on a small scale, you don't need this book. But if you're a developer building a production-grade ML application that needs quick response times, reliability, and good user experience, this is the book for you. It collects principles and practices of machine learning systems that are dramatically easier to run and maintain, and that are reliably better for users. About the Book Machine Learning Systems: Designs that scale teaches you to design and implement production-ready ML systems. You'll learn the principles of reactive design as you build pipelines with Spark, create highly scalable services with Akka, and use powerful machine learning libraries like MLlib on massive datasets. The examples use the Scala language, but the same ideas and tools work in Java, as well. What's Inside Working with Spark, MLlib, and Akka Reactive design patterns Monitoring and maintaining a large-scale system Futures, actors, and supervision About the Reader Readers need intermediate skills in Java or Scala. No prior machine learning experience is assumed. About the Author Jeff Smith builds powerful machine learning systems. For the past decade, he has been working on building data science applications, teams, and companies as part of various teams in New York, San Francisco, and Hong Kong. He blogs (<https://medium.com/@jeffksmithjr>), tweets (@jeffksmithjr), and speaks (www.jeffsmith.tech/speaking) about various aspects of building real-world machine learning systems. Table of Contents PART 1 - FUNDAMENTALS OF REACTIVE MACHINE LEARNING Learning reactive machine learning Using reactive tools PART 2 - BUILDING A REACTIVE MACHINE LEARNING SYSTEM Collecting data Generating features Learning models Evaluating models Publishing models Responding PART 3 - OPERATING A MACHINE LEARNING SYSTEM Delivering Evolving intelligence Upgrade your machine learning models with graph-based algorithms, the perfect structure for complex and interlinked data. Summary In Graph-Powered Machine Learning, you will learn: The lifecycle of a machine learning project Graphs in big data platforms Data source modeling using graphs Graph-based natural language processing, recommendations, and fraud detection techniques Graph algorithms Working with Neo4J Graph-Powered Machine Learning teaches to use graph-based algorithms and data organization strategies to develop superior machine learning applications. You'll dive into the role of graphs in machine learning and big data platforms, and take an in-depth look at data source modeling, algorithm design, recommendations, and fraud detection. Explore end-to-end projects that illustrate architectures and help you optimize with best design practices. Author Alessandro Negro's extensive experience shines through in every chapter, as you learn from examples and concrete scenarios based on his work with real clients! Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Identifying relationships is the foundation of machine learning. By recognizing and analyzing the connections in your data, graph-centric algorithms like K-nearest neighbor or PageRank radically improve the effectiveness of ML applications. Graph-based machine learning techniques offer a powerful new perspective for machine learning in social networking, fraud detection, natural language processing, and recommendation systems. About the book Graph-Powered Machine Learning teaches you how to exploit the natural relationships in structured and unstructured datasets using graph-oriented machine learning algorithms and tools. In this authoritative book, you'll master the architectures and design practices of graphs, and avoid common pitfalls. Author Alessandro Negro explores examples from real-world applications that connect GraphML concepts to real world tasks. What's inside Graphs in big data platforms Recommendations, natural language processing, fraud detection Graph algorithms Working with the Neo4J graph database About the reader For readers comfortable with machine learning basics. About the author Alessandro Negro is Chief Scientist at GraphAware. He has been a speaker at many conferences, and holds a PhD in Computer Science. Table of Contents PART 1 INTRODUCTION 1 Machine learning and graphs: An introduction 2 Graph data engineering 3 Graphs in machine learning applications PART 2 RECOMMENDATIONS 4 Content-based recommendations 5 Collaborative filtering 6 Session-based recommendations 7 Context-aware and hybrid recommendations PART 3 FIGHTING FRAUD 8 Basic approaches to graph-powered fraud detection 9 Proximity-based algorithms 10 Social network analysis against fraud PART 4 TAMING TEXT WITH GRAPHS 11 Graph-based natural language processing 12 Knowledge graphs Graph Databases in Action introduces you to graph database concepts by comparing them with relational database constructs. You'll learn just enough theory to get started, then progress to hands-on development. Discover use cases involving social networking, recommendation engines, and personalization. Summary Relationships in data often look far more like a web than an orderly set of rows and columns. Graph databases shine when it comes to revealing valuable insights within complex, interconnected data such as demographics, financial records, or computer networks. In Graph Databases in Action, experts Dave Bechberger and Josh Perryman illuminate the design and implementation of graph databases in real-world applications. You'll learn how to choose the right database solutions for your tasks, and how to use your new knowledge to build agile, flexible, and high-performing graph-powered applications! Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Isolated data is a thing of the past! Now, data is connected, and graph databases—like Amazon Neptune, Microsoft Cosmos DB, and Neo4j—are the essential tools of this new reality. Graph databases represent relationships naturally, speeding the discovery of insights and driving business value. About the book Graph Databases in Action introduces you to graph database concepts by comparing them with relational database constructs. You'll learn just enough theory to get started, then progress to

hands-on development. Discover use cases involving social networking, recommendation engines, and personalization. What's inside Graph databases vs. relational databases Systematic graph data modeling Querying and navigating a graph Graph patterns Pitfalls and antipatterns About the reader For software developers. No experience with graph databases required. About the author Dave Bechberger and Josh Perryman have decades of experience building complex data-driven systems and have worked with graph databases since 2014. Table of Contents PART 1 - GETTING STARTED WITH GRAPH DATABASES 1 Introduction to graphs 2 Graph data modeling 3 Running basic and recursive traversals 4 Pathfinding traversals and mutating graphs 5 Formatting results 6 Developing an application PART 2 - BUILDING ON GRAPH DATABASES 7 Advanced data modeling techniques 8 Building traversals using known walks 9 Working with subgraphs PART 3 - MOVING BEYOND THE BASICS 10 Performance, pitfalls, and anti-patterns 11 What's next: Graph analytics, machine learning, and resources Updated with new code, new projects, and new chapters, Machine Learning with TensorFlow, Second Edition gives readers a solid foundation in machine-learning concepts and the TensorFlow library. Summary Updated with new code, new projects, and new chapters, Machine Learning with TensorFlow, Second Edition gives readers a solid foundation in machine-learning concepts and the TensorFlow library. Written by NASA JPL Deputy CTO and Principal Data Scientist Chris Mattmann, all examples are accompanied by downloadable Jupyter Notebooks for a hands-on experience coding TensorFlow with Python. New and revised content expands coverage of core machine learning algorithms, and advancements in neural networks such as VGG-Face facial identification classifiers and deep speech classifiers. