

Advanced Analytics With Spark Patterns For Learning From Data At Scale

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

Data is bigger, arrives faster, and comes in a variety of formats—and it all needs to be processed at scale for analytics or machine learning. But how can you process such varied workloads efficiently? Enter Apache Spark. Updated to include Spark 3.0, this second edition shows data engineers and data scientists why structure and unification in Spark matters. Specifically, this book explains how to perform simple

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

and complex data analytics and employ machine learning algorithms. Through step-by-step walk-throughs, code snippets, and notebooks, you'll be able to: Learn Python, SQL, Scala, or Java high-level Structured APIs Understand Spark operations and SQL Engine Inspect, tune, and debug Spark operations with Spark configurations and Spark UI Connect to data sources: JSON, Parquet, CSV, Avro, ORC, Hive, S3, or Kafka Perform analytics on batch and streaming data using Structured Streaming Build reliable data pipelines with open source Delta Lake and Spark Develop machine learning pipelines with MLlib and productionize models using MLflow

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,

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

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.

A handy reference guide for data analysts and data scientists to help to obtain value from big data analytics using Spark on Hadoop clusters About This Book This book is based on the latest 2.0 version of Apache Spark and 2.7 version of Hadoop integrated with most commonly used tools. Learn all Spark stack components including latest topics such as DataFrames, DataSets, GraphFrames, Structured Streaming, DataFrame based ML Pipelines and SparkR. Integrations with frameworks such as HDFS, YARN and tools such as Jupyter, Zeppelin, NiFi, Mahout, HBase Spark Connector, GraphFrames, H2O and Hivemall. Who This Book Is For Though this book is primarily aimed at data analysts and data scientists, it will also help architects, programmers, and practitioners. Knowledge of either Spark or Hadoop would be beneficial. It is assumed that you have basic programming background in Scala,

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

Python, SQL, or R programming with basic Linux experience. Working experience within big data environments is not mandatory. What You Will Learn Find out and implement the tools and techniques of big data analytics using Spark on Hadoop clusters with wide variety of tools used with Spark and Hadoop Understand all the Hadoop and Spark ecosystem components Get to know all the Spark components: Spark Core, Spark SQL, DataFrames, DataSets, Conventional and Structured Streaming, MLLib, ML Pipelines and Graphx See batch and real-time data analytics using Spark Core, Spark SQL, and Conventional and Structured Streaming Get to grips with data science and machine learning using MLLib, ML Pipelines, H2O, Hivemall, Graphx, SparkR and Hivemall. In Detail Big Data Analytics book aims at providing the fundamentals of Apache Spark and Hadoop. All Spark components – Spark Core, Spark SQL, DataFrames, Data sets, Conventional Streaming, Structured Streaming, MLLib, Graphx and Hadoop core components – HDFS, MapReduce and Yarn are explored in greater depth with implementation examples on Spark + Hadoop clusters. It is moving away from MapReduce to Spark. So, advantages of Spark over MapReduce are explained at great depth to reap benefits of in-memory speeds. DataFrames API, Data Sources API and new Data set API are explained for building Big Data analytical applications. Real-time data analytics using Spark Streaming with Apache Kafka and HBase is covered to help building streaming applications. New Structured streaming concept is explained with an IOT (Internet of Things) use case. Machine learning techniques are covered using MLLib, ML Pipelines and SparkR and Graph Analytics are covered with GraphX and GraphFrames components of Spark. Readers will also get an opportunity to get started with web based notebooks such as Jupyter, Apache Zeppelin and data flow tool Apache NiFi to analyze and visualize data. Style and approach This step-by-step pragmatic guide will make life easy no matter what your level of experience. You will deep dive into Apache Spark on Hadoop clusters through ample exciting real-life examples. Practical tutorial explains

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

data science in simple terms to help programmers and data analysts get started with Data Science

Data Analytics with Spark Using Python

Big Data Analytics with Java

Big Data Processing with Apache Spark

Mastering Machine Learning with Spark 2.x

Advanced Analytics with Pyspark

Hands-On Big Data Analytics with PySpark

Best Practices for Scaling and Optimizing Apache Spark

Access real-world documentation and examples for the Spark platform for building large-scale, enterprise-grade machine learning applications. The past decade has seen an astonishing series of advances in machine learning. These breakthroughs are disrupting our everyday life and making an impact across every industry. Next-Generation Machine Learning with Spark provides a gentle introduction to Spark and Spark MLlib and advances to more powerful, third-party machine learning algorithms and libraries beyond what is available in the standard Spark MLlib library. By the end of this book, you will be able to apply your knowledge to real-world use cases through dozens of practical examples and insightful explanations. What You Will Learn Be introduced to machine learning, Spark, and Spark MLlib 2.4.x Achieve lightning-fast gradient boosting on Spark with the XGBoost4J-Spark and LightGBM libraries Detect anomalies with the Isolation Forest algorithm for Spark Use the Spark NLP and Stanford CoreNLP libraries that support multiple languages Optimize your ML workload with the Alluxio in-memory data accelerator for Spark Use GraphX and GraphFrames for Graph Analysis Perform image recognition using convolutional neural networks Utilize the Keras

