

Practical Graph Mining With R Chapman Hallcrc Data Mining And Knowledge Discovery Series Published By Chapman And Hallcrc 2013

Data mining is the art and science of intelligent data analysis. By building knowledge from information, data mining adds considerable value to the ever increasing stores of electronic data that abound today. In performing data mining many decisions need to be made regarding the choice of methodology, the choice of data, the choice of tools, and the choice of algorithms. Throughout this book the reader is introduced to the basic concepts and some of the more popular algorithms of data mining. With a focus on the hands-on end-to-end process for data mining, Williams guides the reader through various capabilities of the easy to use, free, and open source Rattle Data Mining Software built on the sophisticated R Statistical Software. The focus on doing data mining rather than just reading about data mining is refreshing. The book covers data understanding, data preparation, data refinement, model building, model evaluation, and practical deployment. The reader will learn to rapidly deliver a data mining project using software easily installed for free from the Internet. Coupling Rattle with R delivers a very sophisticated data mining environment with all the power, and more, of the many commercial offerings.

What does the Web look like? How can we find patterns, communities, outliers, in a social network? Which are the most central nodes in a network? These are the questions that motivate this work. Networks and graphs appear in many diverse settings, for example in social networks, computer-communication networks (intrusion detection, traffic management), protein-protein interaction networks in biology, document-text bipartite graphs in text retrieval, person-account graphs in financial fraud detection, and others. In this work, first we list several surprising patterns that real graphs tend to follow. Then we give a detailed list of generators that try to mirror these patterns. Generators are important, because they can help with "what if" scenarios, extrapolations, and anonymization. Then we provide a list of powerful tools for graph analysis, and specifically spectral methods (Singular Value Decomposition (SVD)), tensors, and case studies like the famous "pageRank" algorithm and the "HITS" algorithm for ranking web search results. Finally, we conclude with a survey of tools and observations from related fields like sociology, which provide complementary viewpoints. Table of Contents: Introduction / Patterns in Static Graphs / Patterns in Evolving Graphs / Patterns in Weighted Graphs / Discussion: The Structure of Specific Graphs / Discussion: Power Laws and Deviations / Summary of Patterns / Graph Generators / Preferential Attachment and Variants / Incorporating Geographical Information / The RMat / Graph Generation by Kronecker Multiplication / Summary and Practitioner's Guide / SVD, Random Walks, and Tensors / Tensors / Community Detection / Influence/Virus Propagation and Immunization / Case Studies / Social Networks / Other Related Work / Conclusions

R and Data Mining introduces researchers, post-graduate students, and analysts to data mining using R, a free software environment for statistical computing and graphics. The book provides practical methods for using R in applications from academia to industry to extract knowledge from vast amounts of data. Readers will find this book a valuable guide to the use of R in tasks such as classification and prediction, clustering, outlier detection, association rules, sequence analysis, text mining, social network analysis, sentiment analysis, and more. Data mining techniques are growing in popularity in a broad range of areas, from banking to insurance, retail, telecom, medicine, research, and government. This book focuses on the modeling phase of the data mining process, also addressing data exploration and model evaluation. With three in-depth case studies, a quick reference guide, bibliography, and links to a wealth of online resources, R and Data Mining is a valuable, practical guide to a powerful method of analysis. Presents an introduction into using R for data mining applications, covering most popular data mining techniques Provides code examples and data so that readers can easily learn the techniques Features case studies in real-world applications to help readers apply the techniques in their work

Data Mining with R: Learning with Case Studies, Second Edition uses practical examples to illustrate the power of R and data mining. Providing an extensive update to the best-selling first edition, this new edition is divided into two parts. The first part will feature introductory material, including a new chapter that provides an introduction to data mining, to complement the already existing introduction to R. The second part includes case studies, and the new edition strongly revises the R code of the case studies making it more up-to-date with recent packages that have emerged in R. The book does not assume any prior knowledge about R. Readers who are new to R and data mining should be able to follow the case studies, and they are designed to be self-contained so the reader can start anywhere in the document. The book is accompanied by a set of freely available R source files that can be obtained at the book's web site. These files include all the code used in the case studies, and they facilitate the "do-it-yourself" approach followed in the book. Designed for users of data analysis tools, as well as researchers and developers, the book should be useful for anyone interested in entering the "world" of R and data mining. About the Author Luis Torgo is an associate professor in the Department of Computer Science at the University of Porto in Portugal. He teaches Data Mining in R in the NYU Stern School of Business' MS in Business Analytics program. An active researcher in machine learning and data mining for more than 20 years, Dr. Torgo is also a researcher in the Laboratory of Artificial Intelligence and Data Analysis (LIAAD) of INESC Porto LA.

