

## Practical Data Science With Hadoop And Spark

A practical guide to obtaining, transforming, exploring, and analyzing data using Python, MongoDB, and Apache Spark About This Book Learn to use various data analysis tools and algorithms to classify, cluster, visualize, simulate, and forecast your data Apply Machine Learning algorithms to different kinds of data such as social networks, time series, and images A hands-on guide to understanding the nature of data and how to turn it into insight Who This Book Is For This book is for developers who want to implement data analysis and data-driven algorithms in a practical way. It is also suitable for those without a background in data analysis or data processing. Basic knowledge of Python programming, statistics, and linear algebra is assumed. What You Will Learn Acquire, format, and visualize your data Build an image-similarity search engine Generate meaningful visualizations anyone can understand Get started with analyzing social network graphs Find out how to implement sentiment text analysis Install data analysis tools such as Pandas, MongoDB, and Apache Spark Get to grips with Apache Spark Implement machine learning algorithms such as classification or forecasting In Detail Beyond buzzwords like Big Data or Data Science, there are a great opportunities to innovate in many businesses using data analysis to get data-driven products. Data analysis involves asking many questions about data in order to discover insights and generate value for a product or a service. This book explains the basic data algorithms without the theoretical jargon, and you'll get hands-on turning data into insights using machine learning techniques. We will perform data-driven innovation processing for several types of data such as text, Images, social network graphs, documents, and time series, showing you how to implement large data processing with MongoDB and Apache Spark. Style and approach This is a hands-on guide to data analysis and data processing. The concrete examples are explained with simple code and accessible data.

Get expert guidance on architecting end-to-end data management solutions with Apache Hadoop. While many sources explain how to use various components in the Hadoop ecosystem, this practical book takes you through architectural considerations necessary to tie those components together into a complete tailored application, based on your particular use case. To reinforce those lessons, the book's second section provides detailed examples of architectures used in some of the most commonly found Hadoop applications. Whether you're designing a new Hadoop application, or planning to integrate Hadoop into your existing data infrastructure, Hadoop Application Architectures will skillfully guide you through the process. This book covers: Factors to consider when using Hadoop to store and model data Best practices for moving data in and out of the system Data processing frameworks, including MapReduce, Spark, and Hive Common Hadoop processing patterns, such as removing duplicate records and using windowing analytics Giraph, GraphX, and other tools for large graph processing on Hadoop Using workflow orchestration and scheduling tools such as Apache Oozie Near-real-time stream processing with Apache Storm, Apache Spark Streaming, and Apache Flume Architecture examples for clickstream analysis, fraud detection, and data warehousing Let Hadoop For Dummies help harness the power of your data and rein in the information overload Big data has become big

business, and companies and organizations of all sizes are struggling to find ways to retrieve valuable information from their massive data sets with becoming overwhelmed. Enter Hadoop and this easy-to-understand For Dummies guide. Hadoop For Dummies helps readers understand the value of big data, make a business case for using Hadoop, navigate the Hadoop ecosystem, and build and manage Hadoop applications and clusters. Explains the origins of Hadoop, its economic benefits, and its functionality and practical applications Helps you find your way around the Hadoop ecosystem, program MapReduce, utilize design patterns, and get your Hadoop cluster up and running quickly and easily Details how to use Hadoop applications for data mining, web analytics and personalization, large-scale text processing, data science, and problem-solving Shows you how to improve the value of your Hadoop cluster, maximize your investment in Hadoop, and avoid common pitfalls when building your Hadoop cluster From programmers challenged with building and maintaining affordable, scaleable data systems to administrators who must deal with huge volumes of information effectively and efficiently, this how-to has something to help you with Hadoop. "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"--

Get command of your organizational Big Data using the power of data science and analytics Key Features A perfect companion to boost your Big Data storing, processing, analyzing skills to help you take informed business decisions Work with the best tools such as Apache Hadoop, R, Python, and Spark for NoSQL platforms to perform massive online analyses Get expert tips on statistical inference, machine learning, mathematical modeling, and data visualization for Big Data Book Description Big Data analytics relates to the strategies used by organizations to collect, organize and analyze large amounts of data to uncover valuable business insights that otherwise cannot be analyzed through traditional systems. Crafting an enterprise-scale cost-efficient Big Data and machine learning solution to uncover insights and value from your organization's data is a challenge. Today, with hundreds of new Big Data systems, machine learning packages and BI Tools, selecting the right combination of technologies is an even greater challenge. This book will help you do that. With the help of this guide, you will be able to bridge the gap between the theoretical world of technology with the practical ground reality of building corporate Big Data and data science platforms. You will get hands-on exposure to Hadoop and Spark, build machine learning dashboards using R and R Shiny, create web-based apps using NoSQL databases such as MongoDB and even learn how to write R code for neural networks. By the end of the book, you will have a very clear and concrete understanding of what Big Data analytics means, how it drives revenues for organizations, and how you can develop your own Big Data analytics solution using different tools and methods articulated in this book. What you will learn - Get a 360-degree view into the world of Big Data, data science and machine learning - Broad range of technical and business Big Data analytics topics that caters to the interests of the technical experts as well as corporate IT executives - Get hands-on experience with industry-standard Big Data and machine learning tools such as Hadoop, Spark, MongoDB, KDB+ and R - Create production-grade machine learning BI Dashboards using R and R Shiny with step-by-step instructions - Learn how to combine open-source Big Data, machine learning and BI Tools to create low-cost business analytics applications - Understand

