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

Geographic Data Mining and Knowledge DiscoveryComputational Methods of Feature SelectionMining Text DataMachine Learning Models and Algorithms for Big Data ClassificationAdvances in Environmental Remote SensingData ClassificationPractical Machine Learning for Data Analysis Using PythonThe Top Ten Algorithms in Data MiningData StreamsOptimization of Complex Systems: Theory, Models, Algorithms and ApplicationsData Mining With Decision Trees: Theory And Applications (2nd Edition)Data MiningCondition Monitoring with Vibration SignalsData ClassificationUnderstanding Machine LearningData ClusteringMachine Learning for Spatial Environmental DataAutomatic Modulation ClassificationData Science for HealthcareProgramming Collective IntelligenceData Mining Algorithms in C++Big Data Analytics: Systems, Algorithms, ApplicationsTemporal Data MiningClassification and Learning Using Genetic AlgorithmsContrast Data MiningIntroduction to Algorithms for Data Mining and Machine LearningData MiningHandbook of Neural ComputationData ClassificationData Mining Applications with RHandbook of Statistical Analysis and Data Mining ApplicationsKnowledge Discovery from Data StreamsFuzzy Machine Learning Algorithms for Remote Sensing Image ClassificationData

MiningClassification, Clustering, and Data AnalysisData ClusteringData Mining: Concepts and TechniquesEvaluating Learning AlgorithmsText MiningLearning from Imbalanced Data Sets

Geographic Data Mining and Knowledge Discovery

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.

Computational Methods of Feature Selection

Temporal data mining deals with the harvesting of useful information from temporal data. New initiatives in health care and business organizations have increased the importance of temporal information in data today. From basic data mining concepts to state-of-the-art advances, Temporal Data Mining covers the theory of this subject as well as its application in a variety of fields. It discusses the incorporation of temporality in databases as well as temporal data representation, similarity computation, data classification, clustering, pattern discovery, and prediction. The book also explores the use of temporal data mining in medicine and biomedical informatics, business and industrial applications, web usage mining, and spatiotemporal data mining. Along with various state-of-the-art algorithms, each chapter includes detailed references and short descriptions of relevant algorithms and techniques described in other references. In the appendices, the author explains how data mining fits the overall goal of an organization and how these data can be interpreted for the purpose of characterizing a population. She also provides programs written in the Java language that implement some of the algorithms presented in

the first chapter. Check out the author's blog at http://theophanomitsa.wordpress.com/

Mining Text Data

Practical Machine Learning for Data Analysis Using Python is a problem solver's guide for creating real-world intelligent systems. It provides a comprehensive approach with concepts, practices, hands-on examples, and sample code. The book teaches readers the vital skills required to understand and solve different problems with machine learning. It teaches machine learning techniques necessary to become a successful practitioner, through the presentation of real-world case studies in Python machine learning ecosystems. The book also focuses on building a foundation of machine learning knowledge to solve different real-world case studies across various fields, including biomedical signal analysis, healthcare, security, economics, and finance. Moreover, it covers a wide range of machine learning models, including regression, classification, and forecasting. The goal of the book is to help a broad range of readers, including IT professionals, analysts, developers, data scientists, engineers, and graduate students, to solve their own real-world problems. Offers a comprehensive overview of the application of machine learning tools in data analysis across a wide range of subject areas Teaches readers how to apply machine learning techniques to biomedical signals, financial data, and healthcare data Explores important classification and regression algorithms as well as other machine learning

techniques Explains how to use Python to handle data extraction, manipulation, and exploration techniques, as well as how to visualize data spread across multiple dimensions and extract useful features

Machine Learning Models and Algorithms for Big Data Classification

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.

Advances in Environmental Remote Sensing

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.

Data Classification

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

Practical Machine Learning for Data Analysis Using Python

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.

The Top Ten Algorithms in Data Mining

Due to increasing demands for dimensionality reduction, research on feature selection has deeply and widely expanded into many fields, including computational statistics, pattern recognition, machine learning, data mining, and knowledge discovery. Highlighting current research issues, Computational Methods of Feature Selection introduces the basic concepts and principles, state-of-the-art algorithms, and novel applications of this tool. The book begins by exploring unsupervised, randomized, and causal feature selection. It then reports on some recent results of empowering feature selection, including active feature selection, decision-border estimate, the use of ensembles with independent probes, and incremental feature selection. This is followed by discussions of weighting and local methods, such as the ReliefF family, k-means clustering, local feature relevance, and a new interpretation of Relief. The book subsequently covers text classification, a new feature selection score, and both constraint-guided and aggressive feature selection. The final section examines applications of feature selection in

bioinformatics, including feature construction as well as redundancy-, ensemble-, and penalty-based feature selection. Through a clear, concise, and coherent presentation of topics, this volume systematically covers the key concepts, underlying principles, and inventive applications of feature selection, illustrating how this powerful tool can efficiently harness massive, high-dimensional data and turn it into valuable, reliable information.