Practical Statistics for Data Scientists

Algorithm, Security and Application

An Introduction

Mastering Social Media Mining with R

Machine Learning and Knowledge Discovery in Databases

European Conference, Antwerp, Belgium, September 15-19, 2008, Proceedings, Part I

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

Much of the data available today is unstructured and text-heavy, making it challenging for analysts to apply their usual data wrangling and visualization tools. With this practical book, you'll explore text-mining techniques with tidytext, a package that authors Julia Silge and David Robinson developed using the tidy principles behind R packages like ggraph and dplyr. You'll learn how tidytext and other tidy tools in R can make text analysis easier and more effective. The authors demonstrate how treating text as data frames enables you to manipulate, summarize, and visualize characteristics of text. You'll also learn how to integrate natural language processing (NLP) into effective workflows. Practical code examples and data explorations will help you generate real insights from literature, news, and social media. Learn how to apply the tidy text format to NLP Use sentiment analysis to mine the emotional content of text Identify a document's most important terms with frequency measurements Explore relationships and connections between words with the ggraph and widyr packages Convert back and forth between R's tidy and non-tidy text formats Use topic modeling to classify document collections into natural groups Examine case studies that compare Twitter archives, dig into NASA metadata, and analyze thousands of Usenet messages

Summary R in Action, Second Edition presents both the R language and the examples that make it so useful for business developers. Focusing on practical solutions, the book offers a crash course in statistics and covers elegant methods for dealing with messy and incomplete data that are difficult to analyze using traditional methods. You'll also master R's extensive graphical capabilities for exploring and presenting data visually. And this expanded second edition includes new chapters on time series analysis, cluster analysis, and classification methodologies, including decision trees, random forests, and support vector machines. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Business pros and researchers thrive on data, and R speaks the language of data analysis. R is a powerful programming language for statistical computing. Unlike general-purpose tools, R provides thousands of modules for solving just about any data-crunching or presentation challenge you're likely to face. R runs on all important platforms and is used by thousands of major corporations and institutions worldwide. About the Book R in Action, Second Edition teaches you how to use the R language by presenting examples relevant to scientific, technical, and business developers. Focusing on practical solutions, the book offers a crash course in statistics, including elegant methods for dealing with messy and incomplete data. You'll also master R's extensive graphical capabilities for exploring and presenting data visually. And this expanded second edition includes new chapters on forecasting, data mining, and dynamic report writing. What's Inside Complete R language tutorial Using R to manage, analyze, and visualize data Techniques for debugging programs and creating packages OOP in R Over 160 graphs About the Author Dr. Rob Kabacoff is a seasoned researcher and teacher who specializes in data analysis. He also maintains the popular Quick-R website at statmethods.net. Table of Contents PART 1 GETTING STARTED Introduction to R Creating a dataset Getting started with graphs Basic data management Advanced data management PART 2 BASIC METHODS Basic graphs Basic statistics PART 3 INTERMEDIATE METHODS Regression Analysis of variance Power analysis Intermediate graphs Resampling statistics and bootstrapping PART 4 ADVANCED METHODS Generalized linear models Principal components and factor analysis Time series Cluster analysis Classification Advanced methods for missing data PART 5 EXPANDING YOUR SKILLS Advanced graphics with ggplot2 Advanced programming Creating a package Creating dynamic reports Advanced graphics with the lattice package available online only from manning.com/kabacoff2

Graph data is powerful, thanks to its ability to model arbitrary relationship between objects and is encountered in a range of real-world applications in fields such as bioinformatics, traffic network, scientific collaboration, world wide web and social networks. Graph data mining is used to discover useful information and knowledge from graph data. The complications of nodes, links and the semi-structure form present challenges in terms of the computation tasks, e.g., node classification, link prediction, and graph classification. In this context, various advanced techniques, including graph embedding and graph neural networks, have recently been proposed to improve the performance of graph data mining. This book provides a state-of-the-art review of graph data mining methods. It addresses a current hot topic—the security of graph data mining—and proposes a series of detection methods to identify adversarial samples in graph data. In addition, it introduces readers to graph augmentation and subgraph networks to further enhance the models, i.e., improve their accuracy and robustness. Lastly, the book describes the applications of these advanced techniques in various scenarios, such as traffic networks, social and technical networks, and blockchains. .

Practical Guide To Principal Component Methods in R

Data Mining with Rattle and R

Data analysis and graphics with R

Graph Representation Learning

Graph Mining

Unsupervised Machine Learning

Although there are several good books on principal component methods (PCMs) and related topics, we felt that many of them are either too theoretical or too advanced. This book provides a solid practical guidance to summarize, visualize and interpret the most important information in a large multivariate data sets, using principal component methods in R. The visualization is based on the factoextra R package that we developed for creating easily beautiful ggplot2-based graphs from the output of PCMs. This book contains 4 parts. Part I provides a quick introduction to R and presents the key features of FactoMineR and factoextra. Part II describes classical principal component methods to analyze data sets containing, predominantly, either continuous or categorical variables. These methods include: Principal Component Analysis (PCA, for continuous variables), simple correspondence analysis (CA, for large contingency tables formed by two categorical variables) and Multiple CA (MCA, for a data set with more than 2 categorical variables). In Part III, you'll learn advanced methods for analyzing a data set containing a mix of variables (continuous and categorical) structured or not into groups: Factor Analysis of Mixed Data (FAMD) and Multiple Factor Analysis (MFA). Part IV covers hierarchical clustering on principal components (HCPC), which is useful for performing clustering with a data set containing only categorical variables or with a mixed data of categorical and continuous variables.

