

Machine Learning Algorithmic Perspective Recognition

Yeah, reviewing a ebook Machine Learning Algorithmic Perspective Recognition could go to your close connections listings. This is just one of the solutions for you to be successful. As understood, completion does not recommend that you have fabulous points.

Comprehending as skillfully as arrangement even more than extra will have the funds for each success. next to, the proclamation as without difficulty as sharpness of this Machine Learning Algorithmic Perspective Recognition can be taken as well as picked to act.

Pattern Recognition M. Narasimha Murty 2011-05-25 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.

Just Enough R! Richard J. Roiger 2020-05-20 Just Enough R! An Interactive Approach to Machine Learning and Analytics presents just enough of the R language, machine learning algorithms, statistical methodology, and analytics for the reader to learn how to find interesting structure in data. The approach might be called "seeing then doing" as it first gives step-by-step explanations using simple, understandable examples of how the various machine learning algorithms work independent of any programming language. This is followed by detailed scripts written in R that apply the algorithms to solve nontrivial problems with real data. The script code is provided, allowing the reader to execute the scripts as they study the explanations given in the text. Features Gets you quickly using R as a problem-solving tool Uses RStudio's integrated development environment Shows how to interface R with SQLite Includes examples using R's Rattle graphical user interface Requires no prior knowledge of R, machine learning, or computer programming Offers over 50 scripts written in R, including several problem-solving templates that, with slight modification, can be used again and again Covers the most popular machine learning techniques, including ensemble-based methods and logistic regression Includes end-of-chapter exercises, many of which can be solved by modifying existing scripts Includes datasets from several areas, including business, health and medicine, and science About the Author Richard J. Roiger is a professor emeritus at Minnesota State University, Mankato, where he taught and performed research in the Computer and Information Science Department for over 30 years.

Data Science and Machine Learning Dirk P. Kroese 2019-11-20 "This textbook is a well-rounded, rigorous, and informative work presenting the mathematics behind modern machine learning techniques. It hits all the right notes: the choice of topics is up-to-date and perfect for a course on data science for mathematics students at the advanced undergraduate or early graduate level. This book fills a sorely-needed gap in the existing literature by not sacrificing depth for breadth, presenting proofs of major theorems and subsequent derivations, as well as providing a copious amount of Python code. I only wish a book like this had been around when I first began my journey!" -Nicholas Hoell, University of Toronto "This is a well-written book that provides a deeper dive into data-scientific methods than many introductory texts. The writing is clear, and the text logically builds up regularization, classification, and decision trees. Compared to its probable competitors, it carves out a unique niche. -Adam Loy, Carleton College The purpose of Data Science and Machine Learning: Mathematical and Statistical Methods is to provide an accessible, yet comprehensive textbook intended for students interested in gaining a better understanding of the mathematics and statistics that underpin the rich variety of ideas and machine learning algorithms in data science. Key Features: Focuses on mathematical understanding. Presentation is self-contained, accessible, and comprehensive. Extensive list of exercises and worked-out examples. Many concrete algorithms with Python code. Full color throughout. Further Resources can be found on the authors website: <https://github.com/DSML-book/Lectures>

Introduction to Machine Learning Ethem Alpaydin 2014-08-22 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.

Intelligent Video Surveillance Systems Maheshkumar H Kolekar 2018-06-27 This book will provide an overview of techniques for visual monitoring including video surveillance and human activity understanding. It will present the basic techniques of processing video from static cameras, starting with object detection and tracking. The author will introduce further video analytic modules including face detection, trajectory analysis and object classification. Examining system design and specific problems in visual surveillance, such as the use of multiple cameras and moving cameras, the author will elaborate on privacy issues focusing on approaches where automatic processing can help protect privacy.

Data Mining Richard J. Roiger 2017-01-06 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.

