

Machine Learning Foundations A Case Study Approach

Yeah, reviewing a books **Machine Learning Foundations A Case Study Approach** could add your near links listings. This is just one of the solutions for you to be successful. As understood, success does not recommend that you have wonderful points.

Comprehending as well as covenant even more than other will manage to pay for each success. bordering to, the declaration as skillfully as sharpness of this Machine Learning Foundations A Case Study Approach can be taken as skillfully as picked to act.

Personalized Machine Learning - Julian McAuley

2022-01-31

Every day we interact with machine learning systems offering individualized predictions for our entertainment, social connections, purchases, or health. These involve several modalities of data, from sequences of clicks to text, images, and social interactions. This book introduces common principles and methods that underpin the design of personalized predictive models for a variety of settings and modalities. The book begins by revising 'traditional' machine learning models, focusing on adapting them to settings involving user data, then presents techniques based on advanced principles such as matrix factorization, deep learning, and generative modeling, and concludes with a detailed study of the consequences and risks of deploying personalized predictive systems. A series of case studies in domains ranging from e-commerce to health plus hands-on projects and code examples will give readers understanding and experience with large-scale real-world datasets and the ability to design models and systems for a wide range of applications.

Handbook of Research on Machine Learning - Monika Mangla

2022-08-04

This volume takes the reader on a technological voyage of machine learning advancements, highlighting the systematic changes in algorithms, challenges, and constraints. The technological advancements in the ML arena have transformed and revolutionized several fields, including transportation, agriculture, finance, weather monitoring, and others. This book brings together researchers, authors, industrialists, and academicians to cover a vast selection of topics in ML, starting with the rudiments of machine learning approaches and going on to specific applications in healthcare and industrial automation. The book begins with an overview of the ethics, security and privacy issues, future directions, and challenges in machine learning as well as a systematic review of deep learning techniques and provides an understanding of building generative adversarial networks. Chapters explore predictive data analytics for health issues. The book also adds a macro dimension by highlighting the industrial applications of machine learning, such as in the steel industry, for urban information retrieval, in garbage detection, in measuring air pollution, for stock market predictions, for underwater fish detection, as a fake news predictor, and more.

Distributed Optimization and Statistical Learning Via the Alternating Direction Method of Multipliers -

Stephen Boyd 2011

Surveys the theory and history of the alternating direction method of multipliers, and discusses its applications to a wide variety of statistical and machine learning problems of recent interest, including the lasso, sparse logistic regression, basis pursuit, covariance selection, support vector machines, and many others.

Foundations of Qualitative Research - Jerry W. Willis

2007-02-02

This text introduces key theoretical and epistemological concepts in an accessible style together with historical and current real-world examples employed to bring these otherwise difficult concepts to life.

Machine Learning - Peter Flach 2012-09-20

Covering all the main approaches in state-of-the-art machine learning research, this will set a new standard as an introductory textbook.

Deep Learning for Coders with fastai and PyTorch -

Jeremy Howard 2020-06-29

Deep learning is often viewed as the exclusive domain of math PhDs and big tech companies. But as this hands-on guide demonstrates, programmers comfortable with Python can achieve impressive results in deep learning with

little math background, small amounts of data, and minimal code. How? With fastai, the first library to provide a consistent interface to the most frequently used deep learning applications. Authors Jeremy Howard and Sylvain Gugger, the creators of fastai, show you how to train a model on a wide range of tasks using fastai and PyTorch. You'll also dive progressively further into deep learning theory to gain a complete understanding of the algorithms behind the scenes. Train models in computer vision, natural language processing, tabular data, and collaborative filtering Learn the latest deep learning techniques that matter most in practice Improve accuracy, speed, and reliability by understanding how deep learning models work Discover how to turn your models into web applications Implement deep learning algorithms from scratch Consider the ethical implications of your work Gain insight from the foreword by PyTorch cofounder, Soumith Chintala

Machine Learning for Hackers - Drew Conway 2012-02-13

If you're an experienced programmer interested in crunching data, this book will get you started with machine learning—a toolkit of algorithms that enables computers to train themselves to automate useful tasks. Authors Drew Conway and John Myles White help you understand machine learning and statistics tools through a series of hands-on case studies, instead of a traditional math-heavy presentation. Each chapter focuses on a specific problem in machine learning, such as classification, prediction, optimization, and recommendation. Using the R programming language, you'll learn how to analyze sample datasets and write simple machine learning algorithms. Machine Learning for Hackers is ideal for programmers from any background, including business, government, and academic research. Develop a naïve Bayesian classifier to determine if an email is spam, based only on its text Use linear regression to predict the number of page views for the top 1,000 websites Learn optimization techniques by attempting to break a simple letter cipher Compare and contrast U.S. Senators statistically, based on their voting records Build a "whom to follow" recommendation system from Twitter data

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.