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Supercharge your data analysis with machine learning! ML algorithms automatically improve as they process data, so results get better over time. You don't have to be a mathematician to use ML: Tools like Google's TensorFlow library help with complex calculations so you can focus on getting the answers you need. About the book Machine Learning with TensorFlow, Second Edition is a fully revised guide to building machine learning models using Python and TensorFlow. You'll apply core ML concepts to real-world challenges, such as sentiment analysis, text classification, and image recognition. Hands-on examples illustrate neural network techniques for deep speech processing, facial identification, and auto-encoding with CIFAR-10. What's inside Machine Learning with TensorFlow Choosing the best ML approaches Visualizing algorithms with TensorBoard Sharing results with collaborators Running models in Docker About the reader Requires intermediate Python skills and knowledge of general algebraic concepts like vectors and matrices. Examples use the super-stable 1.15.x branch of TensorFlow and TensorFlow 2.x. About the author Chris Mattmann is the Division Manager of the Artificial Intelligence, Analytics, and Innovation Organization at NASA Jet Propulsion Lab. The first edition of this book was written by Nishant Shukla with Kenneth Fricklas. Table of Contents PART 1 - YOUR MACHINE-LEARNING RIG 1 A machine-learning odyssey 2 TensorFlow essentials PART 2 - CORE LEARNING ALGORITHMS 3 Linear regression and beyond 4 Using regression for call-center volume prediction 5 A gentle introduction to classification 6 Sentiment classification: Large movie-review dataset 7 Automatically clustering data 8 Inferring user activity from Android accelerometer data 9 Hidden Markov models 10 Part-of-speech tagging and word-sense disambiguation PART 3 - THE NEURAL NETWORK PARADIGM 11 A peek into autoencoders 12 Applying autoencoders: The CIFAR-10 image dataset 13 Reinforcement learning 14 Convolutional neural networks 15 Building a real-world CNN: VGG-Face ad VGG-Face Lite 16 Recurrent neural networks 17 LSTMs and automatic speech recognition 18 Sequence-to-sequence models for chatbots 19 Utility landscape Learn data science with Python by building five real-world projects! Experiment with card game predictions, tracking disease outbreaks, and more, as you build a flexible and intuitive understanding of data science. In Data Science Bookcamp you will learn: - Techniques for computing and plotting probabilities - Statistical analysis using Scipy - How to organize datasets with clustering algorithms - How to visualize complex multi-variable datasets - How to train a decision tree machine learning algorithm In Data Science Bookcamp you'll test and build your knowledge of Python with the kind of open-ended problems that professional data scientists work on every day. Downloadable data sets and thoroughly-explained solutions help you lock in what you've learned, building your confidence and making you ready for an exciting new data science career. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology A data science project has a lot of moving parts, and it takes practice and skill to get all the code, algorithms, datasets, formats, and visualizations working together harmoniously. This unique book guides you through five realistic projects, including tracking disease outbreaks from news headlines, analyzing social networks, and finding relevant patterns in ad click data. About the book Data Science Bookcamp doesn't stop with surface-level theory and toy examples. As you work through each project, you'll learn how to troubleshoot common problems like missing data, messy data, and algorithms that don't quite fit the model you're building. You'll appreciate the detailed setup instructions and the fully explained solutions that highlight common failure points. In the end, you'll be confident in your skills because you can see the results. What's inside - Web scraping - Organize datasets with clustering algorithms - Visualize complex multi-variable datasets - Train a decision tree machine learning algorithm About the reader For readers who know the basics of Python. No prior data science or machine learning skills required. About the author Leonard Apeltsin is the Head of Data Science at Anomaly, where his team applies advanced analytics to uncover healthcare fraud, waste, and abuse. Table of Contents CASE STUDY 1 FINDING THE WINNING STRATEGY IN A CARD GAME 1 Computing probabilities using Python 2 Plotting probabilities using Matplotlib 3 Running random simulations in NumPy 4 Case study 1 solution CASE STUDY 2 ASSESSING ONLINE AD CLICKS FOR SIGNIFICANCE 5 Basic probability and statistical analysis using SciPy 6 Making predictions using the central limit theorem and SciPy 7 Statistical hypothesis testing 8 Analyzing tables using Pandas 9 Case study 2 solution CASE STUDY 3 TRACKING DISEASE OUTBREAKS USING NEWS HEADLINES 10 Clustering data into groups 11 Geographic location visualization and analysis 12 Case study 3 solution CASE STUDY 4 USING ONLINE JOB POSTINGS TO IMPROVE YOUR DATA SCIENCE RESUME 13 Measuring text similarities 14 Dimension reduction of matrix data 15 NLP analysis of large text datasets 16 Extracting text from web pages 17 Case study 4 solution CASE STUDY 5 PREDICTING FUTURE FRIENDSHIPS FROM SOCIAL NETWORK DATA 18 An introduction to graph theory and network analysis 19 Dynamic graph theory techniques for node ranking and social network analysis 20 Network-driven supervised machine learning 21 Training linear classifiers with logistic regression 22 Training nonlinear classifiers with decision tree techniques 23 Case study 5 solution Distributed across servers, difficult to test, and resistant to modification--modern software is complex. Grokking Simplicity is a friendly, practical guide that will change the way you approach software design and development. It introduces a unique approach to functional programming that explains why certain features of software are prone to complexity, and teaches you the functional techniques you can use to simplify these systems so that they're easier to test and debug. Available in PDF (ePub, kindle, and liveBook formats coming soon). about the technology Even experienced developers struggle with software systems that sprawl across distributed servers and APIs, are filled with redundant code, and are difficult to reliably test and modify. Adopting ways of thinking derived from functional programming can help you design and refactor your codebase in ways that reduce complexity, rather than encouraging it. Grokking Simplicity lays out how to use functional programming in a professional environment to write a codebase that's easier to test and reuse, has fewer bugs, and is better at handling the asynchronous nature of distributed systems. about the book In Grokking Simplicity, you'll learn techniques and, more importantly, a mindset that will help you tackle common problems that arise when software gets complex. Veteran functional programmer Eric Normand guides you to a crystal-clear understanding of why certain features of modern software are so prone to complexity and introduces you to the functional techniques you can use to simplify these systems so that they're easier to read, test, and debug. Through hands-on examples, exercises, and numerous self-assessments, you'll learn to organize your code for maximum reusability and internalize methods to keep unwanted complexity out of your codebase. Regardless of the language you're using, the ways of thinking in this book will help recognize problematic code and tame even the most complex software. what's inside Apply functional programming principles to reduce codebase complexity Work with data transformation pipelines for code that's easier to test and reuse Tools for modeling time to simplify asynchrony 60 exercises and 100 questions to test your knowledge about the reader For experienced programmers. Examples are in JavaScript. about the author Eric Normand has been a functional programmer since 2001 and has been teaching functional programming online and in person since 2007. Visit LispCast.com to see more of his credentials. Think big about your data! PySpark brings the powerful Spark big data processing engine to the Python ecosystem, letting you seamlessly scale up your data tasks and create lightning-fast pipelines. In Data Analysis with Python and PySpark you will learn how to: Manage your data as it scales across multiple machines Scale up your data programs with full confidence Read and write data to and from a variety of sources and formats Deal with messy data with PySpark's data manipulation functionality Discover new data sets and perform exploratory data analysis Build automated data pipelines that transform, summarize, and get insights from data Troubleshoot common PySpark errors Creating reliable long-running jobs Data Analysis with Python and PySpark is your guide to delivering successful Python-driven data projects. Packed with relevant examples and essential techniques, this practical book teaches you to build pipelines for reporting, machine learning, and