

Data Mining Concepts And Techniques The Morgan Kaufmann

This is likewise one of the factors by obtaining the soft documents of this **Data Mining Concepts And Techniques The Morgan Kaufmann** by online. You might not require more era to spend to go to the book commencement as skillfully as search for them. In some cases, you likewise attain not discover the revelation Data Mining Concepts And Techniques The Morgan Kaufmann that you are looking for. It will completely squander the time.

However below, afterward you visit this web page, it will be for that reason completely simple to acquire as competently as download lead Data Mining Concepts And Techniques The Morgan Kaufmann

It will not agree to many period as we notify before. You can accomplish it even if acquit yourself something else at house and even in your workplace. for that reason easy! So, are you question? Just exercise just what we give under as with ease as review **Data Mining Concepts And Techniques The Morgan Kaufmann** what you in imitation of to read!

Discrimination and Privacy in the Information Society - Bart Custers 2012-08-11
Vast amounts of data are nowadays collected, stored and processed, in an effort to assist in making a variety of administrative and governmental decisions. These innovative steps considerably improve the speed, effectiveness and quality of decisions. Analyses are increasingly performed by data mining and profiling technologies that statistically and automatically determine patterns and trends. However, when such practices lead to unwanted or unjustified selections, they may result in unacceptable forms of discrimination. Processing vast amounts of data may lead to situations in which data controllers know many of the characteristics, behaviors and whereabouts of people. In some cases, analysts might know more about individuals than these individuals know about themselves. Judging people by their digital identities sheds a different light on our views of privacy and data protection. This book discusses discrimination and privacy issues related to data mining and profiling practices. It provides technological and regulatory solutions, to problems which arise in these innovative contexts. The book explains that common measures for mitigating privacy and discrimination, such as access controls and anonymity, fail to properly resolve privacy and discrimination concerns. Therefore, new solutions, focusing on technology design, transparency and accountability are called for and set forth.

Data Mining - Charu C. Aggarwal 2015-04-13

This textbook explores the different aspects of data mining from the fundamentals to the complex data types and their applications, capturing the wide diversity of problem domains for data mining issues. It goes beyond the traditional focus on data mining problems to introduce advanced data types such as text, time series, discrete sequences, spatial data, graph data, and social networks. Until now, no single book has addressed all these topics in a comprehensive and integrated way. The chapters of this book fall into one of three categories: Fundamental chapters: Data mining has four main problems, which correspond to clustering, classification, association pattern mining, and outlier analysis. These chapters comprehensively discuss a wide variety of methods for these problems. Domain chapters: These chapters discuss the specific methods used for different domains of data such as text data, time-series data, sequence data, graph data, and spatial data. Application chapters: These chapters study important applications such as stream mining, Web mining, ranking, recommendations, social networks, and privacy preservation. The domain chapters also have an applied flavor. Appropriate

for both introductory and advanced data mining courses, Data Mining: The Textbook balances mathematical details and intuition. It contains the necessary mathematical details for professors and researchers, but it is presented in a simple and intuitive style to improve accessibility for students and industrial practitioners (including those with a limited mathematical background). Numerous illustrations, examples, and exercises are included, with an emphasis on semantically interpretable examples. Praise for Data Mining: The Textbook - "As I read through this book, I have already decided to use it in my classes. This is a book written by an outstanding researcher who has made fundamental contributions to data mining, in a way that is both accessible and up to date. The book is complete with theory and practical use cases. It's a must-have for students and professors alike!" -- Qiang Yang, Chair of Computer Science and Engineering at Hong Kong University of Science and Technology "This is the most amazing and comprehensive text book on data mining. It covers not only the fundamental problems, such as clustering, classification, outliers and frequent patterns, and different data types, including text, time series, sequences, spatial data and graphs, but also various applications, such as recommenders, Web, social network and privacy. It is a great book for graduate students and researchers as well as practitioners." -- Philip S. Yu, UIC Distinguished Professor and Wexler Chair in Information Technology at University of Illinois at Chicago

Data Warehousing And Business Intelligence For e-Commerce - Alan R. Simon 2001-05-23