Learning Analytics: Fundamentals, Applications, and Trends - Alejandro Peña-Ayala 2017-02-17

This book provides a conceptual and empirical perspective on learning analytics, its goal being to disseminate the core concepts, research, and outcomes of this emergent field. Divided into nine chapters, it offers reviews oriented on selected topics, recent advances, and innovative applications. It presents the broad learning analytics landscape and in-depth studies on higher education, adaptive assessment, teaching and learning. In addition, it discusses valuable approaches to coping with personalization and huge data, as well as conceptual topics and specialized applications that have shaped the current state of the art. By identifying fundamentals, highlighting applications, and pointing out current trends, the book offers an essential overview of learning analytics to enhance learning achievement in diverse educational settings. As such, it represents a valuable resource for researchers, practitioners, and students interested in updating their knowledge and finding inspirations for their future work.

Fundamentals of Machine Learning for Predictive Data Analytics, second edition - John D. Kelleher 2020-10-20

The second edition of a comprehensive introduction to machine learning approaches used in predictive data

analytics, covering both theory and practice. Machine learning is often used to build predictive models by extracting patterns from large datasets. These models are used in predictive data analytics applications including price prediction, risk assessment, predicting customer behavior, and document classification. This introductory textbook offers a detailed and focused treatment of the most important machine learning approaches used in predictive data analytics, covering both theoretical concepts and practical applications. Technical and mathematical material is augmented with explanatory worked examples, and case studies illustrate the application of these models in the broader business context. This second edition covers recent developments in machine learning, especially in a new chapter on deep learning, and two new chapters that go beyond predictive analytics to cover unsupervised learning and reinforcement learning.

Deep Learning on Graphs - Yao Ma 2021-09-23

A comprehensive text on foundations and techniques of graph neural networks with applications in NLP, data mining, vision and healthcare.

Deep Learning Illustrated - Jon Krohn 2019-08-05

"The authors' clear visual style provides a comprehensive look at what's currently possible with artificial neural networks as well as a glimpse of the magic that's to come." -Tim Urban, author of *Wait But Why Fully Practical, Insightful Guide to Modern Deep Learning* Deep learning is transforming software, facilitating powerful new artificial intelligence capabilities, and driving unprecedented algorithm performance. *Deep Learning Illustrated* is uniquely intuitive and offers a complete introduction to the discipline's techniques. Packed with full-color figures and easy-to-follow code, it sweeps away the complexity of building deep learning models, making the subject approachable and fun to learn. World-class instructor and practitioner Jon Krohn-with visionary content from Grant Beyleveld and beautiful illustrations by Aglaé Bassens-presents straightforward analogies to explain what deep learning is, why it has become so popular, and how it relates to other machine learning approaches. Krohn has created a practical reference and tutorial for developers, data scientists, researchers, analysts, and students who want to start applying it. He illuminates theory with hands-on Python code in accompanying Jupyter notebooks. To help you progress quickly, he focuses on the versatile deep learning library Keras to nimbly construct efficient TensorFlow models; PyTorch, the leading alternative library, is also covered. You'll gain a pragmatic understanding of all major deep learning approaches and their uses in applications ranging from machine vision and natural language processing to image generation and game-playing algorithms. Discover what makes deep learning systems unique, and the implications for practitioners Explore new tools that make deep learning models easier to build, use, and improve Master essential theory: artificial neurons, training, optimization, convolutional nets, recurrent nets, generative adversarial networks (GANs), deep reinforcement learning, and more Walk through building interactive deep learning applications, and move forward with your own artificial intelligence projects Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book for details.