other data-centric tasks. Quick exercises in every chapter help you practice what you've learned, and rapidly start implementing PySpark into your data systems. No previous knowledge of Spark is required. About the technology The Spark data processing engine is an amazing analytics factory: raw data comes in, insight comes out. PySpark wraps Spark's core engine with a Python-based API. It helps simplify Spark's steep learning curve and makes this powerful tool available to anyone working in the Python data ecosystem. About the book Data Analysis with Python and PySpark helps you solve the daily challenges of data science with PySpark. You'll learn how to scale your processing capabilities across multiple machines while ingesting data from any source—whether that's Hadoop clusters, cloud data storage, or local data files. Once you've covered the fundamentals, you'll explore the full versatility of PySpark by building machine learning pipelines, and blending Python, pandas, and PySpark code. What's inside Organizing your PySpark code Managing your data, no matter the size Scale up your data programs with full confidence Troubleshooting common data pipeline problems Creating reliable long-running jobs About the reader Written for data scientists and data engineers comfortable with Python. About the author As a ML director for a data-driven software company, Jonathan Rioux uses PySpark daily. He teaches the software to data scientists, engineers, and data-savvy business analysts. Table of Contents 1 Introduction PART 1 GET ACQUAINTED: FIRST STEPS IN PYSARK 2 Your first data program in PySpark 3 Submitting and scaling your first PySpark program 4 Analyzing tabular data with pyspark.sql 5 Data frame gymnastics: Joining and grouping PART 2 GET PROFICIENT: TRANSLATE YOUR IDEAS INTO CODE 6 Multidimensional data frames: Using PySpark with JSON data 7 Bilingual PySpark: Blending Python and SQL code 8 Extending PySpark with Python: RDD and UDFs 9 Big data is just a lot of small data: Using pandas UDFs 10 Your data under a different lens: Window functions 11 Faster PySpark: Understanding Spark's query planning PART 3 GET CONFIDENT: USING MACHINE

LEARNING WITH PYSPARK 12 Setting the stage: Preparing features for machine learning 13 Robust machine learning with ML Pipelines 14 Building custom ML transformers and estimators Introduces the programming language's syntax, control flow, and basic data structures and covers its interaction with applications and management of large collections of code. Grokking Machine Learning presents machine learning algorithms and techniques in a way that anyone can understand. This book skips the confused academic jargon and offers clear explanations that require only basic algebra. As you go, you'll build interesting projects with Python, including models for spam detection and image recognition. You'll also pick up practical skills for cleaning and preparing data. 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 Summary Deep Learning and the Game of Go teaches you how to apply the power of deep learning to complex reasoning tasks by building a Go-playing AI. After exposing you to the foundations of machine and deep learning, you'll use Python to build a bot and then teach it the rules of the game. Foreword by Thore Graepel, DeepMind Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology The ancient strategy game of Go is an incredible case study for AI. In 2016, a deep learning-based system shocked the Go world by defeating a world champion. Shortly after that, the upgraded AlphaGo Zero crushed the original bot by using deep reinforcement learning to master the game. Now, you can learn those same deep learning techniques by building your own Go bot! About the Book Deep Learning and the Game of Go introduces deep learning by teaching you to build a Go-winning bot. As you progress, you'll apply increasingly complex training techniques and strategies using the Python deep learning library Keras. You'll enjoy watching your bot master the game of Go, and along the way, you'll discover how to apply your new deep learning skills to a wide range of other scenarios! What's inside Build and teach a self-improving game AI Enhance classical game AI systems with deep learning Implement neural networks for deep learning About the Reader All you need are basic Python skills and high school-level math. No deep learning experience required. About the Author Max Pumperla and Kevin Ferguson are experienced deep learning specialists skilled in distributed systems and data science. Together, Max and Kevin built the open source bot BetaGo. Table of Contents PART 1 - FOUNDATIONS Toward deep learning: a machine-learning introduction Go as a machine-learning problem Implementing your first Go bot PART 2 - MACHINE LEARNING AND GAME AI Playing games with tree search Getting started with neural networks Designing a neural network for Go data Learning from data: a deep-learning bot Deploying bots in the wild Learning by practice: reinforcement learning Reinforcement learning with policy gradients Reinforcement learning with value methods Reinforcement learning with actor-critic methods PART 3 - GREATER THAN THE SUM OF ITS PARTS AlphaGo: Bringing it all together AlphaGo Zero: Integrating tree search with reinforcement learning Time to flex your machine learning muscles! Take on the carefully designed challenges of the Machine Learning Bookcamp and master essential ML techniques through practical application. Summary In Machine Learning Bookcamp you will: Collect and clean data for training models Use popular Python tools, including NumPy, Scikit-Learn, and TensorFlow Apply ML to complex datasets with images Deploy ML models to a production-ready environment The only way to learn is to practice! In Machine Learning Bookcamp, you'll create and deploy Python-based machine learning models for a variety of increasingly challenging projects. Taking you from the basics of machine learning to complex applications such as image analysis, each new project builds on what you've learned in previous chapters. You'll build a portfolio of business-relevant machine learning projects that hiring managers will be excited to see. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Master key machine learning concepts as you build actual projects! Machine learning is what you need for analyzing customer behavior, predicting price trends, evaluating risk, and much more. To master ML, you need great examples, clear explanations, and lots of practice. This book delivers all three! About the book Machine Learning Bookcamp presents realistic, practical machine learning scenarios, along with crystal-clear coverage of key concepts. In it, you'll complete engaging projects, such as creating a car price predictor using linear regression and deploying a churn prediction service. You'll go beyond the algorithms and explore important techniques like deploying ML applications on serverless systems and serving models with Kubernetes and Kubeflow. Dig in, get your hands dirty, and have fun building your ML skills! What's inside Collect and clean data for training models Use popular Python tools, including NumPy, Scikit-Learn, and TensorFlow Deploy ML models to a production-ready environment About the reader Python programming skills assumed. No previous machine learning knowledge is required. About the author Alexey Grigorev is a principal data scientist at OLX Group. He runs DataTalks.Club, a community of people who love data. Table of Contents 1 Introduction to machine learning 2 Machine learning for regression 3 Machine learning for classification 4 Evaluation metrics for classification 5 Deploying machine learning models 6 Decision trees and ensemble learning 7 Neural networks and deep learning 8 Serverless deep learning 9 Serving models with Kubernetes and Kubeflow “We finally have the definitive treatise on PyTorch! It covers the basics and abstractions in great detail. I hope this book becomes your extended reference document.” —Soumith Chintala, co-creator of PyTorch Key Features Written by PyTorch's creator and key contributors Develop deep learning models in a familiar Pythonic way Use PyTorch to build an image classifier for cancer detection Diagnose problems with your neural network and improve training with data augmentation Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About The Book Every other day we hear about new ways to put deep learning to good use: improved medical imaging, accurate credit card fraud detection, long range weather forecasting, and more. PyTorch puts these superpowers