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

framework and distributed deep learning libraries with Spark Who This Book Is For Data scientists and machine learning engineers who want to take their knowledge to the next level and use Spark and more powerful, next-generation algorithms and libraries beyond what is available in the standard Spark MLlib library; also serves as a primer for aspiring data scientists and engineers who need an introduction to machine learning, Spark, and Spark MLlib.

Large data sets arriving at every increasing speeds require a new set of efficient data analysis techniques. Data analytics are becoming an essential component for every organization and technologies such as health care, financial trading, Internet of Things, Smart Cities or Cyber Physical Systems. However, these diverse application domains give rise to new research challenges. In this context, the book provides a broad picture on the concepts, techniques, applications, and open research directions in this area. In addition, it serves as a single source of reference for acquiring the knowledge on emerging Big Data Analytics technologies.

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

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

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

Solve Data Analytics Problems with Spark, PySpark, and Related Open Source Tools Spark is at the heart of today ' s Big Data revolution, helping data professionals supercharge efficiency and performance in a wide range of data processing and analytics tasks. In this guide, Big Data expert Jeffrey Aven covers all you need to know to leverage Spark, together with its extensions, subprojects, and wider ecosystem. Aven combines a language-agnostic introduction to foundational Spark concepts with extensive programming examples utilizing the popular and intuitive PySpark development environment. This guide ' s focus on Python makes it widely accessible to large audiences of data professionals, analysts, and developers—even those with little Hadoop or Spark experience. Aven ' s broad coverage ranges from basic to advanced Spark programming, and Spark SQL to machine learning. You ' ll learn how to efficiently manage all forms of data with Spark: streaming, structured, semi-structured, and unstructured. Throughout, concise topic overviews quickly get you up to speed, and extensive hands-on exercises prepare you to solve real problems. Coverage includes:

- Understand Spark ' s evolving role in the Big Data and Hadoop ecosystems
- Create Spark clusters using various deployment modes
- Control and optimize the operation of Spark clusters and applications
- Master Spark Core RDD API programming techniques
- Extend, accelerate, and optimize Spark routines with advanced API platform constructs, including shared variables,

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

RDD storage, and partitioning • Efficiently integrate Spark with both SQL and nonrelational data stores • Perform stream processing and messaging with Spark Streaming and Apache Kafka • Implement predictive modeling with SparkR and Spark MLlib

Mastering Azure Analytics

Big Data Analytics Beyond Hadoop

Advanced Data Analytics Using Python

Real-Time Applications with Storm, Spark, and More Hadoop Alternatives

Patterns for Learning from Data at Scale

Architecting in the Cloud with Azure Data Lake, HDInsight, and Spark

Big Data Analytics with R

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

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

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 Advanced Analytics with SparkPatterns for Learning from Data at Scale"O'Reilly Media, Inc."

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

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

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. Master alternative Big Data technologies that can do what Hadoop can't: real-time analytics and iterative machine learning. When most technical professionals think of Big Data analytics today, they think of Hadoop. But there are many cutting-edge applications that Hadoop isn't well suited for, especially real-time analytics and contexts requiring the use of iterative machine learning algorithms. Fortunately, several powerful new technologies have been developed specifically for use cases such as these. Big Data Analytics Beyond Hadoop is the first guide specifically designed to help you take the next steps beyond Hadoop. Dr. Vijay Srinivas Agneeswaran introduces the breakthrough Berkeley Data Analysis Stack (BDAS) in detail, including its motivation, design, architecture, Mesos cluster management, performance, and more. He presents realistic use cases and up-to-date example code for: Spark, the next generation in-memory computing technology from UC Berkeley Storm, the parallel real-time Big Data analytics technology from Twitter GraphLab, the next-generation graph processing paradigm from CMU and the University of Washington (with comparisons to alternatives such as Pregel and Piccolo) Halo also offers architectural and design guidance and code sketches for scaling machine learning algorithms to Big Data, and then realizing them in real-time. He concludes by previewing emerging trends, including real-time video analytics, SDNs, and even Big Data governance, security, and privacy issues. He

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

identifies intriguing startups and new research possibilities, including BDAS extensions and cutting-edge model-driven analytics. Big Data Analytics Beyond Hadoop is an indispensable resource for everyone who wants to reach the cutting edge of Big Data analytics, and stay there: practitioners, architects, programmers, data scientists, researchers, startup entrepreneurs, and advanced students.