Machine Learning Techniques for Text - Nikos Tsourakis 2022-10-31

Take your Python text processing skills to another level by learning about the latest natural language processing and machine learning techniques with this full color guide Key Features Learn how to acquire and process textual data and visualize the key findings Obtain deeper insight into the most commonly used algorithms and techniques and understand their tradeoffs Implement models for solving real-world problems and evaluate their performance Book Description With the ever-increasing demand for machine learning and programming professionals, it's prime time to invest in the field. This book will help you in this endeavor, focusing specifically on text data and human language by steering a middle path among the various textbooks that present complicated theoretical concepts or focus disproportionately on Python code. A good metaphor this work builds upon is the relationship between an experienced craftsman and their trainee. Based on the

current problem, the former picks a tool from the toolbox, explains its utility, and puts it into action. This approach will help you to identify at least one practical use for each method or technique presented. The content unfolds in ten chapters, each discussing one specific case study. For this reason, the book is solution-oriented. It's accompanied by Python code in the form of Jupyter notebooks to help you obtain hands-on experience. A recurring pattern in the chapters of this book is helping you get some intuition on the data and then implement and contrast various solutions. By the end of this book, you'll be able to understand and apply various techniques with Python for text preprocessing, text representation, dimensionality reduction, machine learning, language modeling, visualization, and evaluation. What you will learn Understand fundamental concepts of machine learning for text Discover how text data can be represented and build language models Perform exploratory data analysis on text corpora Use text preprocessing techniques and understand their trade-offs Apply dimensionality reduction for visualization and classification Incorporate and fine-tune algorithms and models for machine learning Evaluate the performance of the implemented systems Know the tools for retrieving text data and visualizing the machine learning workflow Who this book is for This book is for professionals in the area of computer science, programming, data science, informatics, business analytics, statistics, language technology, and more who aim for a gentle career shift in machine learning for text. Students in relevant disciplines that seek a textbook in the field will benefit from the practical aspects of the content and how the theory is presented. Finally, professors teaching a similar course will be able to pick pertinent topics in terms of content and difficulty. Beginner-level knowledge of Python programming is needed to get started with this book.

Machine Learning Refined - Jeremy Watt 2016-09-08

Providing a unique approach to machine learning, this text contains fresh and intuitive, yet rigorous, descriptions of all fundamental concepts necessary to conduct research, build products, tinker, and play. By prioritizing geometric intuition, algorithmic thinking, and practical real world applications in disciplines including computer vision, natural language processing, economics, neuroscience, recommender systems, physics, and biology, this text provides readers with both a lucid understanding of foundational material as well as the practical tools needed to solve real-world problems. With in-depth Python and MATLAB/OCTAVE-based computational exercises and a complete treatment of cutting edge numerical optimization techniques, this is an essential resource for students and an ideal reference for researchers and practitioners working in machine learning, computer science, electrical engineering, signal processing, and numerical optimization.

Machine Learning for Email - Drew Conway 2011-10-25

If you're an experienced programmer willing to crunch data, this concise guide will show you how to use machine learning to work with email. You'll learn how to write algorithms that automatically sort and redirect email based on statistical patterns. Authors Drew Conway and John Myles White approach the process in a practical fashion, using a case-study driven approach rather than a traditional math-heavy presentation. This book also includes a short tutorial on using the popular R language to manipulate and analyze data. You'll get clear examples for analyzing sample data and writing machine learning programs with R. Mine email content with R functions, using a collection of sample files Analyze the data and use the results to write a Bayesian spam classifier Rank email by importance, using factors such as thread activity Use your email ranking analysis to write a priority inbox program Test your classifier and priority inbox with a separate email sample set