Deep Learning in Biometrics Mayank Vatsa 2018-03-05 Deep Learning is now synonymous with applied machine learning. Many technology giants (e.g. Google, Microsoft, Apple, IBM) as well as start-ups are focusing on deep learning-based techniques for data analytics and artificial intelligence. This technology applies quite strongly to biometrics. This book covers topics in deep learning, namely convolutional neural networks, deep belief network and stacked autoencoders. The focus is also on the application of these techniques to various biometric modalities: face, iris, palmprint, and fingerprints, while examining the future trends in deep learning and biometric research. Contains chapters written by authors who are leading researchers in biometrics. Presents a comprehensive overview on the internal mechanisms of deep learning. Discusses the latest developments in biometric research. Examines future trends in deep learning and biometric research. Provides extensive references at the end of each chapter to enhance further study.

An Algorithmic Perspective on Imitation Learning Takayuki Osa 2018-03-27 Familiarizes machine learning experts with imitation learning, statistical supervised learning theory, and reinforcement learning. It also roboticists and experts in applied artificial intelligence with a broader appreciation for the frameworks and tools available for imitation learning.

Applied Natural Language Processing with Python Taweh Beysolow II 2018-09-11 Learn to harness the power of AI for natural language processing, performing tasks such as spell check, text summarization, document classification, and natural language generation. Along the way, you will learn the skills to implement these methods in larger infrastructures to replace existing code or create new algorithms. Applied Natural Language Processing with Python starts with reviewing the necessary machine learning concepts before moving onto discussing various NLP problems. After reading this book, you will have the skills to apply these concepts in your own professional environment. What You Will Learn Utilize various machine learning and natural language processing libraries such as TensorFlow, Keras, NLTK, and Gensim Manipulate and preprocess raw text data in formats such as .txt and .pdf Strengthen your skills in data science by learning both the theory and the application of various algorithms Who This Book Is For You should be at least a beginner in ML to get the most out of this text, but you needn't feel that you need be an expert to understand the content.

Chinese Handwriting Recognition: An Algorithmic Perspective Tonghua Su 2013-01-11 Designing machines that can read handwriting like human beings has been an ambitious goal for more than half a century, driving talented researchers to explore diverse approaches. Obstacles have often been encountered that at first appeared insurmountable but were indeed overcome before long. Yet some open issues remain to be solved. As an indispensable branch, Chinese handwriting recognition has been termed as one of the most difficult Pattern Recognition tasks. Chinese handwriting recognition poses its own unique challenges, such as huge variations in strokes, diversity of writing styles, and a large set of confusable categories. With ever-increasing training data, researchers have pursued elaborate algorithms to discern characters from different categories and compensate for the sample variations within the same category. As a result, Chinese handwriting recognition has evolved substantially and amazing achievements can be seen. This book introduces integral algorithms used in Chinese handwriting recognition and the applications of Chinese handwriting recognizers. The first part of the book covers both widespread canonical algorithms to a reliable recognizer and newly developed scalable methods in Chinese handwriting recognition. The recognition of Chinese handwritten text is presented systematically, including instructive guidelines for collecting samples, novel recognition paradigms, distributed discriminative learning of appearance models and distributed estimation of contextual models for large categories, in addition to celebrated methods, e.g. Gradient features, MQDF and HMMs. In the second part of this book, endeavors are made to create a friendlier human-machine interface through application of Chinese handwriting recognition. Four scenarios are exemplified: grid-assisted input, shortest moving input, handwritten micro-blog, and instant handwriting messenger. All the while, the book moves from basic to more complex approaches, also providing a list for further reading with literature comments.

Pattern Recognition and Machine Learning Christopher M. Bishop 2016-08-23 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.

Understanding Machine Learning Shai Shalev-Shwartz 2014-05-19 Introduces machine learning and its algorithmic paradigms, explaining the principles behind automated learning approaches and the considerations underlying their usage.

