

Data Classification Algorithms And Applications Chapman Hallcrc Data Mining And Knowledge Discovery Series

This book provides a unified framework that describes how genetic learning can be used to design pattern recognition and learning systems. It examines how a search technique, the genetic algorithm, can be used for pattern classification mainly through approximating decision boundaries. Coverage also demonstrates the effectiveness of the genetic classifiers vis-à-vis several widely used classifiers, including neural networks.

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

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.

This book provides a general and comprehensible overview of imbalanced learning. It contains a formal description of a problem, and focuses on its main features, and the most relevant proposed solutions. Additionally, it considers the different scenarios in Data Science for which the imbalanced classification can create a real challenge. This book stresses the gap with standard classification tasks by reviewing the case studies and ad-hoc performance metrics that are applied in this area. It also covers the different approaches that have been traditionally applied to address the binary skewed class distribution. Specifically, it reviews cost-sensitive learning, data-level preprocessing methods and algorithm-level solutions, taking also into account those ensemble-learning solutions that embed any of the former alternatives. Furthermore, it focuses on the extension of the problem for multi-class problems, where the former classical methods are no longer to be applied in a straightforward way. This book also focuses on the data intrinsic characteristics that are the main causes which, added to the uneven class distribution, truly hinders the performance of classification algorithms in this scenario. Then, some notes on data reduction are provided in order to understand the advantages related to the use of this type of approaches. Finally this book introduces some novel areas of study that are gathering a deeper attention on the imbalanced data issue. Specifically, it considers the classification of data streams, non-classical classification problems, and the scalability related to Big Data. Examples of software libraries and modules to address imbalanced classification are provided. This book is highly suitable for technical professionals, senior undergraduate and graduate students in the areas of data science, computer science and engineering. It will also be useful for scientists and researchers to gain insight on the current developments in this area of study, as well as future research directions. Providing a unique approach to machine learning, this text contains fresh and intuitive, yet rigorous, descriptions of all fundamental concepts necessary to conduct research, build products, tinker, and play. By prioritizing geometric intuition, algorithmic thinking, and practical real world applications in disciplines including computer vision, natural language processing, economics, neuroscience, recommender systems, physics, and biology, this text provides readers with both a lucid understanding of foundational material as well as the practical tools needed to solve real-world problems. With in-depth Python and MATLAB/OCTAVE-based computational exercises and a complete treatment of cutting edge numerical optimization techniques, this is an essential resource for students and an ideal reference for researchers and practitioners working in machine learning, computer science, electrical engineering, signal processing, and numerical optimization.

"This book investigates machine learning (ML), one of the most fruitful fields of current research, both in the proposal of new techniques and theoretic algorithms and in their application to real-life problems"--Provided by publisher.

Identifying some of the most influential algorithms that are widely used in the data mining community, The Top Ten Algorithms in Data Mining provides a description of each algorithm, discusses its impact, and reviews current and future research. Thoroughly evaluated by independent reviewers, each chapter focuses on a particular algorithm and is written by either the original authors of the algorithm or world-class researchers who have extensively studied the respective algorithm. The book concentrates on the following important algorithms: C4.5, k-Means, SVM, Apriori, EM, PageRank, AdaBoost, kNN, Naive Bayes, and CART. Examples illustrate how each algorithm works and highlight its overall

performance in a real-world application. The text covers key topics—including classification, clustering, statistical learning, association analysis, and link mining—in data mining research and development as well as in data mining, machine learning, and artificial intelligence courses. By naming the leading algorithms in this field, this book encourages the use of data mining techniques in a broader realm of real-world applications. It should inspire more data mining researchers to further explore the impact and novel research issues of these algorithms.

Big data consists of data sets that are too large and complex for traditional data processing and data management applications. Therefore, to obtain the valuable information within the data, one must use a variety of innovative analytical methods, such as web analytics, machine learning, and network analytics. As the study of big data becomes more popular, there is an urgent demand for studies on high-level computational intelligence and computing services for analyzing this significant area of information science. Big Data Analytics for Sustainable Computing is a collection of innovative research that focuses on new computing and system development issues in emerging sustainable applications. Featuring coverage on a wide range of topics such as data filtering, knowledge engineering, and cognitive analytics, this publication is ideally designed for data scientists, IT specialists, computer science practitioners, computer engineers, academicians, professionals, and students seeking current research on emerging analytical techniques and data processing software.