corporate strategies for successful Big Data and data science projects - Go beyond general-purpose analytics to develop cutting-edge Big Data applications using emerging technologies Who this book is for The book is intended for existing and aspiring Big Data professionals who wish to become the go-to person in their organization when it comes to Big Data architecture, analytics, and governance. While no prior knowledge of Big Data or related technologies is assumed, it will be helpful to have some programming experience.

Big Data Analytics with R and Hadoop is a tutorial style book that focuses on all the powerful big data tasks that can be achieved by integrating R and Hadoop. This book is ideal for R developers who are looking for a way to perform big data analytics with Hadoop. This book is also aimed at those who know Hadoop and want to build some intelligent applications over Big data with R packages. It would be helpful if readers have basic knowledge of R.

Summary Hadoop in Practice, Second Edition provides over 100 tested, instantly useful techniques that will help you conquer big data, using Hadoop. This revised new edition covers changes and new features in the Hadoop core architecture, including MapReduce 2. Brand new chapters cover YARN and integrating Kafka, Impala, and Spark SQL with Hadoop. You'll also get new and updated techniques for Flume, Sqoop, and Mahout, all of which have seen major new versions recently. In short, this is the most practical, up-to-date coverage of Hadoop available anywhere. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Book It's always a good time to upgrade your Hadoop skills!

Hadoop in Practice, Second Edition provides a collection of 104 tested, instantly useful techniques for analyzing real-time streams, moving data securely, machine learning, managing large-scale clusters, and taming big data using Hadoop. This completely revised edition covers changes and new features in Hadoop core, including MapReduce 2 and YARN. You'll pick up hands-on best practices for integrating Spark, Kafka, and Impala with Hadoop, and get new and updated techniques for the latest versions of Flume, Sqoop, and Mahout. In short, this is the most practical, up-to-date coverage of Hadoop available. Readers need to know a programming language like Java and have basic familiarity with Hadoop. What's Inside Thoroughly updated for Hadoop 2 How to write YARN applications Integrate real-time technologies like Storm, Impala, and Spark Predictive analytics using Mahout and RR Readers need to know a programming language like Java and have basic familiarity with Hadoop. About the Author Alex Holmes works on tough big-data problems. He is a software engineer, author, speaker, and blogger specializing in large-scale Hadoop projects. Table of Contents PART 1 BACKGROUND AND FUNDAMENTALS Hadoop in a heartbeat Introduction to YARN PART 2 DATA LOGISTICS Data serialization—working with text and beyond Organizing and optimizing data in HDFS Moving data into and out of Hadoop PART 3 BIG DATA PATTERNS Applying MapReduce patterns to big data Utilizing data structures and algorithms at scale Tuning, debugging, and testing PART 4 BEYOND MAPREDUCE SQL on Hadoop Writing a YARN application

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

The book provides a one-stop solution for getting into data science with Python and teaches how to extract insights from data.

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

The Example-Rich, Hands-On Guide to Data Munging with Apache Hadoop™ Data scientists spend much of their time “munging” data: handling day-to-day tasks such as data cleansing, normalization, aggregation, sampling, and transformation. These tasks are both critical and surprisingly interesting. Most important, they deepen your understanding of your data's structure and limitations: crucial insight for improving accuracy and mitigating risk in any analytical project. Now, two leading Hortonworks data scientists, Ofer Mendelevitch and Casey Stella, bring together powerful, practical insights for effective Hadoop-based data munging of large datasets. Drawing on extensive experience with advanced analytics, the authors offer realistic examples that address the common issues you're most likely to face. They describe each task in detail, presenting example code based on widely used tools such as Pig, Hive, and Spark. This concise, hands-on eBook is valuable for every data scientist, data engineer, and architect who wants to master data munging: not just in theory, but in practice with the field's #1 platform—Hadoop. Coverage includes A framework for understanding the various types of data quality checks, including cell-based rules, distribution validation, and outlier analysis Assessing tradeoffs in common approaches to imputing missing values Implementing quality checks with Pig or Hive UDFs Transforming raw data into “feature matrix” format for machine learning algorithms Choosing features and



instances Implementing text features via “bag-of-words” and NLP techniques Handling time-series data via frequency- or time-domain methods Manipulating feature values to prepare for modeling Data Munging with Hadoop is part of a larger, forthcoming work entitled Data Science Using Hadoop. To be notified when the larger work is available, register your purchase of Data Munging with Hadoop at [informit.com/register](http://informit.com/register) and check the box “I would like to hear from InformIT and its family of brands about products and special offers.”