Introduction to Pattern Recognition and Machine Learning M Narasimha Murty 2015-04-22 This book adopts a detailed and methodological algorithmic approach to explain the concepts of pattern recognition. While the text provides a systematic account of its major topics such as pattern representation and nearest neighbour based classifiers, current topics — neural networks, support vector machines and decision trees — attributed to the recent vast progress in this field are also dealt with. Introduction to Pattern Recognition and Machine Learning will equip readers, especially senior computer science undergraduates, with a deeper understanding of the subject matter. Contents: Introduction Types of Data Feature Extraction and Feature Selection Bayesian Learning Classification Classification Using Soft Computing Techniques Data Clustering Soft Clustering Application — Social and Information Networks Readership: Academics and working professionals in computer science. Key Features: The algorithmic approach taken and the practical issues dealt with will aid the reader in writing programs and implementing methods Covers recent and advanced topics by providing working exercises, examples and illustrations in each chapter Provides the reader with a deeper understanding of the subject matter Keywords: Clustering; Classification; Supervised Learning; Soft Computing

Machine Learning Kevin P. Murphy 2012-08-24 A comprehensive introduction to machine learning that uses probabilistic models and inference as a unifying approach. Today's Web-enabled deluge of electronic data calls for automated methods of data analysis. Machine learning provides these, developing methods that can automatically detect patterns in data and then use the uncovered patterns to predict future data. This textbook offers a comprehensive and self-contained introduction to the field of machine learning, based on a unified, probabilistic approach. The coverage combines breadth and depth, offering necessary background material on such topics as probability, optimization, and linear algebra as well as discussion of recent developments in the field, including conditional random fields, L1 regularization, and deep learning. The book is written in an informal, accessible style, complete with pseudo-code for the most important algorithms. All topics are copiously illustrated with color images and worked examples drawn from such application domains as biology, text processing, computer vision, and robotics. Rather than providing a cookbook of different heuristic methods, the book stresses a principled model-based approach, often using the language of graphical models to specify models in a concise and intuitive way. Almost all the models described have been implemented in a MATLAB software package—PMTK (probabilistic modeling toolkit)—that is freely available online. The book is suitable for upper-level undergraduates with an introductory-level college math background and beginning graduate students.

The Digital Mind Arlindo Oliveira 2017-03-17 The Red Queen's race -- The exponential nature of technology -- From Maxwell to the Internet -- The universal machine -- The quest for intelligent machines -- Cells, bodies, and brains -- Biology meets computation -- How the brain works -- Understanding the brain -- Brains, minds, and machines -- Challenges and promises -- Speculations

Algorithmic Aspects of Machine Learning Ankur Moitra 2018-09-27 Introduces cutting-edge research on machine learning theory and practice, providing an accessible, modern algorithmic toolkit.

Chinese Handwriting Recognition: An Algorithmic Perspective Tonghua Su 2013-01-11 Designing machines that can read handwriting like human beings has been an ambitious goal for more than half a century, driving talented researchers to explore diverse approaches. Obstacles have often been encountered that at first appeared insurmountable but were indeed overcome before long. Yet some open issues remain to be solved. As an indispensable branch, Chinese handwriting recognition has been termed as one of the most difficult Pattern Recognition tasks. Chinese handwriting recognition poses its own unique challenges, such as huge variations in strokes, diversity of writing styles, and a large set of confusable categories. With ever-increasing training data, researchers have pursued elaborate algorithms to discern characters from different categories and compensate for the sample variations within the same category. As a result, Chinese handwriting recognition has evolved substantially and amazing achievements can be seen. This book introduces integral algorithms used in Chinese handwriting recognition and the applications of Chinese handwriting recognizers. The first part of the book covers both widespread canonical algorithms to a reliable recognizer and newly developed scalable methods in Chinese handwriting recognition. The recognition of Chinese handwritten text is presented systematically, including instructive guidelines for collecting samples, novel recognition paradigms, distributed discriminative learning of appearance models and distributed estimation of contextual models for large categories, in addition to celebrated methods, e.g. Gradient features, MQDF and HMMs. In the second part of this book, endeavors are made to create a friendlier human-machine interface through application of Chinese handwriting recognition. Four scenarios are exemplified: grid-assisted input, shortest moving input, handwritten micro-blog, and instant handwriting messenger. All the while, the book moves from basic to more complex approaches, also providing a list for further reading with literature comments.

