Download Ebook Beginning Apache Pig Big Data Processing Made Easy Read Pdf Free

Spark: The Definitive Guide *Beginning Apache Pig* **Data Processing Learning Spark** Learning Spark *Hadoop: The Definitive Guide* R for Data Science **Fast Data Processing with Spark 2** *Big Data Processing with Apache Spark* **Quantitative Data Processing in Scanning Probe Microscopy** **Mastering Spark with R** *Python for Data Analysis* **Data Processing Handbook for Complex Biological Data Sources** Data Analysis with Python and PySpark **Health Data Processing** **The Enterprise Big Data Lake** *Data Processing and Modeling with Hadoop* Streaming Systems Stream Processing with Apache Spark *Doing Data Science* *The Self-Service Data Roadmap* *Kafka: The Definitive Guide* **Data Analytics with Hadoop** Environmental Data Analysis with MatLab Data Science for Business *Head First Data Analysis* **Tensors for Data Processing** *Spark in Action* *High Performance Spark* **Apache Spark Implementation on IBM z/OS** **Handbook of Big Geospatial Data** *Data Processing and Reconciliation for Chemical Process Operations* **PySpark Cookbook** **Practical Data Science with Hadoop and Spark** *Frank Kane's Taming Big Data with Apache Spark and Python* **DATA PROCESSING MADE SIMPLE.** **Designing Data-Intensive Applications** **Advanced Analytics with Spark** *Stream Processing with Apache Flink* **Big Data Processing Using Spark in Cloud**

Quantitative Data Processing in Scanning Probe Microscopy May 10 2022 Quantitative Data Processing in Scanning Probe Microscopy: SPM Applications for Nanometrology, Second Edition describes the recommended practices for measurements and data processing for various SPM techniques, also discussing associated numerical techniques and recommendations for further reading for particular physical quantities measurements. Each chapter has been revised and updated for this new edition to reflect the progress that has been made in SPM techniques in recent years. New features for this edition include more step-by-step examples, better sample data and more links to related documentation in open source software. Scanning Probe Microscopy (SPM) techniques have the potential to produce information on various local physical properties. Unfortunately, there is still a large gap between what is measured by commercial devices and what could be considered as a quantitative result. This book determines to educate and close that gap. Associated data sets can be downloaded from http://gwyddion.net/qspm/ Features step-by-step guidance to aid readers in progressing from a general understanding of SPM principles to a greater mastery of complex data measurement techniques Includes a focus on metrology aspects of measurements, arming readers with a solid grasp of instrumentation and measuring methods accuracy Worked examples show quantitative data processing for different SPM analytical techniques

Learning Spark Nov 16 2022 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 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

Stream Processing with Apache Flink Nov 11 2019 Get started with Apache Flink, the open source framework that powers some of the world's largest stream processing applications. With this practical book, you'll explore the fundamental concepts of parallel stream processing and discover how this technology differs from traditional batch data processing. Longtime Apache Flink committers Fabian Hueske and Vasia Kalavri show you how to implement scalable streaming applications with Flink's DataStream API and continuously run and maintain these applications in operational environments. Stream processing is ideal for many use cases, including low-latency ETL, streaming analytics, and real-time dashboards as well as fraud detection, anomaly detection, and alerting. You can process continuous data of any kind, including user interactions, financial transactions, and IoT data, as soon as you generate them. Learn concepts and challenges of distributed stateful stream processing Explore Flink's system architecture, including its event-time processing mode and fault-tolerance model Understand the fundamentals and building blocks of the DataStream API, including its time-based and statefuloperators Read data from and write data to external systems with exactly-once consistency Deploy and configure Flink clusters Operate continuously running streaming applications

Mastering Spark with R Apr 09 2022 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

Streaming Systems Sep 02 2021 Streaming data is a big deal in big data these days. As more and more businesses seek to tame the massive unbounded data sets that pervade our world, streaming systems have finally reached a level of maturity sufficient for mainstream adoption. With this practical guide, data engineers, data scientists, and developers will learn how to work with streaming data in a conceptual and platform-agnostic way. Expanded from Tyler Akidau's popular blog posts "Streaming 101" and "Streaming 102", this book takes you from an introductory level to a nuanced understanding of the what, where, when, and how of processing real-time data streams. You'll also dive deep into watermarks and exactly-once processing with co-authors Slava Chernyak and Reuven Lax. You'll explore: How streaming and batch data processing patterns compare The core principles and concepts behind robust out-of-order data processing How watermarks track progress and completeness in infinite datasets How exactly-once data processing techniques ensure correctness How the concepts of streams and tables form the foundations of both batch and streaming data processing The practical motivations behind a powerful persistent state mechanism, driven by a real-world example How time-varying relations provide a link between stream processing and the world of SQL and relational algebra