You go online to buy a digital camera. Soon, you realize you've bought a more expensive camera than intended, along with extra batteries, charger, and graphics software-all at the prompting of the retailer. Happy with your purchases? The retailer certainly is, and if you are too, you both can be said to be the beneficiaries of "customer intimacy" achieved through the transformation of data collected during this visit or stored from previous visits into real business intelligence that can be exercised in real time. Data Warehousing and Business Intelligence for e-Commerce is a practical exploration of the technological innovations through which traditional data warehousing is brought to bear on this and other less modest e-commerce applications, such as those at work in B2B, G2C, B2G, and B2E models. The authors examine the core technologies and commercial products in use today, providing a nuts-and-bolts understanding of how you can deploy customer and product data in ways that meet the unique requirements of the

online marketplace-particularly if you are part of a brick-and-mortar company with specific online aspirations. In so doing, they build a powerful case for investment in and aggressive development of these approaches, which are likely to separate winners from losers as e-commerce grows and matures. * Includes the latest from successful data warehousing consultants whose work has encouraged the field's new focus on e-commerce. * Presents information that is written for both consultants and practitioners in companies of all sizes. * Emphasizes the special needs and opportunities of traditional brick-and-mortar businesses that are going online or participating in B2B supply chains or e-marketplaces. * Explains how long-standing assumptions about data warehousing have to be rethought in light of emerging business models that depend on customer intimacy. * Provides advice on maintaining data quality and integrity in environments marked by extensive customer self-input. * Advocates careful planning that will help both old economy and new economy companies develop long-lived and successful e-commerce strategies. * Focuses on data warehousing for emerging e-commerce areas such as e-government and B2E environments.

Data Mining - Ian H. Witten 2005-07-13

Data Mining, Second Edition, describes data mining techniques and shows how they work. The book is a major revision of the first edition that appeared in 1999. While the basic core remains the same, it has been updated to reflect the changes that have taken place over five years, and now has nearly double the references. The highlights of this new edition include thirty new technique sections; an enhanced Weka machine learning workbench, which now features an interactive interface; comprehensive information on neural networks; a new section on Bayesian networks; and much more. This text is designed for information systems practitioners, programmers, consultants, developers, information technology managers, specification writers as well as professors and students of graduate-level data mining and machine learning courses. Algorithmic methods at the heart of successful data mining—including tried and true techniques as well as leading edge methods Performance improvement techniques that work by transforming the input or output

Predictive Analytics and Data Mining - Vijay Kotu 2014-11-27

Put Predictive Analytics into Action Learn the basics of Predictive Analysis and Data Mining through an easy to understand conceptual framework and immediately practice the concepts learned using the open source RapidMiner tool. Whether you are brand new to Data Mining or working on your tenth project, this book will show you how to analyze data, uncover hidden patterns and relationships to aid important decisions and predictions. Data Mining has become an essential tool for any enterprise that collects, stores and processes data as part of its operations. This book is ideal for business users, data analysts, business analysts, business intelligence and data warehousing professionals and for anyone who wants to learn Data Mining. You'll be able to: 1. Gain the necessary knowledge of different data mining techniques, so that you can select the right technique for a given data problem and create a general purpose analytics process. 2. Get up and running fast with more than two dozen commonly used powerful algorithms for predictive analytics using practical use cases. 3. Implement a simple step-by-step process for predicting an outcome or discovering hidden relationships from the data using RapidMiner, an open source GUI based data mining tool Predictive analytics and Data Mining techniques covered: Exploratory Data Analysis, Visualization, Decision trees, Rule induction, k-Nearest Neighbors, Naïve Bayesian, Artificial Neural Networks, Support Vector machines, Ensemble models, Bagging, Boosting, Random

Forests, Linear regression, Logistic regression, Association analysis using Apriori and FP Growth, K-Means clustering, Density based clustering, Self Organizing Maps, Text Mining, Time series forecasting, Anomaly detection and Feature selection. Implementation files can be downloaded from the book companion site at www.LearnPredictiveAnalytics.com Demystifies data mining concepts with easy to understand language Shows how to get up and running fast with 20 commonly used powerful techniques for predictive analysis Explains the process of using open source RapidMiner tools Discusses a simple 5 step process for implementing algorithms that can be used for performing predictive analytics Includes practical use cases and examples

Data Mining - Ian H. Witten 2016-10-01

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 *Data Preparation for Data Mining Using SAS* - Mamdouh Refaat 2010-07-27 Are you a data mining analyst, who spends up to 80% of your time assuring data quality, then preparing that data for developing and deploying predictive models? And do you find lots of literature on data mining theory and concepts, but when it comes to practical advice on developing good mining views find little "how to information? And are you, like most analysts, preparing the data in SAS? This book is intended to fill this gap as your source of practical recipes. It introduces a framework for the process of data preparation for data mining, and presents the detailed implementation of each step in SAS. In addition, business applications of data mining modeling require you to deal with a large number of variables, typically hundreds if not thousands. Therefore, the book devotes several chapters to the methods of data transformation and variable selection. A complete framework for the data preparation process, including implementation details for each step. The complete SAS implementation code, which is readily usable by professional

analysts and data miners. A unique and comprehensive approach for the treatment of missing values, optimal binning, and cardinality reduction. Assumes minimal proficiency in SAS and includes a quick-start chapter on writing SAS macros.

