

Advanced Analytics With Spark Patterns For Learning From Data At Scale

Naked Statistics: Stripping the Dread from the Data
 R for Everyone
 Real-Time Big Data Analytics
 Hands-On Big Data Analytics with PySpark
 Spark: The Definitive Guide
 The Data Warehouse Toolkit
 Advanced Analytics with Pyspark
 Learning Spark
 Advanced Analytics with Spark
 Advanced Analytics with Spark
 Learning Spark
 Data Algorithms
 Graph Algorithms
 Big Data Analytics with Java
 Practical Data Science with Hadoop and Spark
 Mastering Spark with R
 Advanced Analytics with PySpark
 Scala for Data Science
 Machine Learning with Apache Spark Quick Start Guide
 Data Analysis with Python and PySpark
 Learning PySpark
 Data Analytics with Spark Using Python
 Mastering Machine Learning with Spark 2.x
 Mastering Spark for Data Science
 Stream Processing with Apache Spark
 Advanced Data Analytics Using Python
 Data Algorithms with Spark
 Modern Data Engineering with Apache Spark
 Big Data Analytics with Spark
 Practical Big Data Analytics
 Scala and Spark for Big Data Analytics
 Big Data Analytics with R
 Advanced Analytics with Spark
 High Performance Spark
 Big Data Analytics Beyond Hadoop
 Advanced Analytics with Spark
 Data Engineering with Apache Spark, Delta Lake, and Lakehouse
 Data Analytics with Hadoop
 Big Data Processing with Apache Spark
 Big Data Analytics

*Advanced Analytics With Spark
 Patterns For Learning From Data At
 Scale* Downloaded from archive.imba.com by
 guest

GUERRA AMIR

Naked Statistics: Stripping the Dread from the Data "O'Reilly Media, Inc."

Apache Spark's speed, ease of use, sophisticated analytics, and multilanguage support makes practical knowledge of this cluster-computing framework a required skill for data engineers and data scientists. With this hands-on guide, anyone looking for an introduction to Spark will learn practical algorithms and examples using PySpark. In each chapter, author Mahmoud Parsian shows you how to solve a data problem with a set of Spark transformations and algorithms. You'll learn how to tackle problems involving ETL, design patterns, machine learning algorithms, data partitioning, and genomics analysis. Each detailed recipe includes PySpark algorithms using the PySpark driver and shell script. With this book, you will: Learn how to select Spark transformations for optimized solutions Explore powerful transformations and reductions including `reduceByKey()`, `combineByKey()`, and `mapPartitions()` Understand data partitioning for optimized queries Build and apply a model using PySpark design patterns Apply motif-finding algorithms to graph data Analyze graph data by using the GraphFrames API Apply PySpark algorithms to clinical and genomics data Learn how to use and apply feature engineering in ML algorithms Understand and use practical and pragmatic data design patterns

R for Everyone Addison-Wesley Professional

In the second edition of this practical book, four Cloudera data scientists present a set of self-contained patterns for performing large-scale data analysis with Spark. The authors bring Spark, statistical methods, and real-world data sets together to teach you how to approach analytics problems by example. Updated for Spark 2.1, this edition acts as an introduction to these techniques and other best practices in Spark programming. You'll start with an introduction to Spark and its ecosystem, and then dive into patterns that apply common techniques—including classification, clustering, collaborative filtering, and anomaly detection—to fields such as genomics, security, and finance. If you have an entry-level understanding of machine learning and statistics, and you program in Java, Python, or Scala, you'll find the book's patterns useful for working on your own data applications. With this book, you will: Familiarize yourself with the Spark programming model Become comfortable within the Spark ecosystem Learn general approaches in data science Examine complete implementations that analyze large public data sets Discover which machine learning tools make sense for particular

problems Acquire code that can be adapted to many uses *Real-Time Big Data Analytics* "O'Reilly Media, Inc."