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 Data science libraries, frameworks, modules, and toolkits are great for doing data science, but they’re also a good way to dive into the discipline without actually understanding data science. In this book, you’ll learn how many of the most fundamental data science tools and algorithms work by implementing them from scratch. If you have an aptitude for mathematics and some programming skills, author Joel Grus will help you get comfortable with the math and statistics at the core of data science, and with hacking skills you need to get started as a data scientist. Today’s messy glut of data holds answers to questions no one’s even thought to ask. This book provides you with the know-how to dig those answers out. Get a crash course in Python Learn the basics of linear algebra, statistics, and probability—and understand how and when they’re used in data science Collect, explore, clean, munge, and manipulate data Dive into the fundamentals of machine learning Implement models such as k-nearest Neighbors, Naive Bayes, linear and logistic regression, decision trees, neural networks, and clustering Explore recommender systems, natural language processing, network analysis, MapReduce, and databases

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

The definitive guide to successfully integrating social, mobile, Big-Data analytics, cloud and IoT principles and technologies The main goal of

## Download File PDF Practical Data Science With Hadoop And Spark

this book is to spur the development of effective big-data computing operations on smart clouds that are fully supported by IoT sensing, machine learning and analytics systems. To that end, the authors draw upon their original research and proven track record in the field to describe a practical approach integrating big-data theories, cloud design principles, Internet of Things (IoT) sensing, machine learning, data analytics and Hadoop and Spark programming. Part 1 focuses on data science, the roles of clouds and IoT devices and frameworks for big-data computing. Big data analytics and cognitive machine learning, as well as cloud architecture, IoT and cognitive systems are explored, and mobile cloud-IoT-interaction frameworks are illustrated with concrete system design examples. Part 2 is devoted to the principles of and algorithms for machine learning, data analytics and deep learning in big data applications. Part 3 concentrates on cloud programming software libraries from MapReduce to Hadoop, Spark and TensorFlow and describes business, educational, healthcare and social media applications for those tools. The first book describing a practical approach to integrating social, mobile, analytics, cloud and IoT (SMACT) principles and technologies Covers theory and computing techniques and technologies, making it suitable for use in both computer science and electrical engineering programs Offers an extremely well-informed vision of future intelligent and cognitive computing environments integrating SMACT technologies Fully illustrated throughout with examples, figures and approximately 150 problems to support and reinforce learning Features a companion website with an instructor manual and PowerPoint slides [www.wiley.com/go/hwangIoT](http://www.wiley.com/go/hwangIoT) Big-Data Analytics for Cloud, IoT and Cognitive Computing satisfies the demand among university faculty and students for cutting-edge information on emerging intelligent and cognitive computing systems and technologies. Professionals working in data science, cloud computing and IoT applications will also find this book to be an extremely useful working resource.

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

Hadoop in Action teaches readers how to use Hadoop and write MapReduce programs. The intended readers are programmers, architects, and project managers who have to process large amounts of data offline. Hadoop in Action will lead the reader from obtaining a copy of Hadoop to setting it up in a cluster and writing data analytic programs. The book begins by making the basic idea of Hadoop and MapReduce easier to grasp by applying the default Hadoop installation to a few easy-to-follow tasks, such as analyzing changes in word frequency across a body of documents. The book continues through the basic concepts of MapReduce applications developed using Hadoop, including a close look at framework components, use of Hadoop for a variety of data analysis tasks, and numerous examples of Hadoop in action. Hadoop in Action will explain how to use Hadoop and present design patterns and practices of programming MapReduce. MapReduce is a complex idea both conceptually and in its implementation,

and Hadoop users are challenged to learn all the knobs and levers for running Hadoop. This book takes you beyond the mechanics of running Hadoop, teaching you to write meaningful programs in a MapReduce framework. This book assumes the reader will have a basic familiarity with Java, as most code examples will be written in Java. Familiarity with basic statistical concepts (e.g. histogram, correlation) will help the reader appreciate the more advanced data processing examples. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book. 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