This book presents machine learning models and algorithms to address big data classification problems. Existing machine learning techniques like the decision tree (a hierarchical approach), random forest (an ensemble hierarchical approach), and deep learning (a layered approach) are highly suitable for the system that can handle such problems. This book helps readers, especially students and newcomers to the field of big data and machine learning, to gain a quick understanding of the techniques and technologies; therefore, the theory, examples, and programs (Matlab and R) presented in this book have been simplified, hardcoded, repeated, or spaced for improvements. They provide vehicles to test and understand the complicated concepts of various topics in the field. It is expected that the readers adopt these programs to experiment with the examples, and then modify or write their own programs toward advancing their knowledge for solving more complex and challenging problems. The presentation format of this book focuses on simplicity, readability, and dependability so that both undergraduate and graduate students as well as new researchers, developers, and practitioners in this field can easily trust and grasp the concepts, and learn them effectively. It has been written to reduce the mathematical complexity and help the vast majority of readers to understand the topics and get interested in the field. This book consists of four parts, with the total of 14 chapters. The first part mainly focuses on the topics that are needed to help analyze and understand data and big data. The second part covers the topics that can explain the systems required for processing big data. The third part presents the topics required to understand and select machine learning techniques to classify big data. Finally, the fourth part concentrates on the topics that explain the scaling-up machine learning, an important solution for modern big data problems.

At the intersection of computer science and healthcare, data analytics has emerged as a promising tool for solving problems across many healthcare-related disciplines. Supplying a comprehensive overview of recent healthcare analytics research, Healthcare Data Analytics provides a clear understanding of the analytical techniques currently available to solve healthcare problems. The book details novel techniques for acquiring, handling, retrieving, and making best use of healthcare data. It analyzes recent developments in healthcare computing and discusses emerging technologies that can help improve the health and well-being of patients. Written by prominent researchers and experts working in the healthcare domain, the book sheds light on many of the computational challenges in the field of medical informatics.

Each chapter in the book is structured as a "survey-style" article discussing the prominent research issues and the advances made on that research topic. The book is divided into three major categories: Healthcare Data Sources and Basic Analytics - details the various healthcare data sources and analytical techniques used in the processing and analysis of such data Advanced Data Analytics for Healthcare - covers advanced analytical methods, including clinical prediction models, temporal pattern mining methods, and visual analytics Applications and Practical Systems for Healthcare - covers the applications of data analytics to pervasive healthcare, fraud detection, and drug discovery along with systems for medical imaging and decision support Computer scientists are usually not trained in domain-specific medical concepts, whereas medical practitioners and researchers have limited exposure to the data analytics area. The contents of this book will help to bring together these diverse communities by carefully and comprehensively discussing the most relevant contributions from each domain.

"This edited book will start with an introduction to feature engineering and then move onto recent concepts, methods and applications with the use of various data types that includes: text, image, streaming data, social network data, financial data, biomedical data, bioinformatics etc. to help readers gain insight into how features can be extracted and transformed from raw data"--

Data Mining: A Tutorial-Based Primer, Second Edition provides a comprehensive introduction to data mining with a focus on model building and testing, as well as on interpreting and validating results. The text guides students to understand how data mining can be employed to solve real problems and recognize whether a data mining solution is a feasible alternative for a specific problem. Fundamental data mining strategies, techniques, and evaluation methods are presented and implemented with the help of two well-known software tools. Several new topics have been added to the second edition including an introduction to Big Data and data analytics, ROC curves, Pareto lift charts, methods for handling large-sized, streaming and imbalanced data, support vector machines, and extended coverage of textual data mining. The second edition contains tutorials for attribute selection, dealing with imbalanced data, outlier analysis, time series analysis, mining textual data, and more. The text provides in-depth coverage of RapidMiner Studio and Weka's Explorer interface. Both software tools are used for stepping students through the tutorials depicting the knowledge

discovery process. This allows the reader maximum flexibility for their hands-on data mining experience.

Data Classification Algorithms and Applications CRC Press

The aim of this book is to provide insight into Data Science and Artificial Learning Techniques based on Industry 4.0, conveys how Machine Learning & Data Science are becoming an essential part of industrial and academic research. Varying from healthcare to social networking and everywhere hybrid models for Data Science, AI, and Machine Learning are being used. The book describes different theoretical and practical aspects and highlights how new systems are being developed. Along with focusing on the research trends, challenges and future of AI in Data Science, the book explores the potential for integration of advanced AI algorithms, addresses the challenges of Data Science for Industry 4.0, covers different security issues, includes qualitative and quantitative research, and offers case studies with working models. This book also provides an overview of AI and Data Science algorithms for readers who do not have a strong mathematical background. Undergraduates, postgraduates, academicians, researchers, and industry professionals will benefit from this book and use it as a guide.