Master the techniques and sophisticated analytics used to construct Spark-based solutions that scale to deliver production-grade data science products About This Book Develop and apply advanced analytical techniques with Spark Learn how to tell a compelling story with data science using Spark's ecosystem Explore data at scale and work with cutting edge data science methods Who This Book Is For This book is for those who have beginner-level familiarity with the Spark architecture and data science applications, especially those who are looking for a challenge and want to learn cutting edge techniques. This book assumes working knowledge of data science, common machine learning methods, and popular data science tools, and assumes you have previously run proof of concept studies and built prototypes. What You Will Learn Learn the design patterns that integrate Spark into industrialized data science pipelines See how commercial data scientists design scalable code and reusable code for data science services Explore cutting edge data science methods so that you can study trends and causality Discover advanced programming techniques using RDD and the DataFrame and Dataset APIs Find out how Spark can be used as a universal ingestion engine tool and as a web scraper Practice the implementation of advanced topics in graph processing, such as community detection and contact chaining Get to know the best practices when performing Extended Exploratory Data Analysis, commonly used in commercial data science teams Study advanced Spark concepts, solution design patterns, and integration architectures Demonstrate powerful data science pipelines In Detail Data science seeks to transform the world using data, and this is typically achieved through disrupting and changing real processes in real industries. In order to operate at this level you need to build data science solutions of substance -solutions that solve real problems. Spark has emerged as the big data platform of choice for data scientists due to its speed, scalability, and easy-to-use APIs. This book deep dives into using Spark to deliver production-grade data science solutions. This process is demonstrated by exploring the construction of a sophisticated global news analysis service that uses Spark to generate continuous geopolitical and current affairs insights. You will learn all about the core Spark APIs and take a comprehensive tour of advanced libraries, including Spark SQL, Spark Streaming, MLlib, and more. You will be introduced to advanced techniques and methods that will help you to construct commercial-grade data products. Focusing on a sequence of tutorials that deliver a working news intelligence service, you will learn about advanced Spark architectures, how to work with geographic data in Spark, and how to tune Spark algorithms so they scale linearly. Style and

approach This is an advanced guide for those with beginner-level familiarity with the Spark architecture and working with Data Science applications. Mastering Spark for Data Science is a practical tutorial that uses core Spark APIs and takes a deep dive into advanced libraries including: Spark SQL, visual streaming, and MLlib. This book expands on titles like: Machine Learning with Spark and Learning Spark. It is the next learning curve for those comfortable with Spark and looking to improve their skills. *Hands-On Big Data Analytics with PySpark* "O'Reilly Media, Inc." The amount of data being generated today is staggering and growing. Apache Spark has emerged as the de facto tool to analyze big data and is now a critical part of the data science toolbox. Updated for Spark 3.0, this practical guide brings together Spark, statistical methods, and real-world datasets to teach you how to approach analytics problems using PySpark, Spark's Python API, and other best practices in Spark programming. Data scientists Akash Tandon, Sandy Ryza, Uri Laserson, Sean Owen, and Josh Wills offer an introduction to the Spark ecosystem, then dive into patterns that apply common techniques-including classification, clustering, collaborative filtering, and anomaly detection, to fields such as genomics, security, and finance. This updated edition also covers NLP and image processing. If you have a basic understanding of machine learning and statistics and you program in Python, this book will get you started with large-scale data analysis. Familiarize yourself with Spark's programming model and ecosystem Learn general approaches in data science Examine complete implementations that analyze large public datasets Discover which machine learning tools make sense for particular problems Explore code that can be adapted to many uses *Spark: The Definitive Guide* "O'Reilly Media, Inc." Build data-intensive applications locally and deploy at scale using the combined powers of Python and Spark 2.0 About This Book Learn why and how you can efficiently use Python to process data and build machine learning models in Apache Spark 2.0 Develop and deploy efficient, scalable real-time Spark solutions Take your understanding of using Spark with Python to the next level with this jump start guide Who This Book Is For If you are a Python developer who wants to learn about the Apache Spark 2.0 ecosystem, this book is for you. A firm understanding of Python is expected to get the best out of the book. Familiarity with Spark would be useful, but is not mandatory. What You Will Learn Learn about Apache Spark and the Spark 2.0 architecture Build and interact with Spark DataFrames using Spark SQL Learn how to solve graph and deep learning problems using GraphFrames and TensorFrames respectively Read, transform, and understand data and use it to train machine learning models Build machine learning models with MLlib and ML Learn how to submit your