Spark: The Definitive Guide

Advanced Analytics with Spark, 2nd Edition

High Performance Spark

Advanced Analytics with PySpark

Create scalable pipelines that ingest, curate, and aggregate complex data in a timely and secure way

Discovering, Analyzing, Visualizing and Presenting Data

Machine Learning with Apache Spark Quick Start Guide

Big Data Analytics with Spark is a step-by-step guide for learning Spark, which is an open-source fast and general-purpose cluster computing framework for large-scale data analysis. You will learn how to use Spark for different types of big data analytics projects, including batch, interactive, graph, and stream data analysis as well as machine learning. In addition, this book will help you

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

become a much sought-after Spark expert. Spark is one of the hottest Big Data technologies. The amount of data generated today by devices, applications and users is exploding. Therefore, there is a critical need for tools that can analyze large-scale data and unlock value from it. Spark is a powerful technology that meets that need. You can, for example, use Spark to perform low latency computations through the use of efficient caching and iterative algorithms; leverage the features of its shell for easy and interactive Data analysis; employ its fast batch processing and low latency features to process your real time data streams and so on. As a result, adoption of Spark is rapidly growing and is replacing Hadoop MapReduce as the technology of choice for big data analytics. This book provides an introduction to Spark and related big-data technologies. It covers Spark core and its add-on libraries, including Spark SQL, Spark Streaming, GraphX, and MLlib. Big Data Analytics with Spark is therefore written for busy professionals who prefer learning a new technology from a consolidated source

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

instead of spending countless hours on the Internet trying to pick bits and pieces from different sources. The book also provides a chapter on Scala, the hottest functional programming language, and the program that underlies Spark. You'll learn the basics of functional programming in Scala, so that you can write Spark applications in it. What's more, Big Data Analytics with Spark provides an introduction to other big data technologies that are commonly used along with Spark, like Hive, Avro, Kafka and so on. So the book is self-sufficient; all the technologies that you need to know to use Spark are covered. The only thing that you are expected to know is programming in any language. There is a critical shortage of people with big data expertise, so companies are willing to pay top dollar for people with skills in areas like Spark and Scala. So reading this book and absorbing its principles will provide a boost—possibly a big boost—to your career.

Data Science and Big Data Analytics is about harnessing the power of data for new insights. The book covers the breadth

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

of activities and methods and tools that Data Scientists use. The content focuses on concepts, principles and practical applications that are applicable to any industry and technology environment, and the learning is supported and explained with examples that you can replicate using open-source software. This book will help you: Become a contributor on a data science team Deploy a structured lifecycle approach to data analytics problems Apply appropriate analytic techniques and tools to analyzing big data Learn how to tell a compelling story with data to drive business action Prepare for EMC Proven Professional Data Science Certification Corresponding data sets are available from the book's page at Wiley which you can find on the Wiley site by searching for the ISBN 9781118876138. Get started discovering, analyzing, visualizing, and presenting data in a meaningful way today!

You'll start with an introduction to Spark and its ecosystem, and then dive into patterns that apply common techniques including classification, clustering,

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

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.

Summary Spark in Action teaches you the theory and skills you need to effectively handle batch and streaming data using Spark. Fully updated for Spark 2.0. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Big data systems distribute datasets across clusters of machines, making it a challenge to efficiently query, stream, and interpret them. Spark can help. It is a processing system designed specifically for distributed data. It provides easy-to-use interfaces, along with the performance you need for production-quality analytics and machine learning. Spark 2 also adds improved programming APIs, better performance, and countless other upgrades.