This compendium is a completely revised version of an earlier book, *Data Mining in Time Series Databases*, by the same editors. It provides a unique collection of new articles written by leading experts that account for the latest developments in the field of time series and data stream mining. The emerging topics covered by the book include weightless neural modeling for mining data streams, using ensemble classifiers for imbalanced and evolving data streams, document stream mining with active learning, and many more. In particular, it addresses the domain of streaming data, which has recently become one of the emerging topics in Data Science, Big Data, and related areas. Existing titles do not provide sufficient information on this topic. Contents: Streaming Data Mining with Massive Online Analytics (MOA) (Albert Bifet, Jesse Read, Geoff Holmes and Bernhard Pfahringer) Weightless Neural Modeling for Mining Data Streams (Douglas O Cardoso, João Gama and Felipe França) Ensemble Classifiers for Imbalanced and Evolving Data Streams (Dariusz Brzezinski and Jerzy Stefanowski) Consensus Learning for Sequence Data (Andreas Nienkötter and Xiaoyi Jiang) Clustering-Based Classification of Document Streams with Active Learning (Mark Last, Maxim Stolar and Menahem Friedman) Supporting the Mining of Big Data by Means of Domain Knowledge During the Pre-mining Phases (Rémon Cornelisse and Sunil Choenni) Data Analytics: Industrial Perspective & Solutions for Streaming Data (Mohsin Munir, Sebastian Baumbach, Ying Gu, Andreas Dengel and Sheraz Ahmed) Readership: Researchers, academics, professionals and graduate students in artificial intelligence, machine learning, databases, and information science. Keywords: Time Series; Data Streams; Big Data; Internet of Things; Concept Drift; Sequence Mining; Episode Mining; Incremental Learning; Active Learning Review: 0

Data mining aims at finding interesting, useful or profitable information in very large databases. The enormous increase in the size of available scientific and commercial databases (data avalanche) as well as the continuing and exponential growth in performance of present day computers make data mining a very active field. In many cases, the burgeoning volume of data sets has grown so large that it threatens to overwhelm rather than enlighten scientists. Therefore, traditional methods are revised and streamlined, complemented by many new methods to address challenging new problems. Mathematical Programming plays a key role in this endeavor. It helps us to formulate precise objectives (e.g., a clustering criterion or a measure of discrimination) as well as the constraints imposed on the solution (e.g., find a partition, a covering or a hierarchy in clustering). It also provides powerful mathematical tools to build highly performing exact or approximate algorithms. This book is based on lectures presented at the workshop on "Data Mining and Mathematical Programming" (October 10-13, 2006, Montreal) and will be a valuable scientific source of information to faculty, students, and researchers in optimization, data analysis and data mining, as well as people working in computer science, engineering and applied mathematics.

Comprehensive Coverage of the Entire Area of Classification Research on the problem of classification tends to be fragmented across such areas as pattern recognition, database, data mining, and machine learning. Addressing the work of these different communities in a unified way, *Data Classification: Algorithms and Applications* explores the underlying 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.

This book seeks to promote the exploitation of data science in healthcare systems. The focus is on advancing the automated analytical methods used to extract new knowledge from data for healthcare applications. To do so, the book draws on several interrelated disciplines, including machine learning, big data analytics, statistics, pattern recognition, computer vision, and Semantic Web technologies, and focuses on their direct application to healthcare. Building on three tutorial-like chapters on data science in healthcare, the following eleven chapters highlight success stories on the application of data science in healthcare, where data science and artificial intelligence technologies have proven to be

very promising. This book is primarily intended for data scientists involved in the healthcare or medical sector. By reading this book, they will gain essential insights into the modern data science technologies needed to advance innovation for both healthcare businesses and patients. A basic grasp of data science is recommended in order to fully benefit from this book.

