



Yu, UIC Distinguished Professor and Wexler Chair in Information Technology at University of Illinois at Chicago

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](http://www.mkp.com/datamining2e) companion site

Written in lucid language, this valuable textbook brings together fundamental concepts of data mining and data warehousing in a single volume. Important topics including information theory, decision tree, Naïve Bayes classifier, distance metrics, partitioning clustering, associate mining, data marts and operational data store are discussed comprehensively. The textbook is written to cater to the needs of undergraduate students of computer science, engineering and information technology for a course on data mining and data warehousing. The text simplifies the understanding of the concepts through exercises and practical examples. Chapters such as classification, associate mining and cluster analysis are discussed in detail with their practical implementation using Weka and R language data mining tools. Advanced topics including big data analytics, relational data models and NoSQL are discussed in detail. Pedagogical features including unsolved problems and multiple-choice questions are interspersed throughout the book for better understanding.

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

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.

"This book addresses existing solutions for data mining, with particular emphasis on potential real-world applications. It captures defining research on topics such as fuzzy set theory, clustering algorithms, semi-supervised clustering, modeling and managing data mining patterns, and sequence motif mining"--Provided by publisher.

Challenging the focus on great powers in the international debate, this book explores how rising middle power states are engaging with emerging major military innovations and analyses how this will affect the stability and security of the Indo Pacific. Presenting a data-based analysis of how middle power actors in the Indo-Pacific are responding to the emergence of military Artificial Intelligence and Killer Robots, the book asserts that continuing to exclude non-great power actors from our thinking in this field enables the dangerous diffusion of Lethal Autonomous Weapon Systems (LAWS) to smaller states and terrorist groups, and demonstrates the disruptive effects of these military innovations on the balance of power in the Indo-Pacific. Offering a detailed analysis of the resource capacities of China, United States, Singapore and Indonesia, it shows how major military innovation acts as a circuit breaker between competitor states disrupting the conventional superiority of the dominant hegemonic state and giving a successful adopter a distinct advantage over their opponent. This book will appeal to researchers, end-users in the military and law enforcement communities, and policymakers. It will also be a valuable resource for researchers interested in strategic stability for the broader Asia-Pacific and the role of middle power states in hegemonic power transition and conflict.

With an emphasis on social science applications, Event History Analysis with R presents an introduction to survival and event history analysis using real-life examples. Keeping mathematical details to a minimum, the book covers key topics, including both discrete and continuous time data, parametric proportional hazards, and accelerated failure times. Features Introduces parametric proportional hazards models with baseline distributions like the Weibull, Gompertz, Lognormal, and Piecewise constant hazard distributions, in addition to traditional Cox regression Presents mathematical details as well as technical material in an appendix Includes real examples with applications in demography, econometrics, and epidemiology Provides a dedicated R package, eha, containing special treatments, including making cuts in the Lexis diagram, creating communal covariates, and creating period statistics A much-needed primer, Event History Analysis with R is a didactically excellent resource for students and practitioners of applied event history and survival analysis.

A Fruitful Field for Researching Data Mining Methodology and for Solving Real-Life Problems Contrast Data Mining: Concepts, Algorithms, and Applications collects recent results from this specialized area of data mining that have previously been scattered in the literature, making them more accessible to researchers and developers in data mining and other fields. The book not only presents concepts and techniques for contrast data mining, but also explores the use of contrast mining to solve challenging problems in various scientific, medical, and business domains. Learn from Real Case Studies of Contrast

Mining Applications In this volume, researchers from around the world specializing in architecture engineering, bioinformatics, computer science, medicine, and systems engineering focus on the mining and use of contrast patterns. They demonstrate many useful and powerful capabilities of a variety of contrast mining techniques and algorithms, including tree-based structures, zero-suppressed binary decision diagrams, data cube representations, and clustering algorithms. They also examine how contrast mining is used in leukemia characterization, discriminative gene transfer and microarray analysis, computational toxicology, spatial and image data classification, voting analysis, heart disease prediction, crime analysis, understanding customer behavior, genetic algorithms, and network security. 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.

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."

For the first time, a reference on the most relevant applications of adaptive filtering techniques. Top researchers in the field contributed chapters addressing applications in acoustics, speech, wireless and networking, where research is still very active and open.

Created with the input of a distinguished International Board of the foremost authorities in data mining from academia and industry, The Handbook of Data Mining presents comprehensive coverage of data mining concepts and techniques. Algorithms, methodologies, management issues, and tools are all illustrated through engaging examples and real-world

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.

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

This book reviews state-of-the-art methodologies and techniques for analyzing enormous quantities of raw data in high-dimensional data spaces, to extract new information for decision making. The goal of this book is to provide a single introductory source, organized in a systematic way, in which we could direct the readers in analysis of large data sets, through the explanation of basic concepts, models and methodologies developed in recent decades. If you are an instructor or professor and would like to obtain instructor's materials, please visit <http://booksupport.wiley.com> If you are an instructor or professor and would like to obtain a solutions manual, please send an email to: [pressbooks@ieee.org](mailto:pressbooks@ieee.org)

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 study, and an extended list of problems and questions that relate to each chapter. This third edition presents new and expanded information that: • Explores big data and cloud computing • Examines deep learning • Includes information on convolutional neural networks (CNN) • Offers reinforcement learning • Contains semi-supervised learning and S3VM • Reviews model evaluation for unbalanced data Written for graduate students in computer science, computer engineers, and computer information systems professionals, the updated third edition of Data Mining continues to provide an essential guide to the basic principles of the technology and the most recent developments in the field.