Machine Learning and Big Data - Uma N. Dulhare 2020-09-01

This book is intended for academic and industrial developers, exploring and developing applications in the area of big data and machine learning, including those that are solving technology requirements, evaluation of methodology advances and algorithm demonstrations. The intent of this book is to provide awareness of algorithms used for machine learning and big data in the academic and professional community. The 17 chapters are

divided into 5 sections: Theoretical Fundamentals; Big Data and Pattern Recognition; Machine Learning: Algorithms & Applications; Machine Learning's Next Frontier and Hands-On and Case Study. While it dwells on the foundations of machine learning and big data as a part of analytics, it also focuses on contemporary topics for research and development. In this regard, the book covers machine learning algorithms and their modern applications in developing automated systems. Subjects covered in detail include: Mathematical foundations of machine learning with various examples. An empirical study of supervised learning algorithms like Naïve Bayes, KNN and semi-supervised learning algorithms viz. S3VM, Graph-Based, Multiview. Precise study on unsupervised learning algorithms like GMM, K-mean clustering, Dritchlet process mixture model, X-means and Reinforcement learning algorithm with Q learning, R learning, TD learning, SARSA Learning, and so forth. Hands-on machine learning open source tools viz. Apache Mahout, H2O. Case studies for readers to analyze the prescribed cases and present their solutions or interpretations with intrusion detection in MANETS using machine learning. Showcase on novel user-cases: Implications of Electronic Governance as well as Pragmatic Study of BD/ML technologies for agriculture, healthcare, social media, industry, banking, insurance and so on.

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.

Reinforcement Learning, second edition - Richard S. Sutton
2018-11-13

The significantly expanded and updated new edition of a widely used text on reinforcement learning, one of the most active research areas in artificial intelligence. Reinforcement learning, one of the most active research areas in artificial intelligence, is a computational approach to learning whereby an agent tries to maximize the total amount of reward it receives while interacting with a complex, uncertain environment. In Reinforcement Learning, Richard Sutton and Andrew Barto provide a clear and simple account of the field's key ideas and algorithms. This second edition has been significantly expanded and updated, presenting new topics and updating coverage of other topics. Like the first edition, this second edition focuses on core online learning algorithms, with the more mathematical material set off in shaded boxes. Part I covers as much of reinforcement learning as possible without going beyond the tabular case for which exact solutions can be found. Many algorithms presented in this part are new to the second edition, including UCB, Expected Sarsa, and Double Learning. Part II extends these ideas to function approximation, with new sections on such topics as artificial neural networks and the Fourier basis, and offers expanded treatment of off-policy learning and policy-gradient methods. Part III has new chapters on reinforcement learning's relationships to psychology and neuroscience, as well as an updated case-studies chapter including AlphaGo and AlphaGo Zero, Atari game playing, and IBM Watson's wagering strategy. The final chapter discusses the future societal impacts of reinforcement learning.

Encyclopedia of Machine Learning - Claude Sammut
2011-03-28

This comprehensive encyclopedia, in A-Z format, provides

easy access to relevant information for those seeking entry into any aspect within the broad field of Machine Learning. Most of the entries in this preeminent work include useful literature references.
California Preschool Learning Foundations: Visual and performing arts. Physical development. Health - Faye Ong
2008

Fundamentals of Programming in SAS - James Blum
2019-07-27

Unlock the essentials of SAS programming! Fundamentals of Programming in SAS: A Case Studies Approach gives a complete introduction to SAS programming. Perfect for students, novice SAS users, and programmers studying for their Base SAS certification, this book covers all the basics, including: working with data creating visualizations data validation good programming practices Experienced programmers know that real-world scenarios require practical solutions. Designed for use in the classroom and for self-guided learners, this book takes a novel approach to learning SAS programming by following a single case study throughout the text and circling back to previous concepts to reinforce material. Readers will benefit from the variety of exercises, including both multiple choice questions and in-depth case studies. Additional case studies are also provided online for extra practice. This approach mirrors the way good SAS programmers develop their skills—through hands-on work with an eye toward developing the knowledge necessary to tackle more difficult tasks. After reading this book, you will gain the skills and confidence to take on larger challenges with the power of SAS.

Deep Learning - Li Deng 2014

Provides an overview of general deep learning methodology and its applications to a variety of signal and information processing tasks

Smart Monitoring of Rotating Machinery for Industry 4.0 - Fakhre Chaari 2021-08-20

This book offers an overview of current methods for the intelligent monitoring of rotating machines. It describes the foundations of smart monitoring, guiding readers to develop appropriate machine learning and statistical models for answering important challenges, such as the management and analysis of a large volume of data. It also discusses real-world case studies, highlighting some practical issues and proposing solutions to them. The book offers extensive information on research trends, and innovative strategies to solve emerging, practical issues. It addresses both academics and professionals dealing with condition monitoring, and mechanical and production engineering issues, in the era of industry 4.0.

