

Download Ebook Manual Solution Pattern Classification Duda Read Pdf Free

Pattern Classification Pattern Classification 2nd Edition with Computer Manual 2nd Edition Set Pattern Recognition Pattern Classification and Scene Analysis Understanding Machine Learning Pattern Recognition and Neural Networks Pattern Classification Introduction to Pattern Recognition Pattern Classification Pattern Recognition and Machine Learning Patterns, Predictions, and Actions: Foundations of Machine Learning Statistical Pattern Recognition Structural, Syntactic, and Statistical Pattern Recognition Neural Networks for Pattern Recognition Pattern Classification and Scene Analysis A Probabilistic Theory of Pattern Recognition Introduction to Machine Learning

Image Processing and Pattern Recognition Pattern Recognition Optical Pattern Recognition PATTERN RECOGNITION: STATISTICAL, STRUCTURAL AND NEURAL APPROACHES Combining Pattern Classifiers Pattern Classification and Scene Analysis Deep Learning Correlation Pattern Recognition Advanced Lectures on Machine Learning Computer Methods in Image Analysis Pattern Classification and Scene Analysis Machine Learning in Action Computer Manual in Mathematica to Accompany Pattern Classification Machine Learning in Document Analysis and Recognition Materials Design Inspired by Nature Machine Learning and Data Mining Hands-On Pattern Recognition

Pattern Recognition Introduction to Statistical
Machine Learning Games of No Chance
Introduction to Statistical Pattern Recognition
Encyclopedia of Remote Sensing The Data
Science Design Manual

The objective of Document Analysis and Recognition (DAR) is to recognize the text and graphical components of a document and to extract information. This book is a collection of research papers and state-of-the-art reviews by leading researchers all over the world. It includes pointers to challenges and opportunities for future research directions. The main goal of the book is to identify good practices for the use of learning strategies in DAR. Observing the environment and recognising patterns for the purpose of decision making is fundamental to human nature. This book deals with the scientific discipline that enables similar perception in machines through pattern recognition (PR), which has application

in diverse technology areas. This book is an exposition of principal topics in PR using an algorithmic approach. It provides a thorough introduction to the concepts of PR and a systematic account of the major topics in PR besides reviewing the vast progress made in the field in recent times. It includes basic techniques of PR, neural networks, support vector machines and decision trees. While theoretical aspects have been given due coverage, the emphasis is more on the practical. The book is replete with examples and illustrations and includes chapter-end exercises. It is designed to meet the needs of senior undergraduate and postgraduate students of computer science and allied disciplines. A comprehensive guide to the essential principles of image processing and pattern recognition Techniques and applications in the areas of image processing and pattern recognition are growing at an unprecedented rate. Containing the latest state-of-the-art developments in the field, Image Processing and

Pattern Recognition presents clear explanations of the fundamentals as well as the most recent applications. It explains the essential principles so readers will not only be able to easily implement the algorithms and techniques, but also lead themselves to discover new problems and applications. Unlike other books on the subject, this volume presents numerous fundamental and advanced image processing algorithms and pattern recognition techniques to illustrate the framework. Scores of graphs and examples, technical assistance, and practical tools illustrate the basic principles and help simplify the problems, allowing students as well as professionals to easily grasp even complicated theories. It also features unique coverage of the most interesting developments and updated techniques, such as image watermarking, digital steganography, document processing and classification, solar image processing and event classification, 3-D Euclidean distance transformation, shortest path

planning, soft morphology, recursive morphology, regulated morphology, and sweep morphology. Additional topics include enhancement and segmentation techniques, active learning, feature extraction, neural networks, and fuzzy logic. Featuring supplemental materials for instructors and students, Image Processing and Pattern Recognition is designed for undergraduate seniors and graduate students, engineering and scientific researchers, and professionals who work in signal processing, image processing, pattern recognition, information security, document processing, multimedia systems, and solar physics. Introduction -- Supervised learning -- Bayesian decision theory -- Parametric methods -- Multivariate methods -- Dimensionality reduction -- Clustering -- Nonparametric methods -- Decision trees -- Linear discrimination -- Multilayer perceptrons -- Local models -- Kernel machines -- Graphical models -- Brief contents -- Hidden markov