Data Mining for Business Analytics: Concepts, Techniques, and Applications in Python presents an applied approach to data mining concepts and methods, using Python software for illustration Readers will learn how to implement a variety of popular data mining algorithms in Python (a free and open-source software) to tackle business problems and opportunities. This is the sixth version of this successful text, and the first using Python. It covers both statistical and machine learning algorithms for prediction, classification, visualization, dimension reduction, recommender systems, clustering, text mining and network analysis. It also includes: A new co-author, Peter Gedeck, who brings both experience teaching business analytics courses using Python, and expertise in the application of machine learning methods to the drug-discovery process A new section on ethical issues in data mining Updates and new material based on feedback from instructors teaching MBA, undergraduate, diploma and executive courses, and from their students More than a dozen case studies demonstrating applications for the data mining techniques described End-of-chapter exercises that help readers gauge and expand their comprehension and competency of the material presented A companion website with more than two dozen data sets, and instructor materials including exercise solutions, PowerPoint slides, and case solutions Data Mining for Business Analytics: Concepts, Techniques, and Applications in Python is an ideal textbook for graduate and upper-undergraduate level courses in data mining, predictive analytics, and business analytics.

This new edition is also an excellent reference for analysts, researchers, and practitioners working with quantitative methods in the fields of business, finance, marketing, computer science, and information technology. "This book has by far the most comprehensive review of business analytics methods that I have ever seen, covering everything from classical approaches such as linear and logistic regression, through to modern methods like neural networks, bagging and boosting, and even much more business specific procedures such as social network analysis and text mining. If not the bible, it is at the least a

definitive manual on the subject.” —Gareth M. James, University of Southern California and co-author (with Witten, Hastie and Tibshirani) of the best-selling book *An Introduction to Statistical Learning*, with Applications in R

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.

The knowledge discovery process is as old as *Homo sapiens*. Until some time ago this process was solely based on the 'natural personal' computer provided by Mother Nature. Fortunately, in recent decades the problem has begun to be solved based on the development of the Data mining technology, aided by the huge computational power of the 'artificial' computers. Digging intelligently in different large databases, data mining aims to extract implicit, previously unknown and potentially useful information from data, since “knowledge is power”. The goal of this book is to provide, in a friendly way, both theoretical concepts and, especially, practical techniques of this exciting field, ready to be applied in real-world situations. Accordingly, it is meant for all those who wish to learn how to explore and analysis of large quantities of data in order to discover the hidden nugget of information.

Thought-provoking and accessible in approach, this updated and expanded second edition of the *Data Mining: Concepts and Techniques: Concepts and Techniques* (The Morgan Kaufma provides a user-friendly introduction to the subject, Taking a clear structural framework, it guides the reader through the subject's core elements. A flowing writing style combines with the use of illustrations and diagrams throughout the text to ensure the reader understands even the most complex of concepts. This succinct and enlightening overview is a required reading for advanced graduate-level students. We hope you find this book useful in shaping your future career. Feel free to send us your enquiries related to our publications to [info@risepress.pw](mailto:info@risepress.pw) Rise Press

Many companies have invested in building large databases and data warehouses capable of storing vast amounts of information. This book offers business, sales and marketing managers a practical guide to accessing such information.

This book covers the fundamental concepts of data mining, to demonstrate the potential of gathering large sets of data, and analyzing these data sets to gain useful business understanding. The book is organized in three parts. Part I introduces concepts. Part II describes and demonstrates basic data mining algorithms. It also contains chapters on a number of different techniques often used in data mining. Part III focuses on business applications of data mining.

*Data Mining* introduces in clear and simple ways how to use existing data mining methods to obtain effective solutions for a variety of management and engineering design problems. *Data Mining* is organised into two parts: the first provides a focused introduction to data mining and the second goes into greater depth on subjects such as customer analysis. It covers almost all managerial activities of a company, including: • supply chain design, • product development, • manufacturing system design, • product quality control, and • preservation of privacy.

Incorporating recent developments of data mining that have made it possible to deal with management and engineering design problems with greater efficiency and efficacy, *Data Mining* presents a number of state-of-the-art topics. It will be an informative source of information for researchers, but will also be a useful reference work for industrial and managerial practitioners. In recent years, the science of managing and analyzing large datasets has emerged as a critical area of research. In the race to answer vital questions and make knowledgeable decisions, impressive amounts of data are now being generated at a rapid pace, increasing the opportunities and challenges associated with the ability to effectively analyze this data.

*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](http://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.

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.

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: 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 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](http://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

[Copyright: 622e475c03fa12daa5558de4182f3deb](http://www.LearnPredictiveAnalytics.com)