applications programmatically using spark-submit Deploy locally built applications to a cluster In Detail Apache Spark is an open source framework for efficient cluster computing with a strong interface for data parallelism and fault tolerance. This book will show you how to leverage the power of Python and put it to use in the Spark ecosystem. You will start by getting a firm understanding of the Spark 2.0 architecture and how to set up a Python environment for Spark. You will get familiar with the modules available in PySpark. You will learn how to abstract data with RDDs and DataFrames and understand the streaming capabilities of PySpark. Also, you will get a thorough overview of machine learning capabilities of PySpark using ML and MLlib, graph processing using GraphFrames, and polyglot persistence using Blaze. Finally, you will learn how to deploy your applications to the cloud using the spark-submit command. By the end of this book, you will have established a firm understanding of the Spark Python API and how it can be used to build data-intensive applications. Style and approach This book takes a very comprehensive, step-by-step approach so you understand how the Spark ecosystem can be used with Python to develop efficient, scalable solutions. Every chapter is standalone and written in a very easy-to-understand manner, with a focus on both the hows and the whys of each concept.

The Data Warehouse Toolkit Apress

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. *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.

Advanced Analytics with Pyspark Packt Publishing Ltd

Harness the power of Scala to program Spark and analyze tonnes of data in the blink of an eye! About This Book Learn Scala's sophisticated type system that combines Functional Programming and object-oriented concepts Work on a wide array of applications, from simple batch jobs to stream processing and machine learning Explore the most common as well as some complex use-cases to perform large-scale data analysis with Spark Who This Book Is For Anyone who wishes to learn how to perform data analysis by harnessing the power of Spark will find this book extremely useful. No knowledge of Spark or Scala is assumed, although prior programming experience (especially with other JVM languages) will be useful to pick up concepts quicker. What You Will Learn Understand object-oriented & functional programming concepts of Scala In-depth understanding of Scala collection APIs Work with RDD and DataFrame to learn Spark's core abstractions Analysing structured and unstructured data using SparkSQL and GraphX Scalable and fault-tolerant streaming application development using Spark structured streaming Learn machine-learning best practices for classification, regression, dimensionality reduction, and recommendation system to build predictive models with widely used algorithms in Spark MLlib & ML Build clustering models to cluster a vast amount of data Understand tuning, debugging, and monitoring Spark applications Deploy Spark applications on real clusters in Standalone, Mesos, and YARN In Detail Scala has been observing wide adoption over the past few years, especially in the field of data science and analytics. Spark, built on Scala, has gained a lot of recognition and is being used widely in productions. Thus, if you want to leverage the power of Scala and Spark to make sense of big data, this book is for you. The first part introduces you to Scala, helping you understand the object-oriented and functional programming concepts needed for Spark application development. It then moves on to Spark to cover the basic abstractions using RDD and DataFrame. This will help you develop scalable and fault-tolerant streaming applications by analyzing structured and unstructured data using SparkSQL, GraphX, and Spark structured streaming. Finally, the book moves on to some advanced topics, such as monitoring, configuration, debugging, testing, and deployment. You will also learn how to develop Spark applications using SparkR and PySpark APIs, interactive data analytics using Zeppelin, and in-memory data processing with Alluxio. By the end

of this book, you will have a thorough understanding of Spark, and you will be able to perform full-stack data analytics with a feel that no amount of data is too big. Style and approach Filled with practical examples and use cases, this book will not only help you get up and running with Spark, but will also take you farther down the road to becoming a data scientist.

Learning Spark "O'Reilly Media, Inc."

In this practical book, four Cloudera data scientists present a set of self-contained patterns for performing large-scale data analysis with Spark. The authors bring Spark, statistical methods, and real-world data sets together to teach you how to approach analytics problems by example. You'll start with an introduction to Spark and its ecosystem, and then dive into patterns that apply common techniques—classification, collaborative filtering, and anomaly detection among others—to fields such as genomics, security, and finance. If you have an entry-level understanding of machine learning and statistics, and you program in Java, Python, or Scala, you'll find these patterns useful for working on your own data applications. Patterns include: Recommending music and the Audioscrobbler data set Predicting forest cover with decision trees Anomaly detection in network traffic with K-means clustering Understanding Wikipedia with Latent Semantic Analysis Analyzing co-occurrence networks with GraphX Geospatial and temporal data analysis on the New York City Taxi Trips data Estimating financial risk through Monte Carlo simulation Analyzing genomics data and the BDG project Analyzing neuroimaging data with PySpark and Thunder