Ensemble Methods Zhi-Hua Zhou 2012-06-06 An up-to-date, self-contained introduction to a state-of-the-art machine learning approach, Ensemble Methods: Foundations and Algorithms shows how these accurate methods are used in real-world tasks. It gives you the necessary groundwork to carry out further research in this evolving field. After presenting background and terminology, the book covers the main algorithms and theories, including Boosting, Bagging, Random Forest, averaging and voting schemes, the Stacking method, mixture of experts, and diversity measures. It also discusses multiclass extension, noise tolerance, error-ambiguity and bias-variance decompositions, and recent progress in information theoretic diversity. Moving on to more advanced topics, the author explains how to achieve better performance through ensemble pruning and how to generate better clustering results by combining multiple clusterings. In addition, he describes developments of ensemble methods in semi-supervised learning, active learning, cost-sensitive learning, class-imbalance learning, and comprehensibility enhancement.

Port Security Management Kenneth Christopher 2009-03-24 The term homeland security hardly existed before September 11, 2001, yet today it dominates public policy and the economic agendas of world governments. The transportation industries have been subjected to unprecedented scrutiny and regulatory mandates in recent years, and the port and maritime sector are no exception. Port Security Management reflects this altered landscape of the post-9/11 era, providing real-world guidelines for strategic security planning and implementation processes. Balance security with business needs The book begins with a historical and organizational perspective on maritime and port security. It then discusses the management of risk assessment, presenting it within the context of the unique vulnerabilities within the maritime and port environments. The important relationships between risk analysis, facility security planning, and coordination among port stakeholders—including the public and private sector businesses—provide the framework for understanding the pivotal role of security managers, security personnel, and law enforcement in ensuring the safety and security of port users and their interests. Work cohesively with governmental and private entities The text also addresses the ground-level issues, tasks, and responsibilities that must be managed by the security manager in concert with the port director and federal and local law enforcement agencies. The author explores the growth of multiuse port facilities for recreation, hospitality, and external business and commercial interests and offers perspectives on the role of technology in security. Finally, the book examines the need to develop contingency and emergency operations plans and work effectively with federal, state, local, and private enterprises in coordinating both routine and emergency response mechanisms.

A Concise Introduction to Machine Learning A.C. Faul 2019-08-01 The emphasis of the book is on the question of Why - only if why an algorithm is successful is understood, can it be properly applied, and the results trusted. Algorithms are often taught side by side without showing the similarities and differences between them. This book addresses the commonalities, and aims to give a thorough and in-depth treatment and develop intuition, while remaining concise. This useful reference should be an essential on the bookshelves of anyone employing machine learning techniques.

MACHINE LEARNING STEPHEN. MARSLAND 2023

Machine Learning Sergios Theodoridis 2020-02-19 Machine Learning: A Bayesian and Optimization Perspective, 2nd edition, gives a unified perspective on machine learning by covering both pillars of supervised learning, namely regression and classification. The book starts with the basics, including mean square, least squares and maximum likelihood methods, ridge regression, Bayesian decision theory classification, logistic regression, and decision trees. It then progresses to more recent techniques, covering sparse modelling methods, learning in reproducing kernel Hilbert spaces and support vector machines, Bayesian inference with a focus on the EM algorithm and its approximate inference variational versions, Monte Carlo methods, probabilistic graphical models focusing on Bayesian networks, hidden Markov models and particle filtering. Dimensionality reduction and latent variables modelling are also considered in depth. This palette of techniques concludes with an extended chapter on neural networks and deep learning architectures. The book also covers the fundamentals of statistical parameter estimation, Wiener and Kalman filtering, convexity and convex optimization, including a chapter on stochastic approximation and the gradient descent family of algorithms, presenting related online learning techniques as well as concepts and algorithmic versions for distributed optimization. Focusing on the physical reasoning behind the mathematics, without sacrificing rigor, all the various methods and techniques are explained in depth, supported by examples and problems, giving an invaluable resource to the student and researcher for understanding and applying machine learning concepts. Most of the chapters include typical case studies and computer exercises, both in MATLAB and Python. The chapters are written to be as self-contained as possible, making the text suitable for different courses: pattern recognition, statistical/adaptive signal processing, statistical/Bayesian learning, as well as courses on sparse modeling, deep learning, and probabilistic graphical models. New to this edition: Complete re-write of the chapter on Neural Networks and Deep Learning to reflect the latest advances since the 1st edition. The chapter, starting from the basic perceptron and feed-forward neural networks concepts, now presents an in depth treatment of deep networks, including recent optimization algorithms, batch normalization, regularization techniques such as the dropout method, convolutional neural networks, recurrent neural networks, attention mechanisms, adversarial examples and training, capsule networks and generative architectures, such as restricted Boltzmann machines (RBMs), variational autoencoders and generative adversarial networks (GANs). Expanded treatment of Bayesian learning to include nonparametric Bayesian methods, with a focus on the Chinese restaurant and the Indian buffet processes. Presents the physical reasoning, mathematical modeling and algorithmic implementation of each method Updates on the latest trends, including sparsity, convex analysis and optimization, online distributed algorithms, learning in RKH spaces, Bayesian inference, graphical and hidden Markov models, particle filtering, deep learning, dictionary learning and latent variables modeling Provides case studies on a variety of topics, including protein folding prediction, optical character recognition, text authorship identification, fMRI data analysis, change point