High Performance Spark Sep 21 2020 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 hours. 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 for functionality and performance when applying suggested improvements Using Spark MLlib and Spark ML machine learning libraries Spark's Streaming components and external community packages

Data Processing and Reconciliation for Chemical Process Operations Jun 18 2020 Computer techniques have made online measurements available at every sampling period in a chemical process. However, measurement errors are introduced that require suitable techniques for data reconciliation and improvements in accuracy. Reconciliation of process data and reliable monitoring are essential to decisions about possible system modifications (optimization and control procedures), analysis of equipment performance, design of the monitoring system itself, and general management planning. While the reconciliation of the process data has been studied for more than 20 years, there is no single source providing a unified approach to the area with instructions on implementation. Data Processing and Reconciliation for Chemical Process Operations is that source. Competitiveness on the world market as well as increasingly stringent environmental and product safety regulations have increased the need for the chemical industry to introduce such fast and low cost improvements in process operations. Introduces the first unified approach to this important field Bridges theory and practice through numerous worked examples and industrial case studies Provides a highly readable account of all aspects of data classification and reconciliation Presents the reader with material, problems, and directions for further study

Data Processing and Modeling with Hadoop Oct 03 2021 Understand data in a simple way using a data lake. KEY FEATURES ? In-depth practical demonstration of Hadoop/Yarn concepts with numerous examples. ? Includes graphical illustrations and visual explanations for Hadoop commands and parameters. ? Includes details of dimensional modeling and Data Vault modeling. ? Includes details of how to create and define a structure to a data lake. DESCRIPTION The book 'Data Processing and Modeling with Hadoop' explains how a distributed system works and its benefits in the big data era in a straightforward and clear manner. After reading the book, you will be able to plan and organize projects involving a massive amount of data. The book describes the standards and technologies that aid in data management and compares them to other technology business standards. The reader receives practical guidance on how to segregate and separate data into zones, as well as how to develop a model that can aid in data evolution. It discusses security and the measures that are utilized to reduce the impact of security. Self-service analytics, Data Lake, Data Vault 2.0, and Data Mesh are discussed in the book. After reading this book, the reader will have a thorough understanding of how to structure a data lake, as well as the ability to plan, organize, and carry out the implementation of a data-driven business with full governance and security. WHAT YOU WILL LEARN ? Learn the basics of components to the Hadoop Ecosystem. ? Understand the structure, files, and zones of a Data Lake. ? Learn to implement the security part of the Hadoop Ecosystem. ? Learn to work with the Data Vault 2.0 modeling. ? Learn to develop a strategy to define good governance. ? Learn new tools to work with Data and Big Data WHO THIS BOOK IS FOR This book caters to big data developers, technical specialists, consultants, and students who want to build good proficiency in big data. Knowing basic SQL concepts, modeling, and development would be good, although not mandatory. TABLE OF CONTENTS 1. Understanding the Current Moment 2. Defining the Zones 3. The Importance of Modeling 4. Massive Parallel Processing 5. Doing ETL/ELT 6. A Little Governance 7. Talking About Security 8. What Are the Next Steps?

Health Data Processing Dec 05 2021 Health Data Processing: Systemic Approaches focuses on the design of health information systems and touches on the main themes of medical informatics and public health. The book is written for health professionals in practice or training, and is especially useful for decision-makers or future decision-makers in the field of health information systems. Users will find sections on the question of reusing data for other purposes, protection of individual liberties that this data and technologies make more acute, and the irruption of large masses of genetic data and its related problems. This book develops the methodological and conceptual aspects related to these issues. Proposes a methodology for the development of health information systems for the better use of digital technologies Illustrates a systemic, transversal, conceptual vision that supports the complex reality of the healthcare world, where the interoperability of agents (professionals and software) is central Discusses the reuse of resources of data for knowledge improvement, health security and public health

Data Analysis with Python and PySpark Jan 06 2022 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.

Hadoop: The Definitive Guide Sep 14 2022 Ready to unlock the power of your data? With this comprehensive guide, you'll learn how to build and maintain reliable, scalable, distributed systems with Apache Hadoop. This book is ideal for programmers looking to analyze datasets of any size, and for administrators who want to set up and run Hadoop clusters. You'll find illuminating case studies that demonstrate how Hadoop is used to solve specific problems. This third edition covers recent changes to Hadoop, including material on the new MapReduce API, as well as MapReduce 2 and its more flexible execution model (YARN). Store large datasets with the Hadoop Distributed File System (HDFS) Run distributed computations with MapReduce Use Hadoop's data and I/O building blocks for compression, data integrity, serialization (including Avro), and persistence Discover common pitfalls and advanced features for writing real-world MapReduce programs Design, build, and administer a dedicated Hadoop cluster—or run Hadoop in the cloud Load data from relational databases into HDFS, using Sqoop Perform large-scale data processing with the Pig query language Analyze datasets with Hive, Hadoop's data warehousing system Take advantage of HBase for structured and semi-structured data, and ZooKeeper for building distributed systems