The main goal of this book is to explain the core ideas of process mining, and to demonstrate how they can be implemented using just some basic tools that are available to any computer scientist or data scientist. It describes how to analyze event logs in order to discover the behavior of real-world business processes. The end result can often be visualized as a graph, and the book explains how to use Python and Graphviz to render these graphs intuitively. Overall, it enables the reader to implement process mining techniques on his or her own, independently of any specific process mining tool. An introduction to two popular process mining tools, namely Disco and ProM, is also provided. In this second edition the code snippets have been updated to Python 3, and some smaller errors have been corrected. The book will be especially valuable for self-study or as a precursor to a more advanced text. Practitioners and students will be able to follow along on their own, even if they have no prior knowledge of the topic. After reading this book, they will be able to more confidently proceed to the research literature if needed.

Discover Novel and Insightful Knowledge from Data Represented as a GraphPractical Graph Mining with R presents a "do-it-yourself" approach to extracting interesting patterns from graph data. It covers many basic and advanced techniques for the identification of anomalous or frequently recurring patterns in a graph, the discovery of groups or cluste

Practical Graph Mining with RCRC Press

Text Mining with R

Mining User Generated Content

Laws, Tools, and Case Studies

A Primer on Process Mining

Practical Guide to Cluster Analysis in R

Managing and Mining Graph Data

Data Mining: Concepts and Techniques provides the concepts and techniques in processing gathered data or information, which will be used in various applications. Specifically, it explains data mining and the tools used in discovering knowledge from the collected data. This book is referred as the knowledge discovery from data (KDD). It focuses on the feasibility, usefulness, effectiveness, and scalability of techniques of large data sets. After describing data mining, this edition explains the methods of knowing, preprocessing, processing, and warehousing data. It then presents information about data warehouses, online analytical processing (OLAP), and data cube technology. Then, the methods involved in mining frequent patterns, associations, and correlations for large data sets are described. The book details the methods for data classification and introduces the concepts and methods for data clustering. The remaining chapters discuss the outlier detection and the trends, applications, and research frontiers in data mining. This book is intended for Computer Science students, application developers, business professionals, and researchers who seek information on data mining. Presents dozens of algorithms and implementation examples, all in pseudo-code and suitable for use in real-world, large-scale data mining projects Addresses advanced topics such as mining object-relational databases, spatial databases, multimedia databases, time-series databases, text databases, the World Wide Web, and applications in several fields Provides a comprehensive, practical look at the concepts and techniques you need to get the most out of your data

This book presents a perspective of network analysis as a tool to find and quantify significant structures in the interaction patterns between different types of entities. Moreover, network analysis provides the basic means to relate these structures to properties of the entities. It has proven itself to be useful for the analysis of biological and social networks, but also for networks describing complex systems in economy, psychology, geography, and various other fields. Today, network analysis packages in the open-source platform R and other open-source software projects enable scientists from all fields to quickly apply network analytic methods to their data sets. Altogether, these applications offer such a wealth of network analytic methods that it can be overwhelming for someone just entering this field. This book provides a road map through this jungle of network analytic methods, offers advice on how to pick the best method for a given network analytic project, and how to avoid common pitfalls. It introduces the methods which are most often used to analyze complex networks, e.g., different global network measures, types of random graph models, centrality indices, and networks motifs. In addition to introducing these methods, the central focus is on network analysis literacy – the competence to decide when to use which of these methods for which type of question. Furthermore, the book intends to increase the reader's competence to read original literature on network analysis by providing a glossary and intensive translation of formal notation and mathematical symbols in everyday speech. Different aspects of network analysis literacy – understanding formal definitions, programming tasks, or the analysis of structural measures and their interpretation – are deepened in various exercises with provided solutions. This text is an excellent, if not the best starting point for all scientists who want to harness the power of network analysis for their field of expertise.