models -- Bayesian estimation -- Combining multiple learners -- Reinforcement learning -- Design and analysis of machine learning experiments. Recently organized competitions have been instrumental in pushing the state-of-the-art in machine learning, establishing benchmarks to fairly evaluate methods, and identifying techniques that really work. This volume in the Challenges in Machine Learning series harvests three years of effort of hundreds of researchers who have participated in three competitions organized around five datasets from various application domains, designed to explore issues of data representation, model selection, and performance prediction. Summary Machine Learning in Action is unique book that blends the foundational theories of machine learning with the practical realities of building tools for everyday data analysis. You'll use the flexible Python programming language to build programs that implement algorithms for data classification, forecasting, recommendations,

and higher-level features like summarization and simplification. About the Book A machine is said to learn when its performance improves with experience. Learning requires algorithms and programs that capture data and ferret out the interesting or useful patterns. Once the specialized domain of analysts and mathematicians, machine learning is becoming a skill needed by many. Machine Learning in Action is a clearly written tutorial for developers. It avoids academic language and takes you straight to the techniques you'll use in your day-to-day work. Many (Python) examples present the core algorithms of statistical data processing, data analysis, and data visualization in code you can reuse. You'll understand the concepts and how they fit in with tactical tasks like classification, forecasting, recommendations, and higher-level features like summarization and simplification. Readers need no prior experience with machine learning or statistical processing. Familiarity with Python is

offsite.creighton.edu

helpful. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book. What's Inside A no-nonsense introduction Examples showing common ML tasks Everyday data analysis Implementing classic algorithms like Apriori and Adaboos Table of Contents PART 1 CLASSIFICATION Machine learning basics Classifying with k-Nearest Neighbors Splitting datasets one feature at a time: decision trees Classifying with probability theory: naïve Bayes Logistic regression Support vector machines Improving classification with the AdaBoost meta algorithm PART 2 FORECASTING NUMERIC VALUES WITH REGRESSION Predicting numeric values: regression Tree-based regression PART 3 UNSUPERVISED LEARNING Grouping unlabeled items using k-means clustering Association analysis with the Apriori algorithm Efficiently finding frequent itemsets with FP-growth PART 4 ADDITIONAL TOOLS Using principal component analysis to simplify

data Simplifying data with the singular value decomposition Big data and MapReduce This engaging and clearly written textbook/reference provides a must-have introduction to the rapidly emerging interdisciplinary field of data science. It focuses on the principles fundamental to becoming a good data scientist and the key skills needed to build systems for collecting, analyzing, and interpreting data. The Data Science Design Manual is a source of practical insights that highlights what really matters in analyzing data, and provides an intuitive understanding of how these core concepts can be used. The book does not emphasize any particular programming language or suite of data-analysis tools, focusing instead on high-level discussion of important design principles. This easy-to-read text ideally serves the needs of undergraduate and early graduate students embarking on an "Introduction to Data Science" course. It reveals how this discipline sits at the intersection of statistics, computer science, and

offsite.creighton.edu

machine learning, with a distinct heft and character of its own. Practitioners in these and related fields will find this book perfect for self-study as well. Additional learning tools: Contains “War Stories,” offering perspectives on how data science applies in the real world Includes “Homework Problems,” providing a wide range of exercises and projects for self-study Provides a complete set of lecture slides and online video lectures at www.data-manual.com Provides “Take-Home Lessons,” emphasizing the big-picture concepts to learn from each chapter Recommends exciting “Kaggle Challenges” from the online platform Kaggle Highlights “False Starts,” revealing the subtle reasons why certain approaches fail Offers examples taken from the data science television show “The Quant Shop” (www.quant-shop.com) Computer Manual (MATHEMATICA version) to accompany Pattern Classification, Third Edition, and its associated software contains all the MATHEMATICA code for the dynamic figures in the electronic version

offsite.creighton.edu

of PATTERN CLASSIFICATION, 3/e, (Duda, Hart, and Stork) as well as for core algorithms in pattern classification, clustering, and feature extraction described in the text. The code is cross-referenced with the material from the textbook, and uses the same terminology and symbols, so that the conceptual link from course material to working code is tight. Introduction to Mathematical Techniques in Pattern Recognition by Harry C. Andrews This volume is one of the first cohesive treatments of the use of mathematics for studying interactions between various recognition environments. It brings together techniques previously scattered throughout the literature and provides a concise common notation that will facilitate the understanding and comparison of the many aspects of mathematical pattern recognition. The contents of this volume are divided into five interrelated subject areas: Feature Selection, Distribution Free Classification, Statistical Classification, Nonsupervised Learning, and