Advanced Analytics with Spark Packt Publishing Ltd

Data in all domains is getting bigger. How can you work with it efficiently? Recently updated for Spark 1.3, this book introduces Apache Spark, the open source cluster computing system that makes data analytics fast to write and fast to run. With Spark, you can tackle big datasets quickly through simple APIs in Python, Java, and Scala. This edition includes new information on Spark SQL, Spark Streaming, setup, and Maven coordinates. Written by the developers of Spark, this book will have data scientists and engineers up and running in no time. You'll learn how to express parallel jobs with just a few lines of code, and cover applications from simple batch jobs to stream processing and machine learning. Quickly dive into Spark capabilities such as distributed datasets, in-memory caching, and the interactive shell Leverage Spark's powerful built-in libraries, including Spark SQL, Spark Streaming, and MLlib Use one programming paradigm instead of mixing and matching tools like Hive, Hadoop, Mahout, and Storm Learn how to deploy interactive, batch, and streaming applications Connect to data sources including HDFS, Hive, JSON, and S3 Master advanced topics like data partitioning and shared variables

Advanced Analytics with Spark "O'Reilly Media, Inc."

If you are ready to dive into the MapReduce framework for processing large datasets, this practical book takes you step by step through the algorithms and tools you need to build distributed MapReduce applications with Apache Hadoop or Apache Spark. Each chapter provides a recipe for solving a massive computational problem, such as building a recommendation system. You'll learn how to implement the appropriate MapReduce solution with code that you can use in your projects. Dr. Mahmoud Parsian covers basic design patterns, optimization techniques, and data mining and machine learning solutions for problems in bioinformatics, genomics, statistics, and social network analysis. This book also includes an overview of MapReduce, Hadoop, and Spark. Topics include: Market basket analysis for a large set of transactions Data mining algorithms (K-means, KNN, and Naive Bayes) Using huge genomic data to sequence DNA and RNA Naive Bayes theorem and Markov chains for data and market prediction Recommendation algorithms and pairwise document similarity Linear regression, Cox regression, and Pearson correlation Allelic frequency and mining DNA Social network analysis (recommendation systems, counting triangles, sentiment analysis)

Learning Spark Apress

Gain a broad foundation of advanced data analytics concepts and discover the recent revolution in databases such as Neo4j, Elasticsearch, and MongoDB. This book discusses how to implement ETL techniques including topical crawling, which is applied in domains such as high-frequency algorithmic trading and goal-oriented dialog systems. You'll also see examples of machine learning concepts such as semi-supervised learning, deep learning, and NLP. *Advanced Data Analytics Using Python* also covers important traditional data analysis techniques such as time series and principal component analysis. After reading this book you will have experience of every technical aspect of an analytics project. You'll get to know the concepts using Python code, giving you samples to use in your own projects. What You Will Learn Work with data analysis techniques such as classification, clustering, regression, and forecasting Handle structured and unstructured data, ETL techniques, and different kinds of databases such as Neo4j, Elasticsearch, MongoDB, and MySQL Examine the different big data frameworks, including Hadoop and Spark Discover advanced machine learning concepts such as semi-supervised learning, deep learning, and NLP Who This Book Is For Data scientists and software developers

interested in the field of data analytics.

Data Algorithms Packt Publishing Ltd

Updated new edition of Ralph Kimball's groundbreaking book on dimensional modeling for data warehousing and business intelligence! The first edition of Ralph Kimball's *The Data Warehouse Toolkit* introduced the industry to dimensional modeling, and now his books are considered the most authoritative guides in this space. This new third edition is a complete library of updated dimensional modeling techniques, the most comprehensive collection ever. It covers new and enhanced star schema dimensional modeling patterns, adds two new chapters on ETL techniques, includes new and expanded business matrices for 12 case studies, and more. Authored by Ralph Kimball and Margy Ross, known worldwide as educators, consultants, and influential thought leaders in data warehousing and business intelligence Begins with fundamental design recommendations and progresses through increasingly complex scenarios Presents unique modeling techniques for business applications such as inventory management, procurement, invoicing, accounting, customer relationship management, big data analytics, and more Draws real-world case studies from a variety of industries, including retail sales, financial services, telecommunications, education, health care, insurance, e-commerce, and more Design dimensional databases that are easy to understand and provide fast query response with *The Data Warehouse Toolkit: The Definitive Guide to Dimensional Modeling*, 3rd Edition.