In the data stream scenario, input arrives very rapidly and there is limited memory to store the input. Algorithms have to work with one or few passes over the data, space less than linear in the input size or time significantly less than the input size. In the past few years, a new theory has emerged for reasoning about algorithms that work within these constraints on space, time, and number of passes. Some of the methods rely on metric embeddings, pseudo-random computations, sparse approximation theory and communication complexity. The applications for this scenario include IP network traffic analysis, mining text message streams and processing massive data sets in general. Researchers in Theoretical Computer Science, Databases, IP Networking and Computer Systems are working on the data stream challenges. Entropies and entropy-like quantities play an increasing role in modern non-linear data analysis. Fields that benefit from this application range from biosignal analysis to econophysics and engineering. This issue is a collection of papers touching on different aspects of entropy measures in data analysis, as well as theoretical and computational analyses. The relevant topics include the difficulty to achieve adequate application of entropy measures and the acceptable parameter choices for those entropy measures, entropy-based coupling, and similarity analysis, along with the utilization of entropy measures as features in automatic learning and classification. Various real data applications are given. Build strong foundation for entering the world of Machine Learning and data science with the help of this comprehensive guide About This Book Get started in the field of Machine Learning with the help of this solid, concept-rich, yet highly practical guide. Your one-stop solution for everything that matters in mastering the whats and whys of Machine Learning algorithms and their implementation. Get a solid foundation for your entry into Machine Learning by strengthening your roots (algorithms) with this comprehensive guide. Who This Book Is For This book is for IT professionals who want to enter the field of data science and are very new to Machine Learning. Familiarity with languages such as R and Python will be invaluable here. What You Will Learn Acquaint yourself with important elements of Machine Learning Understand the feature selection and feature engineering process Assess performance and error trade-offs for Linear Regression Build a data model and understand how it works by using different types of algorithm Learn to tune the parameters of Support Vector machines Implement clusters to a dataset Explore the concept of Natural Processing Language and Recommendation Systems Create a ML architecture from scratch. In Detail As the amount of data continues to grow at an almost incomprehensible rate, being able to understand and process data is becoming a key differentiator for competitive organizations. Machine learning applications are everywhere, from self-driving cars, spam detection, document search, and trading strategies, to speech recognition. This makes machine learning well-suited to the present-day era of Big Data and Data Science. The main challenge is how to transform data into actionable knowledge. In this book you will learn all the important Machine Learning algorithms that are commonly used in the field of data science. These algorithms can be used for supervised as well as unsupervised learning, reinforcement learning, and semi-supervised learning. A few famous algorithms that are covered in this book are Linear regression, Logistic Regression, SVM, Naive Bayes, K-Means, Random Forest, TensorFlow, and Feature engineering. In this book you will also learn how these algorithms work and their practical implementation to resolve your problems. This book will also introduce you to the Natural Processing Language and Recommendation systems, which help you run multiple algorithms simultaneously. On completion of the book you will have mastered selecting Machine Learning algorithms for clustering, classification, or regression based on for your problem. Style and approach An easy-to-follow, step-by-step guide that will help you get to grips with real -world applications of Algorithms for Machine Learning.

A new starting-point and a new method are requisite, to insure a complete [classification of the Steiner triple systems of order 15]. This method was furnished, and its tedious and difficult execution undertaken, by Mr. Cole. F. N. Cole, L. D. Cummings, and H. S. White (1917) [129] The history of classifying combinatorial objects is as old as the history of the objects themselves. In the mid-19th century, Kirkman, Steiner, and others became the fathers of modern combinatorics, and their work – on various objects, including (what became later known as) Steiner triple systems – led to several classification results. Almost a century earlier, in 1782, Euler [180] published some results on classifying small Latin squares, but for the first few steps in this direction one should actually go at least as far back as ancient Greece and the proof that there are exactly five Platonic solids. One of the most remarkable achievements in the early, pre-computer era is the classification of the Steiner triple systems of order 15, quoted above. An onerous task that, today, no sensible person would attempt by hand calculation. Because, with the exception of occasional parameters for which combinatorial arguments are effective (often to prove nonexistence or uniqueness), classification in general is about algorithms and computation.