Machine Learning and Knowledge Extraction - Andreas Holzinger 2019-08-22

This book constitutes the refereed proceedings of the IFIP TC 5, TC 12, WG 8.4, 8.9, 12.9 International Cross-Domain Conference for Machine Learning and Knowledge Extraction, CD-MAKE 2019, held in Canterbury, UK, in August 2019. The 25 revised full papers presented were carefully reviewed and selected from 45 submissions. The cross-domain integration and appraisal of different fields provides an atmosphere to foster different perspectives and opinions; it will offer a platform for novel ideas and a fresh look on the methodologies to put these ideas into business for the benefit of humanity.

Innovations in Multi-Agent Systems and Application - 1 - Dipti Srinivasan 2010-07-17

In today's world, the increasing requirement for emulating the behavior of real-world applications for achieving effective management and control has necessitated the usage of advanced computational techniques. Computational intelligence-based techniques that combine a variety of problem solvers are becoming increasingly pervasive. The ability of these methods to adapt to the dynamically changing environment and learn in an online manner has increased their usefulness in simulating intelligent behaviors as observed in humans. These intelligent systems are able to handle the stochastic and uncertain nature of the real-world problems. Application domains requiring interaction of people or organizations with different, even possibly conflicting goals and proprietary information handling are growing exponentially. To efficiently handle these types of complex interactions, distributed problem solving systems like multiagent systems have become a necessity. The rapid advancements in network

communication technologies have provided the platform for successful implementation of such intelligent agent-based problem solvers. An agent can be viewed as a self-contained, concurrently executing thread of control that encapsulates some state and communicates with its environment, and possibly other agents via message passing. Agent-based systems offer advantages when independently developed components must interoperate in a heterogeneous environment. Such agent-based systems are increasingly being applied in a wide range of areas including telecommunications, Business process modeling, computer games, distributed system control and robot systems.

Machine Learning - Marco Gori 2017-11-13

Machine Learning: A Constraint-Based Approach provides readers with a refreshing look at the basic models and algorithms of machine learning, with an emphasis on current topics of interest that includes neural networks and kernel machines. The book presents the information in a truly unified manner that is based on the notion of learning from environmental constraints. While regarding symbolic knowledge bases as a collection of constraints, the book draws a path towards a deep integration with machine learning that relies on the idea of adopting multivalued logic formalisms, like in fuzzy systems. A special attention is reserved to deep learning, which nicely fits the constrained-based approach followed in this book. This book presents a simpler unified notion of regularization, which is strictly connected with the parsimony principle, and includes many solved exercises that are classified according to the Donald Knuth ranking of difficulty, which essentially consists of a mix of warm-up exercises that lead to deeper research problems. A software simulator is also included.

Clustering and Information Retrieval - Weili Wu 2013-12-01

Clustering is an important technique for discovering relatively dense sub-regions or sub-spaces of a multi-dimension data distribution. Clustering has been used in information retrieval for many different purposes, such as query expansion, document grouping, document indexing, and visualization of search results. In this book, we address issues of clustering algorithms, evaluation methodologies, applications, and architectures for information retrieval. The first two chapters discuss clustering algorithms. The chapter from Baeza-Yates et al. describes a clustering method for a general metric space which is a common model of data relevant to information retrieval. The chapter by Guha, Rastogi, and Shim presents a survey as well as detailed discussion of two clustering algorithms: CURE and ROCK for numeric data and categorical data respectively. Evaluation methodologies are addressed in the next two chapters. Ertoz et al. demonstrate the use of text retrieval benchmarks, such as TRECS, to evaluate clustering algorithms. He et al. provide objective measures of clustering quality in their chapter. Applications of clustering methods to information retrieval is addressed in the next four chapters. Chu et al. and Noel et al. explore feature selection using word stems, phrases, and link associations for document clustering and indexing. Wen et al. and Sung et al. discuss applications of clustering to user queries and data cleansing. Finally, we consider the problem of designing architectures for information retrieval. Crichton, Hughes, and Kelly elaborate on the development of a scientific data system architecture for information retrieval.