Graph Algorithms Packt Publishing Ltd

Discover how graph algorithms can help you leverage the relationships within your data to develop more intelligent solutions and enhance your machine learning models. You'll learn how graph analytics are uniquely suited to unfold complex structures and reveal difficult-to-find patterns lurking in your data. Whether you are trying to build dynamic network models or forecast real-world behavior, this book illustrates how graph algorithms deliver value—from finding vulnerabilities and bottlenecks to detecting communities and improving machine learning predictions. This practical book walks you through hands-on examples of how to use graph algorithms in Apache Spark and Neo4j—two of the most common choices for graph analytics. Also included: sample code and tips for over 20 practical graph algorithms that cover optimal pathfinding, importance through centrality, and community detection. Learn how graph analytics vary from conventional statistical analysis Understand how classic graph algorithms work, and how they are applied Get guidance on which algorithms to use for different types of questions Explore algorithm examples with working code and sample datasets from Spark and Neo4j See how connected feature extraction can increase machine learning accuracy and precision Walk through creating an ML workflow for link prediction combining Neo4j and Spark

Big Data Analytics with Java Apress

The amount of data being generated today is staggering—and growing. Apache Spark has emerged as the de facto tool to analyze big data and is now a critical part of the data science toolbox. Updated for Spark 3.0, this practical guide brings together Spark, statistical methods, and real-world datasets to teach you how to approach analytics problems using PySpark, Spark's Python API, and other best practices in Spark programming. Data scientists Akash Tandon, Sandy Ryza, Uri Laserson, Sean Owen, and Josh Wills offer an introduction to the Spark ecosystem, then dive into patterns that apply common techniques—including classification, clustering, collaborative filtering, and anomaly detection—to fields such as genomics, security, and finance. This updated edition also covers NLP and image processing. If you have a basic understanding of machine learning and statistics and you program in Python, this book will get you started with large-scale data analysis. Familiarize yourself with Spark's programming model and ecosystem Learn general approaches in data science Examine complete implementations that analyze large public datasets Discover which machine learning tools make sense for particular problems Explore code that can be adapted to many uses

Practical Data Science with Hadoop and Spark Addison-Wesley Professional

Utilize R to uncover hidden patterns in your Big Data About This Book Perform computational analyses on Big Data to generate meaningful results Get a practical knowledge of R programming language while working on Big Data platforms like Hadoop, Spark, H2O and SQL/NoSQL databases, Explore fast, streaming, and scalable data analysis with the most cutting-edge technologies in the market Who This Book Is For This book is intended for Data Analysts, Scientists, Data Engineers, Statisticians, Researchers, who want to integrate R with their current or future Big Data workflows. It is assumed that readers have some experience in data analysis and understanding of data management and algorithmic processing of large quantities of data, however they may lack specific skills related to R. What You Will Learn Learn about current state of Big Data processing using R programming language and its powerful statistical capabilities Deploy Big Data analytics platforms with selected Big Data tools supported by R in a cost-effective and time-saving manner Apply the R language to

real-world Big Data problems on a multi-node Hadoop cluster, e.g. electricity consumption across various socio-demographic indicators and bike share scheme usage. Explore the compatibility of R with Hadoop, Spark, SQL and NoSQL databases, and H2O platform. In Detail Big Data analytics is the process of examining large and complex data sets that often exceed the computational capabilities. R is a leading programming language of data science, consisting of powerful functions to tackle all problems related to Big Data processing. The book will begin with a brief introduction to the Big Data world and its current industry standards. With introduction to the R language and presenting its development, structure, applications in real world, and its shortcomings. Book will progress towards revision of major R functions for data management and transformations. Readers will be introduced to Cloud based Big Data solutions (e.g. Amazon EC2 instances and Amazon RDS, Microsoft Azure and its HDInsight clusters) and also provide guidance on R connectivity with relational and non-relational databases such as MongoDB and HBase etc. It will further expand to include Big Data tools such as Apache Hadoop ecosystem, HDFS and MapReduce frameworks. Also other R compatible tools such as Apache Spark, its machine learning library Spark MLlib, as well as H2O. Style and approach This book will serve as a practical guide to tackling Big Data problems using R programming language and its statistical environment. Each section of the book will present you with concise and easy-to-follow steps on how to process, transform and analyse large data sets.