Data Streams

Decision trees have become one of the most powerful and popular approaches in knowledge discovery and data mining; it is the science of exploring large and complex bodies of data in order to discover useful patterns. Decision tree learning continues to evolve over time. Existing methods are constantly being improved and new methods introduced.This 2nd Edition is dedicated entirely to the field of decision trees in data mining; to cover all aspects of this important technique, as well as improved or new methods and techniques developed after the publication of our first edition. In this new edition, all chapters have been revised and new topics brought in. New topics include Cost-Sensitive Active Learning, Learning with Uncertain and Imbalanced Data, Using Decision Trees beyond Classification Tasks, Privacy Preserving Decision Tree Learning, Lessons Learned from Comparative Studies, and Learning Decision Trees for Big Data. A walk-through guide to existing open-source data mining software is also included in this edition.This book invites readers to explore the

many benefits in data mining that decision trees offer:

Optimization of Complex Systems: Theory, Models, Algorithms and Applications

Handbook of Neural Computation explores neural computation applications, ranging from conventional fields of mechanical and civil engineering, to electronics, electrical engineering and computer science. This book covers the numerous applications of artificial and deep neural networks and their uses in learning machines, including image and speech recognition, natural language processing and risk analysis. Edited by renowned authorities in this field, this work is comprised of articles from reputable industry and academic scholars and experts from around the world. Each contributor presents a specific research issue with its recent and future trends. As the demand rises in the engineering and medical industries for neural networks and other machine learning methods to solve different types of operations, such as data prediction, classification of images, analysis of big data, and intelligent decision-making, this book provides readers with the latest, cutting-edge research in one comprehensive text. Features high-quality research articles on multivariate adaptive regression splines, the minimax probability machine, and more Discusses machine learning techniques, including classification, clustering, regression, web mining, information retrieval and natural language processing Covers supervised, unsupervised, reinforced, ensemble, and nature-

Data Mining With Decision Trees: Theory And Applications (2nd Edition)

Research on the problem of clustering tends to be fragmented across the pattern recognition, database, data mining, and machine learning communities. Addressing this problem in a unified way, Data Clustering: Algorithms and Applications provides complete coverage of the entire area of clustering, from basic methods to more refined and complex data clustering approaches. It pays special attention to recent issues in graphs, social networks, and other domains. The book focuses on three primary aspects of data clustering: Methods, describing key techniques commonly used for clustering, such as feature selection, agglomerative clustering, partitional clustering, density-based clustering, probabilistic clustering, grid-based clustering, spectral clustering, and nonnegative matrix factorization Domains, covering methods used for different domains of data, such as categorical data, text data, multimedia data, graph data, biological data, stream data, uncertain data, time series clustering, high-dimensional clustering, and big data Variations and Insights, discussing important variations of the clustering process, such as semisupervised clustering, interactive clustering, multiview clustering, cluster ensembles, and cluster validation In this book, top researchers from around the world explore the characteristics of clustering problems in a variety of application areas. They also explain how to glean

detailed insight from the clustering process—including how to verify the quality of the underlying clusters—through supervision, human intervention, or the automated generation of alternative clusters.

Data Mining

Introduces machine learning and its algorithmic paradigms, explaining the principles behind automated learning approaches and the considerations underlying their usage.

Condition Monitoring with Vibration Signals

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.

Data Classification

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.