Automatic Modulation Classification (AMC) has been a key technology in many military, security, and civilian telecommunication applications for decades. In military and security applications, modulation often serves as another level of encryption; in modern civilian applications, multiple modulation types can be employed by a signal transmitter to control the data rate and link reliability. This book offers comprehensive documentation of AMC models, algorithms and implementations for successful modulation recognition. It provides an invaluable theoretical and numerical comparison of AMC algorithms, as well as guidance on state-of-the-art classification designs with specific military and civilian applications in mind. Key Features: Provides an important collection of AMC algorithms in five major categories, from likelihood-based classifiers and distribution-test-based classifiers to feature-based classifiers, machine learning assisted classifiers and blind modulation classifiers Lists detailed implementation for each algorithm based on a unified theoretical background and a comprehensive theoretical and numerical performance comparison Gives clear guidance for the

design of specific automatic modulation classifiers for different practical applications in both civilian and military communication systems Includes a MATLAB toolbox on a companion website offering the implementation of a selection of methods discussed in the book

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

A hands-on approach to tasks and techniques in data stream mining and real-time analytics, with examples in MOA, a popular freely available open-source software framework. Today many information sources—including sensor networks, financial markets, social networks, and healthcare monitoring—are so-called data streams, arriving sequentially and at high speed. Analysis must take place in real time, with partial data and without the capacity to store the entire data set. This book presents algorithms and techniques used in data stream mining and real-time analytics. Taking a hands-on approach, the book demonstrates the techniques using MOA (Massive Online Analysis), a popular, freely available open-source software framework, allowing readers to try out the techniques after reading the explanations. The book first offers a brief introduction to the topic, covering big data mining, basic methodologies for mining data streams, and a simple example of MOA. More detailed discussions follow, with chapters on sketching techniques, change, classification, ensemble methods, regression, clustering, and frequent pattern mining. Most of these chapters include exercises, an MOA-based lab session, or both. Finally, the book discusses the MOA software, covering the MOA graphical user interface, the command line, use of its API, and the development of new methods within MOA. The book will be an essential reference for readers who want to use data stream mining as a tool, researchers in innovation or data stream mining, and programmers who want to create new algorithms for MOA.

The field of machine learning has matured to the point where many sophisticated learning approaches can be applied to practical applications. Thus it is of critical importance that researchers have the proper tools to evaluate learning approaches and understand the underlying issues. This book examines various aspects of the evaluation process with an emphasis on classification algorithms. The authors describe several techniques for classifier performance assessment, error estimation and resampling, obtaining statistical significance as well as selecting appropriate domains for evaluation. They also present a unified evaluation framework and highlight how different components of evaluation are both significantly interrelated and interdependent. The techniques presented in the book are illustrated using R and WEKA, facilitating better practical insight as well as implementation. Aimed at researchers in the theory and applications of machine learning, this book offers a solid basis for conducting performance evaluations of algorithms in practical settings. The two volumes LNCS 10199 and 10200 constitute the refereed conference proceedings of the 20th European Conference on the Applications of Evolutionary Computation, EvoApplications 2017, held in Amsterdam, The Netherlands, in April 2017, colocated with the Evo* 2016 events EuroGP, EvoCOP, and EvoMUSART. The 46 revised full papers presented together with 26 poster papers were carefully reviewed and selected from 108 submissions. EvoApplications 2016 consisted of the following 13 tracks: EvoBAFIN (natural computing methods in business analytics and finance), EvoBIO (evolutionary computation, machine learning and data mining in computational biology), EvoCOMNET (nature-inspired techniques for telecommunication networks and other parallel and distributed systems), EvoCOMPLEX (evolutionary algorithms and complex systems), EvoENERGY (evolutionary computation in energy applications), EvoGAMES (bio-inspired algorithms in games), EvoIASP (evolutionary computation in image analysis, signal processing, and pattern recognition), EvoINDUSTRY (nature-inspired techniques in industrial settings), EvoKNOW (knowledge incorporation in evolutionary computation), EvoNUM (bio-inspired algorithms for continuous parameter optimization), EvoPAR (parallel implementation of evolutionary algorithms), EvoROBOT (evolutionary robotics), EvoSET (nature-inspired algorithms in software engineering and testing), and EvoSTOC (evolutionary algorithms in stochastic and dynamic environments).

Comprehensive Coverage of the Entire Area of Classification Research on the problem of classification tends to be fragmented across such areas as pattern recognition, database, data mining, and machine learning. Addressing the work of these different communities in a unified way, Data Classification: Algorithms and Applications explores the underlying algorithms of classification as well as applications of classification in a variety of problem domains, including text, multimedia, social network, and biological data. This comprehensive book focuses on three primary aspects of data classification: Methods-The book first describes common techniques used for classification, including probabilistic methods, decision trees, rule-based methods, instance-based methods, support vector machine methods, and neural networks. Domains-The book then examines specific methods used for data domains such as multimedia, text, time-series, network, discrete sequence, and uncertain data. It also covers large data sets and data streams due to the recent importance of the big data paradigm. Variations-The book concludes with insight on variations of the classification process. It discusses ensembles, rare-class learning, distance function learning, active learning, visual learning, transfer learning, and semi-supervised learning as well as evaluation aspects of classifiers.