in your hands. Instantly familiar to anyone who knows Python data tools like NumPy and Scikit-learn, PyTorch simplifies deep learning without sacrificing advanced features. It's great for building quick models, and it scales smoothly from laptop to enterprise. Deep Learning with PyTorch teaches you to create deep learning and neural network systems with PyTorch. This practical book gets you to work right away building a tumor image classifier from scratch. After covering the basics, you'll learn best practices for the entire deep learning pipeline, tackling advanced projects as your PyTorch skills become more sophisticated. All code samples are easy to explore in downloadable Jupyter notebooks. What You Will Learn Understanding deep learning data structures such as tensors and neural networks Best practices for the PyTorch Tensor API, loading data in Python, and visualizing results Implementing modules and loss functions Utilizing pretrained models from PyTorch Hub Methods for training networks with limited inputs Sifting through unreliable results to diagnose and fix problems in your neural network Improve your results with augmented data, better model architecture, and fine tuning This Book Is Written For For Python programmers with an interest in machine learning. No experience with PyTorch or other deep learning frameworks is required. About The Authors Eli Stevens has worked in Silicon Valley for the past 15 years as a software engineer, and the past 7 years as Chief Technical Officer of a startup making medical device software. Luca Antiga is co-founder and CEO of an AI engineering company located in Bergamo, Italy, and a regular contributor to PyTorch. Thomas Viehmann is a Machine Learning and PyTorch speciality trainer and consultant based in Munich, Germany and a PyTorch core developer. Table of Contents PART 1 - CORE PYTORCH 1 Introducing deep learning and the PyTorch Library 2 Pretrained networks 3 It starts with a tensor 4 Real-world data representation using tensors 5 The mechanics of learning 6 Using a neural network to fit the data 7 Telling birds from airplanes: Learning from images 8 Using convolutions to generalize PART 2 - LEARNING FROM IMAGES IN THE REAL WORLD: EARLY DETECTION OF LUNG CANCER 9 Using PyTorch to fight cancer 10 Combining data sources into a unified dataset 11 Training a classification model to detect suspected tumors 12 Improving training with metrics and augmentation 13 Using segmentation to find suspected nodules 14 End-to-end nodule analysis, and where to go next PART 3 - DEPLOYMENT 15 Deploying to production Summary Deep Learning for Search teaches you how to improve the effectiveness of your search by implementing neural network-based techniques. By the time you're finished with the book, you'll be ready to build amazing search engines that deliver the results your users need and that get better as time goes on! Foreword by Chris Mattmann. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Deep learning handles the toughest search challenges, including imprecise search terms, badly indexed data, and retrieving images with minimal metadata. And with modern tools like DL4J and TensorFlow, you can apply powerful DL techniques without a deep background in data science or natural language processing (NLP). This book will show you how. About the Book Deep Learning for Search teaches you to improve your search results with neural networks. You'll review how DL relates to search basics like indexing and ranking. Then, you'll walk through in-depth examples to upgrade your search with DL techniques using Apache Lucene and Deeplearning4j. As the book progresses, you'll explore advanced topics like searching through images, translating user queries, and designing search engines that improve as they learn! What's inside Accurate and relevant rankings Searching across languages Content-based image search Search with

recommendations About the Reader For developers comfortable with Java or a similar language and search basics. No experience with deep learning or NLP needed.

About the Author Tommaso Teofili is a software engineer with a passion for open source and machine learning. As a member of the Apache Software Foundation, he contributes to a number of open source projects, ranging from topics like information retrieval (such as Lucene and Solr) to natural language processing and machine translation (including OpenNLP, Joshua, and UIMA). He currently works at Adobe, developing search and indexing infrastructure components, and researching the areas of natural language processing, information retrieval, and deep learning. He has presented search and machine learning talks at conferences including BerlinBuzzwords, International Conference on Computational Science, ApacheCon, EclipseCon, and others. You can find him on Twitter at @tteofili.

Table of Contents PART 1 - SEARCH MEETS DEEP LEARNING Neural search Generating synonyms PART 2 - THROWING NEURAL NETS AT A SEARCH ENGINE From plain retrieval to text generation More-sensitive query suggestions Ranking search results with word embeddings Document embeddings for rankings and recommendations PART 3 - ONE STEP BEYOND Searching across languages Content-based image search A peek at performance Summary Machine Learning in Action is unique book that blends the foundational theories of machine learning with the practical realities of building tools for everyday data analysis. You'll use the flexible Python programming language to build programs that implement algorithms for data classification, forecasting, recommendations, and higher-level features like summarization and simplification. About the Book A machine is said to learn when its performance improves with experience. Learning requires algorithms and programs that capture data and ferret out the interesting or useful patterns. Once the specialized domain of analysts and mathematicians, machine learning is becoming a skill needed by many. Machine Learning in Action is a clearly written tutorial for developers. It avoids academic language and takes you straight to the techniques you'll use in your day-to-day work. Many (Python) examples present the core algorithms of statistical data processing, data analysis, and data visualization in code you can reuse. You'll understand the concepts and how they fit in with tactical tasks like classification, forecasting, recommendations, and higher-level features like summarization and simplification. Readers need no prior experience with machine learning or statistical processing. Familiarity with Python is helpful. 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. What's Inside A no-nonsense introduction Examples showing common ML tasks Everyday data analysis Implementing classic algorithms like Apriori and Adaboos Table of Contents PART 1 CLASSIFICATION Machine learning basics Classifying with k-Nearest Neighbors Splitting datasets one feature at a time: decision trees Classifying with probability theory: naïve Bayes Logistic regression Support vector machines Improving classification with the AdaBoost meta algorithm PART 2 FORECASTING NUMERIC VALUES WITH REGRESSION Predicting numeric values: regression Tree-based regression PART 3 UNSUPERVISED LEARNING Grouping unlabeled items using k-means clustering Association analysis with the Apriori algorithm Efficiently finding frequent itemsets with FP-growth PART 4 ADDITIONAL TOOLS Using principal component analysis to simplify data Simplifying data with the singular value decomposition Big data and MapReduce This book teaches you how to build and maintain effective data pipelines. You'll explore the most common usage patterns, including aggregating multiple data sources, connecting to and from data lakes, and cloud deployment. -- Summary HTTP/2 in Action is a complete guide to HTTP/2, one of the core protocols of the web. Because HTTP/2 has been designed to be easy to transition to, including keeping it backwards compatible, adoption is rapid and expected to increase over the next few years. Concentrating on practical matters, this interesting book presents key HTTP/2 concepts such as frames, streams, and multiplexing and explores how they affect the performance and behavior of your websites. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology HTTP—Hypertext Transfer Protocol—is the standard for exchanging messages between websites and browsers. And after 20 years, it's gotten a much-needed upgrade. With support for streams, server push, header compression, and prioritization, HTTP/2 delivers vast improvements in speed, security, and efficiency. About the Book HTTP/2 in Action teaches you everything you need to know to use HTTP/2 effectively. You'll learn how to optimize web performance with new features like frames, multiplexing, and push. You'll also explore real-world examples on advanced topics like flow control and dependencies. With ready-to-implement tips and best practices, this practical guide is sure to get you—and your websites—up to speed! What's Inside HTTP/2 for web developers Upgrading and troubleshooting Real-world examples and case studies QUIC and HTTP/3 About the Reader Written for web developers and site administrators. About the Authors Barry Pollard is a professional developer with two decades of experience developing, supporting, and tuning software and infrastructure. Table of Contents PART 1 MOVING TO HTTP/2 Web technologies and HTTP The road to HTTP/2 Upgrading to HTTP/2 PART 2 USING HTTP/2 HTTP/2 protocol basics Implementing HTTP/2 push Optimizing for HTTP/2 PART 3 ADVANCED HTTP/2 Advanced HTTP/2 concepts HPACK header compression PART 4 THE FUTURE OF HTTP TCP, QUIC, and HTTP/3 Where HTTP goes from here Summary Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Machine learning has made remarkable progress in recent years. We went from near-unusable speech and image recognition, to near-human accuracy. We went from machines that couldn't beat a serious Go player, to defeating a world champion. Behind this progress is deep learning—a combination of engineering advances, best practices, and theory that enables a wealth of previously impossible smart applications. About the Book Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. You'll explore challenging concepts and practice with applications in computer vision, natural-language processing, and generative models. By the time you finish, you'll have the knowledge and hands-on skills to apply deep learning in your own projects. What's Inside Deep learning from first principles Setting up your own deep-learning environment Image-classification models Deep learning for text and sequences Neural style transfer, text generation, and image generation About the Reader Readers need intermediate Python skills. No previous experience with Keras, TensorFlow, or machine learning is required. About the Author François Chollet works on deep learning at Google in Mountain View, CA. He is the creator of the Keras deep-learning library, as well as a contributor to the TensorFlow machine-learning framework. He also does deep-learning research, with a focus on computer vision and the application of machine learning to formal reasoning. His papers have been published at major conferences in the field, including the Conference on Computer Vision and Pattern Recognition (CVPR), the Conference and Workshop on Neural Information Processing Systems (NIPS), the International Conference on Learning Representations (ICLR), and others. Table of Contents PART 1 - FUNDAMENTALS OF DEEP LEARNING What is deep learning? Before we begin: the mathematical building blocks of neural networks Getting started with neural networks Fundamentals of machine learning PART 2 - DEEP LEARNING IN PRACTICE Deep learning for computer vision Deep learning for text and sequences Advanced deep-learning best practices Generative deep learning Conclusions appendix A - Installing Keras and its dependencies on Ubuntu appendix B - Running Jupyter notebooks on an EC2 GPU instance Summary Machine Learning with TensorFlow gives readers a solid foundation in machine-learning concepts plus hands-on experience coding TensorFlow with Python. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology TensorFlow, Google's library for large-scale machine learning, simplifies often-complex computations by representing them as graphs and efficiently mapping parts of the graphs to machines in a cluster or to the processors of a single machine. About the Book Machine Learning with TensorFlow gives readers a solid foundation in machine-learning concepts plus hands-on experience coding TensorFlow with Python. You'll learn the basics by working with classic prediction, classification, and clustering algorithms. Then, you'll move on to the money chapters: exploration of deep-learning concepts like autoencoders, recurrent neural networks, and reinforcement learning. Digest this book and you will be ready to use TensorFlow for machine-learning and deep-learning applications of your own. What's Inside Matching your tasks to the right machine-learning and deep-learning approaches Visualizing algorithms with TensorBoard Understanding and using neural networks About the Reader Written for developers experienced with Python and algebraic concepts like vectors and matrices. About the Author Author Nishant Shukla is a computer vision researcher focused on applying machine-learning techniques in robotics. Senior technical editor, Kenneth Fricklas, is a seasoned developer, author, and machine-learning practitioner. Table of Contents PART 1 - YOUR MACHINE-LEARNING RIG A machine-learning odyssey TensorFlow essentials PART 2 - CORE LEARNING ALGORITHMS Linear regression and beyond A gentle introduction to classification Automatically clustering data Hidden Markov models PART 3 - THE NEURAL NETWORK PARADIGM A peek into autoencoders Reinforcement learning Convolutional neural networks Recurrent neural networks Sequence-to-sequence models for chatbots Utility landscape Summary Machine Learning in Action is unique book that blends the foundational theories of machine learning with the practical realities of building tools for everyday data analysis. You'll use the flexible Python programming language to build programs that implement algorithms for data classification, forecasting, recommendations, and higher-level features like summarization and simplification. About the Book A machine is said to learn when its performance improves with experience. Learning requires algorithms and programs that capture data and ferret out the interesting or useful patterns. Once the specialized domain of analysts and mathematicians, machine learning is becoming a skill needed by many. Machine Learning in Action is a clearly written tutorial for developers. It avoids academic language and takes you straight to the techniques you'll use in your day-to-day work. Many (Python) examples present the core algorithms of statistical data processing, data analysis, and data visualization in code you can reuse. You'll understand the concepts and how they fit in with tactical tasks like classification, forecasting, recommendations, and higher-level features like summarization and simplification. Readers need no prior experience with machine learning or statistical processing. Familiarity with Python is helpful. 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. What's Inside A no-nonsense introduction Examples showing common ML tasks Everyday data analysis Implementing classic algorithms like Apriori and Adaboos Table of Contents PART 1 CLASSIFICATION Machine learning basics Classifying with k-Nearest Neighbors Splitting datasets one feature at a time: decision trees Classifying with probability theory: naïve Bayes Logistic regression Support vector machines Improving classification with the AdaBoost meta algorithm PART 2 FORECASTING NUMERIC VALUES WITH REGRESSION Predicting numeric values: regression Tree-based regression PART 3 UNSUPERVISED LEARNING Grouping unlabeled items using k-means clustering Association analysis with the Apriori algorithm Efficiently finding frequent itemsets with FP-growth PART 4 ADDITIONAL TOOLS Using principal component analysis to simplify data Simplifying data with the singular value decomposition Big data and MapReduce Summary Deep Learning with R introduces the world of deep learning using the