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

About the Book Spark in Action teaches you the theory and skills you need to effectively handle batch and streaming data using Spark. You'll get comfortable with the Spark CLI as you work through a few introductory examples. Then, you'll start programming Spark using its core APIs. Along the way, you'll work with structured data using Spark SQL, process near-real-time streaming data, apply machine learning algorithms, and munge graph data using Spark GraphX. For a zero-effort startup, you can download the preconfigured virtual machine ready for you to try the book's code. What's Inside Updated for Spark 2.0 Real-life case studies Spark DevOps with Docker Examples in Scala, and online in Java and Python About the Reader Written for experienced programmers with some background in big data or machine learning. About the Authors Petar Zečević and Marko Bonaći are seasoned developers heavily involved in the Spark community. Table of Contents PART 1 - FIRST STEPS Introduction to Apache Spark Spark fundamentals Writing Spark applications The Spark API in depth PART 2 - MEET THE

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

SPARK FAMILY Sparkling queries with Spark SQL Ingesting data with Spark Streaming Getting smart with MLlib ML:

classification and clustering Connecting the dots with

GraphX PART 3 – SPARK OPS Running Spark Running on a Spark standalone cluster Running on YARN and Mesos PART 4 –

BRINGING IT TOGETHER Case study: real-time dashboard Deep learning on Spark with H2O

Distributed Computing and Event Processing using Apache Spark, Flink, Storm, and Kafka

Practical Real-time Data Processing and Analytics

Data Science and Big Data Analytics

Next-Generation Machine Learning with Spark

R for Everyone

Data Analytics Basics

Mastering Spark with R

Get started with distributed computing using PySpark, a single unified framework to solve end-to-end data analytics at scale Key Features

Discover how to convert huge amounts of raw data into meaningful and actionable insights Use Spark's unified analytics engine for end-to-end

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

analytics, from data preparation to predictive analytics Perform data ingestion, cleansing, and integration for ML, data analytics, and data visualization Book Description Apache Spark is a unified data analytics engine designed to process huge volumes of data quickly and efficiently. PySpark is Apache Spark's Python language API, which offers Python developers an easy-to-use scalable data analytics framework. Essential PySpark for Scalable Data Analytics starts by exploring the distributed computing paradigm and provides a high-level overview of Apache Spark. You'll begin your analytics journey with the data engineering process, learning how to perform data ingestion, cleansing, and integration at scale. This book helps you build real-time analytics pipelines that help you gain insights faster. You'll then discover methods for building cloud-based data lakes, and explore Delta Lake, which brings reliability to data lakes. The book also covers Data Lakehouse, an emerging paradigm, which combines the structure and performance of a data warehouse with the scalability of cloud-based data lakes. Later, you'll perform scalable data science and machine learning tasks using PySpark, such as data preparation, feature engineering, and model training and productionization. Finally, you'll learn ways to scale out standard Python ML libraries along with a new pandas API on top of PySpark called Koalas. By the end of this PySpark book, you'll

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

be able to harness the power of PySpark to solve business problems. What you will learn

- Understand the role of distributed computing in the world of big data
- Gain an appreciation for Apache Spark as the de facto go-to for big data processing
- Scale out your data analytics process using Apache Spark
- Build data pipelines using data lakes, and perform data visualization with PySpark and Spark SQL
- Leverage the cloud to build truly scalable and real-time data analytics applications
- Explore the applications of data science and scalable machine learning with PySpark
- Integrate your clean and curated data with BI and SQL analysis tools

Who this book is for This book is for practicing data engineers, data scientists, data analysts, and data enthusiasts who are already using data analytics to explore distributed and scalable data analytics. Basic to intermediate knowledge of the disciplines of data engineering, data science, and SQL analytics is expected. General proficiency in using any programming language, especially Python, and working knowledge of performing data analytics using frameworks such as pandas and SQL will help you to get the most out of this book.

Learn how to use, deploy, and maintain Apache Spark with this comprehensive guide, written by the creators of the open-source cluster-computing framework. With an emphasis on improvements and new features in Spark 2.0, authors Bill Chambers and Matei Zaharia break down

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

Spark topics into distinct sections, each with unique goals. You'll explore the basic operations and common functions of Spark's structured APIs, as well as Structured Streaming, a new high-level API for building end-to-end streaming applications. Developers and system administrators will learn the fundamentals of monitoring, tuning, and debugging Spark, and explore machine learning techniques and scenarios for employing MLlib, Spark's scalable machine-learning library. Get a gentle overview of big data and Spark Learn about DataFrames, SQL, and Datasets—Spark's core APIs—through worked examples Dive into Spark's low-level APIs, RDDs, and execution of SQL and DataFrames Understand how Spark runs on a cluster Debug, monitor, and tune Spark clusters and applications Learn the power of Structured Streaming, Spark's stream-processing engine Learn how you can apply MLlib to a variety of problems, including classification or recommendation

Data science teams looking to turn research into useful analytics applications require not only the right tools, but also the right approach if they're to succeed. With the revised second edition of this hands-on guide, up-and-coming data scientists will learn how to use the Agile Data Science development methodology to build data applications with Python, Apache Spark, Kafka, and other tools. Author Russell Journey demonstrates how to