Fuzzy Modeling and Genetic Algorithms for Data Mining and Exploration - Earl Cox 2005-02

Foundations and ideas -- Principal model types -- Approaches to model building -- Fundamental concepts of fuzzy logic -- Fundamental concepts of fuzzy systems -- Fuzzy SQL and intelligent queries -- Fuzzy clustering -- Fuzzy rule induction -- Fundamental concepts of genetic algorithms -- Genetic resource scheduling optimization -- Genetic tuning of fuzzy models.

Data Mining - Jiawei Han 2001

Data Mining - Jiawei Han 2011

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.

Predictive Data Mining - Sholom M. Weiss 1998

This book is the first technical guide to provide a complete, generalized road map for developing data-mining applications, together with advice on performing these large-scale, open-ended analyses for real-world data warehouses.

Data Mining - Ian H. Witten 2000

This book offers a thorough grounding in machine learning concepts combined with practical advice on applying machine learning tools and techniques in real-world data mining situations. Clearly written and effectively illustrated, this book is ideal for anyone involved at any level in the work of extracting usable knowledge from large collections of data. Complementing the book's instruction is fully functional machine learning software.

Formal Concept Analysis - Bernhard Ganter 2012-12-06

This first textbook on formal concept analysis gives a systematic presentation of the mathematical foundations and their relations to applications in computer science, especially in data analysis and knowledge processing. Above all, it presents graphical methods for representing conceptual systems that have proved

themselves in communicating knowledge. The mathematical foundations are treated thoroughly and are illuminated by means of numerous examples, making the basic theory readily accessible in compact form.

Algorithms and Theory of Computation Handbook, Second Edition, Volume 1 - Mikhail J. Atallah 2009-11-20

Algorithms and Theory of Computation Handbook, Second Edition: General Concepts and Techniques provides an up-to-date compendium of fundamental computer science topics and techniques. It also illustrates how the topics and techniques come together to deliver efficient solutions to important practical problems. Along with updating and revising many of the existing chapters, this second edition contains four new chapters that cover external memory and parameterized algorithms as well as computational number theory and algorithmic coding theory. This best-selling handbook continues to help computer professionals and engineers find significant information on various algorithmic topics. The expert contributors clearly define the terminology, present basic results and techniques, and offer a number of current references to the in-depth literature. They also provide a glimpse of the major research issues concerning the relevant topics.

Data Mining Techniques - Arun K. Pujari 2001

This Book Addresses All The Major And Latest Techniques Of Data Mining And Data Warehousing. It Deals With The Latest Algorithms For Discussing Association Rules, Decision Trees, Clustering, Neural Networks And Genetic Algorithms. The Book Also Discusses The Mining Of Web Data, Temporal And Text Data. It Can Serve As A Textbook For Students Of Computer Science, Mathematical Science And Management Science, And Also Be An Excellent Handbook For Researchers In The Area Of Data Mining And Warehousing.

Data Mining - Jiawei Han 2022-10-01

Data Mining: Concepts and Techniques, Fourth Edition provides the theories and methods for processing data or information used in various applications. Specifically, it explains data mining and the tools used in discovering knowledge from collected data, known as KDD. The book focuses on the feasibility, usefulness, effectiveness and scalability of techniques of large datasets. After describing data mining, the authors explain the methods of knowing, preprocessing, processing and warehousing data. They then present 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. Users from computer science students, application developers, business professionals, and researchers who seek information on data mining will find this resource very helpful. 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 needed to get the most out of your data

Data Warehousing - Paul Westerman 2001

What is data warehousing? -- Project planning -- Business exploration -- Business case study and ROI analysis -- Organizational integration -- Technology -- Database maintenance -- Technical construction of the Wal-Mart data warehouse --

Postimplementation of the Wal-Mart data warehouse -- Store operations sample analyses -- Merchandising sample analyses.