Introduces professionals and scientists to statistics and machine learning using the programming language R Written by and for practitioners, this book provides an overall introduction to R, focusing on tools and methods commonly used in data science, and placing emphasis on practice and business use. It covers a wide range of topics in a single volume, including big data, databases, statistical machine learning, data wrangling, data visualization, and the reporting of results. The topics covered are all important for someone with a science/math background that is looking to quickly learn several practical technologies to enter or transition to the growing field of data science. The Big R-Book for Professionals: From Data Science to Learning Machines and Reporting with R includes nine parts, starting with an introduction to the subject and followed by an overview of R and elements of statistics. The third part revolves around data, while the fourth focuses on data wrangling. Part 5 teaches readers about exploring data. In Part 6 we learn to build models, Part 7 introduces the reader to the reality in companies, Part 8 covers reports and interactive applications and finally Part 9 introduces the reader to big data and performance computing. It also includes some helpful appendices. Provides a practical guide for non-experts with a focus on business users Contains a unique combination of topics including an introduction to R, machine learning, mathematical models, data wrangling, and reporting Uses a practical tone and integrates multiple topics in a coherent framework Demystifies the hype around machine learning and AI by enabling readers to understand the provided models and program them in R Shows readers how to visualize results in static and interactive reports Supplementary materials includes PDF slides based on the book's content, as well as all the extracted R-code and is available to everyone on a Wiley Book Companion Site The Big R-Book is an excellent guide for science technology, engineering, or

mathematics students who wish to make a successful transition from the academic world to the professional. It will also appeal to all young data scientists, quantitative analysts, and analytics professionals, as well as those who make mathematical models. Learn how to fuse today's data science tools and techniques with your SAP enterprise resource planning (ERP) system. With this practical guide, SAP veterans Greg Foss and Paul Modderman demonstrate how to use several data analysis tools to solve interesting problems with your SAP data. Data engineers and scientists will explore ways to add SAP data to their analysis processes, while SAP business analysts will learn practical methods for answering questions about the business. By focusing on grounded explanations of both SAP processes and data science tools, this book gives data scientists and business analysts powerful methods for discovering deep data truths. You'll explore:

- Examples of how data analysis can help you solve several SAP challenges
- Natural language processing for unlocking the secrets in text
- Data science techniques for data clustering and segmentation
- Methods for detecting anomalies in your SAP data
- Data visualization techniques for making your data come to life

Primer into the multidisciplinary world of Data Science

### KEY FEATURES

- Explore and use the key concepts of Statistics required to solve data science problems
- Use Docker, Jenkins, and Git for Continuous Development and Continuous Integration of your web app
- Learn how to build Data Science solutions with GCP and AWS

### DESCRIPTION

The book will initially explain the What-Why of Data Science and the process of solving a Data Science problem. The fundamental concepts of Data Science, such as Statistics, Machine Learning, Business Intelligence, Data pipeline, and Cloud Computing, will also be discussed. All the topics will be explained with an example problem and will show how the industry approaches to solve such a problem. The book will pose questions to the learners to solve the problems and build the problem-solving aptitude and effectively learn. The book uses Mathematics wherever necessary and will show you how it is implemented using Python with the help of an example dataset.

### WHAT WILL YOU LEARN

- Understand the multi-disciplinary nature of Data Science
- Get familiar with the key concepts in Mathematics and Statistics
- Explore a few key ML algorithms and their use cases
- Learn how to implement the basics of Data Pipelines
- Get an overview of Cloud Computing & DevOps
- Learn how to create visualizations using Tableau

### WHO THIS BOOK IS FOR

This book is ideal for Data Science enthusiasts who want to explore various aspects of Data Science. Useful for Academicians, Business owners, and Researchers for a quick reference on industrial practices in Data Science.

### TABLE OF CONTENTS

1. Data Science in Practice
2. Mathematics Essentials
3. Statistics Essentials
4. Exploratory Data Analysis
5. Data preprocessing
6. Feature Engineering
7. Machine learning algorithms
8. Productionizing ML models
9. Data Flows in Enterprises
10. Introduction to Databases
11. Introduction to Big Data
12. DevOps for Data Science
13. Introduction to Cloud Computing
14. Deploy Model to Cloud
15. Introduction to Business Intelligence
16. Data Visualization Tools
17. Industry Use Case 1 – FormAssist
18. Industry Use Case 2 – PeopleReporter
19. Data Science Learning Resources
20. Do It Your Self Challenges
21. MCQs for Assessments

Explore big data concepts, platforms, analytics, and their applications using the power of Hadoop

### 3 Key Features

Learn Hadoop 3 to build effective big data analytics solutions on-premise and on cloud