Apache Spark Implementation on IBM z/OS Aug 21 2020 The term big data refers to extremely large sets of data that are analyzed to reveal insights, such as patterns, trends, and associations. The algorithms that analyze this data to provide these insights must extract value from a wide range of data sources, including business data and live, streaming, social media data. However, the real value of these insights comes from their timeliness. Rapid delivery of insights enables anyone (not only data scientists) to make effective decisions, applying deep intelligence to every enterprise application. Apache Spark is an integrated analytics framework and runtime to accelerate and simplify algorithm development, depoyment, and realization of business insight from analytics. Apache Spark on IBM® z/OS® puts the open source engine, augmented with unique differentiated features, built specifically for data science, where big data resides. This IBM Redbooks® publication describes the installation and configuration of IBM z/OS Platform for Apache Spark for field teams and clients. Additionally, it includes examples of business analytics scenarios.

Python for Data Analysis Mar 08 2022 Get complete instructions for manipulating, processing, cleaning, and crunching datasets in Python. Updated for Python 3.6, the second edition of this hands-on guide is packed with practical case studies that show you how to solve a broad set of data analysis problems effectively. You'll learn the latest versions of pandas, NumPy, IPython, and Jupyter in the process. Written by Wes McKinney, the creator of the Python pandas project, this book is a practical, modern introduction to data science tools in Python. It's ideal for analysts new to Python and for Python programmers new to data science and scientific computing. Data files and related material are available on GitHub. Use the IPython shell and Jupyter notebook for exploratory computing Learn basic and advanced features in NumPy (Numerical Python) Get started with data analysis tools in the pandas library Use flexible tools to load, clean, transform, merge, and reshape data Create informative visualizations with matplotlib Apply the pandas groupby facility to slice, dice, and summarize datasets Analyze and manipulate regular and irregular time series data Learn how to solve real-world data analysis problems with thorough, detailed examples

Advanced Analytics with Spark Dec 13 2019 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

Frank Kane's Taming Big Data with Apache Spark and Python Mar 16 2020 Frank Kane's hands-on Spark training course, based on his bestselling Taming Big Data with Apache Spark and Python video, now available in a book. Understand and analyze large data sets using Spark on a single system or on a cluster. About This Book Understand how Spark can be distributed across computing clusters Develop and run Spark jobs efficiently using Python A hands-on tutorial by Frank Kane with over 15 real-world examples teaching you Big Data processing with Spark Who This Book Is For If you are a data scientist or data analyst who wants to learn Big Data processing using Apache Spark and Python, this book is for you. If you have some programming experience in Python, and want to learn how to process large amounts of data using Apache Spark, Frank Kane's Taming Big Data with Apache Spark and Python will also help you. What You Will Learn Find out how you can identify Big Data problems as Spark problems Install and run Apache Spark on your computer or on a cluster Analyze large data sets across many CPUs using Spark's Resilient Distributed Datasets Implement machine learning on Spark using the MLlib library Process continuous streams of data in real time using the Spark streaming module Perform complex network analysis using Spark's GraphX library Use Amazon's Elastic MapReduce service to run your Spark jobs on a cluster In Detail Frank Kane's Taming Big Data with Apache Spark and Python is your companion to learning Apache Spark in a hands-on manner. Frank will start you off by teaching you how to set up Spark on a single system or on a cluster, and you'll soon move on to analyzing large data sets using Spark RDD, and developing and running effective Spark jobs quickly using Python. Apache Spark has emerged as the next big thing in the Big Data domain – quickly rising from an ascending technology to an established superstar in just a matter of years. Spark allows you to quickly extract actionable insights from large amounts of data, on a real-time basis, making it an essential tool in many modern businesses. Frank has packed this book with over 15 interactive, fun-filled examples relevant to the real world, and he will empower you to understand the Spark ecosystem and implement production-grade real-time Spark projects with ease. Style and approach Frank Kane's Taming Big Data with Apache Spark and Python is a hands-on tutorial with over 15 real-world examples carefully explained by Frank in a step-by-step manner. The examples vary in complexity, and you can move through them at your own pace.