Data Mining and Knowledge Discovery Handbook - Oded Maimon 2006-05-28

Data Mining and Knowledge Discovery Handbook organizes all major concepts, theories, methodologies, trends, challenges and applications of data mining (DM) and knowledge discovery in databases (KDD) into a coherent and unified repository. This book first surveys, then provides comprehensive yet concise algorithmic descriptions of methods, including classic methods plus the extensions and novel methods developed recently. This volume concludes with in-depth descriptions of data mining applications in various interdisciplinary industries including finance, marketing, medicine, biology, engineering, telecommunications, software, and security. Data Mining and Knowledge Discovery Handbook is designed for research scientists and graduate-level students in computer science and engineering. This book is also suitable for professionals in fields such as computing applications, information systems management, and strategic research management.

R for Data Science - Hadley Wickham 2016-12-12

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

Data Mining - Jiawei Han 2012

Mining of Data with Complex Structures explores nature of data with complex structure including sequences, trees and graphs. Readers will find a detailed description of the state-of-the-art of sequence mining, tree mining and graph mining, and more.

Data Mining: Concepts And Techniques, 2E - Jiawei Han 2010-01-01

Data Mining: Know It All - Soumen Chakrabarti 2008-10-31

This book brings all of the elements of data mining together in a single volume, saving the reader the time and expense of making multiple purchases. It consolidates both introductory and advanced topics, thereby covering the gamut of data mining and machine learning tactics ? from data integration and pre-processing, to fundamental algorithms, to optimization techniques and web mining methodology. The proposed book expertly combines the finest data mining material from the Morgan Kaufmann portfolio. Individual chapters are derived from a select group of MK books authored by the best and brightest in the field. These chapters are combined into one comprehensive volume in a way that allows it to be used as a reference work for those interested in new and developing aspects of data mining. This book represents a quick and efficient way to unite valuable content from

leading data mining experts, thereby creating a definitive, one-stop-shopping opportunity for customers to receive the information they would otherwise need to round up from separate sources. Chapters contributed by various recognized experts in the field let the reader remain up to date and fully informed from multiple viewpoints. Presents multiple methods of analysis and algorithmic problem-solving techniques, enhancing the reader's technical expertise and ability to implement practical solutions. Coverage of both theory and practice brings all of the elements of data mining together in a single volume, saving the reader the time and expense of making multiple purchases.

Data Mining, Southeast Asia Edition - Jiawei Han 2006-04-06

Our ability to generate and collect data has been increasing rapidly. Not only are all of our business, scientific, and government transactions now computerized, but the widespread use of digital cameras, publication tools, and bar codes also generate data. On the collection side, scanned text and image platforms, satellite remote sensing systems, and the World Wide Web have flooded us with a tremendous amount of data. This explosive growth has generated an even more urgent need for new techniques and automated tools that can help us transform this data into useful information and knowledge. Like the first edition, voted the most popular data mining book by KD Nuggets readers, this book explores concepts and techniques for the discovery of patterns hidden in large data sets, focusing on issues relating to their feasibility, usefulness, effectiveness, and scalability. However, since the publication of the first edition, great progress has been made in the development of new data mining methods, systems, and applications. This new edition substantially enhances the first edition, and new chapters have been added to address recent developments on mining complex types of data— including stream data, sequence data, graph structured data, social network data, and multi-relational data. A comprehensive, practical look at the concepts and techniques you need to know to get the most out of real business data Updates that incorporate input from readers, changes in the field, and more material on statistics and machine learning Dozens of algorithms and implementation examples, all in easily understood pseudo-code and suitable for use in real-world, large-scale data mining projects Complete classroom support for instructors at www.mkp.com/datamining2e companion site

Advanced Engineering Mathematics, 22e - Dass H.K.

"Advanced Engineering Mathematics" is written for the students of all engineering disciplines. Topics such as Partial Differentiation, Differential Equations, Complex Numbers, Statistics, Probability, Fuzzy Sets and Linear Programming which are an important part of all major universities have been well-explained. Filled with examples and in-text exercises, the book successfully helps the student to practice and retain the understanding of otherwise difficult concepts.

Data Preparation for Data Mining - Dorian Pyle 1999-03-22

This book focuses on the importance of clean, well-structured data as the first step to successful data mining. It shows how data should be prepared prior to mining in order to maximize mining performance.