Understanding Machine Learning

Data Mining Applications with R is a great resource for researchers and professionals to understand the wide use of R, a free software environment for statistical computing and graphics, in solving different problems in industry. R is widely used in leveraging data mining techniques across many different industries, including government, finance, insurance, medicine, scientific research and more. This book presents 15 different real-world case studies illustrating various techniques in rapidly growing areas. It is an ideal companion for data mining researchers in academia and industry looking for ways to turn this versatile software into a powerful analytic tool. R code, Data and color figures for the book are provided at the RDataMining.com

website. Helps data miners to learn to use R in their specific area of work and see how R can apply in different industries Presents various case studies in real-world applications, which will help readers to apply the techniques in their work Provides code examples and sample data for readers to easily learn the techniques by running the code by themselves

Data Clustering

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

Machine Learning for Spatial Environmental Data

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

Automatic Modulation Classification

This book contains 112 papers selected from about 250 submissions to the 6th World Congress on Global Optimization (WCGO 2019) which takes place on July 8–10, 2019 at University of Lorraine, Metz, France. The book covers both theoretical and algorithmic aspects of Nonconvex Optimization, as well as its applications to modeling and solving decision problems in various domains. It is composed of 10 parts, each of them deals with either the theory and/or methods in a branch of optimization such as Continuous optimization, DC Programming and DCA, Discrete optimization & Network optimization, Multiobjective programming, Optimization under uncertainty, or models and optimization methods in a specific application area including Data science, Economics & Finance, Energy & Water management, Engineering systems, Transportation, Logistics, Resource allocation & Production management. The researchers and practitioners working in Nonconvex Optimization and several application areas can find here many inspiring ideas and useful tools & techniques for their works.

Data Science for Healthcare

The Definitive Volume on Cutting-Edge Exploratory Analysis of Massive Spatial and Spatiotemporal DatabasesSince the publication of the first edition of Geographic Data Mining and Knowledge Discovery, new techniques for geographic data warehousing (GDW), spatial data mining, and geovisualization (GVis) have been developed. In addition, there has bee

Programming Collective Intelligence

This book provides a comprehensive survey of techniques, technologies and applications of Big Data and its analysis. The Big Data phenomenon is increasingly impacting all sectors of business and industry, producing an emerging new information ecosystem. On the applications front, the book offers detailed descriptions of various application areas for Big Data Analytics in the important domains of Social Semantic Web Mining, Banking and Financial Services, Capital Markets, Insurance, Advertisement, Recommendation Systems, Bio-Informatics, the IoT and Fog Computing, before delving into issues of security and privacy. With regard to machine learning techniques, the book presents all the standard algorithms for learning – including supervised, semi-supervised and unsupervised techniques such as clustering and reinforcement learning techniques to perform collective Deep Learning. Multi-layered and nonlinear learning for Big Data are also covered. In turn, the book highlights real-life case studies on successful implementations of Big Data Analytics at large IT companies such as Google, Facebook, LinkedIn and Microsoft. Multi-sectorial case studies on domain-based companies such as Deutsche Bank, the power provider Opower, Delta Airlines and a Chinese City Transportation application represent a valuable addition. Given its comprehensive coverage of Big Data Analytics, the book offers a unique resource for undergraduate and graduate students, researchers, educators and IT professionals alike.

Data Mining Algorithms in C++

Provides an extensive, up-to-date treatment of techniques used for machine condition monitoring Clear and concise throughout, this accessible book is the first to be wholly devoted to the field of condition monitoring for rotating machines using vibration signals. It covers various feature extraction, feature selection, and classification methods as well as their applications to machine vibration datasets. It also presents new methods including machine learning and compressive sampling, which help to improve safety, reliability, and performance. Condition Monitoring with Vibration Signals: Compressive Sampling and Learning Algorithms for Rotating Machines starts by introducing readers to Vibration Analysis Techniques and Machine Condition Monitoring (MCM). It then offers readers sections covering: Rotating Machine Condition Monitoring using Learning Algorithms; Classification Algorithms; and New Fault Diagnosis Frameworks designed for MCM. Readers will learn signal processing in the time-frequency domain, methods for linear subspace learning, and the basic principles of the learning method Artificial Neural Network (ANN). They will also discover recent trends of deep learning in the field of machine condition monitoring, new feature learning frameworks based on compressive sampling, subspace learning techniques for machine condition monitoring, and much more. Covers the fundamental as well as the state-of-the-art approaches to machine condition monitoringguiding readers from the basics of rotating machines to the generation of knowledge

using vibration signals Provides new methods, including machine learning and compressive sampling, which offer significant improvements in accuracy with reduced computational costs Features learning algorithms that can be used for fault diagnosis and prognosis Includes previously and recently developed dimensionality reduction techniques and classification algorithms Condition Monitoring with Vibration Signals: Compressive Sampling and Learning Algorithms for Rotating Machines is an excellent book for research students, postgraduate students, industrial practitioners, and researchers.