Data Processing Handbook for Complex Biological Data Sources Feb 07 2022 Data Processing Handbook for Complex Biological Data provides relevant and to the point content for those who need to understand the different types of biological data and the techniques to process and interpret them. The book includes feedback the editor received from students studying at both undergraduate and graduate levels, and from her peers. In order to succeed in data processing for biological data sources, it is necessary to master the type of data and general methods and tools for modern data processing. For instance, many labs follow the path of interdisciplinary studies and get their data validated by several methods. Researchers at those labs may not perform all the techniques themselves, but either in collaboration or through outsourcing, they make use of a range of them, because, in the absence of cross validation using different techniques, the chances for acceptance of an article for publication in high profile journals is weakened. Explains how to interpret enormous amounts of data generated using several experimental approaches in simple terms, thus relating biology and physics at the atomic level Presents sample data files and explains the usage of equations and web servers cited in research articles to extract useful information from their own biological data Discusses, in detail, raw data files, data processing strategies, and the web based sources relevant for data processing

Learning Spark Oct 15 2022 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

Environmental Data Analysis with MatLab Feb 24 2021 "Environmental Data Analysis with MatLab" is for students and researchers working to analyze real data sets in the environmental sciences. One only has to consider the global warming debate to realize how critically important it is to be able to derive clear conclusions from often-noisy data drawn from a broad range of sources. This book teaches the basics of the underlying theory of data analysis, and then reinforces that knowledge with carefully chosen, realistic scenarios. MatLab, a commercial data processing environment, is used in these scenarios; significant content is devoted to teaching how it can be effectively used in an environmental data analysis setting. The book, though written in a self-contained way, is supplemented with data sets and MatLab scripts that can be used as a data analysis tutorial. It is well written and outlines a clear learning path for researchers and students. It uses real world environmental examples and case studies. It has MatLab software for application in a readily-available software environment. Homework problems help user follow up upon case studies with homework that expands them.

Stream Processing with Apache Spark Aug 01 2021 Before you can build analytics tools to gain quick insights, you first need to know how to process data in real time. With this practical guide, developers familiar with Apache Spark will learn how to put this in-memory framework to use for streaming data. You'll discover how Spark enables you to write streaming jobs in almost the same way you write batch jobs. Authors Gerard Maas and François Garillot help you explore the theoretical underpinnings of Apache Spark. This comprehensive guide features two sections that compare and contrast the streaming APIs Spark now supports: the original Spark Streaming library and the newer Structured Streaming API. Learn fundamental stream processing concepts and examine different streaming architectures Explore Structured Streaming through practical examples; learn different aspects of stream processing in detail Create and operate streaming jobs and applications with Spark Streaming; integrate Spark Streaming with other Spark APIs Learn advanced Spark Streaming techniques, including approximation algorithms and machine learning algorithms Compare Apache Spark to other stream processing projects, including Apache Storm, Apache Flink, and Apache Kafka Streams

R for Data Science Aug 13 2022 Learn how to use R to turn raw data into insight, knowledge, and understanding. This book introduces you to R, RStudio, and the tidyverse, a collection of R packages designed to work together to make data science fast, fluent, and fun. Suitable for readers with no previous programming experience, R for Data Science is designed to get you doing data science as quickly as possible. Authors Hadley Wickham and Garrett Grolemund guide you through the steps of importing, wrangling, exploring, and modeling your data and communicating the results. You'll get a complete, big-picture understanding of the data science cycle, along with basic tools you need to manage the details. Each section of the book is paired with exercises to help you practice what you've learned along the way. You'll learn how to: Wrangle—transform your datasets into a form convenient for analysis Program—learn powerful R tools for solving data problems with greater clarity and ease Explore—examine your data, generate hypotheses, and quickly test them Model—provide a low-dimensional summary that captures true "signals" in your dataset Communicate—learn R Markdown for integrating prose, code, and results

Handbook of Big Geospatial Data Jul 20 2020 This handbook covers a wide range of topics related to the collection, processing, analysis, and use of geospatial data in their various forms. This handbook provides an overview of how spatial computing technologies for big data can be organized and implemented to solve real-world problems. Diverse subdomains ranging from indoor mapping and navigation over trajectory computing to earth observation from space, are also present in this handbook. It combines fundamental contributions focusing on spatio-textual analysis, uncertain databases, and spatial statistics with application examples such as road network detection or colocation detection using GPUs. In summary, this handbook gives an essential introduction and overview of the rich field of spatial information science and big geospatial data. It introduces three different perspectives, which together define the field of big geospatial data: a societal, governmental, and governance perspective. It discusses questions of how the acquisition, distribution and exploitation of big geospatial data must be organized both on the scale of companies and countries. A second perspective is a theory-oriented set of contributions on arbitrary spatial data with contributions introducing into the exciting field of spatial statistics or into uncertain databases. A third perspective is taking a very practical perspective to big geospatial data, ranging from chapters that describe how big geospatial data infrastructures can be implemented and how specific applications can be implemented on top of big geospatial data. This would include for example, research in historic map data, road network extraction, damage estimation from remote sensing imagery, or the analysis of spatio-textual collections and social media. This multi-disciplinary approach makes the book unique. This handbook can be used as a reference for undergraduate students, graduate students and researchers focused on big geospatial data. Professionals can use this book, as well as practitioners facing big collections of geospatial data.