Machine Learning for Factor Investing: R Version - Guillaume Coqueret 2020-08-31

Machine learning (ML) is progressively reshaping the fields of quantitative finance and algorithmic trading. ML tools are increasingly adopted by hedge funds and asset managers, notably for alpha signal generation and stocks selection. The technicality of the subject can make it hard for non-specialists to join the bandwagon, as the jargon and coding requirements may seem out of reach. *Machine Learning for Factor Investing: R Version* bridges this gap. It provides a comprehensive tour of modern ML-based investment strategies that rely on firm characteristics. The book covers a wide array of subjects which range from economic rationales to rigorous portfolio back-testing and encompass both data processing and model interpretability. Common supervised learning algorithms such as tree models and neural networks are explained in the context of style investing and the reader can also dig into more complex techniques

like autoencoder asset returns, Bayesian additive trees, and causal models. All topics are illustrated with self-contained R code samples and snippets that are applied to a large public dataset that contains over 90 predictors. The material, along with the content of the book, is available online so that readers can reproduce and enhance the examples at their convenience. If you have even a basic knowledge of quantitative finance, this combination of theoretical concepts and practical illustrations will help you learn quickly and deepen your financial and technical expertise.

Machine Learning - Ryszard S. Michalski 1994-02-09

Multistrategy learning is one of the newest and most promising research directions in the development of machine learning systems. The objectives of research in this area are to study trade-offs between different learning strategies and to develop learning systems that employ multiple types of inference or computational paradigms in a learning process. Multistrategy systems offer significant advantages over monostrategy systems. They are more flexible in the type of input they can learn from and the type of knowledge they can acquire. As a consequence, multistrategy systems have the potential to be applicable to a wide range of practical problems. This volume is the first book in this fast growing field. It contains a selection of contributions by leading researchers specializing in this area. See below for earlier volumes in the series.

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

Foundations of Computational Intelligence - Aboul-Ella Hassanien 2009-05-02

Foundations of Computational Intelligence Volume 1: Learning and Approximation: Theoretical Foundations and Applications Learning methods and approximation algorithms are fundamental tools that deal with computationally hard problems and problems in which the input is gradually disclosed over time. Both kinds of problems have a large number of applications arising from a variety of fields, such as algorithmic game theory, approximation classes, coloring and partitioning, competitive analysis, computational finance, cuts and connectivity, inapproximability results, mechanism design, network design, packing and covering, paradigms for design and analysis of approximation and online algorithms, randomization techniques, real-world applications, scheduling problems and so on. The past years have witnessed a large number of interesting applications using various techniques of Computational Intelligence such as rough sets, connectionist learning; fuzzy logic; evolutionary computing; artificial immune systems; swarm intelligence; reinforcement learning, intelligent multimedia processing etc. . In spite of numerous successful applications of Computational Intelligence in business and industry, it is sometimes difficult to explain the performance of these techniques and algorithms from a theoretical perspective. Therefore, we encouraged authors to present original ideas dealing with the incorporation of different mechanisms of Computational Intelligent dealing with Learning and Approximation algorithms and underlying processes. This edited volume comprises 15 chapters, including an overview chapter, which provides an up-to-date and state-of-the art research on the application of Computational Intelligence for learning and approximation.

Foundations of Intelligent Systems - Troels Andreasen 2014-06-03

This book constitutes the refereed proceedings of the 21st International Symposium on Methodologies for Intelligent Systems, ISMIS 2014, held in Roskilde, Denmark, in June 2014. The 61 revised full papers were carefully reviewed and selected from 111 submissions. The papers are organized in topical sections on complex networks and data stream mining; data mining methods; intelligent systems applications; knowledge representation in databases and systems; textual data analysis and mining; special session: challenges in text mining and semantic information retrieval; special session: warehousing and OLAPing complex, spatial and spatio-temporal data; ISMIS posters.