Sequential Learning. Appendices describing specific aspects of feature selection and extensive reference and bibliographies are included. 1972 253 pp. Threshold Logic and its Applications by Saburo Muroga This is the first in-depth exposition of threshold logic and its applications using linear programming and integer programming as optimization tools. It presents threshold logic as a unified theory of conventional simple gates, threshold gates and their networks. This unified viewpoint explicitly reveals many important properties that were formerly concealed in the framework of conventional switching theory (based essentially on and, or and not gates). 1971 478 pp. Knowing and Guessing A Quantitative Study of Inference and Information By Satosi Watanabe This volume presents a coherent theoretical view of a field now split into different disciplines: philosophy, information science, cybernetics, psychology, electrical engineering, and physics. The target of investigation is the cognitive

process of knowing and guessing. In contrast to traditional philosophy, the approach is quantitative rather than qualitative. The study is formal in the sense that the author is not interested in the contents of knowledge or the physiological mechanism of the process of knowing. "The author's style is lucid, his comments are illuminating. The result is a fascinating book, which will be of interest to scientists in many different fields." — Nature 1969 592 pp. Statistical pattern recognition; Probability density estimation; Single-layer networks; The multi-layer perceptron; Radial basis functions; Error functions; Parameter optimization algorithms; Pre-processing and feature extraction; Learning and generalization; Bayesian techniques; Appendix; References; Index. This book constitutes the refereed proceedings of the 12th International Workshop on Structural and Syntactic Pattern Recognition, SSPR 2008 and the 7th International Workshop on Statistical Techniques in Pattern Recognition,

SPR 2008, held jointly in Orlando, FL, USA, in December 2008 as a satellite event of the 19th International Conference of Pattern Recognition, ICPR 2008. The 56 revised full papers and 42 revised poster papers presented together with the abstracts of 4 invited papers were carefully reviewed and selected from 175 submissions. The papers are organized in topical sections on graph-based methods, probabilistic and stochastic structural models for PR, image and video analysis, shape analysis, kernel methods, recognition and classification, applications, ensemble methods, feature selection, density estimation and clustering, computer vision and biometrics, pattern recognition and applications, pattern recognition, as well as feature selection and clustering. Good data mining practice for business intelligence (the art of turning raw software into meaningful information) is demonstrated by the many new techniques and developments in the conversion of fresh scientific discovery into widely accessible

software solutions. Written as an introduction to the main issues associated with the basics of machine learning and the algorithms used in data mining, this text is suitable for advanced undergraduates, postgraduates and tutors in a wide area of computer science and technology, as well as researchers looking to adapt various algorithms for particular data mining tasks. A valuable addition to libraries and bookshelves of the many companies who are using the principles of data mining to effectively deliver solid business and industry solutions. Is Nine-Men Morris, in the hands of perfect players, a win for white or for black - or a draw? Can king, rook, and knight always defeat king and two knights in chess? What can Go players learn from economists? What are nimbers, tinies, switches and minies? This book deals with combinatorial games, that is, games not involving chance or hidden information. Their study is at once old and young: though some games, such as chess, have been analyzed for

centuries, the first full analysis of a nontrivial combinatorial game (Nim) only appeared in 1902. The first part of this book will be accessible to anyone, regardless of background: it contains introductory expositions, reports of unusual tournaments, and a fascinating article by John H. Conway on the possibly everlasting contest between an angel and a devil. For those who want to delve more deeply, the book also contains combinatorial studies of chess and Go; reports on computer advances such as the solution of Nine-Men Morris and Pentominoes; and theoretical approaches to such problems as games with many players. If you have read and enjoyed Martin Gardner, or if you like to learn and analyze new games, this book is for you. This 1996 book explains the statistical framework for pattern recognition and machine learning, now in paperback. Covering pattern classification methods, *Combining Classifiers: Ideas and Methods* focuses on the important and widely studied issue of how to combine several

classifiers together in order to achieve improved recognition performance. It is one of the first books to provide unified, coherent, and expansive coverage of the topic and as such will be welcomed by those involved in the area. With case studies that bring the text alive and demonstrate 'real-world' applications it is destined to become essential reading. *PATTERN CLASSIFICATION* a unified view of statistical and neural approaches The product of years of research and practical experience in pattern classification, this book offers a theory-based engineering perspective on neural networks and statistical pattern classification. *Pattern Classification* sheds new light on the relationship between seemingly unrelated approaches to pattern recognition, including statistical methods, polynomial regression, multilayer perceptron, and radial basis functions. Important topics such as feature selection, reject criteria, classifier performance measurement, and classifier combinations are fully covered, as