Mastering Spark with R "O'Reilly Media, Inc."

Learn the basics of analytics on big data using Java, machine learning and other big data tools. About This Book Acquire level data science applications. Surpasses the barrier of other languages in data science and learn to create useful object-oriented codes. Extensive use of Java compliant big data tools like Apache Spark, Hadoop, etc. Who This Book Is For This book is for Java developers who are looking to perform data analysis in production environment. Those who wish to implement data analysis in their Big Data applications will find this book helpful. What You Will Learn Start from simple analytic tasks on big data. Get into more complex tasks with predictive analytics on big data using machine learning. Learn real-time analytic tasks. Understand the concepts with examples and case studies. Prepare and refine data for analysis. Create charts in order to understand the data. See various real-world datasets. In Detail This book covers case studies such as sentiment analysis on a tweet dataset, recommendations on a movielens dataset, customer segmentation on an ecommerce dataset, and graph analysis on actual flights dataset. This book is an end-to-end guide to implement analytics on big data with Java. Java is the de facto language for major big data environments, including Hadoop. This book will teach you how to perform analytics on big data with production-friendly Java. This book is basically divided into two sections. The first part is an introduction that will help the readers get acquainted with big data environments, whereas the second part will contain a hardcore discussion on all the concepts in analytics on big data. It will take you from data analysis and data visualization to the core concepts and advantages of machine learning, real-life usage of regression and classification using Naive Bayes, a deep discussion on the concepts of clustering, and a review of simple neural networks on big data using deepLearning4j or plain Java Spark code. This book is a must-have book for Java developers who want to start learning big data analytics and want to use it in the real world. Style and approach The approach of the book is to deliver practical

learning modules in manageable content. Each chapter is a self-contained unit of a concept in big data analytics. Book will step by step build the competency in the area of big data analytics. Examples using real world case studies to give ideas of real applications and how to use the techniques mentioned. The examples and case studies will be shown using both theory and code.

Advanced Analytics with PySpark O'Reilly Media

If you're like most R users, you have deep knowledge and love for statistics. But as your organization continues to collect huge amounts of data, adding tools such as Apache Spark makes a lot of sense. With this practical book, data scientists and professionals working with large-scale data applications will learn how to use Spark from R to tackle big data and big compute problems. Authors Javier Luraschi, Kevin Kuo, and Edgar Ruiz show you how to use R with Spark to solve different data analysis problems. This book covers relevant data science topics, cluster computing, and issues that should interest even the most advanced users. Analyze, explore, transform, and visualize data in Apache Spark with R. Create statistical models to extract information and predict outcomes; automate the process in production-ready workflows. Perform analysis and modeling across many machines using distributed computing techniques. Use large-scale data from multiple sources and different formats with ease from within Spark. Learn about alternative modeling frameworks for graph processing, geospatial analysis, and genomics at scale. Dive into advanced topics including custom transformations, real-time data processing, and creating custom Spark extensions.