Integrate Hadoop with other big data tools such as R,



Python, Apache Spark, and Apache Flink Exploit big data using Hadoop 3 with real-world examples Book Description Apache Hadoop is the most popular platform for big data processing, and can be combined with a host of other big data tools to build powerful analytics solutions. Big Data Analytics with Hadoop 3 shows you how to do just that, by providing insights into the software as well as its benefits with the help of practical examples. Once you have taken a tour of Hadoop 3's latest features, you will get an overview of HDFS, MapReduce, and YARN, and how they enable faster, more efficient big data processing. You will then move on to learning how to integrate Hadoop with the open source tools, such as Python and R, to analyze and visualize data and perform statistical computing on big data. As you get acquainted with all this, you will explore how to use Hadoop 3 with Apache Spark and Apache Flink for real-time data analytics and stream processing. In addition to this, you will understand how to use Hadoop to build analytics solutions on the cloud and an end-to-end pipeline to perform big data analysis using practical use cases. By the end of this book, you will be well-versed with the analytical capabilities of the Hadoop ecosystem. You will be able to build powerful solutions to perform big data analytics and get insight effortlessly. What you will learn Explore the new features of Hadoop 3 along with HDFS, YARN, and MapReduce Get well-versed with the analytical capabilities of Hadoop ecosystem using practical examples Integrate Hadoop with R and Python for more efficient big data processing Learn to use Hadoop with Apache Spark and Apache Flink for real-time data analytics Set up a Hadoop cluster on AWS cloud Perform big data analytics on AWS using Elastic Map Reduce Who this book is for Big Data Analytics with Hadoop 3 is for you if you are looking to build high-performance analytics solutions for your enterprise or business using Hadoop 3's powerful features, or you're new to big data analytics. A basic understanding of the Java programming language is required.

This textbook on practical data analytics unites fundamental principles, algorithms, and data. Algorithms are the keystone of data analytics and the focal point of this textbook. Clear and intuitive explanations of the mathematical and statistical foundations make the algorithms transparent. But practical data analytics requires more than just the foundations. Problems and data are enormously variable and only the most elementary of algorithms can be used without modification. Programming fluency and experience with real and challenging data is indispensable and so the reader is immersed in Python and R and real data analysis. By the end of the book, the reader will have gained the ability to adapt algorithms to new problems and carry out innovative analyses. This book has three parts:(a) Data Reduction: Begins with the concepts of data reduction, data maps, and information extraction. The second chapter introduces associative statistics, the mathematical foundation of scalable algorithms and distributed computing. Practical aspects of distributed computing is the subject of the Hadoop and MapReduce chapter.(b) Extracting Information from Data: Linear regression and data visualization are the principal topics of Part II. The authors dedicate a chapter to the critical domain of Healthcare Analytics for an extended example of practical data analytics. The algorithms and analytics will be of much interest to practitioners interested in utilizing the large and unwieldy data sets of the Centers for Disease Control and Prevention's Behavioral Risk Factor Surveillance System.(c) Predictive Analytics Two foundational and widely used algorithms, k-nearest neighbors and naive Bayes, are developed in detail. A chapter is dedicated to forecasting. The last chapter

focuses on streaming data and uses publicly accessible data streams originating from the Twitter API and the NASDAQ stock market in the tutorials. This book is intended for a one- or two-semester course in data analytics for upper-division undergraduate and graduate students in mathematics, statistics, and computer science. The prerequisites are kept low, and students with one or two courses in probability or statistics, an exposure to vectors and matrices, and a programming course will have no difficulty. The core material of every chapter is accessible to all with these prerequisites. The chapters often expand at the close with innovations of interest to practitioners of data science. Each chapter includes exercises of varying levels of difficulty. The text is eminently suitable for self-study and an exceptional resource for practitioners.

Learn how to process and analysis data using Python KEY FEATURES - The book has theories explained elaborately along with Python code and corresponding output to support the theoretical explanations. The Python codes are provided with step-by-step comments to explain each instruction of the code. - The book is not just dealing with the background mathematics alone or only the programs but beautifully correlates the background mathematics to the theory and then finally translating it into the programs. - A rich set of chapter-end exercises are provided, consisting of both short-answer questions and long-answer questions.

DESCRIPTION This book introduces the fundamental concepts of Data Science, which has proved to be a major game-changer in business solving problems. Topics covered in the book include fundamentals of Data Science, data preprocessing, data plotting and visualization, statistical data analysis, machine learning for data analysis, time-series analysis, deep learning for Data Science, social media analytics, business analytics, and Big Data analytics. The content of the book describes the fundamentals of each of the Data Science related topics together with illustrative examples as to how various data analysis techniques can be implemented using different tools and libraries of Python programming language. Each chapter contains numerous examples and illustrative output to explain the important basic concepts. An appropriate number of questions is presented at the end of each chapter for self-assessing the conceptual understanding. The references presented at the end of every chapter will help the readers to explore more on a given topic.

WHAT WILL YOU LEARN Perform processing on data for making it ready for visual plot and understand the pattern in data over time. Understand what machine learning is and how learning can be incorporated into a program. Know how tools can be used to perform analysis on big data using python and other standard tools. Perform social media analytics, business analytics, and data analytics on any data of a company or organization.