offsite.creighton.edu

well as material on techniques that, until now, would have required an extensive literature search to locate. A full program of illustrations, graphs, and examples helps make the operations and general properties of different classification approaches intuitively understandable. Offering a lucid presentation of complex applications and their algorithms, Pattern Classification is an invaluable resource for researchers, engineers, and graduate students in this rapidly developing field. Machine Learning has become a key enabling technology for many engineering applications, investigating scientific questions and theoretical problems alike. To stimulate discussions and to disseminate new results, a summer school series was started in February 2002, the documentation of which is published as LNAI 2600. This book presents revised lectures of two subsequent summer schools held in 2003 in Canberra, Australia, and in Tübingen, Germany. The tutorial lectures included are devoted to statistical learning theory,

unsupervised learning, Bayesian inference, and applications in pattern recognition; they provide in-depth overviews of exciting new developments and contain a large number of references. Graduate students, lecturers, researchers and professionals alike will find this book a useful resource in learning and teaching machine learning. About The Book: This book explores the heart of pattern recognition concepts, methods and applications using statistical, syntactic and neural approaches. Divided into four sections, it clearly demonstrates the similarities and differences among the three approaches. The second part deals with the statistical pattern recognition approach, starting with a simple example and finishing with unsupervised learning through clustering. Section three discusses the syntactic approach and explores such topics as the capabilities of string grammars and parsing; higher dimensional representations and graphical approaches. Part four presents an excellent

overview of the emerging neural approach including an examination of pattern associations and feedforward nets. Along with examples, each chapter provides the reader with pertinent literature for a more in-depth study of specific topics. An introduction to a broad range of topics in deep learning, covering mathematical and conceptual background, deep learning techniques used in industry, and research perspectives. “Written by three experts in the field, Deep Learning is the only comprehensive book on the subject.” —Elon Musk, cochair of OpenAI; cofounder and CEO of Tesla and SpaceX

Deep learning is a form of machine learning that enables computers to learn from experience and understand the world in terms of a hierarchy of concepts. Because the computer gathers knowledge from experience, there is no need for a human computer operator to formally specify all the knowledge that the computer needs. The hierarchy of concepts allows the computer to learn complicated concepts by building them out

of simpler ones; a graph of these hierarchies would be many layers deep. This book introduces a broad range of topics in deep learning. The text offers mathematical and conceptual background, covering relevant concepts in linear algebra, probability theory and information theory, numerical computation, and machine learning. It describes deep learning techniques used by practitioners in industry, including deep feedforward networks, regularization, optimization algorithms, convolutional networks, sequence modeling, and practical methodology; and it surveys such applications as natural language processing, speech recognition, computer vision, online recommendation systems, bioinformatics, and videogames. Finally, the book offers research perspectives, covering such theoretical topics as linear factor models, autoencoders, representation learning, structured probabilistic models, Monte Carlo methods, the partition function, approximate inference, and deep

generative models. Deep Learning can be used by undergraduate or graduate students planning careers in either industry or research, and by software engineers who want to begin using deep learning in their products or platforms. A website offers supplementary material for both readers and instructors. Machine learning allows computers to learn and discern patterns without actually being programmed. When Statistical techniques and machine learning are combined together they are a powerful tool for analysing various kinds of data in many computer science/engineering areas including, image processing, speech processing, natural language processing, robot control, as well as in fundamental sciences such as biology, medicine, astronomy, physics, and materials. Introduction to Statistical Machine Learning provides a general introduction to machine learning that covers a wide range of topics concisely and will help you bridge the gap between theory and practice. Part I discusses the fundamental

concepts of statistics and probability that are used in describing machine learning algorithms. Part II and Part III explain the two major approaches of machine learning techniques; generative methods and discriminative methods. While Part III provides an in-depth look at advanced topics that play essential roles in making machine learning algorithms more useful in practice. The accompanying MATLAB/Octave programs provide you with the necessary practical skills needed to accomplish a wide range of data analysis tasks. Provides the necessary background material to understand machine learning such as statistics, probability, linear algebra, and calculus Complete coverage of the generative approach to statistical pattern recognition and the discriminative approach to statistical machine learning Includes MATLAB/Octave programs so that readers can test the algorithms numerically and acquire both mathematical and practical skills in a wide range of data analysis tasks Discusses a wide range of