Data Mining and Analysis - Mohammed J. Zaki 2014-05-12

A comprehensive overview of data mining from an algorithmic perspective, integrating related concepts from machine learning and statistics.

Applications of Data Mining in Computer Security - Daniel Barbará 2012-12-06

Data mining is becoming a pervasive technology in activities as diverse as using historical data to predict the success of a marketing campaign, looking for patterns in financial transactions to discover illegal activities or analyzing

genome sequences. From this perspective, it was just a matter of time for the discipline to reach the important area of computer security. Applications Of Data Mining In Computer Security presents a collection of research efforts on the use of data mining in computer security. Applications Of Data Mining In Computer Security concentrates heavily on the use of data mining in the area of intrusion detection. The reason for this is twofold. First, the volume of data dealing with both network and host activity is so large that it makes it an ideal candidate for using data mining techniques. Second, intrusion detection is an extremely critical activity. This book also addresses the application of data mining to computer forensics. This is a crucial area that seeks to address the needs of law enforcement in analyzing the digital evidence.

Information Visualization in Data Mining and Knowledge Discovery - Usama M. Fayyad 2002

This text surveys research from the fields of data mining and information visualisation and presents a case for techniques by which information visualisation can be used to uncover real knowledge hidden away in large databases.

Data Mining: Concepts and Techniques - Jiawei Han 2011-06-09

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

Data Mining and Predictive Analytics - Daniel T. Larose 2015-02-19

Learn methods of data analysis and their application to real-world data sets This updated second edition serves as an introduction to data mining methods and models, including association rules, clustering, neural networks, logistic regression, and multivariate analysis. The authors apply a unified "white box" approach to data mining methods and models. This approach is designed to walk readers through the operations and nuances of the various methods, using small data sets, so readers can gain an insight into the inner workings of the method under review. Chapters provide readers with hands-on analysis problems, representing an opportunity for readers to apply their newly-acquired data mining expertise to solving real problems using large, real-world data sets. Data Mining

and Predictive Analytics: Offers comprehensive coverage of association rules, clustering, neural networks, logistic regression, multivariate analysis, and R statistical programming language Features over 750 chapter exercises, allowing readers to assess their understanding of the new material Provides a detailed case study that brings together the lessons learned in the book Includes access to the companion website, www.dataminingconsultant.com, with exclusive password-protected instructor content Data Mining and Predictive Analytics will appeal to computer science and statistic students, as well as students in MBA programs, and chief executives.

Hadoop in Action - Chuck Lam 2010-11-30

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.

Frequent Pattern Mining - Charu C. Aggarwal 2014-08-29

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.

Java Data Mining: Strategy, Standard, and Practice - Mark F. Hornick 2010-07-26

Whether you are a software developer, systems architect, data analyst, or business analyst, if you want to take advantage of data mining in the development of advanced analytic applications, Java Data Mining, JDM, the new standard now implemented in core DBMS and data mining/analysis software, is a key solution component. This book is the essential guide to the usage of the JDM standard interface, written by contributors to the JDM standard. Data mining introduction - an overview of data mining and the problems it can address across industries; JDM's place in strategic solutions to data mining-related problems JDM essentials - concepts, design approach and design issues, with detailed code examples in Java; a Web Services interface to enable JDM functionality in an SOA environment;

and illustration of JDM XML Schema for JDM objects JDM in practice - the use of JDM from vendor implementations and approaches to customer applications, integration, and usage; impact of data mining on IT infrastructure; a how-to guide for building applications that use the JDM API Free, downloadable KJDM source code referenced in the book available here

Data Mining - Vikram Pudi 2009

Designed to serve as a textbook for undergraduate computer science engineering and MCA students, *Data Mining: Concepts and Techniques* imparts a clear understanding of the algorithms and techniques that can be used to structure large databases and then extract interesting patterns from them.

