

Han Kamber Data Mining Concepts 5th Edition

Data Mining: Practical Machine Learning Tools and Techniques, Third Edition, offers a thorough grounding in machine learning concepts as well as practical advice on applying machine learning tools and techniques in real-world data mining situations. This highly anticipated third edition of the most acclaimed work on data mining and machine learning will teach you everything you need to know about preparing inputs, interpreting outputs, evaluating results, and the algorithmic methods at the heart of successful data mining. Thorough updates reflect the technical changes and modernizations that have taken place in the field since the last edition, including new material on Data Transformations, Ensemble Learning, Massive Data Sets, Multi-instance Learning, plus a new version of the popular Weka machine learning software developed by the authors. Witten, Frank, and Hall include both tried-and-true techniques of today as well as methods at the leading edge of contemporary research. The book is targeted at information systems practitioners, programmers, consultants, developers, information technology managers, specification writers, data analysts, data modelers, database R&D professionals, data warehouse engineers, data mining professionals. The book will also be useful for professors and students of upper-level undergraduate and graduate-level data mining and machine learning courses who want to incorporate data mining as part of their data management knowledge base and expertise. Provides a thorough grounding in machine learning concepts as well as practical advice on applying the tools and techniques to your data mining projects Offers concrete tips and techniques for performance improvement that work by transforming the input or output in machine learning methods Includes downloadable Weka software toolkit, a collection of machine learning algorithms for data mining tasks—in an updated, interactive interface. Algorithms in toolkit cover: data pre-processing, classification, regression, clustering, association rules, visualization 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 for Business Analytics: Concepts, Techniques, and Applications in XLMiner®, Third Edition presents an applied approach to data mining and predictive analytics with clear exposition, hands-on exercises, and real-life case studies. Readers will work with all of the standard data mining methods using the Microsoft® Office Excel® add-in XLMiner® to develop predictive models and learn how to obtain business value from Big Data. Featuring updated topical coverage on text mining, social network analysis, collaborative filtering, ensemble methods, uplift modeling and more, the Third Edition also includes: Real-world examples to build a theoretical and practical understanding of key data mining methods End-of-chapter exercises that help readers better understand the presented material Data-rich case studies to illustrate various applications of data mining techniques Completely new chapters on social network analysis and text mining A companion site with additional data sets, instructors material that include solutions to exercises and case studies, and Microsoft PowerPoint® slides <https://www.dataminingbook.com> Free 140-day license to use XLMiner for Education software Data Mining for Business Analytics: Concepts, Techniques, and Applications in XLMiner®, Third Edition is an ideal textbook for upper-

undergraduate and graduate-level courses as well as professional programs on data mining, predictive modeling, and Big Data analytics. The new edition is also a unique reference for analysts, researchers, and practitioners working with predictive analytics in the fields of business, finance, marketing, computer science, and information technology. Praise for the Second Edition "...full of vivid and thought-provoking anecdotes... needs to be read by anyone with a serious interest in research and marketing." – Research Magazine "Shmueli et al. have done a wonderful job in presenting the field of data mining - a welcome addition to the literature." – ComputingReviews.com "Excellent choice for business analysts...The book is a perfect fit for its intended audience." – Keith McCormick, Consultant and Author of SPSS Statistics For Dummies, Third Edition and SPSS Statistics for Data Analysis and Visualization Galit Shmueli, PhD, is Distinguished Professor at National Tsing Hua University's Institute of Service Science. She has designed and instructed data mining courses since 2004 at University of Maryland, Statistics.com, The Indian School of Business, and National Tsing Hua University, Taiwan. Professor Shmueli is known for her research and teaching in business analytics, with a focus on statistical and data mining methods in information systems and healthcare. She has authored over 70 journal articles, books, textbooks and book chapters. Peter C. Bruce is President and Founder of the Institute for Statistics Education at www.statistics.com. He has written multiple journal articles and is the developer of Resampling Stats software. He is the author of *Introductory Statistics and Analytics: A Resampling Perspective*, also published by Wiley. Nitin R. Patel, PhD, is Chairman and cofounder of Cytel, Inc., based in Cambridge, Massachusetts. A Fellow of the American Statistical Association, Dr. Patel has also served as a Visiting Professor at the Massachusetts Institute of Technology and at Harvard University. He is a Fellow of the Computer Society of India and was a professor at the Indian Institute of Management, Ahmedabad for 15 years. This book offers detailed surveys and systematic discussion of models, algorithms and applications for link mining, focusing on theory and technique, and related applications: text mining, social network analysis, collaborative filtering and bioinformatics. This book constitutes the refereed proceedings of the 8th Pacific-Asia Conference on Knowledge Discovery and Data mining, PAKDD 2004, held in Sydney, Australia in May 2004. The 50 revised full papers and 31 revised short papers presented were carefully reviewed and selected from a total of 238 submissions. The papers are organized in topical sections on classification; clustering; association rules; novel algorithms; event mining, anomaly detection, and intrusion detection; ensemble learning; Bayesian network and graph mining; text mining; multimedia mining; text mining and Web mining; statistical methods, sequential data mining, and time series mining; and biomedical data mining. *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.