Data Mining: Practical Machine Learning Tools and Techniques, Fourth Edition, offers a thorough grounding in machine learning concepts, along with practical advice on applying these tools and techniques in real-world data mining situations. This highly anticipated fourth edition of the most acclaimed work on data mining and machine learning teaches readers everything they need to know to get going, from preparing inputs, interpreting outputs, evaluating results, to the algorithmic methods at the heart of successful data mining approaches. Extensive updates reflect the technical changes and modernizations that have taken place in the field since the last edition, including substantial new chapters on probabilistic methods and on deep learning. Accompanying the book is a new version of the popular WEKA machine learning software from the University of Waikato. Authors Witten, Frank, Hall, and Pal include today's techniques coupled with the methods at the leading edge of contemporary research. Please visit the book companion website at http://www.cs.waikato.ac.nz/ml/weka/book.html It contains Powerpoint slides for Chapters 1-12. This is a very comprehensive teaching resource, with many PPT slides covering each chapter of the book Online Appendix on the Weka workbench; again a very comprehensive learning aid for the open source software that goes with the book Table of contents, highlighting the many new sections in the 4th edition, along with reviews of the 1st edition, errata, etc. Provides a thorough grounding in machine learning concepts, as well as practical advice on applying the tools and techniques to data mining projects Presents concrete tips and techniques for performance improvement that work by transforming the input or output in machine learning methods Includes a downloadable WEKA software toolkit, a comprehensive collection of machine learning algorithms for data mining tasks-in an easy-to-use interactive interface Includes open-access online courses that introduce practical applications of the material in the book

The Book of R is a comprehensive, beginner-friendly guide to R, the world's most popular programming language for statistical analysis. Even if you have no programming experience and little more than a grounding in the basics of mathematics, you'll find everything you need to begin using R effectively for statistical analysis. You'll start with the basics, like how to handle data and write simple programs, before moving on to more advanced topics, like producing statistical summaries of your data and performing statistical tests and modeling. You'll even learn how to create impressive data visualizations with R's basic graphics tools and contributed packages, like ggplot2 and gvis, as well as interactive 3D visualizations using the rgl package. Dozens of hands-on exercises (with downloadable solutions) take you from theory to practice, as you learn: –The fundamentals of programming in R, including how to write data frames, create functions, and use variables, statements, and loops –Statistical concepts like exploratory data analysis, probabilities, hypothesis tests, and regression modeling, and how to execute them in R –How to access R's thousands of functions, libraries, and data sets –How to draw valid and useful conclusions from your data –How to create publication-quality graphics of your results Combining detailed explanations with real-world examples and exercises, this book will provide you with a solid understanding of both statistics and the depth of R's functionality. Make The Book of R your doorway into the growing world of data analysis.

Practical Machine Learning Tools and Techniques

Practical Graph Mining with R

Graph Databases

Business Analytics Using R - A Practical Approach

A Practical Perspective

Graph Algorithms

Handbook of Statistical Analysis and Data Mining Applications, Second Edition, is a comprehensive professional reference book that guides business analysts, scientists, engineers and researchers, both academic and industrial, through all stages of data analysis, model building and implementation. The

handbook helps users discern technical and business problems, understand the strengths and weaknesses of modern data mining algorithms and employ the right statistical methods for practical application. This book is an ideal reference for users who want to address massive and complex datasets with novel statistical approaches and be able to objectively evaluate analyses and solutions. It has clear, intuitive explanations of the principles and tools for solving problems using modern analytic techniques and discusses their application to real problems in ways accessible and beneficial to practitioners across several areas—from science and engineering, to medicine, academia and commerce. Includes input by practitioners for practitioners Includes tutorials in numerous fields of study that provide step-by-step instruction on how to use supplied tools to build models Contains practical advice from successful real-world implementations Brings together, in a single resource, all the information a beginner needs to understand the tools and issues in data mining to build successful data mining solutions Features clear, intuitive explanations of novel analytical tools and techniques, and their practical applications

Presenting a comprehensive resource for the mastery of network analysis in R, the goal of Network Analysis with R is to introduce modern network analysis techniques in R to social, physical, and health scientists. The mathematical foundations of network analysis are emphasized in an accessible way and readers are guided through the basic steps of network studies: network conceptualization, data collection and management, network description, visualization, and building and testing statistical models of networks. As with all of the books in the Use R! series, each chapter contains extensive R code and detailed visualizations of datasets. Appendices will describe the R network packages and the datasets used in the book. An R package developed specifically for the book, available to readers on GitHub, contains relevant code and real-world network datasets as well.

Delve into your data for the key to success Data mining is quickly becoming integral to creating value and business momentum. The ability to detect unseen patterns hidden in the numbers exhaustively generated by day-to-day operations allows savvy decision-makers to exploit every tool at their disposal in the pursuit of better business. By creating models and testing whether patterns hold up, it is possible to discover new intelligence that could change your business's entire paradigm for a more successful outcome. Data Mining for Dummies shows you why it doesn't take a data scientist to gain this advantage, and empowers average business people to start shaping a process relevant to their business's needs. In this book, you'll learn the hows and whys of mining to the depths of your data, and how to make the case for heavier investment into data mining capabilities. The book explains the details of the knowledge discovery process including: Model creation, validity testing, and interpretation Effective communication of findings Available tools, both paid and open-source Data selection, transformation, and evaluation Data Mining for Dummies takes you step-by-step through a real-world data-mining project using open-source tools that allow you to get immediate hands-on experience working with large amounts of data. You'll gain the confidence you need to start making data mining practices a routine part of your successful business. If you're serious about doing everything you can to push your company to the top, Data Mining for Dummies is your ticket to effective data mining.