Spark in Action Oct 23 2020 Summary The Spark distributed data processing platform provides an easy-to-implement tool for ingesting, streaming, and processing data from any source. In Spark in Action, Second Edition, you'll learn to take advantage of Spark's core features and incredible processing speed, with applications including real-time computation, delayed evaluation, and machine learning. Spark skills are a hot commodity in enterprises worldwide, and with Spark's powerful and flexible Java APIs, you can reap all the benefits without first learning Scala or Hadoop. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Analyzing enterprise data starts by reading, filtering, and merging files and streams from many sources. The Spark data processing engine handles this varied volume like a champ, delivering speeds 100 times faster than Hadoop systems. Thanks to SQL support, an intuitive interface, and a straightforward multilanguage API, you can use Spark without learning a complex new ecosystem. About the book Spark in Action, Second Edition, teaches you to create end-to-end analytics applications. In this entirely new book, you'll learn from interesting Java-based examples, including a complete data pipeline for processing NASA satellite data. And you'll discover Java, Python, and Scala code samples hosted on GitHub that you can explore and adapt, plus appendixes that give you a cheat sheet for installing tools and understanding Spark-specific terms. What's inside Writing Spark applications in Java Spark application architecture Ingestion through files, databases, streaming, and Elasticsearch Querying distributed datasets with Spark SQL About the reader This book does not assume previous experience with Spark, Scala, or Hadoop. About the author Jean-Georges Perrin is an experienced data and software architect. He is France's first IBM Champion and has been honored for 12 consecutive years. Table of Contents PART 1 - THE THEORY CRIPPLED BY AWESOME EXAMPLES 1 So, what is Spark, anyway? 2 Architecture and flow 3 The majestic role of the dataframe 4 Fundamentally lazy 5 Building a simple app for deployment 6 Deploying your simple app PART 2 - INGESTION 7 Ingestion from files 8 Ingestion from databases 9 Advanced ingestion: finding data sources and building your own 10 Ingestion through structured streaming PART 3 - TRANSFORMING YOUR DATA 11 Working with SQL 12 Transforming your data 13 Transforming entire documents 14 Extending transformations with user-defined functions 15 Aggregating your data PART 4 - GOING FURTHER 16 Cache and checkpoint: Enhancing Spark's performances 17 Exporting data and building full data pipelines 18 Exploring deployment

PySpark Cookbook May 18 2020 Combine the power of Apache Spark and Python to build effective big data applications Key Features Perform effective data processing, machine learning, and analytics using PySpark Overcome challenges in developing and deploying Spark solutions using Python Explore recipes for efficiently combining Python and Apache Spark to process data Book Description Apache Spark is an open source framework for efficient cluster computing with a strong interface for data parallelism and fault tolerance. The PySpark Cookbook presents effective and time-saving recipes for leveraging the power of Python and putting it to use in the Spark ecosystem. You'll start by learning the Apache Spark architecture and how to set up a Python environment for Spark. You'll then get familiar with the modules available in PySpark and start using them effortlessly. In addition to this, you'll discover how to abstract data with RDDs and DataFrames, and understand the streaming capabilities of PySpark. You'll then move on to using ML and MLlib in order to solve any problems related to the machine learning capabilities of PySpark and use GraphFrames to solve graph-processing problems. Finally, you will explore how to deploy your applications to the cloud using the spark-submit command. By the end of this book, you will be able to use the Python API for Apache Spark to solve any problems associated with building data-intensive applications. What you will learn Configure a local instance of PySpark in a virtual environment Install and configure Jupyter in local and multi-node environments Create DataFrames from JSON and a dictionary using pyspark.sql Explore regression and clustering models available in the ML module Use DataFrames to transform data used for modeling Connect to PubNub and perform aggregations on streams Who this book is for The PySpark Cookbook is for you if you are a Python developer looking for hands-on recipes for using the Apache Spark 2.x ecosystem in the best possible way. A thorough understanding of Python (and some familiarity with Spark) will help you get the best out of the book.

Doing Data Science Jun 30 2021 Now that people are aware that data can make the difference in an election or a business model, data science as an occupation is gaining ground. But how can you get started working in a wide-ranging, interdisciplinary field that's so clouded in hype? This insightful book, based on Columbia University's Introduction to Data Science class, tells you what you need to know. In many of these chapter-long lectures, data scientists from companies such as Google, Microsoft, and eBay share new algorithms, methods, and models by presenting case studies and the code they use. If you're familiar with linear algebra, probability, and statistics, and have programming experience, this book is an ideal introduction to data science. Topics include: Statistical inference, exploratory data analysis, and the data science process Algorithms Spam filters, Naive Bayes, and data wrangling Logistic regression Financial modeling Recommendation engines and causality Data visualization Social networks and data journalism Data engineering, MapReduce, Pregel, and Hadoop Doing Data Science is collaboration between course instructor Rachel Schutt, Senior VP of Data Science at News Corp, and data science consultant Cathy O'Neil, a senior data scientist at Johnson Research Labs, who attended and blogged about the course.