This volume contains the Proceedings of ICFCA 2004, the 2nd International Conference on Formal Concept Analysis. The ICFCA conference series aims to be the premier forum for the publication of advances in applied lattice and order theory and in particular scientific advances related to formal concept analysis. Formal concept analysis emerged in the 1980s from efforts to restructure lattice theory to promote better communication between lattice theorists and potential users of lattice theory. Since then, the field has developed into a growing research area in its own right with a thriving theoretical community and an increasing number of applications in data and knowledge processing including data visualization, information retrieval, machine learning, data analysis and knowledge management. In terms of theory, formal concept analysis has been extended into attribute exploration, Boolean judgment, contextual logic and so on to create a powerful general framework for knowledge representation and reasoning. This conference aims to unify theoretical and applied practitioners who use formal concept analysis, drawing on the fields of mathematics, computer and library sciences and software engineering. The theme of the 2004 conference was 'Concept Lattices' to acknowledge the colloquial term used for the line diagrams that appear in almost every paper in this volume. ICFCA 2004 included tutorial sessions, demonstrating the practical benefits of formal concept analysis, and highlighted developments in the foundational theory and standards. The conference showcased the increasing variety of formal concept analysis software and included eight invited lectures from distinguished speakers in the field. Seven of the eight invited speakers submitted accompanying papers and these were reviewed and appear in this volume.

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: 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.

The two-volume set LNAI 8346 and 8347 constitutes the thoroughly refereed proceedings of the 9th International Conference on Advanced Data Mining and Applications, ADMA 2013, held in Hangzhou, China, in December 2013. The 32 regular papers and 64 short papers presented in these two volumes were carefully reviewed and selected from 222 submissions. The papers included in these two volumes cover the following topics: opinion mining, behavior mining, data stream mining, sequential data mining, web mining, image mining, text mining, social network mining, classification, clustering, association rule mining, pattern mining, regression, predication, feature extraction, identification, privacy preservation, applications, and machine learning.

Presents the latest techniques for analyzing and extracting information from large amounts of data in high-dimensional data spaces The revised and updated third edition of Data Mining contains in one volume an introduction to a systematic approach to the analysis of large data sets that integrates results from disciplines such as statistics, artificial intelligence, data bases, pattern recognition, and computer visualization. Advances in deep learning technology have opened an entire new spectrum of applications. The author—a noted expert on the topic—explains the basic concepts, models, and methodologies that have been developed in recent years. This new edition introduces and expands on many topics, as well as providing revised sections on software tools and data mining applications. Additional changes include an updated list of references for further

topical sections on Web information modeling; payment and security; information extraction; advanced applications; performance issues; linkage analysis and document clustering; Web caching and content analysis; XML query processing; Web search and personalization; workflow management and enterprise information systems; business processes; deep Web and dynamic content; Web information systems design; ontologies and applications; multimedia, user interfaces, and languages; and peer-to-peer and grid systems.

Introduction to Data Mining presents fundamental concepts and algorithms for those learning data mining for the first time. Each concept is explored thoroughly and supported with numerous examples. The text requires only a modest background in mathematics. Each major topic is organized into two chapters, beginning with basic concepts that provide necessary background for understanding each data mining technique, followed by more advanced concepts and algorithms.