Originating from Facebook, LinkedIn, Twitter, Instagram, YouTube, and many other networking sites, the social media shared by users and the associated metadata are collectively known as user generated content (UGC). To analyze UGC and glean insight about user behavior, robust techniques are needed to tackle the huge amount of real-time, multimedia, and multilingual data. Researchers must also know how to assess the social aspects of UGC, such as user relations and influential users. Mining User Generated Content is the first focused effort to compile state-of-the-art research and address future directions of UGC. It explains how to collect, index, and analyze UGC to uncover social trends and user habits. Divided into four parts, the book focuses on the mining and applications of UGC. The first part presents an introduction to this new and exciting topic. Covering the mining of UGC of different medium types, the second part discusses the social annotation of UGC, social network graph construction and community mining, mining of UGC to assist in music retrieval, and the popular but difficult topic of UGC sentiment analysis. The third part describes the mining and searching of various types of UGC, including knowledge extraction, search techniques for UGC content, and a specific study on the analysis and annotation of Japanese blogs. The fourth part on applications explores the use of UGC to support question-answering, information summarization, and recommendations.

Handbook of Statistical Analysis and Data Mining Applications

Mining of Massive Datasets

A User's Guide to Network Analysis in R

R in Action

Scientific Data Mining

Graph Data Mining

Now in its second edition, this book focuses on practical algorithms for mining data from even the largest datasets.

Integrates social media, social network analysis, and data mining to provide an understanding of the potentials of social media mining.

The main goal of this book is to explain the core ideas of process mining, and to demonstrate how they can be implemented using just some basic tools that are available to any computer scientist or data scientist. It describes how to analyze event logs in order to discover the behavior of real-world business processes. The end result can often be visualized as a graph, and the book explains how to use Python and Graphviz to render these graphs intuitively. Overall, it enables the reader to implement process mining techniques on his or her own, independently of any specific process mining tool. An introduction to two popular process mining tools, namely Disco and ProM, is also provided. The book will be especially valuable for self-study or as a precursor to a more advanced text. Practitioners and students will be able to follow along on their own, even if they have no prior knowledge of the topic. After reading this book, they will be able to more confidently proceed to the research literature if needed.

Learn the fundamental aspects of the business statistics, data mining, and machine learning techniques required to understand the huge amount of data generated by your organization. This book explains practical business analytics through examples, covers the steps involved in using it correctly, and shows you the context in which a particular technique does not make sense. Further, Practical Business Analytics using R helps you understand specific issues faced by organizations and how the solutions to these issues can be facilitated by business analytics. This book will discuss and explore the following through examples and case studies: An introduction to R: data management and R functions The architecture, framework, and life cycle of a business analytics project Descriptive analytics using R: descriptive statistics and data cleaning Data mining: classification, association rules, and clustering Predictive analytics: simple regression, multiple regression, and logistic regression This book includes case studies on important business analytic techniques, such as classification, association, clustering, and regression. The R language is the statistical tool used to demonstrate the concepts throughout the book. What You Will Learn • Write R programs to handle data • Build analytical models and draw useful inferences from them • Discover the basic concepts of data mining and machine learning • Carry out predictive modeling • Define a business issue as an analytical problem Who This Book Is For Beginners who want to understand and learn the fundamentals of analytics using R. Students, managers, executives, strategy and planning professionals, software professionals, and BI/DW professionals.