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

compose a data platform for building, deploying, and refining analytics applications with Apache Kafka, MongoDB, Elasticsearch, d3.js, scikit-learn, and Apache Airflow. You'll learn an iterative approach that lets you quickly change the kind of analysis you're doing, depending on what the data is telling you. Publish data science work as a web application, and affect meaningful change in your organization. Build value from your data in a series of agile sprints, using the data-value pyramid Extract features for statistical models from a single dataset Visualize data with charts, and expose different aspects through interactive reports Use historical data to predict the future via classification and regression Translate predictions into actions Get feedback from users after each sprint to keep your project on track

Ready to use statistical and machine-learning techniques across large data sets? This practical guide shows you why the Hadoop ecosystem is perfect for the job. Instead of deployment, operations, or software development usually associated with distributed computing, you'll focus on particular analyses you can build, the data warehousing techniques that Hadoop provides, and higher order data workflows this framework can produce. Data scientists and analysts will learn how to perform a wide range of techniques, from writing MapReduce and Spark applications with Python to

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

using advanced modeling and data management with Spark MLlib, Hive, and HBase. You'll also learn about the analytical processes and data systems available to build and empower data products that can handle—and actually require—huge amounts of data. Understand core concepts behind Hadoop and cluster computing Use design patterns and parallel analytical algorithms to create distributed data analysis jobs Learn about data management, mining, and warehousing in a distributed context using Apache Hive and HBase Use Sqoop and Apache Flume to ingest data from relational databases Program complex Hadoop and Spark applications with Apache Pig and Spark DataFrames Perform machine learning techniques such as classification, clustering, and collaborative filtering with Spark's MLlib

Graph Algorithms

Advanced Analytics and Graphics

Explore the concepts of functional programming, data streaming, and machine learning

Agile Data Science 2.0

Big Data Analytics

A Beginner's Guide

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

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

Statistical Computation for Programmers, Scientists, Quants, Excel Users, and Other Professionals Using the open source R language, you can build powerful statistical models to answer many of your most challenging questions. R has traditionally been difficult for non-statisticians to learn, and most R books assume far too much

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

knowledge to be of help. R for Everyone, Second Edition, is the solution. Drawing on his unsurpassed experience teaching new users, professional data scientist Jared P. Lander has written the perfect tutorial for anyone new to statistical programming and modeling. Organized to make learning easy and intuitive, this guide focuses on the 20 percent of R functionality you'll need to accomplish 80 percent of modern data tasks. Lander's self-contained chapters start with the absolute basics, offering extensive hands-on practice and sample code. You'll download and install R; navigate and use the R environment; master basic program control, data import, manipulation, and visualization; and walk through several essential tests. Then, building on this foundation, you'll construct several complete models, both linear and nonlinear, and use some data mining techniques. After all this you'll make your code reproducible with LaTeX, RMarkdown, and Shiny. By the time you're done, you won't just know how to write R programs, you'll be ready to tackle the statistical problems you care about most. Coverage includes Explore R, RStudio, and R packages Use R for math: variable types, vectors, calling functions, and more Exploit data structures, including data.frames, matrices, and lists Read many different types of data Create attractive, intuitive statistical graphics Write user-defined functions Control program flow with if, ifelse, and complex checks Improve program efficiency with group manipulations Combine and reshape multiple datasets Manipulate strings using R's facilities and regular expressions Create normal, binomial, and Poisson probability

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

distributions Build linear, generalized linear, and nonlinear models Program basic statistics: mean, standard deviation, and t-tests Train machine learning models Assess the quality of models and variable selection Prevent overfitting and perform variable selection, using the Elastic Net and Bayesian methods Analyze univariate and multivariate time series data Group data via K-means and hierarchical clustering Prepare reports, slideshows, and web pages with knitr Display interactive data with RMarkdown and htmlwidgets Implement dashboards with Shiny Build reusable R packages with devtools and Rcpp Register your product at informit.com/register for convenient access to downloads, updates, and corrections as they become available. 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

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

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

Helps users understand the breadth of Azure services by organizing them into a reference framework they can use when crafting their own big-data analytics solution.