In this book the author provides the reader with a comprehensive coverage of data mining topics and algorithms. Data base perspective is maintained throughout the book which provides students with a focused discussion of algorithms, data structures, data types and complexity of algorithms and space. It also emphasizes the use of data mining concepts in real-world applications with large database components.

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

Special Features:

- Best-in-class data mining techniques for solving critical problems in all areas of business
- Explains how to pick the right data mining techniques for specific problems
- Shows how to perform analysis and evaluate results
- Features real-world examples from across various industry sectors
- Companion Web site with updates on data mining products and service providers

About The Book: Companies have invested in building data warehouses to capture vast amounts of customer information. The payoff comes with mining or getting access to the data within this information gold mine to make better business decisions.

Readers and reviewers loved Berry and Linoff's first book, Data Mining Techniques, because the authors so clearly illustrate practical techniques with real benefits for improved marketing and sales. Mastering Data Mining takes off from there-assuming readers know the basic techniques covered in the first book, the authors focus on how to best apply these techniques to real business cases. They start with simple applications and work up to the most powerful and sophisticated examples over the course of about 20 cases. (Ralph Kimball used this same approach in his highly successful Data Warehouse Toolkit). As with their first book, Mastering Data Mining is sufficiently technical for database analysts, but is accessible to technically

savvy business and marketing managers. It should also appeal to a new breed of database marketing managers.

This book constitutes the thoroughly refereed post-proceedings of three workshops and an industrial track held in conjunction with the 11th Pacific-Asia Conference on Knowledge Discovery and Data Mining, PAKDD 2007, held in Nanjing, China in May 2007. The 62 revised full papers presented together with an overview article to each workshop were carefully reviewed and selected from 355 submissions.

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

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.

This volume comprises the proceedings of the Industrial Conference on Data Mining (ICDM 2009) held in Leipzig (www.data-mining-forum.de). For this edition the Program Committee received 130 submissions. After the peer-review process, we accepted 32 high-quality papers for oral presentation that are included in this book. The topics range from theoretical aspects of data mining to applications of data mining, such as on multimedia data, in marketing, finance and telecommunication, in medicine and agriculture, and in process control, industry and society. Ten papers were selected for poster presentations that are published in the ICDM Poster Proceedings Volume by ibai-publishing (www.ibai-publishing.org). In conjunction with ICDM two workshops were run focusing on special hot application-oriented topics in data mining. The workshop Data Mining in Marketing DMM 2009 was run for the second time. The papers are published in a separate workshop book “Advances in Data Mining on Marketing” by ibai-

publishing (www.ibai-publishing.org). The Workshop on Case-Based Reasoning for Multimedia Data CBR-MD ran for the second year. The papers are published in a special issue of the International Journal of Transactions on Case-Based Reasoning (www.ibai-publishing.org/journal/cbr).

"This book describes the process of analyzing data. The authors have extensive experience both managing data analysts and conducting their own data analyses, and this book is a distillation of their experience in a format that is applicable to both practitioners and managers in data science."--Leanpub.com.

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

Essay from the year 2012 in the subject Computer Science - Theory, grade: B, Atlantic International University (School of Science and Engineering), course: Doctorate in Information Technology, language: English, abstract: Data mining is an independent science that based on advanced ways for information retrieval. Data mining is dealing with knowledge discovery in data warehouses without predefined hypotheses. So it is quite different from other applications such as decision support systems, OLAP and others which are looking for information on the factors and assumptions that we know it in advance. Data Mining supports multiple algorithms which have the ability to adopt automatic classification of historical data and predict future events. Data mining in the databases is designed to extract the hidden information, and it is a modern technology that imposed itself strongly in the information revolution, in the light of the great technological development and widespread use of data warehouses. Data mining techniques focus on building future forecasts and explore the behavior and trends, allowing a good estimation for right decisions that taken in a timely manner. This paper provides a general definition of data mining science and its most important techniques and algorithms used.

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

This book constitutes the refereed proceedings of the 13th Industrial Conference on Data Mining, ICDM 2013, held in New York, NY, in July 2013. The 22 revised full papers presented were carefully reviewed and selected from 112 submissions. The topics range from theoretical aspects of data mining to applications of data mining, such as in multimedia data, in marketing, finance and telecommunication, in medicine and agriculture, and in process control, industry

and society.