The Enterprise Big Data Lake Nov 04 2021 The data lake is a daring new approach for harnessing the power of big data technology and providing convenient self-service capabilities. But is it right for your company? This book is based on discussions with practitioners and executives from more than a hundred organizations, ranging from data-driven companies such as Google, LinkedIn, and Facebook, to governments and traditional corporate enterprises. You'll learn what a data lake is, why enterprises need one, and how to build one successfully with the best practices in this book. Alex Gorelik, CTO and founder of Waterline Data, explains why old systems and processes can no longer support data needs in the enterprise. Then, in a collection of essays about data lake implementation, you'll examine data lake initiatives, analytic projects, experiences, and best practices from data experts working in various industries. Get a succinct introduction to data warehousing, big data, and data science Learn various paths enterprises take to build a data lake Explore how to build a self-service model and best practices for providing analysts access to the data Use different methods for architecting your data lake Discover ways to implement a data lake from experts in different industries

Tensors for Data Processing Nov 23 2020 Tensors for Data Processing: Theory, Methods and Applications presents both classical and state-of-the-art methods on tensor computation for data processing, covering computation theories, processing methods, computing and engineering applications, with an emphasis on techniques for data processing. This reference is ideal for students, researchers and industry developers who want to understand and use tensor-based data processing theories and methods. As a higher-order generalization of a matrix, tensor-based processing can avoid multi-linear data structure loss that occurs in classical matrix-based data processing methods. This move from matrix to tensors is beneficial for many diverse application areas, including signal processing, computer science, acoustics, neuroscience, communication, medical engineering, seismology, psychometric, chemometrics, biometric, quantum physics and quantum chemistry. Provides a complete reference on classical and state-of-the-art tensor-based methods for data processing Includes a wide range of applications from different disciplines Gives guidance for their application

Big Data Processing Using Spark in Cloud Oct 11 2019 The book describes the emergence of big data technologies and the role of Spark in the entire big data stack. It compares Spark and Hadoop and identifies the shortcomings of Hadoop that have been overcome by Spark. The book mainly focuses on the in-depth architecture of Spark and our understanding of Spark RDDs and how RDD complements big data's immutable nature, and solves it with lazy evaluation, cacheable and type inference. It also addresses advanced topics in Spark, starting with the basics of Scala and the core Spark framework, and exploring Spark data frames, machine learning using Mllib, graph analytics using Graph X and real-time processing with Apache Kafka, AWS Kenisis, and Azure Event Hub. It then goes on to investigate Spark using PySpark and R. Focusing on the current big data stack, the book examines the interaction with current big data tools, with Spark being the core processing layer for all types of data. The book is intended for data engineers and scientists working on massive datasets and big data technologies in the cloud. In addition to industry professionals, it is helpful for aspiring data processing professionals and students working in big data processing and cloud computing environments.

Big Data Processing with Apache Spark Jun 11 2022 Apache Spark is a popular open-source big-data processing framework thatÕs built around speed, ease of use, and unified distributed computing architecture. Not only it supports developing applications in different languages like Java, Scala, Python, and R, itÕs also hundred times faster in memory and ten times faster even when running on disk compared to traditional data processing frameworks. Whether you are currently working on a big data project or interested in learning more about topics like machine learning, streaming data processing, and graph data analytics, this book is for you. You can learn about Apache Spark and develop Spark programs for various use cases in big data analytics using the code examples provided. This book covers all the libraries in Spark ecosystem: Spark Core, Spark SQL, Spark Streaming, Spark ML, and Spark GraphX.

Kafka: The Definitive Guide Apr 28 2021 Every enterprise application creates data, whether it's log messages, metrics, user activity, outgoing messages, or something else. And how to move all of this data becomes nearly as important as the data itself. If you're an application architect, developer, or production engineer new to Apache Kafka, this practical guide shows you how to use this open source streaming platform to handle real-time data feeds. Engineers from Confluent and LinkedIn who are responsible for developing Kafka explain how to deploy production Kafka clusters, write reliable event-driven microservices, and build scalable stream-processing applications with this platform. Through detailed examples, you'll learn Kafka's design principles, reliability guarantees, key APIs, and architecture details, including the replication protocol, the controller, and the storage layer. Understand publish-subscribe messaging and how it fits in the big data ecosystem. Explore Kafka producers and consumers for writing and reading messages Understand Kafka patterns and use-case requirements to ensure reliable data delivery Get best practices for building data pipelines and applications with Kafka Manage Kafka in production, and learn to perform monitoring, tuning, and maintenance tasks Learn the most critical metrics among Kafka's operational measurements Explore how Kafka's stream delivery capabilities make it a perfect source for stream processing systems