detection, hyperspectral image unmixing, target localization, and more

Building Machine Learning Systems with Python - Second Edition Luis Pedro Coelho 2015-03-26 This book primarily targets Python developers who want to learn and use Python's machine learning capabilities and gain valuable insights from data to develop effective solutions for business problems.

Foundations of Data Science Avrim Blum 2020-01-23 This book provides an introduction to the mathematical and algorithmic foundations of data science, including machine learning, high-dimensional geometry, and analysis of large networks. Topics include the counterintuitive nature of data in high dimensions, important linear algebraic techniques such as singular value decomposition, the theory of random walks and Markov chains, the fundamentals of and important algorithms for machine learning, algorithms and analysis for clustering, probabilistic models for large networks, representation learning including topic modelling and non-negative matrix factorization, wavelets and compressed sensing. Important probabilistic techniques are developed including the law of large numbers, tail inequalities, analysis of random projections, generalization guarantees in machine learning, and moment methods for analysis of phase transitions in large random graphs. Additionally, important structural and complexity measures are discussed such as matrix norms and VC-dimension. This book is suitable for both undergraduate and graduate courses in the design and analysis of algorithms for data.

Machine Learning in Action Peter Harrington 2012-04-03 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 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 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: naive 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

Mathematics for Machine Learning Marc Peter Deisenroth 2020-04-23 The fundamental mathematical tools needed to understand machine learning include linear algebra, analytic geometry, matrix decompositions, vector calculus, optimization, probability and statistics. These topics are traditionally taught in disparate courses, making it hard for data science or computer science students, or professionals, to efficiently learn the mathematics. This self-contained textbook bridges the gap between mathematical and machine learning texts, introducing the mathematical concepts with a minimum of prerequisites. It uses these concepts to derive four central machine learning methods: linear regression, principal component analysis, Gaussian mixture models and support vector machines. For students and others with a mathematical background, these derivations provide a starting point to machine learning texts. For those learning the mathematics for the first time, the methods help build intuition and practical experience with applying mathematical concepts. Every chapter includes worked examples and exercises to test understanding. Programming tutorials are offered on the book's web site.

Machine Learning Stephen Marsland 2011-03-23 Traditional books on machine learning can be divided into two groups- those aimed at advanced undergraduates or early postgraduates with reasonable mathematical knowledge and those that are primers on how to code algorithms. The field is ready for a text that not only demonstrates how to use the algorithms that make up machine learning methods, but