Scala for Data Science Packt Publishing Ltd

A New York Times bestseller "Brilliant...the best math teacher you never had." —San Francisco Chronicle. Once considered tedious, the field of statistics is rapidly evolving into a discipline. Hal Varian, chief economist at Google, has actually called "sexy." From batting averages and political polls to game shows and medical research, the real-world application of statistics continues to grow by leaps and bounds. How can we catch schools that cheat on standardized tests? How does Netflix know which movies you'll like? What is causing the rising incidence of autism? As best-selling author Charles Wheelan shows us in *Naked Statistics*, the right data and a few well-chosen statistical tools can help us answer these questions and more. For those who slept through Stats 101, this book is a lifesaver. Wheelan strips away the arcane and technical details and focuses on the underlying intuition that drives statistical analysis. He clarifies key concepts such as inference, correlation, and regression analysis, reveals how biased or careless parties can manipulate or misrepresent data, and shows us how brilliant and creative researchers are exploiting the valuable data from natural experiments to tackle thorny questions. And in Wheelan's trademark style, there's not a dull page in sight. You'll encounter clever Schlitz Beer marketers leveraging basic probability, an International Sausage Festival illuminating the tenets of the central limit theorem, and a head-scratching choice from the famous game show *Let's Make a Deal*—and you'll come away with insights each time. With the wit, accessibility, and sheer fun that turned *Naked Economics* into a bestseller, Wheelan defies the odds yet again by bringing another essential, formerly unglamorous discipline to life.

Machine Learning with Apache Spark Quick Start Guide FT Press

Leverage Apache Spark within a modern data engineering ecosystem. This hands-on guide will teach you how to write fully functional applications, follow industry best practices, and learn

the rationale behind these decisions. With Apache Spark as the foundation, you will follow a step-by-step journey beginning with the basics of data ingestion, processing, and transformation, and ending up with an entire local data platform running Apache Spark, Apache Zeppelin, Apache Kafka, Redis, MySQL, Minio (S3), and Apache Airflow. Apache Spark applications solve a wide range of data problems from traditional data loading and processing to rich SQL-based analysis as well as complex machine learning workloads and even near real-time processing of streaming data. Spark fits well as a central foundation for any data engineering workload. This book will teach you to write interactive Spark applications using Apache Zeppelin notebooks, write and compile reusable applications and modules, and fully test both batch and streaming. You will also learn to containerize your applications using Docker and run and deploy your Spark applications using a variety of tools such as Apache Airflow, Docker and Kubernetes. Reading this book will empower you to take advantage of Apache Spark to optimize your data pipelines and teach you to craft modular and testable Spark applications. You will create and deploy mission-critical streaming spark applications in a low-stress environment that paves the way for your own path to production. What You Will Learn Simplify data transformation with Spark Pipelines and Spark SQL Bridge. Data engineering with machine learning. Architect modular data pipeline applications. Build reusable application components and libraries. Containerize your Spark applications for consistency and reliability. Use Docker and Kubernetes to deploy your Spark applications. Speed up application experimentation using Apache Zeppelin and Docker. Understand serializable structured data and data contracts. Harness effective strategies for optimizing data in your data lakes. Build end-to-end Spark structured streaming applications using Redis and Apache Kafka. Embrace testing for your batch and streaming applications. Deploy and monitor your Spark applications. Who This Book Is For Professional software engineers who want to take their current skills and apply them to new and exciting opportunities within the data ecosystem, practicing data engineers who are looking for a guiding light while traversing the many challenges of moving from batch to streaming modes, data architects who wish to provide clear and concise direction for how best to harness and use Apache Spark within their organization, and those interested in the ins and outs of becoming a modern data engineer in today's fast-paced and data-hungry world.

Data Analysis with Python and PySpark Packt Publishing Ltd

In the second edition of this practical book, four Cloudera data scientists present a set of self-contained patterns for performing large-scale data analysis with Spark. The authors bring Spark, statistical methods, and real-world data sets together to teach you how to approach analytics problems by example. Updated for Spark 2.1, this edition acts as an introduction to these techniques and other best practices in Spark programming. You'll start with an introduction to Spark and its ecosystem, and then dive into patterns that apply common techniques—including classification, clustering, collaborative filtering, and anomaly detection—to fields such as genomics, security, and finance. If you have an entry-level understanding of machine learning and statistics, and you program in Java, Python, or Scala, you'll find the book's patterns useful for working on your own data applications. With this book, you will: Familiarize yourself with the Spark programming model. Become comfortable within the Spark ecosystem. Learn general approaches in data science. Examine complete implementations that analyze large public data sets. Discover which machine learning tools make sense for particular problems. Acquire code that can be adapted to many uses.

Related with [Advanced Analytics With Spark Patterns For Learning From Data At Scale](#):

- Fantasy Alarm Draft Guide : [click here](#)