Practical Data Science with Hadoop and Spark Apr 16 2020 The Complete Guide to Data Science with Hadoop—For Technical Professionals, Businesspeople, and Students Demand is soaring for professionals who can solve real data science problems with Hadoop and Spark. Practical Data Science with Hadoop® and Spark is your complete guide to doing just that. Drawing on immense experience with Hadoop and big data, three leading experts bring together everything you need: high-level concepts, deep-dive techniques, real-world use cases, practical applications, and hands-on tutorials. The authors introduce the essentials of data science and the modern Hadoop ecosystem, explaining how Hadoop and Spark have evolved into an effective platform for solving data science problems at scale. In addition to comprehensive application coverage, the authors also provide useful guidance on the important steps of data ingestion, data munging, and visualization. Once the groundwork is in place, the authors focus on specific applications, including machine learning, predictive modeling for sentiment analysis, clustering for document analysis, anomaly detection, and natural language processing (NLP). This guide provides a strong technical foundation for those who want to do practical data science, and also presents business-driven guidance on how to apply Hadoop and Spark to optimize ROI of data science initiatives. Learn What data science is, how it has evolved, and how to plan a data science career How data volume, variety, and velocity shape data science use cases Hadoop and its ecosystem, including HDFS, MapReduce, YARN, and Spark Data importation with Hive and Spark Data quality, preprocessing, preparation, and modeling Visualization: surfacing insights from huge data sets Machine learning: classification, regression, clustering, and anomaly detection Algorithms and Hadoop tools for predictive modeling Cluster analysis and similarity functions Large-scale anomaly detection NLP: applying data science to human language

Data Science for Business Jan 26 2021 Written by renowned data science experts Foster Provost and Tom Fawcett, Data Science for Business introduces the fundamental principles of data science, and walks you through the "data-analytic thinking" necessary for extracting useful knowledge and business value from the data you collect. This guide also helps you understand the many data-mining techniques in use today. Based on an MBA course Provost has taught at New York University over the past ten years, Data Science for Business provides examples of real-world business problems to illustrate these principles. You'll not only learn how to improve communication between business stakeholders and data scientists, but also how participate intelligently in your company's data science projects. You'll also discover how to think data-analytically, and fully appreciate how data science methods can support business decision-making. Understand how data science fits in your organization—and how you can use it for competitive advantage Treat data as a business asset that requires careful investment if you're to gain real value Approach business problems data-analytically, using the data-mining process to gather good data in the most appropriate way Learn general concepts for actually extracting knowledge from data Apply data science principles when interviewing data science job candidates

Data Processing Dec 17 2022 Data Processing: Made Simple, Second Edition presents discussions of a number of trends and developments in the world of commercial data processing. The book covers the rapid growth of micro- and mini-computers for both home and office use; word processing and the 'automated office'; the advent of distributed data processing; and the continued growth of database-oriented systems. The text also discusses modern digital computers; fundamental computer concepts; information and data processing requirements of commercial organizations; and the historical perspective of the computer industry. The computer hardware and software and the development and implementation of a computer system are considered. The book tackles careers in data processing; the tasks carried out by the data processing department; and the way in which the data processing department fits in with the rest of the organization. The text concludes by examining some of the problems of running a data processing department, and by suggesting some possible solutions. Computer science students will find the book invaluable.

Fast Data Processing with Spark 2 Jul 12 2022 Learn how to use Spark to process big data at speed and scale for sharper analytics. Put the principles into practice for faster, slicker big data projects. About This Book A quick way to get started with Spark – and reap the rewards From analytics to engineering your big data architecture, we've got it covered Bring your Scala and Java knowledge – and put it to work on new and exciting problems Who This Book Is For This book is for developers with little to no knowledge of Spark, but with a background in Scala/Java programming. It's recommended that you have experience in dealing and working with big data and a strong interest in data science. What You Will Learn Install and set up Spark in your cluster Prototype distributed applications with Spark's interactive shell Perform data wrangling using the new DataFrame APIs Get to know the different ways to interact with Spark's distributed representation of data (RDDs) Query Spark with a SQL-like query syntax See how Spark works with big data Implement machine learning systems with highly scalable algorithms Use R, the popular statistical language, to work with Spark Apply interesting graph algorithms and graph processing with GraphX In Detail When people want a way to process big data at speed, Spark is invariably the solution. With its ease of development (in comparison to the relative complexity of Hadoop), it's unsurprising that it's becoming popular with data analysts and engineers everywhere. Beginning with the fundamentals, we'll show you how to get set up with Spark with minimum fuss. You'll then get to grips with some simple APIs before investigating machine learning and graph processing – throughout we'll make sure you know exactly how to apply your knowledge. You will also learn how to use the Spark shell, how to load data before finding out how to build and run your own Spark applications. Discover how to manipulate your RDD and get stuck into a range of DataFrame APIs. As if that's not enough, you'll also learn some useful Machine Learning algorithms with the help of Spark MLlib and integrating Spark with R. We'll also make sure you're confident and prepared for graph processing, as you learn more about the GraphX API. Style and approach This book is a basic, step-by-step tutorial that will help you take advantage of all that Spark has to offer.