Machine Learning for Algorithmic Trading Stefan Jansen 2020-07-31 Leverage machine learning to design and back-test automated trading strategies for real-world markets using pandas, TA-Lib, scikit-learn, LightGBM, SpaCy, Gensim, TensorFlow 2, Zipline, backtrader, Alphalens, and pyfolio. Key Features Design, train, and evaluate machine learning algorithms that underpin automated trading strategies Create a research and strategy development process to apply predictive modeling to trading decisions Leverage NLP and deep learning to extract tradeable signals from market and alternative data Book Description The explosive growth of digital data has boosted the demand for expertise in trading strategies that use machine learning (ML). This revised and expanded second edition enables you to build and evaluate sophisticated supervised, unsupervised, and reinforcement learning models. This book introduces end-to-end machine learning for the trading workflow, from the idea and feature engineering to model optimization, strategy design, and backtesting. It illustrates this by using examples ranging from linear models and tree-based ensembles to deep-learning techniques from cutting edge research. This edition shows how to work with market, fundamental, and alternative data, such as tick data, minute and daily bars, SEC filings, earnings call transcripts, financial news, or satellite images to generate tradeable signals. It illustrates how to engineer financial features or alpha factors that enable an ML model to predict returns from price data for US and international stocks and ETFs. It also shows how to assess the signal content of new features using Alphalens and SHAP values and includes a new appendix with over one hundred alpha factor examples. By the end, you will be proficient in translating ML model predictions into a trading strategy that operates at daily or intraday horizons, and in evaluating its performance. What you will learn Leverage market, fundamental, and alternative text and image data Research and evaluate alpha factors using statistics, Alphalens, and SHAP values Implement machine learning techniques to solve investment and trading problems Backtest and evaluate trading strategies based on machine learning using Zipline and Backtrader Optimize portfolio risk and performance analysis using pandas, NumPy, and pyfolio Create a pairs trading strategy based on cointegration for US equities and ETFs Train a gradient boosting model to predict intraday returns using AlgoSeek's high-quality trades and quotes data Who this book is for If you are a data analyst, data scientist, Python developer, investment analyst, or portfolio manager interested in getting hands-on machine learning knowledge for trading, this book is for you. This book is for you if you want to learn how to extract value from a diverse set of data sources using machine learning to design your own systematic trading strategies. Some understanding of Python and machine learning techniques is required.

Trends in Deep Learning Methodologies Vincenzo Piuri 2020-11-12 Trends in Deep Learning Methodologies: Algorithms, Applications, and Systems covers deep learning approaches such as neural networks, deep belief networks, recurrent neural networks, convolutional neural networks, deep auto-encoder, and deep generative networks, which have emerged as powerful computational models. Chapters elaborate on these models which have shown significant success in dealing with massive data for a large number of applications, given their capacity to extract complex hidden features and learn efficient representation in unsupervised settings. Chapters investigate deep learning-based algorithms in a variety of application, including biomedical and health informatics, computer vision, image processing, and more. In recent years, many powerful algorithms have been developed for matching patterns in data and making predictions about future events. The major advantage of deep learning is to process big data analytics for better analysis and self-adaptive algorithms to handle more data. Deep learning methods can deal with multiple levels of representation in which the system learns to abstract higher level representations of raw data. Earlier, it was a common requirement to have a domain expert to develop a specific model for each specific application, however, recent advancements in representation learning algorithms allow researchers across various subject domains to automatically learn the patterns and representation of the given data for the development of specific models. Provides insights into the theory, algorithms, implementation and the application of deep learning techniques Covers a wide range of applications of deep learning across smart healthcare and smart engineering Investigates the development of new models and how they can be exploited to find appropriate solutions

Information Theory, Inference and Learning Algorithms David J. C. MacKay 2003-09-25 Table of contents

Statistical Machine Learning Richard Golden 2020-06-24 The recent rapid growth in the variety and complexity of new machine learning architectures requires the development of improved methods for designing, analyzing, evaluating, and communicating machine learning technologies. Statistical Machine Learning: A Unified Framework provides students, engineers, and scientists with tools from mathematical statistics and nonlinear optimization theory to become experts in the field of machine learning. In particular, the material in this text directly supports the mathematical analysis and design of old, new, and not-yet-invented nonlinear high-dimensional machine learning algorithms. Features: Unified empirical risk minimization framework supports rigorous mathematical analyses of widely used supervised, unsupervised, and reinforcement machine learning algorithms Matrix calculus methods for supporting machine learning analysis and design applications Explicit conditions for ensuring convergence of adaptive, batch, minibatch, MCEM, and MCMC learning algorithms that minimize both unimodal and multimodal objective functions Explicit conditions for characterizing asymptotic properties of M-estimators and model selection criteria such as AIC and BIC in the presence of possible model misspecification This advanced text is suitable for graduate students or highly motivated undergraduate students in statistics, computer science, electrical engineering, and applied mathematics. The text is self-contained and only assumes knowledge of lower-division linear algebra and upper-division probability theory. Students, professional engineers, and multidisciplinary scientists possessing these minimal prerequisites will find this text challenging yet accessible. About the Author: Richard M. Golden (Ph.D., M.S.E.E., B.S.E.E.) is Professor of Cognitive Science and Participating Faculty Member in Electrical Engineering at the University of Texas at Dallas. Dr. Golden has published articles and given talks at scientific conferences on a wide range of topics in the fields of both statistics and machine learning over the past three decades. His long-term research interests include identifying conditions for the convergence of deterministic and stochastic machine learning algorithms and investigating estimation and inference in the presence of possibly misspecified probability models.