As the first book devoted to relational data mining, this coherently written multi-author monograph provides a thorough introduction and systematic overview of the area. The first part introduces the reader to the basics and principles of classical knowledge discovery in databases and inductive logic programming; subsequent chapters by leading experts assess the techniques in relational data mining in a principled and comprehensive way; finally, three chapters deal with advanced applications in various fields and refer the reader to resources for relational data mining. This book will become a valuable source of reference for R&D professionals active in relational data mining. Students as well as IT professionals and ambitious practitioners interested in learning about relational data mining will appreciate the book as a useful text and gentle introduction to this exciting new field.

Data Preparation for Data Mining addresses an issue unfortunately ignored by most authorities on data mining: data preparation. Thanks largely to its perceived difficulty, data preparation has traditionally taken a backseat to the more alluring question of how best to extract meaningful knowledge. But without adequate preparation of your data, the return on the resources invested in mining is certain to be disappointing. Dorian Pyle corrects this imbalance. A twenty-five-year veteran of what has become the data mining industry, Pyle shares his own successful data preparation methodology, offering both a conceptual overview for managers and complete technical details for IT professionals. Apply his techniques and watch your mining efforts pay off-in the form of improved performance, reduced distortion, and more valuable results. On the enclosed CD-ROM, you'll find a suite of programs as C source code and compiled into a command-line-driven toolkit. This code illustrates how the author's techniques can be applied to arrive at an automated preparation solution that works for you. Also included are demonstration versions of three commercial products that help with data preparation, along with sample data with which you can practice and experiment. * Offers in-depth coverage of an essential but largely ignored subject. * Goes far beyond theory, leading you-step by step-through the author's own data preparation techniques. * Provides practical illustrations of the author's methodology using realistic sample data sets. * Includes algorithms you can apply directly to your own project, along with instructions for understanding when automation is possible and when greater intervention is required. * Explains how to identify and correct data problems that may be present in your application. * Prepares miners, helping them head into preparation with a better understanding of data sets and their limitations.

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."

Real-world physical and abstract data objects are interconnected, forming gigantic, interconnected networks. By structuring these data objects and interactions between these objects into multiple types, such networks become semi-structured heterogeneous information networks. Most real-world applications that handle big data, including interconnected social media and social networks, scientific, engineering, or medical information systems, online e-commerce systems, and most database systems, can be structured into heterogeneous information networks. Therefore, effective analysis of large-scale heterogeneous information networks poses an interesting but critical challenge. In this book, we investigate the principles and methodologies of mining heterogeneous information networks. Departing from many existing network models that view interconnected data as homogeneous graphs or networks, our semi-structured heterogeneous information network model leverages the rich semantics of typed nodes and links in a network and uncovers surprisingly rich knowledge from the network. This semi-structured heterogeneous network modeling leads to a series of new principles and powerful methodologies for mining interconnected data, including: (1) rank-based clustering and classification; (2) meta-path-based similarity search and mining; (3) relation strength-aware mining, and many other potential developments. This book introduces this new research frontier and points out some promising research directions. Table of Contents: Introduction / Ranking-Based Clustering / Classification of Heterogeneous Information Networks / Meta-Path-Based Similarity Search / Meta-Path-Based Relationship Prediction / Relation Strength-Aware Clustering with Incomplete Attributes / User-Guided Clustering

via Meta-Path Selection / Research Frontiers

The fundamental algorithms in data mining and machine learning form the basis of data science, utilizing automated methods to analyze patterns and models for all kinds of data in applications ranging from scientific discovery to business analytics. This textbook for senior undergraduate and graduate courses provides a comprehensive, in-depth overview of data mining, machine learning and statistics, offering solid guidance for students, researchers, and practitioners. The book lays the foundations of data analysis, pattern mining, clustering, classification and regression, with a focus on the algorithms and the underlying algebraic, geometric, and probabilistic concepts. New to this second edition is an entire part devoted to regression methods, including neural networks and deep learning.

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

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 Elsevier

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. Whether you are a seasoned professional or a new student of data mining, this book has much to offer you: * 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.

[Copyright: cbe9f4868d414f0a49fbdf0f3aaeabda](http://www.mkp.com/datamining2e)