Spark: The Definitive Guide Feb 19 2023 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 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

Beginning Apache Pig Jan 18 2023 Learn to use Apache Pig to develop lightweight big data applications easily and quickly. This book shows you many optimization techniques and covers every context where Pig is used in big data analytics. Beginning Apache Pig shows you how Pig is easy to learn and requires relatively little time to develop big data applications.The book is divided into four parts: the complete features of Apache Pig; integration with other tools; how to solve complex business problems; and optimization of tools.You'll discover topics such as MapReduce and why it cannot meet every business need; the features of Pig Latin such as data types for each load, store, joins, groups, and ordering; how Pig workflows can be created; submitting Pig jobs using Hue; and working with Oozie. You'll also see how to extend the framework by writing UDFs and custom load, store, and filter functions. Finally you'll cover different optimization techniques such as gathering statistics about a Pig script, joining strategies, parallelism, and the role of data formats in good performance. What You Will Learn• Use all the features of Apache Pig• Integrate Apache Pig with other tools• Extend Apache Pig• Optimize Pig Latin code• Solve different use cases for Pig LatinWho This Book Is ForAll levels of IT professionals: architects, big data enthusiasts, engineers, developers, and big data administrators

Head First Data Analysis Dec 25 2020 A guide for data managers and analyzers shares guidelines for identifying patterns, predicting future outcomes, and presenting findings to others; drawing on current research in cognitive science and learning theory while covering such additional topics as assessing data quality, handling ambiguous information, and organizing data within market groups. Original.

Designing Data-Intensive Applications Jan 14 2020 Data is at the center of many challenges in system design today. Difficult issues need to be figured out, such as scalability, consistency, reliability, efficiency, and maintainability. In addition, we have an overwhelming variety of tools, including relational databases, NoSQL datastores, stream or batch processors, and message brokers. What are the right choices for your application? How do you make sense of all these buzzwords? In this practical and comprehensive guide, author Martin Kleppmann helps you navigate this diverse landscape by examining the pros and cons of various technologies for processing and storing data. Software keeps changing, but the fundamental principles remain the same. With this book, software engineers and architects will learn how to apply those ideas in practice, and how to make full use of data in modern applications. Peer under the hood of the systems you already use, and learn how to use and operate them more effectively Make informed decisions by identifying the strengths and weaknesses of different tools Navigate the trade-offs around consistency, scalability, fault tolerance, and complexity Understand the distributed systems research upon which modern databases are built Peek behind the scenes of major online services, and learn from their architectures

Data Analytics with Hadoop Mar 28 2021 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 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

The Self-Service Data Roadmap May 30 2021 Data-driven insights are a key competitive advantage for any industry today, but deriving insights from raw data can still take days or weeks. Most organizations can't scale data science teams fast enough to keep up with the growing amounts of data to transform. What's the answer? Self-service data. With this practical book, data engineers, data scientists, and team managers will learn how to build a self-service data science platform that helps anyone in your organization extract insights from data. Sandeep Uttamchandani provides a scorecard to track and address bottlenecks that slow down time to insight across data discovery, transformation, processing, and production. This book bridges the gap between data scientists bottlenecked by engineering realities and data engineers unclear about ways to make self-service work. Build a self-service portal to support data discovery, quality, lineage, and governance Select the best approach for each self-service capability using open source cloud technologies Tailor self-service for the people, processes, and technology maturity of your data platform Implement capabilities to democratize data and reduce time to insight Scale your self-service portal to support a large number of users within your organization

DATA PROCESSING MADE SIMPLE. Feb 13 2020

- Tensors For Data Processing
- Spark In Action
- High Performance Spark
- Apache Spark Implementation On IBM Z OS
- Handbook Of Big Geospatial Data
- Data Processing And Reconciliation For Chemical Process Operations
- PySpark Cookbook
- Practical Data Science With Hadoop And Spark
- Frank Kanes Taming Big Data With Apache Spark And Python
- DATA PROCESSING MADE SIMPLE
- Designing Data Intensive Applications
- Advanced Analytics With Spark
- Stream Processing With Apache Flink
- Big Data Processing Using Spark In Cloud