WHO THIS BOOK IS FOR The book is for readers with basic programming and mathematical skills. The book is for any engineering graduates that wish to apply data science in their projects or wish to build a career in this direction. The book can be read by anyone who has an interest in data analysis and would like to explore more out of interest or to apply it to certain real-life problems.

TABLE OF CONTENTS 1. Fundamentals of Data Science 2. Data Preprocessing 3. Data Plotting and Visualization 4. Statistical Data Analysis 5. Machine Learning for Data Science 6. Time-Series Analysis 7. Deep Learning for Data Science 8. Social Media Analytics 9. Business Analytics 10. Big Data Analytics

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

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, 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

This invaluable addition to any data scientist's library shows you how to apply the R programming language and useful statistical techniques to everyday business situations as well as how to effectively present results to audiences of all levels. To answer the ever-increasing demand for machine learning and analysis, this new edition boasts additional R tools, modeling techniques, and more. Practical Data Science with R, Second Edition takes a practice-oriented approach to explaining basic principles in the ever-expanding field of data science. You'll jump right to real-world use cases as you apply the R programming language and statistical analysis techniques to carefully explained examples based in marketing, business intelligence, and decision support. Purchase of

the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications.

Get Started Fast with Apache Hadoop® 2, YARN, and Today's Hadoop Ecosystem With Hadoop 2.x and YARN, Hadoop moves beyond MapReduce to become practical for virtually any type of data processing. Hadoop 2.x and the Data Lake concept represent a radical shift away from conventional approaches to data usage and storage. Hadoop 2.x installations offer unmatched scalability and breakthrough extensibility that supports new and existing Big Data analytics processing methods and models. Hadoop® 2 Quick-Start Guide is the first easy, accessible guide to Apache Hadoop 2.x, YARN, and the modern Hadoop ecosystem. Building on his unsurpassed experience teaching Hadoop and Big Data, author Douglas Eadline covers all the basics you need to know to install and use Hadoop 2 on personal computers or servers, and to navigate the powerful technologies that complement it. Eadline concisely introduces and explains every key Hadoop 2 concept, tool, and service, illustrating each with a simple “beginning-to-end” example and identifying trustworthy, up-to-date resources for learning more. This guide is ideal if you want to learn about Hadoop 2 without getting mired in technical details. Douglas Eadline will bring you up to speed quickly, whether you're a user, admin, devops specialist, programmer, architect, analyst, or data scientist. Coverage Includes Understanding what Hadoop 2 and YARN do, and how they improve on Hadoop 1 with MapReduce Understanding Hadoop-based Data Lakes versus RDBMS Data Warehouses Installing Hadoop 2 and core services on Linux machines, virtualized sandboxes, or clusters Exploring the Hadoop Distributed File System (HDFS) Understanding the essentials of MapReduce and YARN application programming Simplifying programming and data movement with Apache Pig, Hive, Sqoop, Flume, Oozie, and HBase Observing application progress, controlling jobs, and managing workflows Managing Hadoop efficiently with Apache Ambari—including recipes for HDFS to NFSv3 gateway, HDFS snapshots, and YARN configuration Learning basic Hadoop 2 troubleshooting, and installing Apache Hue and Apache Spark

Learn how to build a data science technology stack and perform good data science with repeatable methods. You will learn how to turn data lakes into business assets. The data science technology stack demonstrated in Practical Data Science is built from components in general use in the industry. Data scientist Andreas Vermeulen demonstrates in detail how to build and provision a technology stack to yield repeatable results. He shows you how to apply practical methods to extract actionable business knowledge from data lakes consisting of data from a polyglot of data types and dimensions. What You'll Learn Become fluent in the essential concepts and terminology of data science and data engineering Build and use a technology stack that meets industry criteria Master the methods for retrieving actionable business knowledge Coordinate the handling of polyglot data types in a data lake for repeatable results Who This Book Is For Data scientists and data engineers who are required to convert data from a data lake into actionable knowledge for their business, and students who aspire to be data scientists and data engineers

This multi-contributed handbook focuses on the latest workings of IoT (internet of Things) and Big Data. As the resources are limited, it's the endeavor of the authors to support and bring the information into one resource. The book is divided into 4 sections that covers IoT and technologies, the future of Big Data, algorithms, and case studies showing IoT and Big Data in various fields



such as health care, manufacturing and automation. Features Focuses on the latest workings of IoT and Big Data Discusses the emerging role of technologies and the fast-growing market of Big Data Covers the movement toward automation with hardware, software, and sensors, and trying to save on energy resources Offers the latest technology on IoT Presents the future horizons on Big Data