Data Mining

Mining Graph Data

Foundations of Data Science

Introduction to Data Mining

Data Mining with R

Social Media Mining

This book concentrates on mining networks, a subfield within data science. Data science uses scientific and computational tools to extract valuable knowledge from large data sets. Once data is processed and cleaned, it is analyzed and presented to support decision-making processes. Data science and machine learning tools have become widely used in companies of all sizes. Networks are often large-scale, decentralized, and evolve dynamically over time. Mining complex networks aim to understand the principles governing the organization and the behavior of such networks is crucial for a broad range of fields of study. Here are a few selected typical applications of mining networks: Community detection (which users on some social media platforms are close friends). Link prediction (who is likely to connect to whom on such platforms). Node attribute prediction (what advertisement should be shown to a given user of a particular platform to match their interests). Influential node detection (which social media users would be the best ambassadors of a specific product). This textbook is suitable for an upper-year undergraduate course or a graduate course in programs such as data science, mathematics, computer science, business, engineering, physics, statistics, and social science. This book can be successfully used by all enthusiasts of data science at various levels of sophistication to expand their knowledge or consider changing their career path. Jupiter notebooks (in Python and Julia) accompany the book and can be accessed on <https://www.ryerson.ca/mining-complex-networks/>. These not only contain all the experiments presented in the book, but also include additional material. Bogumi? Kami?ski is the Chairman of the Scientific Council for the Discipline of Economics and Finance at SGH Warsaw School of Economics. He is also an Adjunct Professor at the Data Science Laboratory at Ryerson University. Bogumi? is an expert in applications of mathematical modeling to solving complex real-life problems. He is also a substantial open-source contributor to the development of the Julia language and its package ecosystem. Pawe? Pra?at is a Professor of Mathematics in Ryerson University, whose main research interests are in random graph theory, especially in modeling and mining complex networks. He is the Director of Fields-CQAM Lab on Computational Methods in Industrial Mathematics in The Fields Institute for Research in Mathematical Sciences and has pursued collaborations with various industry partners as well as the Government of Canada. He has written over 170 papers and three books with 130 plus collaborators. François Th?berge holds a B.Sc. degree in applied mathematics from the University of Ottawa, a M.Sc. in telecommunications from INRS and a PhD in electrical engineering from McGill University. He has been employed by the Government of Canada since 1996 where he was involved in the creation of the data science team as well as the research group now known as the Tutte Institute for Mathematics and Computing. He also holds an adjunct professorial position in the Department of Mathematics and Statistics at the University of Ottawa. His current interests include relational-data mining and deep learning.

Although there are several good books on unsupervised machine learning, we felt that many of them are too theoretical. This book provides practical guide to cluster analysis, elegant visualization and interpretation. It contains 5 parts. Part I provides a quick introduction to R and presents required R packages, as well as, data formats and dissimilarity measures for cluster analysis and visualization. Part II covers partitioning clustering methods, which subdivide the data sets into a set of k groups, where k is the number of groups pre-specified by the analyst. Partitioning clustering approaches include: K-means, K-Medoids (PAM) and CLARA algorithms. In Part III, we consider hierarchical clustering method, which is an alternative approach to partitioning clustering. The result of hierarchical clustering is a tree-based representation of the objects called dendrogram. In this part, we describe how to compute, visualize, interpret and compare dendrograms. Part IV describes clustering validation and evaluation strategies, which consists of measuring the goodness of clustering results. Among the chapters covered here, there are: Assessing clustering tendency, Determining the optimal number of clusters, Cluster validation statistics, Choosing the best clustering algorithms and Computing p-value for hierarchical clustering. Part V presents advanced clustering methods, including: Hierarchical k-means clustering, Fuzzy clustering, Model-based clustering and Density-based clustering.

A comprehensive text on foundations and techniques of graph neural networks with applications in NLP, data mining, vision and healthcare.

This book constitutes the refereed proceedings of the joint conference on Machine Learning and Knowledge Discovery in Databases: ECML PKDD 2008, held in Antwerp, Belgium, in September 2008. The 100 papers presented in two volumes, together with 5 invited talks, were carefully reviewed and selected from 521 submissions. In addition to the regular papers the volume contains 14 abstracts of papers appearing in full version in the Machine Learning Journal and the Knowledge Discovery and Databases Journal of Springer. The conference intends to provide an international forum for the discussion of the latest high quality research results in all areas related to machine learning and knowledge discovery in databases. The topics addressed are application of machine learning and data mining methods to real-world problems, particularly exploratory research that describes novel learning and mining tasks and applications requiring non-standard techniques.

A First Course in Programming and Statistics

Examples and Case Studies

A Tidy Approach

The Book of R

Practical Text Mining and Statistical Analysis for Non-structured Text Data Applications

Mining Complex Networks

Graph-structured data is ubiquitous throughout the natural and social sciences, from telecommunication networks to quantum chemistry. Building relational inductive biases into deep learning architectures is crucial for creating systems that can learn, reason, and generalize from this kind of data. Recent years have seen a surge in research on graph representation learning, including techniques for deep graph embeddings, generalizations of convolutional neural networks to graph-structured data, and neural message-passing approaches inspired by belief propagation. These advances in graph representation learning have led to new state-of-the-art results in numerous domains, including chemical synthesis, 3D vision, recommender systems, question answering, and social network analysis. This book provides a synthesis and overview of graph representation learning. It begins with a discussion of the goals of graph representation learning as well as key methodological foundations in graph theory and network analysis. Following this, the book introduces and reviews methods for learning node embeddings, including random-walk-based methods and applications to knowledge graphs. It then provides a technical synthesis and introduction to the highly successful graph neural network (GNN) formalism, which has become a dominant and fast-growing paradigm for deep learning with graph data. The book concludes with a synthesis of recent advancements in deep generative models for graphs—a nascent but quickly growing subset of graph representation learning.