Multidimensional Particle Swarm Optimization for Machine Learning and Pattern Recognition Serkan Kiranyaz 2013-07-16 For many engineering problems we require optimization processes with dynamic adaptation as we aim to establish the dimension of the search space where the optimum solution resides and develop robust techniques to avoid the local optima usually associated with multimodal problems. This book explores multidimensional particle swarm optimization, a technique developed by the authors that addresses these requirements in a well-defined algorithmic approach. After an introduction to the key optimization techniques, the authors introduce their unified framework and demonstrate its advantages in challenging application domains, focusing on the state of the art of multidimensional extensions such as global convergence in particle swarm optimization, dynamic data clustering, evolutionary neural networks, biomedical applications and personalized ECG classification, content-based image classification and retrieval, and evolutionary feature synthesis. The content is characterized

by strong practical considerations, and the book is supported with fully documented source code for all applications presented, as well as many sample datasets. The book will be of benefit to researchers and practitioners working in the areas of machine intelligence, signal processing, pattern recognition, and data mining, or using principles from these areas in their application domains. It may also be used as a reference text for graduate courses on swarm optimization, data clustering and classification, content-based multimedia search, and biomedical signal processing applications.

Introduction to Statistical Machine Learning Masashi Sugiyama 2015-10-31 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 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.

Artificial Intelligence and Causal Inference Momiao Xiong 2022-02-03 Artificial Intelligence and Causal Inference address the recent development of relationships between artificial intelligence (AI) and causal inference. Despite significant progress in AI, a great challenge in AI development we are still facing is to understand mechanism underlying intelligence, including reasoning, planning and imagination. Understanding, transfer and generalization are major principles that give rise intelligence. One of a key component for understanding is causal inference. Causal inference includes intervention, domain shift learning, temporal structure and counterfactual thinking as major concepts to understand causation and reasoning. Unfortunately, these essential components of the causality are often overlooked by machine learning, which leads to some failure of the deep learning. AI and causal inference involve (1) using AI techniques as major tools for causal analysis and (2) applying the causal concepts and causal analysis methods to solving AI problems. The purpose of this book is to fill the gap between the AI and modern causal analysis for further facilitating the AI revolution. This book is ideal for graduate students and researchers in AI, data science, causal inference, statistics, genomics, bioinformatics and precision medicine. Key Features: Cover three types of neural networks, formulate deep learning as an optimal control problem and use Pontryagin's Maximum Principle for network training. Deep learning for nonlinear mediation and instrumental variable causal analysis. Construction of causal networks is formulated as a continuous optimization problem. Transformer and attention are used to encode-decode graphics. RL is used to infer large causal networks. Use VAE, GAN, neural differential equations, recurrent neural network (RNN) and RL to estimate counterfactual outcomes. AI-based methods for estimation of individualized treatment effect in the presence of network interference.