Statistical Process Monitoring Using Advanced Data-Driven and Deep Learning Approaches - Fouzi Harrou 2020-07-03

Statistical Process Monitoring Using Advanced Data-Driven and Deep Learning Approaches tackles multivariate challenges in process monitoring by merging the advantages of univariate and traditional multivariate techniques to enhance their performance and widen their practical applicability. The book proceeds with merging the desirable properties of shallow learning approaches - such as a one-class support vector machine and k-nearest neighbours and unsupervised deep learning approaches - to develop more sophisticated and efficient monitoring techniques. Finally, the developed approaches are applied to monitor many processes, such as wastewater treatment plants, detection of obstacles in driving environments for autonomous robots and vehicles, robot swarm, chemical processes (continuous stirred tank reactor, plug flow reactor, and distillation columns), ozone pollution, road traffic congestion, and solar photovoltaic systems. Uses a data-driven based approach to fault detection and attribution Provides an in-depth understanding of fault detection and attribution in complex and multivariate systems Familiarises you with the most suitable data-driven based techniques including multivariate statistical techniques and deep learning-

based methods Includes case studies and comparison of different methods

Demystifying Big Data, Machine Learning, and Deep Learning for Healthcare Analytics - Pradeep N 2021-06-25

Demystifying Big Data, Machine Learning, and Deep Learning for Healthcare Analytics presents the changing world of data utilization, especially in clinical healthcare. Various techniques, methodologies, and algorithms are presented in this book to organize data in a structured manner that will assist physicians in the care of patients and help biomedical engineers and computer scientists understand the impact of these techniques on healthcare analytics. The book is divided into two parts: Part 1 covers big data aspects such as healthcare decision support systems and analytics-related topics. Part 2 focuses on the current frameworks and applications of deep learning and machine learning, and provides an outlook on future directions of research and development. The entire book takes a case study approach, providing a wealth of real-world case studies in the application chapters to act as a foundational reference for biomedical engineers, computer scientists, healthcare researchers, and clinicians. Provides a comprehensive reference for biomedical engineers, computer scientists, advanced industry practitioners, researchers, and clinicians to understand and develop healthcare analytics using advanced tools and technologies Includes in-depth illustrations of advanced techniques via dataset samples, statistical tables, and graphs with algorithms and computational methods for developing new applications in healthcare informatics Unique case study approach provides readers with insights for practical clinical implementation

Foundations of Machine Learning, second edition - Mehryar Mohri 2018-12-25

A new edition of a graduate-level machine learning textbook that focuses on the analysis and theory of algorithms. This book is a general introduction to machine learning that can serve as a textbook for graduate students and a reference for researchers. It covers fundamental modern topics in machine learning while providing the theoretical basis and conceptual tools needed for the discussion and justification of algorithms. It also describes several key aspects of the application of these algorithms. The authors aim to present novel theoretical tools and concepts while giving concise proofs even for relatively advanced topics. Foundations of Machine Learning is unique in its focus on the analysis and theory of algorithms. The first four chapters lay the theoretical foundation for what follows; subsequent chapters are mostly self-contained. Topics covered include the Probably Approximately Correct (PAC) learning framework; generalization bounds based on Rademacher complexity and VC-dimension; Support Vector Machines (SVMs); kernel methods; boosting; on-line learning; multi-class classification; ranking; regression; algorithmic stability; dimensionality reduction; learning automata and languages; and reinforcement learning. Each chapter ends with a set of exercises. Appendixes provide additional material including concise probability review. This second edition offers three new chapters, on model selection, maximum entropy models, and conditional entropy models. New material in the appendixes includes a major section on Fenchel duality, expanded coverage of concentration inequalities, and an entirely new entry on information theory. More than half of the exercises are new to this edition.

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.

Introducing Data Science - Davy Cielen 2016-05-02

Summary Introducing Data Science teaches you how to accomplish the fundamental tasks that occupy data scientists. Using the Python language and common Python libraries, you'll experience firsthand the challenges of dealing with data at scale and gain a solid foundation in data science. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Many companies need developers with data science skills to work on projects ranging from social media marketing to machine learning. Discovering what you need to learn to begin a career as a data scientist can seem bewildering. This book is designed to help you get started. About the Book Introducing Data Science explains vital data science concepts and teaches you how to accomplish the fundamental tasks that occupy data scientists. You'll explore data visualization, graph databases, the use of NoSQL, and the data science process. You'll use the Python language and common Python libraries as you experience firsthand the challenges of dealing with data at scale. Discover how Python allows you to gain insights from data sets so big that they need to be stored on multiple machines, or from data moving so quickly that no single machine can handle it. This book gives you hands-on experience with the most popular Python data science libraries, Scikit-learn and StatsModels. After reading this book, you'll have the solid foundation you need to start a career in data science. What's Inside Handling large data Introduction to machine learning Using Python to work with data Writing data science algorithms About the Reader This book assumes you're comfortable reading code in Python or a similar language, such as C, Ruby, or JavaScript. No prior experience with data science is required. About the Authors Davy Cielen, Arno D. B. Meysman, and Mohamed Ali are the founders and managing partners of Optimately and Maiton, where they focus on developing data science projects and solutions in various sectors. Table of Contents Data science in a big data world The data science process Machine learning Handling large data on a single computer First steps in big data Join the NoSQL movement The rise of graph databases Text mining and text analytics Data visualization to the end user *Model-Based Machine Learning* - Taylor & Francis Group 2018-12-07