Data Algorithms

Uncover patterns, derive actionable insights, and learn from big data using MLlib

Advanced Analytics with Spark

Patterns for Learning from Data at Scale Using Python and Spark

Concepts, Techniques, and Applications

The Complete Guide to Large-Scale Analysis and Modeling

Advanced Analytics With Spark

Data analytics is increasingly becoming a key element in shaping a company's business strategy. Today, data influences every decision made by an organization, and this is driving the wide-scale adoption of data analytics, including machine learning technology and artificial intelligence solutions. The heightened focus is propelling a surge in data analytics spending, reflected in various studies conducted by leading market research

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

firms. The field of data analytics offers some amazing salaries and is not only the hottest job, but it is also one of the best-paying jobs in the world. This guide aims at providing readers with everything they need to know about the data analytics field, basic terminologies, key concepts, real-life use cases, skills you must master in order to scale your career, and training and certifications you might need to reach your dream job. Apache Spark is amazing when everything clicks. But if you haven't seen the performance improvements you expected, or still don't feel confident enough to use Spark in production, this practical book is for you. Authors Holden Karau and Rachel Warren demonstrate performance optimizations to help your Spark queries run faster and handle larger data sizes, while using fewer resources. Ideal for software engineers, data engineers, developers, and system administrators working with large-scale data applications, this book describes techniques that can reduce data infrastructure costs and developer headaches. Not only will you gain a more comprehensive understanding of Spark, you'll also learn how to make it sing. With this book, you'll explore: How Spark SQL's new interfaces improve performance over SQL's RDD data structure The choice between data joins in Core Spark and Spark SQL Techniques for getting the most out of standard RDD transformations How to work around performance issues in Spark's key/value pair paradigm Writing high-performance Spark code without Scala or the JVM How to test functionality and performance when applying suggested improvements Using Spark ML and Spark ML machine learning libraries Spark's Streaming components and external

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

community packages

Unlock the complexities of machine learning algorithms in Spark to generate useful data insights through this data analysis tutorial About This Book Process and analyze big data in a distributed and scalable way Write sophisticated Spark pipelines that incorporate elaborate extraction Build and use regression models to predict flight delays Who This Book Is For Are you a developer with a background in machine learning and statistics who is feeling limited by the current slow and “small data” machine learning tools? Then this is the book for you! In this book, you will create scalable machine learning applications to power a modern data-driven business using Spark. We assume that you already know machine learning concepts and algorithms and have Spark up and running (whether on a cluster or locally) and have a basic knowledge of the various libraries contained in Spark What You Will Learn Use Spark streams to cluster tweets online Run the PageRank algorithm to compute user influence Perform complex manipulation of DataFrames using Spark Define Spark pipelines to compose individual data transformations Utilize generated models for off-line/on-line prediction Transfer the learning from an ensemble to a simpler Neural Network Understand basic graph properties and important graph operations Use GraphFrames, an extension of DataFrames to graphs, to study graphs using an elegant query language Use K-means algorithm to cluster movie reviews data In Detail The purpose of machine learning is to build systems that learn from data. Being able to understand trends and patterns in complex data is critical to success; it is one

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

key strategies to unlock growth in the challenging contemporary marketplace today. With the meteoric rise of machine learning, developers are now keen on finding out how can they make their Spark applications smarter. This book gives you access to transform code into actionable knowledge. The book commences by defining machine learning primitive by the MLlib and H2O libraries. You will learn how to use Binary classification to detect the Higgs Boson particle in the huge amount of data produced by CERN particle collider and classify daily health activities using ensemble Methods for Multi-Class Classification. Next, you will solve a typical regression problem involving flight delay predictions and write sophisticated Spark pipelines. You will analyze Twitter data with help of the doc2vec algorithm and K-means clustering. Finally, you will build different pattern mining models using MLlib, perform complex manipulation of DataFrames using Spark and Spark SQL, and deploy your app in a Spark streaming environment. Style and approach This book takes a practical approach to help you get to grips with using Spark for analytics and implement machine learning algorithms. We'll teach you about advanced applications of machine learning through illustrative examples. These examples will equip you to harness the potential of machine learning, through Spark, in a variety of enterprise-grade systems. 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. You'll start with an introduction to Spark and its ecosystem and then dive into patterns that apply common techniques. Discover which machine

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

learning tools make sense for particular problems.