Big Data Analytics: Systems, Algorithms, Applications

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.

Temporal Data Mining

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

Classification and Learning Using Genetic Algorithms

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.

Contrast Data Mining

Accompanying CD-RM contains Machine learning office software, MLO guide (pdf) and examples of data.

Introduction to Algorithms for Data Mining and Machine Learning

New technologies have enabled us to collect massive amounts of data in many fields. However, our pace of discovering useful information and knowledge from

these data falls far behind our pace of collecting the data. Data Mining: Theories, Algorithms, and Examples introduces and explains a comprehensive set of data mining algorithms from various dat

Data Mining

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.

Handbook of Neural 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

Data Classification

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

Data Mining Applications with R

The Definitive Resource on Text Mining Theory and Applications from Foremost Researchers in the Field Giving a broad perspective of the field from numerous vantage points, Text Mining: Classification, Clustering, and Applications focuses on statistical methods for text mining and analysis. It examines methods to automatically cluster and classify text documents and applies these methods in a variety of areas, including adaptive information filtering, information distillation, and text search. The book begins with chapters on the classification of documents into predefined categories. It presents state-of-the-art algorithms and their use in practice. The next chapters describe novel methods for clustering documents into groups that are not predefined. These methods seek to automatically determine topical structures that may exist in a document corpus. The book concludes by discussing various text mining applications that have significant implications for future research and industrial use. There is no doubt that text mining will continue to play a critical role in the development of future information systems and advances in research will be instrumental to their success. This book captures the technical depth and immense practical potential of text mining, guiding readers to a sound appreciation of this burgeoning field.

Handbook of Statistical Analysis and Data Mining Applications

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

Knowledge Discovery from Data Streams

Introduction to Algorithms for Data Mining and Machine Learning introduces the essential ideas behind all key algorithms and techniques for data mining and machine learning, along with optimization techniques. Its strong formal mathematical approach, well selected examples, and practical software recommendations help readers develop confidence in their data modeling skills so they can process and

interpret data for classification, clustering, curve-fitting and predictions. Masterfully balancing theory and practice, it is especially useful for those who need relevant, well explained, but not rigorous (proofs based) background theory and clear guidelines for working with big data. Presents an informal, theorem-free approach with concise, compact coverage of all fundamental topics Includes worked examples that help users increase confidence in their understanding of key algorithms, thus encouraging self-study Provides algorithms and techniques that can be implemented in any programming language, with each chapter including notes about relevant software packages

Fuzzy Machine Learning Algorithms for Remote Sensing Image Classification

The book presents a long list of useful methods for classification, clustering and data analysis. By combining theoretical aspects with practical problems, it is designed for researchers as well as for applied statisticians and will support the fast transfer of new methodological advances to a wide range of applications.

Data Mining

This book covers the state-of-art image classification methods for discrimination of earth objects from remote sensing satellite data with an emphasis on fuzzy machine learning and deep learning algorithms. Both types of algorithms are described in such details

that these can be implemented directly for thematic mapping of multiple-class or specific-class landcover from multispectral optical remote sensing data. These algorithms along with multi-date, multi-sensor remote sensing are capable to monitor specific stage (for e.g., phenology of growing crop) of a particular class also included. With these capabilities fuzzy machine learning algorithms have strong applications in areas like crop insurance, forest fire mapping, stubble burning, post disaster damage mapping etc. It also provides details about the temporal indices database using proposed Class Based Sensor Independent (CBSI) approach supported by practical examples. As well, this book addresses other related algorithms based on distance, kernel based as well as spatial information through Markov Random Field (MRF)/Local convolution methods to handle mixed pixels, non-linearity and noisy pixels. Further, this book covers about techniques for quantiative assessment of soft classified fraction outputs from soft classification and supported by in-house developed tool called sub-pixel multi-spectral image classifier (SMIC). It is aimed at graduate, postgraduate, research scholars and working professionals of different branches such as Geoinformation sciences, Geography, Electrical, Electronics and Computer Sciences etc., working in the fields of earth observation and satellite image processing. Learning algorithms discussed in this book may also be useful in other related fields, for example, in medical imaging. Overall, this book aims to: exclusive focus on using large range of fuzzy classification algorithms for remote sensing images; discuss ANN, CNN, RNN, and hybrid learning

classifiers application on remote sensing images; describe sub-pixel multi-spectral image classifier tool (SMIC) to support discussed fuzzy and learning algorithms; explain how to assess soft classified outputs as fraction images using fuzzy error matrix (FERM) and its advance versions with FERM tool, Entropy, Correlation Coefficient, Root Mean Square Error and Receiver Operating Characteristic (ROC) methods and; combines explanation of the algorithms with case studies and practical applications.