Foundations of Intelligent Systems - Li Chen 2013-12-06

This book constitutes the proceedings of the 20th International Symposium on Methodologies for Intelligent Systems, ISMIS 2012, held in Macau, China, in December 2012. The 42 regular papers and 11 short papers presented were carefully reviewed and selected from 88 submissions. They are organized in topical sections named: knowledge discovery and data mining; intelligent information systems; text mining and language processing; knowledge representation and integration; music information retrieval; recommender systems; technology intelligence and applications; product

configuration; human factors in information retrieval; social recommender systems; and warehousing and OLAPing complex, spatial and spatio-temporal data.

Machine Learning with TensorFlow, Second Edition - Mattmann A. Chris 2021-02-02

Updated with new code, new projects, and new chapters, Machine Learning with TensorFlow, Second Edition gives readers a solid foundation in machine-learning concepts and the TensorFlow library. Summary Updated with new code, new projects, and new chapters, Machine Learning with TensorFlow, Second Edition gives readers a solid foundation in machine-learning concepts and the TensorFlow library. Written by NASA JPL Deputy CTO and Principal Data Scientist Chris Mattmann, all examples are accompanied by downloadable Jupyter Notebooks for a hands-on experience coding TensorFlow with Python. New and revised content expands coverage of core machine learning algorithms, and advancements in neural networks such as VGG-Face facial identification classifiers and deep speech classifiers. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Supercharge your data analysis with machine learning! ML algorithms automatically improve as they process data, so results get better over time. You don't have to be a mathematician to use ML: Tools like Google's TensorFlow library help with complex calculations so you can focus on getting the answers you need. About the book Machine Learning with TensorFlow, Second Edition is a fully revised guide to building machine learning models using Python and TensorFlow. You'll apply core ML concepts to real-world challenges, such as sentiment analysis, text classification, and image recognition. Hands-on examples illustrate neural network techniques for deep speech processing, facial identification, and auto-encoding with CIFAR-10. What's inside Machine Learning with TensorFlow Choosing the best ML approaches Visualizing algorithms with TensorBoard Sharing results with collaborators Running models in Docker About the reader Requires intermediate Python skills and knowledge of general algebraic concepts like vectors and matrices. Examples use the super-stable 1.15.x branch of TensorFlow and TensorFlow 2.x. About the author Chris Mattmann is the Division Manager of the Artificial Intelligence, Analytics, and Innovation Organization at NASA Jet Propulsion Lab. The first edition of this book was written by Nishant Shukla with Kenneth Fricklas. Table of Contents PART 1 - YOUR MACHINE-LEARNING RIG 1 A machine-learning odyssey 2 TensorFlow essentials PART 2 - CORE LEARNING ALGORITHMS 3 Linear regression and beyond 4 Using regression for call-center volume prediction 5 A gentle introduction to classification 6 Sentiment classification: Large movie-review dataset 7 Automatically clustering data 8 Inferring user activity from Android accelerometer data 9 Hidden Markov models 10 Part-of-speech tagging and word-sense disambiguation PART 3 - THE NEURAL NETWORK PARADIGM 11 A peek into autoencoders 12 Applying autoencoders: The CIFAR-10 image dataset 13 Reinforcement learning 14 Convolutional neural networks 15 Building a real-world CNN: VGG-Face ad VGG-Face Lite 16 Recurrent neural networks 17 LSTMs and automatic speech recognition 18 Sequence-to-sequence models for chatbots 19 Utility landscape