powerful Keras library and its R language interface. The book builds your understanding of deep learning through intuitive explanations and practical examples. Continue your journey into the world of deep learning with Deep Learning with R in Motion, a practical, hands-on video course available exclusively at Manning.com (www.manning.com/livevideo/deep-learning-with-r-in-motion). Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Machine learning has made remarkable progress in recent years. Deep-learning systems now enable previously impossible smart applications, revolutionizing image recognition and natural-language processing, and identifying complex patterns in data. The Keras deep-learning library provides data scientists and developers working in R a state-of-the-art toolset for tackling deep-learning tasks. About the Book Deep Learning with R introduces the world of deep learning using the powerful Keras library and its R language interface. Initially written for Python as Deep Learning with Python by Keras creator and Google AI researcher François Chollet and adapted for R by RStudio founder J. J. Allaire, this book builds your understanding of deep learning through intuitive explanations and practical examples. You'll practice your new skills with R-based applications in computer vision, natural-language processing, and generative models. What's Inside Deep learning from first principles Setting up your own deep-learning environment Image classification and generation Deep learning for text and sequences About the Reader You'll need intermediate R programming skills. No previous experience with machine learning or deep learning is assumed. About the Authors François Chollet is a deep-learning researcher at Google and the author of the Keras library. J.J. Allaire is the founder of RStudio and the author of the R interfaces to TensorFlow and Keras. Table of Contents PART 1 - FUNDAMENTALS OF DEEP LEARNING What is deep learning? Before we begin: the mathematical building blocks of neural networks Getting started with neural networks Fundamentals of machine learning PART 2 - DEEP LEARNING IN PRACTICE Deep learning for computer vision Deep learning for text and sequences Advanced deep-learning best practices Generative deep learning Conclusions Summary This third revision of Manning's popular The Quick Python Book offers a clear, crisp updated introduction to the elegant Python programming language and its famously easy-to-read syntax. Written for programmers new to Python, this latest edition includes new exercises throughout. It covers features common to other languages concisely, while introducing Python's comprehensive standard functions library and unique features in detail. Foreword by Nicholas Tollervey, Python Software Foundation. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Initially Guido van Rossum's 1989 holiday project, Python has grown into an amazing computer language. It's a joy to learn and read, and powerful enough to handle everything from low-level system resources to advanced applications like deep learning. Elegantly simple and complete, it also boasts a massive ecosystem of libraries and frameworks. Python programmers are in high demand; you can't afford not to be fluent! About the Book The Quick Python Book, Third Edition is a comprehensive guide to the Python language by a Python authority, Naomi Ceder. With the personal touch of a skilled teacher, she beautifully balances details of the language with the insights and advice you need to handle any task. Extensive, relevant examples and learn-by-doing exercises help you master each important concept the first time through. Whether you're scraping websites or playing around with nested tuples, you'll appreciate this book's clarity, focus, and attention to detail. What's Inside Clear coverage of Python 3 Core libraries, packages, and tools In-depth exercises Five new data science-related chapters About the Reader Written for readers familiar with programming concepts--no Python experience assumed. About the Author Naomi Ceder is chair of the Python Software Foundation. She has been learning, using, and teaching Python since 2001. Table of Contents PART 1 - STARTING OUT 1. About Python 2. Getting started 3. The Quick Python overview PART 2 - THE ESSENTIALS 4. The absolute basics 5. Lists, tuples, and sets 6. Strings 7. Dictionaries 8. Control flow 9. Functions 10. Modules and scoping rules 11. Python programs 12. Using the filesystem 13. Reading and writing files 14. Exceptions PART 3 - ADVANCED LANGUAGE FEATURES 15. Classes and object-oriented programming 16. Regular expressions 17. Data types as objects 18. Packages 19. Using Python libraries PART 4 - WORKING WITH DATA 20. Basic file wrangling 21. Processing data files 22. Data over the network 23. Saving data 24. Exploring data Deliver huge improvements to your machine learning pipelines without spending hours fine-tuning parameters! This book's practical case-studies reveal feature engineering techniques that upgrade your data wrangling—and your ML results. In Feature Engineering Bookcamp you will learn how to: Identify and implement feature transformations for your data Build powerful machine learning pipelines with unstructured data like text and images Quantify and minimize bias in machine learning pipelines at the data level Use feature stores to build real-time feature engineering pipelines Enhance existing machine learning pipelines by manipulating the input data Use state-of-the-art deep learning models to extract hidden patterns in data Feature Engineering Bookcamp guides you through a collection of projects that give you hands-on practice with core feature engineering techniques. You'll work with feature engineering practices that speed up the time it takes to process data and deliver real improvements in your model's performance. This instantly-useful book skips the abstract mathematical theory and minutely-detailed formulas; instead you'll learn through interesting code-driven case studies, including tweet classification, COVID detection, recidivism prediction, stock price movement detection, and more. About the technology Get better output from machine learning pipelines by improving your training data! Use feature engineering, a machine learning technique for designing relevant input variables based on your existing data, to simplify training and enhance model performance. While fine-tuning hyperparameters or tweaking models may give you a minor performance bump, feature engineering delivers dramatic improvements by transforming your data pipeline. About the book Feature Engineering Bookcamp walks you through six hands-on projects where you'll learn to upgrade your training data using feature engineering. Each chapter explores a new code-driven case study, taken from real-world industries like finance and healthcare. You'll practice cleaning and transforming data, mitigating bias, and more. The book is full of performance-enhancing tips for all major ML subdomains—from natural language processing to time-series analysis. What's inside Identify and implement feature transformations Build machine learning pipelines with unstructured data Quantify and minimize bias in ML pipelines Use feature stores to build real-time feature engineering pipelines Enhance existing pipelines by manipulating input data About the reader For experienced machine learning engineers familiar with Python. About the author Sinan Ozdemir is the founder and CTO of Shiba, a former lecturer of Data Science at Johns Hopkins University, and the author of multiple textbooks on data science and machine learning. Table of Contents 1 Introduction to feature engineering 2 The basics of feature engineering 3 Healthcare: Diagnosing COVID-19 4 Bias and fairness: Modeling recidivism 5 Natural language processing: Classifying social media sentiment 6 Computer vision: Object recognition 7 Time series analysis: Day trading with machine learning 8 Feature stores 9 Putting it all together Human-in-the-Loop Machine Learning lays out methods for humans and machines to work together effectively. Summary Most machine learning systems that are deployed in the world today learn from human feedback. However, most machine learning courses focus almost exclusively on the algorithms, not the human-computer interaction part of the systems. This can leave a big knowledge gap for data scientists working in real-world machine learning, where data scientists spend more time on data management than on building algorithms. Human-in-the-Loop Machine Learning is a practical guide to optimizing the entire machine learning process, including techniques for annotation, active learning, transfer learning, and using machine learning to optimize every step of the process. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Machine learning applications perform better with human feedback. Keeping the right people in the loop improves the accuracy of models, reduces errors in data, lowers costs, and helps you ship models faster. About the book Human-in-the-Loop Machine Learning lays out methods for humans and machines to work together effectively. You'll find best practices on selecting sample data for human feedback, quality control for human annotations, and designing annotation interfaces. You'll learn to create training data for labeling, object detection, and semantic segmentation, sequence labeling, and more. The book starts with the basics and progresses to advanced techniques like transfer learning and self-supervision within annotation workflows. What's inside Identifying the right training and evaluation data Finding and managing people to annotate data Selecting annotation quality control strategies Designing interfaces to improve accuracy and efficiency About the author Robert (Munro) Monarch is a data scientist and engineer who has built machine learning data for companies such as Apple, Amazon, Google, and IBM. He holds a PhD from Stanford. Robert holds a PhD from Stanford focused on Human-in-the-Loop machine learning for healthcare and disaster response, and is a disaster response professional in addition to being a machine learning professional. A worked example throughout this text is classifying disaster-related messages from real disasters that Robert has helped respond to in the past. Table of Contents PART 1 - FIRST STEPS 1 Introduction to human-in-the-loop machine learning 2 Getting started with human-in-the-loop machine learning PART 2 - ACTIVE LEARNING 3 Uncertainty sampling 4 Diversity sampling 5 Advanced active learning 6 Applying active learning to different machine learning tasks PART 3 - ANNOTATION 7 Working with the people annotating your data 8 Quality control for data annotation 9 Advanced data annotation and augmentation 10 Annotation quality for different machine learning tasks PART 4 - HUMAN-COMPUTER INTERACTION FOR MACHINE LEARNING 11 Interfaces for data annotation 12 Human-in-the-loop machine learning products Class-tested and coherent, this textbook teaches classical and web information retrieval, including web search and the related areas of text classification and text clustering from basic concepts. It gives an up-to-date treatment of all aspects of the design and implementation of systems for gathering, indexing, and searching documents; methods for evaluating systems; and an introduction to the use of machine learning methods on text collections. All the important ideas are explained using examples and figures, making it perfect for introductory courses in information retrieval for advanced undergraduates and graduate students in computer science. Based on feedback from extensive classroom experience, the book has been carefully structured in order to make teaching more natural and effective. Slides and additional exercises (with solutions for lecturers) are also available through the book's supporting website to help course instructors prepare their lectures. Summary Humans learn best from feedback—we are encouraged to take actions that lead to positive results while deterred by decisions with negative consequences. This reinforcement process can be applied to computer programs allowing them to solve more complex problems that classical programming cannot. Deep Reinforcement Learning in Action teaches you the fundamental concepts and terminology of deep reinforcement learning, along with the practical skills and techniques you'll need to implement it into your own projects. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Deep reinforcement learning AI systems rapidly adapt to new environments, a vast improvement over standard neural networks. A DRL agent learns like people do, taking in raw data such as sensor input and refining its responses and predictions through trial and error. About the book Deep Reinforcement Learning in Action teaches you how to program AI agents that adapt and improve based on direct feedback from their environment. In this example-rich tutorial, you'll master foundational and advanced DRL techniques by taking on interesting challenges like navigating a maze and playing video games. Along the way, you'll work with core algorithms, including deep Q-networks and policy gradients, along with industry-standard tools like

PyTorch and OpenAI Gym. What's inside Building and training DRL networks The most popular DRL algorithms for learning and problem solving Evolutionary algorithms for curiosity and multi-agent learning All examples available as Jupyter Notebooks About the reader For readers with intermediate skills in Python and deep learning. About the author Alexander Zai is a machine learning engineer at Amazon AI. Brandon Brown is a machine learning and data analysis blogger. Table of Contents PART 1 - FOUNDATIONS 1. What is reinforcement learning? 2. Modeling reinforcement learning problems: Markov decision processes 3. Predicting the best states and actions: Deep Q-networks 4. Learning to pick the best policy: Policy gradient methods 5. Tackling more complex problems with actor-critic methods PART 2 - ABOVE AND BEYOND 6. Alternative optimization methods: Evolutionary algorithms 7. Distributional DQN: Getting the full story 8. Curiosity-driven exploration 9. Multi-agent reinforcement learning 10. Interpretable reinforcement learning: Attention and relational models 11. In conclusion: A review and roadmap This book introduces a collection of algorithms for complex programming challenges in data analysis, machine learning, and graph computing. You'll discover cutting-edge approaches to a variety of tricky scenarios. -- Summary The Spark distributed data processing platform provides an easy-to-implement tool for ingesting, streaming, and processing data from any source. In Spark in Action, Second Edition, you'll learn to take advantage of Spark's core features and incredible processing speed, with applications including real-time computation, delayed evaluation, and machine learning. Spark skills are a hot commodity in enterprises worldwide, and with Spark's powerful and flexible Java APIs, you can reap all the benefits without first learning Scala or Hadoop. Foreword by Rob Thomas. About the technology Analyzing enterprise data starts by reading, filtering, and merging files and streams from many sources. The Spark data processing engine handles this varied volume like a champ, delivering speeds 100 times faster than Hadoop systems. Thanks to SQL support, an intuitive interface, and a straightforward multilanguage API, you can use Spark without learning a complex new ecosystem. About the book Spark in Action, Second Edition, teaches you to create end-to-end analytics applications. In this entirely new book, you'll learn from interesting Java-based examples, including a complete data pipeline for processing NASA satellite data. And you'll discover Java, Python, and Scala code samples hosted on GitHub that you can explore and adapt, plus appendixes that give you a cheat sheet for installing tools and understanding Spark-specific terms. What's inside Writing Spark applications in Java Spark application architecture Ingestion through files, databases, streaming, and Elasticsearch Querying distributed datasets with Spark SQL About the reader This book does not assume previous experience with Spark, Scala, or Hadoop. About the author Jean-Georges Perrin is an experienced data and software architect. He is France's first IBM Champion and has been honored for 12 consecutive years. Table of Contents PART 1 - THE THEORY CRIPPLED BY AWESOME EXAMPLES 1 So, what is Spark, anyway? 2 Architecture and flow 3 The majestic role of the dataframe 4 Fundamentally lazy 5 Building a simple app for deployment 6 Deploying your simple app PART 2 - INGESTION 7 Ingestion from files 8 Ingestion from databases 9 Advanced ingestion: finding data sources and building your own 10 Ingestion through structured streaming PART 3 - TRANSFORMING YOUR DATA 11 Working with SQL 12 Transforming your data 13 Transforming entire documents 14 Extending transformations with user-defined functions 15 Aggregating your data PART 4 - GOING FURTHER 16 Cache and checkpoint: Enhancing Spark's performances 17 Exporting data and building full data pipelines 18 Exploring deployment 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 Summary Grokking Deep Learning teaches you to build deep learning neural networks from scratch! In his engaging style, seasoned deep learning expert Andrew Trask shows you the science under the hood, so you grok for yourself every detail of training neural networks. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Deep learning, a branch of artificial intelligence, teaches computers to learn by using neural networks, technology inspired by the human brain. Online text translation, self-driving cars, personalized product recommendations, and virtual voice assistants are just a few of the exciting modern advancements possible thanks to deep learning. About the Book Grokking Deep Learning teaches you to build deep learning neural networks from scratch! In his engaging style, seasoned deep learning expert Andrew Trask shows you the science under the hood, so you grok for yourself every detail of training neural networks. Using only Python and its math-supporting library, NumPy, you'll train your own neural networks to see and understand images, translate text into different languages, and even write like Shakespeare! When you're done, you'll be fully prepared to move on to mastering deep learning frameworks. What's inside The science behind deep learning Building and training your own neural networks Privacy concepts, including federated learning Tips for continuing your pursuit of deep learning About the Reader For readers with high school-level math and intermediate programming skills. About the Author Andrew Trask is a PhD student at Oxford University and a research scientist at DeepMind. Previously, Andrew was a researcher and analytics product manager at Digital Reasoning, where he trained the world's largest artificial neural network and helped guide the analytics roadmap for the Synthesys cognitive computing platform. Table of Contents Introducing deep learning: why you should learn it Fundamental concepts: how do machines learn? Introduction to neural prediction: forward propagation Introduction to neural learning: gradient descent Learning multiple weights at a time: generalizing gradient descent Building your first deep neural network: introduction to backpropagation How to picture neural networks: in your head and on paper Learning signal and ignoring noise: introduction to regularization and batching Modeling probabilities and nonlinearities: activation functions Neural learning about edges and corners: intro to convolutional neural networks Neural networks that understand language: king - man + woman == ? Neural networks that write like Shakespeare: recurrent layers for variable-length data Introducing automatic optimization: let's build a deep learning framework Learning to write like Shakespeare: long short-term memory Deep learning on unseen data: introducing federated learning Where to go from here: a brief guide Summary Hadoop in Practice, Second Edition provides over 100 tested, instantly useful techniques that will help you conquer big data, using Hadoop. This revised new edition covers changes and new features in the Hadoop core architecture, including MapReduce 2. Brand new chapters cover YARN and integrating Kafka, Impala, and Spark SQL with Hadoop. You'll also get new and updated techniques for Flume, Sqoop, and Mahout, all of which have seen major new versions recently. In short, this is the most practical, up-to-date coverage of Hadoop available anywhere. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Book It's always a good time to upgrade your Hadoop skills! Hadoop in Practice, Second Edition provides a collection of 104 tested, instantly useful techniques for analyzing real-time streams, moving data securely, machine learning, managing large-scale clusters, and taming big data using Hadoop. This completely revised edition covers changes and new features in Hadoop core, including MapReduce 2 and YARN. You'll pick up hands-on best practices for integrating Spark, Kafka, and Impala with Hadoop, and get new and updated techniques for the latest versions of Flume, Sqoop, and Mahout. In short, this is the most practical, up-to-date coverage of Hadoop available. Readers need to know a programming language like Java and have basic familiarity with Hadoop. What's Inside Thoroughly updated for Hadoop 2 How to write YARN applications Integrate real-time technologies like Storm, Impala, and Spark Predictive analytics using Mahout and RR Readers need to know a programming language like Java and have basic familiarity with Hadoop. About the Author Alex Holmes works on tough big-data problems. He is a software engineer, author, speaker, and blogger specializing in large-scale Hadoop projects. Table of Contents PART 1 BACKGROUND AND FUNDAMENTALS Hadoop in a heartbeat Introduction to YARN PART 2 DATA LOGISTICS Data serialization—working with text and beyond Organizing and optimizing data in HDFS Moving data into and out of Hadoop PART 3 BIG DATA PATTERNS Applying MapReduce patterns to big data Utilizing data structures and algorithms at scale Tuning, debugging, and testing PART 4 BEYOND MAPREDUCE SQL on Hadoop Writing a YARN application Grokking Deep Reinforcement Learning uses engaging exercises to teach you how to build deep learning systems. This book combines annotated Python code with intuitive explanations to explore DRL techniques. You'll see how algorithms function and learn to develop your own DRL agents using evaluative feedback. Summary We all learn through trial and error. We avoid the things that cause us to experience pain and failure. We embrace and build on the things that give us reward and success. This common pattern is the foundation of deep reinforcement learning: building machine learning systems that explore and learn based on the responses of the environment. Grokking Deep Reinforcement Learning introduces this powerful machine learning approach, using examples, illustrations, exercises, and crystal-clear teaching. You'll love the perfectly paced teaching and the clever, engaging writing style as you dig into this awesome exploration of reinforcement learning fundamentals, effective deep learning techniques, and practical applications in this emerging field. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology We learn by interacting with our environment, and the rewards or punishments we experience guide our future behavior. Deep reinforcement learning brings that same natural process to artificial intelligence, analyzing results to uncover the most efficient ways forward. DRL agents can improve marketing campaigns, predict stock performance, and beat grand masters in Go and chess. About the book Grokking Deep Reinforcement Learning uses engaging exercises to teach you how to build deep learning systems. This book combines annotated Python code with intuitive explanations to explore DRL techniques. You'll see how algorithms function and learn to develop your own DRL agents using evaluative feedback. What's inside An introduction to reinforcement learning DRL agents with human-like behaviors Applying DRL

to complex situations About the reader For developers with basic deep learning experience. About the author Miguel Morales works on reinforcement learning at Lockheed Martin and is an instructor for the Georgia Institute of Technology's Reinforcement Learning and Decision Making course. Table of Contents 1 Introduction to deep reinforcement learning 2 Mathematical foundations of reinforcement learning 3 Balancing immediate and long-term goals 4 Balancing the gathering and use of information 5 Evaluating agents' behaviors 6 Improving agents' behaviors 7 Achieving goals more effectively and efficiently 8 Introduction to value-based deep reinforcement learning 9 More stable value-based methods 10 Sample-efficient value-based methods 11 Policy-gradient and actor-critic methods 12 Advanced actor-critic methods 13 Toward artificial general intelligence

- [Introduction To Information Retrieval](#)
- [Introducing Data Science](#)
- [Data Science Bookcamp](#)
- [Deep Learning With Python](#)
- [Machine Learning In Action](#)
- [Feature Engineering Bookcamp](#)
- [Deep Learning And The Game Of Go](#)
- [Build A Career In Data Science](#)
- [Math For Programmers](#)
- [Spark In Action](#)
- [Grokking Machine Learning](#)
- [Machine Learning In Action](#)
- [Deep Learning For Search](#)
- [Practices Of The Python Pro](#)
- [Grokking Deep Learning](#)
- [Speech Language Processing](#)
- [Deep Learning With R](#)
- [Machine Learning Bookcamp](#)
- [Graph Powered Machine Learning](#)
- [Data Analysis With Python And PySpark](#)
- [Real World Machine Learning](#)
- [Grokking Simplicity](#)
- [The Quick Python Book](#)
- [Ensemble Methods For Machine Learning](#)
- [MLOps Engineering At Scale](#)
- [Graph Databases In Action](#)
- [Machine Learning With TensorFlow Second Edition](#)
- [Machine Learning With TensorFlow](#)
- [Deep Learning With PyTorch](#)
- [Advanced Algorithms And Data Structures](#)
- [Machine Learning Systems](#)
- [Data Pipelines With Apache Airflow](#)
- [Human in the Loop Machine Learning](#)
- [The Quick Python Book](#)
- [Hadoop In Practice](#)
- [Grokking Deep Reinforcement Learning](#)
- [Deep Reinforcement Learning In Action](#)
- [Grokking Artificial Intelligence Algorithms](#)
- [HTTP 2 In Action](#)
- [Natural Language Processing In Action](#)