Data Analytics

Big Data Processing Made Simple

Essential PySpark for Scalable Data Analytics

Building Full-Stack Data Analytics Applications with Spark

Data Engineering with Apache Spark, Delta Lake, and Lakehouse

A beginner's guide to harnessing the power and ease of PySpark 3

Recipes for Scaling Up with Hadoop and Spark

A practical guide to help you tackle different real-time data processing and analytics problems using the best tools for each scenario About This Book Learn about the various challenges in real-time data processing and use the right tools to overcome them This book covers popular tools and frameworks such as Spark, Flink, and Apache Storm to solve all your distributed processing problems A practical guide filled with examples, tips, and tricks to help you perform efficient Big Data processing in real-time Who This Book Is For If you are a Java developer who would like to be equipped with all the tools required to devise an end-to-end practical solution on real-time data streaming, then this book is for you. Basic knowledge of

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

real-time processing would be helpful, and knowing the fundamentals of Maven, Shell, and Eclipse would be great. What You Will Learn Get an introduction to the established real-time stack Understand the key integration of all the components Get a thorough understanding of the basic building blocks for real-time solution designing Garnish the search and visualization aspects for your real-time solution Get conceptually and practically acquainted with real-time analytics Be well equipped to apply the knowledge and create your own solutions In Detail With the rise of Big Data, there is an increasing need to process large amounts of data continuously, with a shorter turnaround time. Real-time data processing involves continuous input, processing and output of data, with the condition that the time required for processing is as short as possible. This book covers the majority of the existing and evolving open source technology stack for real-time processing and analytics. You will get to know about all the real-time solution aspects, from the source to the presentation to persistence. Through this practical book, you'll be equipped with a clear understanding of how to solve challenges on your own. We'll cover topics such as

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

how to set up components, basic executions, integrations, advanced use cases, alerts, and monitoring. You'll be exposed to the popular tools used in real-time processing today such as Apache Spark, Apache Flink, and Storm. Finally, you will put your knowledge to practical use by implementing all of the techniques in the form of a practical, real-world use case. By the end of this book, you will have a solid understanding of all the aspects of real-time data processing and analytics, and will know how to deploy the solutions in production environments in the best possible manner. Style and Approach In this practical guide to real-time analytics, each chapter begins with a basic high-level concept of the topic, followed by a practical, hands-on implementation of each concept, where you can see the working and execution of it. The book is written in a DIY style, with plenty of practical use cases, well-explained code examples, and relevant screenshots and diagrams.

Learn the basics of analytics on big data using Java, machine learning and other big data tools About This Book Acquire real-world set of tools for building enterprise level data science applications Surpasses the barrier of other languages in data

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

science and learn 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 basically divided into two sections.

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

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 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 builds 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.

Use PySpark to easily crush messy data at-scale and discover

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

proven techniques to create testable, immutable, and easily parallelizable Spark jobs

Key Features

- Work with large amounts of agile data using distributed datasets and in-memory caching
- Source data from all popular data hosting platforms, such as HDFS, Hive, JSON, and S3
- Employ the easy-to-use PySpark API to deploy big data Analytics for production

Book Description

Apache Spark is an open source parallel-processing framework that has been around for quite some time now. One of the many uses of Apache Spark is for data analytics applications across clustered computers. In this book, you will not only learn how to use Spark and the Python API to create high-performance analytics with big data, but also discover techniques for testing, immunizing, and parallelizing Spark jobs. You will learn how to source data from all popular data hosting platforms, including HDFS, Hive, JSON, and S3, and deal with large datasets with PySpark to gain practical big data experience. This book will help you work on prototypes on local machines and subsequently go on to handle messy data in production and at scale. This book covers installing and setting up PySpark, RDD operations, big data cleaning and wrangling, and aggregating and summarizing

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

data into useful reports. You will also learn how to implement some practical and proven techniques to improve certain aspects of programming and administration in Apache Spark. By the end of the book, you will be able to build big data analytical solutions using the various PySpark offerings and also optimize them effectively. What you will learn

Get practical big data experience while working on messy datasets

Analyze patterns with Spark SQL to improve your business intelligence

Use PySpark's interactive shell to speed up development time

Create highly concurrent Spark programs by leveraging immutability

Discover ways to avoid the most expensive operation in the Spark API: the shuffle operation

Re-design your jobs to use reduceByKey instead of groupBy

Create robust processing pipelines by testing Apache Spark jobs

Who this book is for

This book is for developers, data scientists, business analysts, or anyone who needs to reliably analyze large amounts of large-scale, real-world data. Whether you're tasked with creating your company's business intelligence function or creating great data platforms for your machine learning models, or are looking to use code to magnify the impact of your business, this book is for you.