The main goal of the new field of data mining is the analysis of large and complex datasets. Some very important datasets may be derived from business and industrial activities. This kind of data is known as OC enterprise dataOCO. The common characteristic of such datasets is that the analyst wishes to analyze them for the purpose of designing a more cost-effective strategy for optimizing some type of performance measure, such as reducing production time, improving quality, eliminating wastes, or maximizing profit. Data in this category may describe different scheduling scenarios in a manufacturing environment, quality control of some process, fault diagnosis in the operation of a machine or process, risk analysis when issuing credit to applicants, management of supply chains in a manufacturing system, or data for business related decision-making. Sample Chapter(s). Foreword (37 KB). Chapter 1: Enterprise Data Mining: A Review and Research Directions (655 KB). Contents: Enterprise Data Mining: A Review and Research Directions (T W

Liao); Application and Comparison of Classification Techniques in Controlling Credit Risk (L Yu et al.); Predictive Classification with Imbalanced Enterprise Data (S Daskalaki et al.); Data Mining Applications of Process Platform Formation for High Variety Production (J Jiao & L Zhang); Multivariate Control Charts from a Data Mining Perspective (G C Porzio & G Ragozini); Maintenance Planning Using Enterprise Data Mining (L P Khoo et al.); Mining Images of Cell-Based Assays (P Perner); Support Vector Machines and Applications (T B Trafalis & O O Oladunni); A Survey of Manifold-Based Learning Methods (X Huo et al.); and other papers. Readership: Graduate students in engineering, computer science, and business schools; researchers and practitioners of data mining with emphasis of enterprise data mining."

Discover hidden relationships among the variables in your data, and learn how to exploit these relationships. This book presents a collection of data-mining algorithms that are effective in a wide variety of prediction and classification applications. All algorithms include an intuitive explanation of operation, essential equations, references to more rigorous theory, and commented C++ source code. Many of these techniques are recent developments, still not in widespread use. Others are standard algorithms given a fresh look. In every case, the focus is on practical applicability, with all code written in such a way that it can easily be included into any program. The Windows-based DATAMINE program lets you experiment with the techniques before incorporating them into your own work. What You'll Learn Use Monte-Carlo permutation tests to provide statistically sound assessments of relationships present in your data Discover how combinatorially symmetric cross validation reveals whether your model has true power or has just learned noise by overfitting the data Work with feature weighting as regularized energy-based learning to rank variables according to their predictive power when there is too little data for traditional methods See how the eigenstructure of a dataset enables clustering of variables into groups that exist only within meaningful subspaces of the data Plot regions of the variable space where there is disagreement between marginal and actual densities, or where contribution to mutual information is high Who This Book Is For Anyone interested in discovering and exploiting relationships among variables. Although all code examples are written in C++, the algorithms are described in sufficient detail that they can easily be programmed in any language.

Text mining applications have experienced tremendous advances because of web 2.0 and social networking applications. Recent advances in hardware and software technology have lead to a number of unique scenarios where text mining algorithms are learned. Mining Text Data introduces an important niche in the text analytics field, and is an edited volume contributed by leading international researchers and practitioners focused on social networks & data mining. This book contains a wide swath in topics across social networks & data mining. Each chapter contains a comprehensive survey including the key research content on the topic, and the future directions of research in the field. There is a special focus on Text Embedded with Heterogeneous and Multimedia Data which makes the mining process much more challenging. A number of methods have been designed such as transfer learning and cross-lingual mining for such cases. Mining Text Data simplifies the content, so that advanced-level students, practitioners and researchers in computer science can benefit from this book. Academic and corporate libraries, as well as ACM, IEEE, and Management Science focused on information security, electronic commerce, databases, data mining, machine learning, and statistics are the primary buyers for this reference book.