Over 85 recipes to help you complete real-world data science projects in R and Python About This Book Tackle every step in the data science pipeline and use it to acquire, clean, analyze, and visualize your data Get beyond the theory and implement real-world projects in data science using R and Python Easy-to-follow recipes will help you understand and implement the numerical computing concepts Who This Book Is For If you are an aspiring data scientist who wants to learn data science and numerical programming concepts through hands-on, real-world project examples, this is the book for you. Whether you are brand new to data science or you are a seasoned expert, you will benefit from learning about the structure of real-world data science projects and the programming examples in R and Python. What You Will Learn Learn and understand the installation procedure and environment required for R and Python on various platforms Prepare data for analysis by implement various data science concepts such as acquisition, cleaning and munging through R and Python Build a predictive model and an exploratory model Analyze the results of your model and create reports on the acquired data Build various tree-based methods and Build random forest In Detail As increasing amounts of data are generated each year, the need to analyze and create value out of it is more important than ever. Companies that know what to do with their data and how to do it well will have a competitive advantage over companies that don't. Because of this, there will be an increasing demand for people that possess both the analytical and technical abilities to extract valuable insights from data and create valuable solutions that put those insights to use. Starting with the basics, this book covers how to set up your numerical programming environment, introduces you to the data science pipeline, and guides you through several data projects in a step-by-step format. By sequentially working through the steps in each chapter, you will quickly familiarize yourself with the process and learn how to apply it to a variety of situations with examples using the two most popular programming languages for data analysis—R and Python. Style and approach This step-by-step guide to data science is full of hands-on examples of real-world data science tasks. Each recipe focuses on a particular task involved in the data science pipeline, ranging from readying the dataset to analytics and visualization

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.

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 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

Tackle the real-world complexities of modern machine learning with innovative, cutting-edge, techniques About This Book Fully-coded working examples using a wide range of machine learning libraries and tools, including Python, R, Julia, and Spark Comprehensive practical solutions taking you into the future of machine learning Go a step further and integrate your machine learning projects with Hadoop Who This Book Is For This book has been created for data scientists who want to see machine learning in action and explore its real-world application. With guidance on everything from the fundamentals of machine learning and predictive analytics to the latest innovations set to lead the big data revolution into the future, this is an unmissable resource for anyone dedicated to tackling current big data challenges. Knowledge of programming (Python and R) and mathematics is advisable if you want to get started immediately. What You Will Learn Implement a wide range of algorithms and techniques for tackling complex data Get to grips with some of the most powerful languages in data science, including R, Python, and Julia Harness the capabilities of Spark and Hadoop to manage and process data successfully Apply the appropriate machine learning technique to address real-world problems Get acquainted with Deep learning and find out how neural networks are being used at the cutting-edge of machine learning Explore the future of machine learning and dive deeper into polyglot persistence, semantic data, and more In Detail Finding meaning in increasingly larger and more complex datasets is a growing demand of the modern world.

Machine learning and predictive analytics have become the most important approaches to uncover data gold mines. Machine learning uses complex algorithms to make improved predictions of outcomes based on historical patterns and the behaviour of data sets. Machine learning can deliver dynamic insights into trends, patterns, and relationships within data, immensely valuable to business growth and development. This book explores an extensive range of machine learning techniques uncovering hidden tricks and tips for several types of data using practical and real-world examples. While machine learning can be highly theoretical, this book offers a refreshing hands-on approach without losing sight of the underlying principles. Inside, a full exploration of the various algorithms gives you high-quality guidance so you can begin to see just how effective machine learning is at tackling contemporary challenges of big data. This is the only book you need to implement a whole suite of open source tools, frameworks, and languages in machine learning. We will cover the leading data science languages, Python and R, and the underrated but powerful Julia, as well as a range of other big data platforms including Spark, Hadoop, and Mahout. Practical Machine Learning is an essential resource for the modern data scientists who want to get to grips with its real-world application. With this book, you will not only learn the fundamentals of machine learning but dive deep into the complexities of real world data before moving on to using Hadoop and its wider ecosystem of tools to process and manage your structured and unstructured data. You will explore different machine learning techniques for both supervised and unsupervised learning; from decision trees to Naive Bayes classifiers and linear and clustering methods, you will learn strategies for a truly advanced approach to the statistical analysis of data. The book also explores the cutting-edge advancements in machine learning, with worked examples and guidance on deep learning and reinforcement learning, providing you with practical demonstrations and samples that help take the theory—and mystery—out of even the most advanced machine learning methodologies. Style and approach A practical data science tutorial designed to give you an insight into the practical application of machine learning, this book takes you through complex concepts and tasks in an accessible way. Featuring information on a wide range of data science techniques, Practical Machine Learning is a comprehensive data science resource.