Classification, Clustering, and Data Analysis

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.

Data Clustering

Cluster analysis is an unsupervised process that divides a set of objects into homogeneous groups. This book starts with basic information on cluster analysis, including the classification of data and the corresponding similarity measures, followed by the presentation of over 50 clustering algorithms in groups according to some specific baseline methodologies such as hierarchical, center-based, and search-based methods. As a result, readers and users can easily identify an appropriate algorithm for their applications and compare novel ideas with existing results. The book also provides examples of clustering applications to illustrate the advantages and shortcomings of different clustering architectures and algorithms. Application areas include pattern recognition, artificial intelligence, information technology, image processing, biology, psychology, and marketing. Readers also learn how to perform cluster analysis with the C/C++ and MATLAB programming languages.

Data Mining: Concepts and Techniques

Want to tap the power behind search rankings, product recommendations, social bookmarking, and online matchmaking? This fascinating book demonstrates how you can build Web 2.0 applications to mine the enormous amount of data created by people on the Internet. With the sophisticated algorithms in this book, you can write smart programs to access interesting datasets from other web sites, collect data from users of your own applications, and analyze and understand the data once you've found it. Programming Collective Intelligence takes you into the world of machine learning and statistics, and explains how to draw conclusions about user experience, marketing, personal tastes, and human behavior in general -- all from information that you and others collect every day. Each algorithm is described clearly and concisely with code that can immediately be used on your web site, blog, Wiki, or specialized application. This book explains: Collaborative filtering techniques that enable online retailers to recommend products or media Methods of clustering to detect groups of similar items in a large dataset Search engine features -- crawlers, indexers, query engines, and the PageRank algorithm Optimization algorithms that search millions of possible solutions to a problem and choose the best one Bayesian filtering, used in spam filters for classifying documents based on word types and other features Using decision trees not only to make predictions, but to model the way decisions are made Predicting numerical values rather than classifications

to build price models Support vector machines to match people in online dating sites Non-negative matrix factorization to find the independent features in a dataset Evolving intelligence for problem solving -- how a computer develops its skill by improving its own code the more it plays a game Each chapter includes exercises for extending the algorithms to make them more powerful. Go beyond simple database-backed applications and put the wealth of Internet data to work for you. "Bravo! I cannot think of a better way for a developer to first learn these algorithms and methods, nor can I think of a better way for me (an old AI dog) to reinvigorate my knowledge of the details." -- Dan Russell, Google "Toby's book does a great job of breaking down the complex subject matter of machine-learning algorithms into practical, easy-to-understand examples that can be directly applied to analysis of social interaction across the Web today. If I had this book two years ago, it would have saved precious time going down some fruitless paths." -- Tim Wolters, CTO, Collective Intellect

Evaluating Learning Algorithms

Since the beginning of the Internet age and the increased use of ubiquitous computing devices, the large volume and continuous flow of distributed data have imposed new constraints on the design of learning algorithms. Exploring how to extract knowledge structures from evolving and time-changing data, Knowledge Discovery from Data Streams presents a coherent overview of state-of-the-

art research in learning from data streams. The book covers the fundamentals that are imperative to understanding data streams and describes important applications, such as TCP/IP traffic, GPS data, sensor networks, and customer click streams. It also addresses several challenges of data mining in the future, when stream mining will be at the core of many applications. These challenges involve designing useful and efficient data mining solutions applicable to real-world problems. In the appendix, the author includes examples of publicly available software and online data sets. This practical, up-to-date book focuses on the new requirements of the next generation of data mining. Although the concepts presented in the text are mainly about data streams, they also are valid for different areas of machine learning and data mining.

Text Mining

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.

Learning from Imbalanced Data Sets

Comprehensive Coverage of the Entire Area of ClassificationResearch 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 underlyi

ROMANCE ACTION & ADVENTURE MYSTERY & THRILLER BIOGRAPHIES & HISTORY CHILDREN'S YOUNG ADULT FANTASY HISTORICAL FICTION HORROR LITERARY FICTION NON-FICTION SCIENCE FICTION