Statistical methods are a key part of data science, yet very few data scientists have any formal statistics training. Courses and books on basic statistics rarely cover the topic from a data science perspective. This practical guide explains how to apply various statistical methods to data science, tells you how to avoid their misuse, and gives you advice on what's important and what's not. Many data science resources incorporate statistical methods but lack a deeper statistical perspective. If you're familiar with the R programming language, and have some exposure to statistics, this quick reference bridges the gap in an accessible, readable format. With this book, you'll learn: Why exploratory data analysis is a key preliminary step in data science How random sampling can reduce bias and yield a higher quality dataset, even with big data How the principles of experimental design yield definitive answers to questions How to use regression to estimate outcomes and detect anomalies Key classification techniques for predicting which categories a record belongs to Statistical machine learning methods that "learn" from data Unsupervised learning methods for extracting meaning from unlabeled data

This comprehensive reference consists of 18 chapters from prominent researchers in the field. Each chapter is self-contained, and synthesizes one aspect of frequent pattern mining. An emphasis is placed on simplifying the content, so that students and practitioners can benefit from the book. Each chapter contains a survey describing key research on the topic, a case study and future directions. Key topics include: Pattern Growth Methods, Frequent Pattern Mining in Data Streams, Mining Graph Patterns, Big Data Frequent Pattern Mining, Algorithms for Data Clustering and more. Advanced-level students in computer science, researchers and practitioners from industry will find this book an invaluable reference.

Extract valuable data from your social media sites and make better business decisions using R About This Book Explore the social media APIs in R to capture data and tame it Employ the machine learning capabilities of R to gain optimal business value A hands-on guide with real-world examples to help you take advantage of the vast opportunities that come with social media data Who This Book Is For If you have basic knowledge of R in terms of its libraries and are aware of different machine learning techniques, this book is for you. Those with experience in data analysis who are interested in mining social media data will find this book useful. What You Will Learn Access APIs of popular social media sites and extract data Perform sentiment analysis and identify trending topics Measure CTR performance for social media campaigns Implement exploratory data analysis and correlation analysis Build a logistic regression model to detect spam messages Construct clusters of pictures using the K-means algorithm and identify popular personalities and destinations Develop recommendation systems using Collaborative Filtering and the Apriori algorithm In Detail With an increase in the number of users on the web, the content generated has increased substantially, bringing in the need to gain insights into the untapped gold mine that is social media data. For computational statistics, R has an advantage over other languages in providing readily-available data extraction and transformation packages, making it easier to carry out your ETL tasks. Along with this, its data visualization packages help users get a better understanding of the underlying data distributions while its range of "standard" statistical packages simplify analysis of the data. This book will teach you how powerful business cases are solved by applying machine learning techniques on social media data. You will learn about important and recent developments in the field of social media, along with a few advanced topics such as Open Authorization (OAuth). Through practical examples, you will access data from R using APIs of various social media sites such as Twitter, Facebook, Instagram, GitHub, Foursquare, LinkedIn, Blogger, and other networks. We will provide you with detailed explanations on the implementation of various use cases using R programming. With this handy guide, you will be ready to embark on your journey as an independent social media analyst. Style and approach This easy-to-follow guide is packed with hands-on, step-by-step examples that will enable you to convert your real-world social media data into useful, practical information.

PCA, M(CA), FAMD, MFA, HCPC, factextra

R and Data Mining

Learning with Case Studies, Second Edition

50 Essential Concepts

Practical Skills with Python and Graphviz

Data Mining: Concepts and Techniques

The world contains an unimaginably vast amount of digital information which is getting ever vaster ever more rapidly. This makes it possible to do many things that previously could not be done: spot business trends, prevent diseases, combat crime and so on. Managed well, the textual data can be used to unlock new sources of economic value, provide fresh insights into science and hold governments to account. As the Internet expands and our natural capacity to process the unstructured text that it contains diminishes, the value of text mining for information retrieval and search will increase dramatically. This comprehensive professional reference brings together all the information, tools and methods a professional will need to efficiently use text mining applications and statistical analysis. The Handbook of Practical Text Mining and Statistical Analysis for Non-structured Text Data Applications presents a comprehensive how- to reference that shows the user how to conduct text mining and statistically analyze results. In addition to providing an in-depth examination of core text mining and link detection tools, methods and operations, the book examines advanced preprocessing techniques, knowledge representation considerations, and visualization approaches. Finally, the book explores current real-world, mission-critical applications of text mining and link detection using real world example tutorials in such varied fields as corporate, finance, business intelligence, genomics research, and counterterrorism activities. -Extensive case studies, most in a tutorial format, allow the reader to 'click through' the example using a software program, thus learning to conduct text mining analyses in the most rapid manner of learning possible -Numerous examples, tutorials, power points and datasets available via companion website on Elsevierdirect.com -Glossary of text mining terms provided in the appendix