offsite.creighton.edu

applications in machine learning and statistics and provides examples drawn from image processing, speech processing, natural language processing, robot control, as well as biology, medicine, astronomy, physics, and materials. This is the first textbook on pattern recognition to present the Bayesian viewpoint. The book presents approximate inference algorithms that permit fast approximate answers in situations where exact answers are not feasible. It uses graphical models to describe probability distributions when no other books apply graphical models to machine learning. No previous knowledge of pattern recognition or machine learning concepts is assumed. Familiarity with multivariate calculus and basic linear algebra is required, and some experience in the use of probabilities would be helpful though not essential as the book includes a self-contained introduction to basic probability theory. The inner architecture of a material can have an astonishing effect on its overall

properties and is vital to understand when designing new materials. Nature is a master at designing hierarchical structures and so researchers are looking at biological examples for inspiration, specifically to understand how nature arranges the inner architectures for a particular function in order to apply these design principles into man-made materials. *Materials Design Inspired by Nature* is the first book to address the relationship between the inner architecture of natural materials and their physical properties for materials design. The book explores examples from plants, the marine world, arthropods and bacteria, where the inner architecture is exploited to obtain specific mechanical, optical or magnetic properties along with how these design principles are used in man-made products. Details of the experimental methods used to investigate hierarchical structures are also given. Written by leading experts in bio-inspired materials research, this is essential reading for anyone developing new

offsite.creighton.edu

materials. This first encyclopaedic reference on remote sensing describes the concepts, techniques, instrumentation, data analysis, interpretation, and applications of remote sensing, both airborne and space-based. Scientists, engineers, academics, and students can quickly access answers to their reference questions and direction for further study. A self-contained and coherent account of probabilistic techniques, covering: distance measures, kernel rules, nearest neighbour rules, Vapnik-Chervonenkis theory, parametric classification, and feature extraction. Each chapter concludes with problems and exercises to further the readers understanding. Both research workers and graduate students will benefit from this wide-ranging and up-to-date account of a fast-moving field. A comprehensive review of optical pattern recognition techniques and implementations, for graduate students and researchers. Statistical pattern recognition is a very active area of study and research, which has

seen many advances in recent years. New and emerging applications - such as data mining, web searching, multimedia data retrieval, face recognition, and cursive handwriting recognition - require robust and efficient pattern recognition techniques. Statistical decision making and estimation are regarded as fundamental to the study of pattern recognition. Statistical Pattern Recognition, Second Edition has been fully updated with new methods, applications and references. It provides a comprehensive introduction to this vibrant area - with material drawn from engineering, statistics, computer science and the social sciences - and covers many application areas, such as database design, artificial neural networks, and decision support systems. * Provides a self-contained introduction to statistical pattern recognition. * Each technique described is illustrated by real examples. * Covers Bayesian methods, neural networks, support vector machines, and unsupervised classification.

* Each section concludes with a description of the applications that have been addressed and with further developments of the theory. * Includes background material on dissimilarity, parameter estimation, data, linear algebra and probability. * Features a variety of exercises, from 'open-book' questions to more lengthy projects. The book is aimed primarily at senior undergraduate and graduate students studying statistical pattern recognition, pattern processing, neural networks, and data mining, in both statistics and engineering departments. It is also an excellent source of reference for technical professionals working in advanced information development environments. For further information on the techniques and applications discussed in this book please visit <http://www.statistical-pattern-recognition.net/> www.statistical-pattern-recognition.net/a Introduction to Pattern Recognition: A Matlab Approach is an accompanying manual to

www.creighton.edu

Theodoridis/Koutroumbas' Pattern Recognition. It includes Matlab code of the most common methods and algorithms in the book, together with a descriptive summary and solved examples, and including real-life data sets in imaging and audio recognition. This text is designed for electronic engineering, computer science, computer engineering, biomedical engineering and applied mathematics students taking graduate courses on pattern recognition and machine learning as well as R&D engineers and university researchers in image and signal processing/analysis, and computer vision. Matlab code and descriptive summary of the most common methods and algorithms in Theodoridis/Koutroumbas, Pattern Recognition, Fourth Edition Solved examples in Matlab, including real-life data sets in imaging and audio recognition Available separately or at a special package price with the main text (ISBN for package: 978-0-12-374491-3) The first edition, published in 1973, has become a classic