Lecture Notes in Data Mining - Michael W. Berry 2006

The continual explosion of information technology and the need for better data collection and management methods has made data mining an even more relevant topic of study. Books on data mining tend to be either broad and introductory or focus on some very specific technical aspect of the field. This book is a series of seventeen edited OC student-authored lecturesOCO which explore in depth the core of data mining (classification, clustering and association rules) by offering overviews that include both analysis and insight. The initial chapters lay a framework of data mining techniques by explaining some of the basics such as applications of Bayes Theorem, similarity measures, and decision trees. Before focusing on the pillars of classification, clustering and association rules, the book also considers alternative candidates such as point estimation and genetic algorithms. The book's discussion of classification includes an introduction to decision tree algorithms, rule-based algorithms (a popular alternative to decision trees) and distance-based algorithms. Five of the lecture-chapters are devoted to the concept of clustering or unsupervised classification. The functionality of hierarchical and partitional clustering algorithms is also covered as well as the efficient and scalable clustering algorithms used in large databases. The concept of association rules in terms of basic algorithms, parallel and distributive algorithms and advanced measures that help determine the value of association rules are discussed. The final chapter discusses algorithms for spatial data mining. Sample Chapter(s). Chapter 1: Point Estimation Algorithms (397 KB).

Contents: Point Estimation Algorithms; Applications of Bayes Theorem; Similarity Measures; Decision Trees; Genetic Algorithms; Classification: Distance Based Algorithms; Decision Tree-Based Algorithms; Covering (Rule-Based) Algorithms; Clustering: An Overview; Clustering Hierarchical Algorithms; Clustering Partitional Algorithms; Clustering: Large Databases; Clustering Categorical Attributes; Association Rules: An Overview; Association Rules: Parallel and Distributed Algorithms; Association Rules: Advanced Techniques and Measures; Spatial Mining: Techniques and Algorithms. Readership: An introductory data mining textbook or a technical data mining book for an upper level undergraduate or graduate level course."

Business Modeling and Data Mining - Dorian Pyle 2003-05-17

Business Modeling and Data Mining demonstrates how real world business problems can be formulated so that data mining can answer them. The concepts and techniques presented in this book are the essential building blocks in understanding what models are and how they can be used practically to reveal hidden assumptions and needs, determine problems, discover data, determine costs, and explore the whole domain of the problem. This book articulately explains how to understand both the strategic and tactical aspects of any business problem, identify where the key

leverage points are and determine where quantitative techniques of analysis -- such as data mining -- can yield most benefit. It addresses techniques for discovering how to turn colloquial expression and vague descriptions of a business problem first into qualitative models and then into well-defined quantitative models (using data mining) that can then be used to find a solution. The book completes the process by illustrating how these findings from data mining can be turned into strategic or tactical implementations. · Teaches how to discover, construct and refine models that are useful in business situations · Teaches how to design, discover and develop the data necessary for mining · Provides a practical approach to mining data for all business situations · Provides a comprehensive, easy-to-use, fully interactive methodology for building models and mining data · Provides pointers to supplemental online resources, including a downloadable version of the methodology and software tools.

Mining of Massive Datasets - Jure Leskovec 2014-11-13

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

Principles of Data Mining - Max Bramer 2016-11-09

This book explains and explores the principal techniques of Data Mining, the automatic extraction of implicit and potentially useful information from data, which is increasingly used in commercial, scientific and other application areas. It focuses on classification, association rule mining and clustering. Each topic is clearly explained, with a focus on algorithms not mathematical formalism, and is illustrated by detailed worked examples. The book is written for readers without a strong background in mathematics or statistics and any formulae used are explained in detail. It can be used as a textbook to support courses at undergraduate or postgraduate levels in a wide range of subjects including Computer Science, Business Studies, Marketing, Artificial Intelligence, Bioinformatics and Forensic Science. As an aid to self study, this book aims to help general readers develop the necessary understanding of what is inside the 'black box' so they can use commercial data mining packages discriminately, as well as enabling advanced readers or academic researchers to understand or contribute to future technical advances in the field. Each chapter has practical exercises to enable readers to check their progress. A full glossary of technical terms used is included. This expanded third edition includes detailed descriptions of algorithms for classifying streaming data, both stationary data, where the underlying model is fixed, and data that is time-dependent, where the underlying model changes from time to time - a phenomenon known as concept drift.

Advances in Electrical Engineering and Computational Science - Len Gelman 2009-04-21

Advances in Electrical Engineering and Computational Science contains sixty-one revised and extended research articles written by prominent researchers participating in the conference. Topics covered include Control Engineering, Network Management, Wireless Networks, Biotechnology, Signal Processing, Computational Intelligence, Computational Statistics, Internet Computing, High Performance Computing, and industrial applications. *Advances in Electrical Engineering and Computational Science* will offer the state of art of tremendous advances in electrical engineering and computational science and also serve as an excellent reference work for researchers and graduate students working with/on electrical engineering and computational science.

Data Mining: Introductory And Advanced Topics - Margaret H Dunham 2006-09