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

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

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

analysis)

Spark in Action

Learning Spark

Data Analytics with Hadoop

Practical Examples in Apache Spark and Neo4j

An Introduction for Data Scientists

Data Algorithms with Spark

Lightning-Fast Big Data Analysis

Understand the complexities of modern-day data engineering platforms and explore strategies to deal with them with the help of use case scenarios led by an industry expert in big data Key Features Become well-versed with the core concepts of Apache Spark and Delta Lake for building data platforms Learn how to ingest, process, and analyze data that can be later used for training machine learning models Understand how to operationalize data models in production using curated data Book Description In the world of ever-changing data and schemas, it is important to build data pipelines that can auto-adjust to changes. This book will help you build scalable data platforms that managers, data scientists, and data analysts can rely on. Starting with an introduction to data engineering, along with its key concepts and architectures, this book will show you how to use Microsoft Azure Cloud services effectively for data

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

engineering. You'll cover data lake design patterns and the different stages through which the data needs to flow in a typical data lake. Once you've explored the main features of Delta Lake to build data lakes with fast performance and governance in mind, you'll advance to implementing the lambda architecture using Delta Lake. Packed with practical examples and code snippets, this book takes you through real-world examples based on production scenarios faced by the author in his 10 years of experience working with big data. Finally, you'll cover data lake deployment strategies that play an important role in provisioning the cloud resources and deploying the data pipelines in a repeatable and continuous way. By the end of this data engineering book, you'll know how to effectively deal with ever-changing data and create scalable data pipelines to streamline data science, ML, and artificial intelligence (AI) tasks. What you will learn Discover the challenges you may face in the data engineering world Add ACID transactions to Apache Spark using Delta Lake Understand effective design strategies to build enterprise-grade data lakes Explore architectural and design patterns for building efficient data ingestion pipelines Orchestrate a data pipeline for preprocessing data using Apache Spark and Delta Lake APIs Automate deployment and monitoring of data pipelines in production Get to grips with securing, monitoring, and managing data pipelines models efficiently Who this book is for This book is

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

for aspiring data engineers and data analysts who are new to the world of data engineering and are looking for a practical guide to building scalable data platforms. If you already work with PySpark and want to use Delta Lake for data engineering, you'll find this book useful. Basic knowledge of Python, Spark, and SQL is expected.

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

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

genomics data Learn how to use and apply feature engineering in ML algorithms Understand and use practical and pragmatic data design patterns 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 Harness the power of Scala to program Spark and analyze tonnes of data in the

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

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

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

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.

A Practitioner's Guide to Using Spark for Large Scale Data Analysis
Covers XGBoost, LightGBM, Spark NLP, Distributed Deep Learning with

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

Keras, and More

Analyze large datasets and discover techniques for testing, immunizing, and parallelizing Spark jobs

Big Data Analytics with Spark

Scala and Spark for Big Data Analytics

With Machine Learning, Deep Learning and NLP Examples

Combine advanced analytics including Machine Learning, Deep Learning

Neural Networks and Natural Language Processing with modern scalable

technologies including Apache Spark to derive actionable insights from Big

Data in real-time Key Features Make a hands-on start in the fields of Big Data,

Distributed Technologies and Machine Learning Learn how to design, develop

and interpret the results of common Machine Learning algorithms Uncover

hidden patterns in your data in order to derive real actionable insights and

business value Book Description Every person and every organization in the

world manages data, whether they realize it or not. Data is used to describe

the world around us and can be used for almost any purpose, from analyzing

consumer habits to fighting disease and serious organized crime. Ultimately,

we manage data in order to derive value from it, and many organizations

around the world have traditionally invested in technology to help process their

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

data faster and more efficiently. But we now live in an interconnected world driven by mass data creation and consumption where data is no longer rows and columns restricted to a spreadsheet, but an organic and evolving asset in its own right. With this realization comes major challenges for organizations: how do we manage the sheer size of data being created every second (think not only spreadsheets and databases, but also social media posts, images, videos, music, blogs and so on)? And once we can manage all of this data, how do we derive real value from it? The focus of Machine Learning with Apache Spark is to help us answer these questions in a hands-on manner. We introduce the latest scalable technologies to help us manage and process big data. We then introduce advanced analytical algorithms applied to real-world use cases in order to uncover patterns, derive actionable insights, and learn from this big data. What you will learn

- Understand how Spark fits in the context of the big data ecosystem
- Understand how to deploy and configure a local development environment using Apache Spark
- Understand how to design supervised and unsupervised learning models
- Build models to perform NLP, deep learning, and cognitive services using Spark ML libraries
- Design real-time machine learning pipelines in Apache Spark
- Become familiar with advanced techniques for processing a large volume of data by applying machine learning

File Type PDF Advanced Analytics With Spark Patterns For Learning From Data At Scale

algorithmsWho this book is for This book is aimed at Business Analysts, Data Analysts and Data Scientists who wish to make a hands-on start in order to take advantage of modern Big Data technologies combined with Advanced Analytics.