Port Security Management Kenneth Christopher 2014-06-20 Sea and freshwater ports are a key component of critical infrastructure and essential for maintaining global and domestic economies. In order to effectively secure a dynamic port facility operation, one must understand the business of maritime commerce. Following in the tradition of its bestselling predecessor, Port Security Management, Second Edit

Machine Learning Steven W. Knox 2018-03-08 AN INTRODUCTION TO MACHINE LEARNING THAT INCLUDES THE FUNDAMENTAL TECHNIQUES, METHODS, AND APPLICATIONS Machine Learning: a Concise Introduction offers a comprehensive introduction to the core concepts, approaches, and applications of machine learning. The author—an expert in the field—presents fundamental ideas, terminology, and techniques for solving applied problems in classification, regression, clustering, density estimation, and dimension reduction. The design principles behind the techniques are emphasized, including the bias-variance trade-off and its influence on the design of ensemble methods. Understanding these principles leads to more flexible and successful applications. Machine Learning: a Concise Introduction also includes methods for optimization, risk estimation, and model selection—essential elements of most applied projects. This important resource: Illustrates many classification methods with a single, running example, highlighting similarities and differences between methods Presents R source code which shows how to apply and interpret many of the techniques covered Includes many thoughtful exercises as an integral part of the text, with an appendix of selected solutions Contains useful information for effectively communicating with clients A volume in the popular Wiley Series in Probability and Statistics, Machine Learning: a Concise Introduction offers the practical information needed for an understanding of the methods and application of machine learning. STEVEN W. KNOX holds a Ph.D. in Mathematics from the University of Illinois and an M.S. in Statistics from Carnegie Mellon University. He has over twenty years' experience in using Machine Learning, Statistics, and Mathematics to solve real-world problems. He currently serves as Technical Director of Mathematics Research and Senior Advocate for Data Science at the National Security Agency.

Introduction to Machine Learning with Applications in Information Security Mark Stamp 2017-09-22 Introduction to Machine Learning with Applications in Information Security provides a class-tested introduction to a wide variety of machine learning algorithms, reinforced through realistic applications. The book is accessible and doesn't prove theorems, or otherwise dwell on mathematical theory. The goal is to present topics at an intuitive level, with just enough detail to clarify the underlying concepts. The book covers core machine learning topics in-depth, including Hidden Markov Models, Principal Component Analysis, Support Vector Machines, and Clustering. It also includes coverage of Nearest Neighbors, Neural Networks, Boosting and AdaBoost, Random Forests, Linear Discriminant Analysis, Vector Quantization, Naive Bayes, Regression Analysis, Conditional Random Fields, and Data Analysis. Most of the examples in the book are drawn from the field of information security, with many of the machine learning applications specifically focused on malware. The applications presented are designed to demystify machine learning techniques by providing straightforward scenarios. Many of the exercises in this book require some programming, and basic computing concepts are assumed in a few of the application sections. However, anyone with a modest amount of programming experience should have no trouble with this aspect of the book. Instructor resources, including PowerPoint slides, lecture videos, and other relevant material are provided on an accompanying website: <http://www.cs.sjsu.edu/~stamp/ML/>. For the reader's benefit, the figures in the book are also available in electronic form, and in color. About the Author Mark Stamp has been a Professor of Computer Science at San Jose State University since 2002. Prior to that, he worked at the National Security Agency (NSA) for seven years, and a Silicon Valley startup company for two years. He received his Ph.D. from Texas Tech University in 1992. His love affair with machine learning began in the early 1990s, when he was working at the NSA, and continues today at SJSU, where he has supervised vast numbers of master's student projects, most of which involve a combination of information security and machine learning.

Machine Learning for Microbial Phenotype Prediction Roman Feldbauer 2016-06-15 This thesis presents a scalable, generic methodology for microbial phenotype prediction based on supervised machine learning, several models for biological and ecological traits of high relevance, and the deployment in metagenomic datasets. The results suggest that the presented prediction tool can be used to automatically annotate phenotypes in near-complete microbial genome sequences, as generated in large numbers in current metagenomic studies. Unraveling relationships between a living organism's genetic information and its observable traits is a central biological problem. Phenotype prediction facilitated by machine learning techniques will be a major step forward to creating biological knowledge from big data.

Bayesian Reasoning and Machine Learning David Barber 2012-02-02 A practical introduction perfect for final-year undergraduate and graduate students without a solid background in linear algebra and calculus.

Log-Linear Models David Knoke 1980-08 Introduces methods for quantitative assessment of relationships among categorical variables in multivariable crosstabulations. Procedures to estimate and interpret effect parameters for hierarchical models are described for both the general loglinear model and its logit version.