Currently many different application areas for Big Data (BD) and Machine Learning (ML) are being explored. These promising application areas for BD/ML are the social sites, search engines, multimedia sharing sites, various stock exchange sites, online gaming, online survey sites and various news sites, and so on. To date, various use-cases for this application area are being researched and developed. Software applications are already being published and used in various settings from education and training to discover useful hidden patterns and other information like customer choices and market trends that can help organizations make more informed and customer-oriented business decisions. Combining BD with ML will provide powerful, largely unexplored application areas that will revolutionize practice in Videos Surveillance, Social Media Services, Email Spam and Malware Filtering, Online Fraud Detection, and so on. It is very important to continuously monitor and understand these effects from safety and societal point of view. Hence, the main purpose of this book is for researchers, software developers and practitioners, academicians and students to showcase novel use-cases and applications, present empirical research results from user-centered qualitative and quantitative experiments of these new applications, and facilitate a discussion forum to explore the latest trends in big data and machine learning by providing algorithms which can be trained to perform interdisciplinary techniques such as statistics, linear algebra, and optimization and also create automated systems that can sift through large volumes of data at high speed to make predictions or decisions without human intervention

This book is a collection of high-quality research work on cutting-edge technologies and the most-happening areas of computational intelligence and data engineering. It includes selected papers from the International Conference on Computational Intelligence and Data Engineering (ICCIDE 2020). It covers various topics, including collective intelligence, intelligent transportation systems, fuzzy systems, Bayesian network, ant colony optimization, data privacy and security, data mining, data warehousing, big data analytics, cloud computing, natural language processing, swarm intelligence and speech processing.

Data clustering, also known as cluster analysis, is an unsupervised process that divides a set of objects into homogeneous groups. Since the publication of the first edition of this monograph in 2007, development in the area has exploded, especially in clustering algorithms for big data and open-source software for cluster analysis. This second edition reflects these new developments, covers the basics of data clustering, includes a list of popular clustering algorithms, and provides program code that helps users implement clustering algorithms. Data Clustering: Theory, Algorithms and Applications, Second Edition will be of interest to researchers, practitioners, and data scientists as well as

undergraduate and graduate students.

Generating a satisfactory classification image from remote sensing data is not a straightforward task. Many factors contribute to this difficulty including the characteristics of a study area, availability of suitable remote sensing data, ancillary and ground reference data, proper use of variables and classification algorithms, and the analyst's e

This book constitutes the refereed proceedings of the 7th International Workshop on Information Security Applications, WISA 2006, held in Jeju Island, Korea in August 2006. Coverage in the 30 revised full papers includes public key crypto applications and virus protection, cyber indication and intrusion detection, biometrics and security trust management, secure software and systems, smart cards and secure hardware, and mobile security.

Data Science and Analytics with Python is designed for practitioners in data science and data analytics in both academic and business environments. The aim is to present the reader with the main concepts used in data science using tools developed in Python, such as SciKit-learn, Pandas, Numpy, and others. The use of Python is of particular interest, given its recent popularity in the data science community. The book can be used by seasoned programmers and newcomers alike. The book is organized in a way that individual chapters are sufficiently independent from each other so that the reader is comfortable using the contents as a reference. The book discusses what data science and analytics are, from the point of view of the process and results obtained. Important features of Python are also covered, including a Python primer. The basic elements of machine learning, pattern recognition, and artificial intelligence that underpin the algorithms and implementations used in the rest of the book also appear in the first part of the book. Regression analysis using Python, clustering techniques, and classification algorithms are covered in the second part of the book. Hierarchical clustering, decision trees, and ensemble techniques are also explored, along with dimensionality reduction techniques and recommendation systems. The support vector machine algorithm and the Kernel trick are discussed in the last part of the book. About the Author Dr. Jesús Rogel-Salazar is a Lead Data scientist with experience in the field working for companies such as AKQA, IBM Data Science Studio, Dow Jones and others. He is a visiting researcher at the Department of Physics at Imperial College London, UK and a member of the School of Physics, Astronomy and Mathematics at the University of Hertfordshire, UK, He obtained his doctorate in physics at Imperial College London for work on quantum atom optics and ultra-cold matter. He has held a position as senior lecturer in mathematics as well as a consultant in the financial industry since 2006. He is the author of the book Essential Matlab and Octave, also published by CRC Press. His interests include mathematical modelling, data science, and optimization in a wide range of applications including optics, quantum mechanics, data journalism, and finance.

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

[Copyright: 4b9072390226f6e1ccf0ba976d9e6694](https://doi.org/10.1007/978-1-4939-9999-9)