Data Science is booming thanks to R and Python, but Java brings the robustness, convenience, and ability to scale critical to today's data science applications. With this practical book, Java software engineers looking to add data science skills will take a logical journey through the data science pipeline. Author Michael Brzustowicz explains the basic math theory behind each step of the data science process, as well as how to apply these concepts with Java. You'll learn the critical roles that data IO, linear algebra, statistics, data operations, learning and prediction, and Hadoop MapReduce play in the process. Throughout this book, you'll find code examples you can use in your applications. Examine methods for obtaining, cleaning, and arranging data into its purest form Understand the matrix structure that

your data should take Learn basic concepts for testing the origin and validity of data Transform your data into stable and usable numerical values Understand supervised and unsupervised learning algorithms, and methods for evaluating their success Get up and running with MapReduce, using customized components suitable for data science algorithms Practical Data Science with Hadoop and Spark Designing and Building Effective Analytics at Scale Addison-Wesley Professional

Dive into the world of SQL on Hadoop and get the most out of your Hive data warehouses. This book is your go-to resource for using Hive: authors Scott Shaw, Ankur Gupta, David Kjerrumgaard, and Andreas Francois Vermeulen take you through learning HiveQL, the SQL-like language specific to Hive, to analyze, export, and massage the data stored across your Hadoop environment. From deploying Hive on your hardware or virtual machine and setting up its initial configuration to learning how Hive interacts with Hadoop, MapReduce, Tez and other big data technologies, Practical Hive gives you a detailed treatment of the software. In addition, this book discusses the value of open source software, Hive performance tuning, and how to leverage semi-structured and unstructured data. What You Will Learn Install and configure Hive for new and existing datasets Perform DDL operations Execute efficient DML operations Use tables, partitions, buckets, and user-defined functions Discover performance tuning tips and Hive best practices Who This Book Is For Developers, companies, and professionals who deal with large amounts of data and could use software that can efficiently manage large volumes of input. It is assumed that readers have the ability to work with SQL.

Summary Introducing Data Science teaches you how to accomplish the fundamental tasks that occupy data scientists. Using the Python language and common Python libraries, you'll experience firsthand the challenges of dealing with data at scale and gain a solid foundation in data science. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Many companies need developers with data science skills to work on projects ranging from social media marketing to machine learning. Discovering what you need to learn to begin a career as a data scientist can seem bewildering. This book is designed to help you get started. About the Book Introducing Data Science Introducing Data Science explains vital data science concepts and teaches you how to accomplish the fundamental tasks that occupy data scientists. You'll explore data visualization, graph databases, the use of NoSQL, and the data science process. You'll use the Python language and common Python libraries as you experience firsthand the challenges of dealing with data at scale. Discover how Python allows you to gain insights from data sets so big that they need to be stored on multiple machines, or from data moving so quickly that no single machine can handle it. This book gives you hands-on experience with the most popular Python data science libraries, Scikit-learn and StatsModels. After reading this book, you'll have the solid foundation you need to start a career in data science.



What's Inside Handling large data Introduction to machine learning Using Python to work with data Writing data science algorithms About the Reader This book assumes you're comfortable reading code in Python or a similar language, such as C, Ruby, or JavaScript. No prior experience with data science is required. About the Authors Davy Clenden, Arno D. B. Meysman, and Mohamed Ali are the founders and managing partners of Optimately and Maiton, where they focus on developing data science projects and solutions in various sectors. Table of Contents Data science in a big data world The data science process Machine learning Handling large data on a single computer First steps in big data Join the NoSQL movement The rise of graph databases Text mining and text analytics Data visualization to the end user

This book describes current problems in data science and Big Data. Key topics are data classification, Graph Cut, the Laplacian Matrix, Google Page Rank, efficient algorithms, hardness of problems, different types of big data, geometric data structures, topological data processing, and various learning methods. For unsolved problems such as incomplete data relation and reconstruction, the book includes possible solutions and both statistical and computational methods for data analysis. Initial chapters focus on exploring the properties of incomplete data sets and partial-connectedness among data points or data sets. Discussions also cover the completion problem of Netflix matrix; machine learning method on massive data sets; image segmentation and video search. This book introduces software tools for data science and Big Data such MapReduce, Hadoop, and Spark. This book contains three parts. The first part explores the fundamental tools of data science. It includes basic graph theoretical methods, statistical and AI methods for massive data sets. In second part, chapters focus on the procedural treatment of data science problems including machine learning methods, mathematical image and video processing, topological data analysis, and statistical methods. The final section provides case studies on special topics in variational learning, manifold learning, business and financial data recovery, geometric search, and computing models. Mathematical Problems in Data Science is a valuable resource for researchers and professionals working in data science, information systems and networks. Advanced-level students studying computer science, electrical engineering and mathematics will also find the content helpful.

[Copyright: c40f330457b429b8036a5299e1c1d0d5](#)