Managing and Mining Graph Data is a comprehensive survey book in graph management and mining. It contains extensive surveys on a variety of important graph topics such as graph languages, indexing, clustering, data generation, pattern mining, classification, keyword search, pattern matching, and privacy. It also studies a number of domain-specific scenarios such as stream mining, web graphs, social networks, chemical and biological data. The chapters are written by well known researchers in the field, and provide a broad perspective of the area. This is the first comprehensive survey book in the emerging topic of graph data processing. Managing and Mining Graph Data is designed for a varied audience composed of professors, researchers and practitioners in industry. This volume is also suitable as a reference book for advanced-level database students in computer science and engineering.

This text takes a focused and comprehensive look at mining data represented as a graph, with the latest findings and applications in both theory and practice provided. Even if you have minimal background in analyzing graph data, with this book you'll be able to represent data as graphs, extract patterns and concepts from the data, and apply the methodologies presented in the text to real datasets. There is a misprint with the link to the accompanying Web page for this book. For those readers who would like to experiment with the techniques found in this book or test their own ideas on graph data, the Web page for the book should be <http://www.eecs.wsu.edu/MGD>.

This book provides an introduction to the mathematical and algorithmic foundations of data science, including machine learning, high-dimensional geometry, and analysis of large networks. Topics include the counterintuitive nature of data in high dimensions, important linear algebraic techniques such as singular value decomposition, the theory of random walks and Markov chains, the fundamentals of and important algorithms for machine learning, algorithms and analysis for clustering, probabilistic models for large networks, representation learning including topic modelling and non-negative matrix factorization, wavelets and compressed sensing. Important probabilistic techniques are developed including the law of large numbers, tail inequalities, analysis of random projections, generalization guarantees in machine learning, and moment methods for analysis of phase transitions in large random graphs. Additionally, important structural and complexity measures are discussed such as matrix norms and VC-dimension. This book is suitable for both undergraduate and graduate courses in the design and analysis of algorithms for data.

The Art of Excavating Data for Knowledge Discovery

Practical Examples in Apache Spark and Neo4j

Frequent Pattern Mining

A Practical Approach to the Analysis of Networks

Deep Learning on Graphs

Chandrika Kamath describes how techniques from the multi-disciplinary field of data mining can be used to address the modern problem of data overload in science and engineering domains. Starting with a survey of analysis problems in different applications, it identifies the common themes across these domains.

Discover how graph databases can help you manage and query highly connected data. With this practical book, you'll learn how to design and implement a graph database that brings the power of graphs to bear on a broad range of problem domains. Whether you want to speed up your response to user queries or build a database that can adapt as your business evolves, this book shows you how to apply the schema-free graph model to real-world problems. Learn how different organizations are using graph databases to outperform their competitors. With this book's data modeling, query, and code examples, you'll quickly be able to implement your own solution. Model data with the Cypher query language and property graph model Learn best practices and common pitfalls when modeling with graphs Plan and implement a graph database solution in test-driven fashion Explore real-world examples to learn how and why organizations use a graph database Understand common patterns and components of graph database architecture Use analytical techniques and algorithms to mine graph database information

Discover Novel and Insightful Knowledge from Data Represented as a Graph Practical Graph Mining with R presents a "do-it-yourself" approach to extracting interesting patterns from graph data. It covers many basic and advanced techniques for the identification of anomalous or frequently recurring patterns in a graph, the discovery of groups or clusters of nodes that share common patterns of attributes and relationships, the extraction of patterns that distinguish one category of graphs from another, and the use of those patterns to predict the category of new graphs. Hands-On Application of Graph Data Mining Each chapter in the

book focuses on a graph mining task, such as link analysis, cluster analysis, and classification. Through applications using real data sets, the book demonstrates how computational techniques can help solve real-world problems. The applications covered include network intrusion detection, tumor cell diagnostics, face recognition, predictive toxicology, mining metabolic and protein-protein interaction networks, and community detection in social networks. Develops Intuition through Easy-to-Follow Examples and Rigorous Mathematical Foundations Every algorithm and example is accompanied with R code. This allows readers to see how the algorithmic techniques correspond to the process of graph data analysis and to use the graph mining techniques in practice. The text also gives a rigorous, formal explanation of the underlying mathematics of each technique. Makes Graph Mining Accessible to Various Levels of Expertise Assuming no prior knowledge of mathematics or data mining, this self-contained book is accessible to students, researchers, and practitioners of graph data mining. It is suitable as a primary textbook for graph mining or as a supplement to a standard data mining course. It can also be used as a reference for researchers in computer, information, and computational science as well as a handy guide for data analytics practitioners.

Data Mining For Dummies
Network Analysis Literacy