reference in the field. Now with the second edition, readers will find information on key new topics such as neural networks and statistical pattern recognition, the theory of machine learning, and the theory of invariances. Also included are worked examples, comparisons between different methods, extensive graphics, expanded exercises and computer project topics. An Instructor's Manual presenting detailed solutions to all the problems in the book is available from the Wiley editorial department. The first edition, published in 1973, has become a classic reference in the field. Now with the second edition, readers will find information on key new topics such as neural networks and statistical pattern recognition, the theory of machine learning, and the theory of invariances. Also included are worked examples, comparisons between different methods, extensive graphics, expanded exercises and computer project topics. An Instructor's Manual presenting detailed solutions to all the problems in the book is

offsite.creighton.edu

available from the Wiley editorial department. Introduces machine learning and its algorithmic paradigms, explaining the principles behind automated learning approaches and the considerations underlying their usage. Correlation is a robust and general technique for pattern recognition and is used in many applications, such as automatic target recognition, biometric recognition and optical character recognition. The design, analysis and use of correlation pattern recognition algorithms requires background information, including linear systems theory, random variables and processes, matrix/vector methods, detection and estimation theory, digital signal processing and optical processing. This book provides a needed review of this diverse background material and develops the signal processing theory, the pattern recognition metrics, and the practical application know-how from basic premises. It shows both digital and optical implementations. It also contains technology presented by the

team that developed it and includes case studies of significant interest, such as face and fingerprint recognition. Suitable for graduate students taking courses in pattern recognition theory, whilst reaching technical levels of interest to the professional practitioner. This book constitutes the proceedings of the 12th Mexican Conference on Pattern Recognition, MCPR 2020, which was due to be held in Morelia, Mexico, in June 2020. The conference was held virtually due to the COVID-19 pandemic. The 31 papers presented in this volume were carefully reviewed and selected from 67 submissions. They were organized in the following topical sections: pattern recognition techniques; image processing and analysis; computer vision; industrial and medical applications of pattern recognition; natural language processing and recognition; artificial intelligence techniques and recognition. This completely revised second edition presents an introduction to statistical pattern recognition.

[offsite.creighton.edu](https://www.creighton.edu)

Pattern recognition in general covers a wide range of problems: it is applied to engineering problems, such as character readers and wave form analysis as well as to brain modeling in biology and psychology. Statistical decision and estimation, which are the main subjects of this book, are regarded as fundamental to the study of pattern recognition. This book is appropriate as a text for introductory courses in pattern recognition and as a reference book for workers in the field. Each chapter contains computer projects as well as exercises. Pattern recognition is a scientific discipline that is becoming increasingly important in the age of automation and information handling and retrieval. Pattern Recognition, 2e covers the entire spectrum of pattern recognition applications, from image analysis to speech recognition and communications. This book presents cutting-edge material on neural networks, - a set of linked microprocessors that can form associations and uses pattern recognition to

"learn" -and enhances student motivation by approaching pattern recognition from the designer's point of view. A direct result of more than 10 years of teaching experience, the text was developed by the authors through use in their own classrooms. *Approaches pattern recognition from the designer's point of view *New edition highlights latest developments in this growing field, including independent components and support vector machines, not available elsewhere *Supplemented by computer examples selected from applications of interest An authoritative, up-to-date graduate textbook on machine learning that highlights its historical context and societal impacts Patterns, Predictions, and Actions introduces graduate students to the essentials of machine learning while offering invaluable perspective on its history and social implications. Beginning with the foundations of decision making, Moritz

Hardt and Benjamin Recht explain how representation, optimization, and generalization are the constituents of supervised learning. They go on to provide self-contained discussions of causality, the practice of causal inference, sequential decision making, and reinforcement learning, equipping readers with the concepts and tools they need to assess the consequences that may arise from acting on statistical decisions. Provides a modern introduction to machine learning, showing how data patterns support predictions and consequential actions Pays special attention to societal impacts and fairness in decision making Traces the development of machine learning from its origins to today Features a novel chapter on machine learning benchmarks and datasets Invites readers from all backgrounds, requiring some experience with probability, calculus, and linear algebra An essential textbook for students and a guide for researchers