
Acces PDF Pdf Yaser Data From Learning

This is likewise one of the factors by obtaining the soft documents of this **Pdf Yaser Data From Learning** by online. You might not require more grow old to spend to go to the books instigation as competently as search for them. In some cases, you likewise reach not discover the notice Pdf Yaser Data From Learning that you are looking for. It will totally squander the time.

However below, past you visit this web page, it will be thus completely simple to get as with ease as download guide Pdf Yaser Data From Learning

It will not undertake many get older as we run by before. You can pull off it even though operate something else at house and even in your workplace. in view of that easy! So, are you question? Just exercise just what we give below as with ease as review **Pdf Yaser Data From Learning** what you with to read!

KEY=FROM - HUGHES COSTA

Learning from Data

A Short Course

Information Theory, Inference and Learning Algorithms

Cambridge University Press Table of contents

Portfolio Management under Stress

A Bayesian-Net Approach to Coherent Asset Allocation

Cambridge University Press A rigorous presentation of a novel methodology for asset allocation in financial portfolios under conditions of market distress.

Machine Learning Refined

Foundations, Algorithms, and Applications

Cambridge University Press 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.

Fundamentals of Clinical Data Science

Springer This open access book comprehensively covers the fundamentals of clinical data science, focusing on data collection, modelling and clinical applications. Topics covered in the first section on data collection include: data sources, data at scale (big data), data stewardship (FAIR data) and related privacy concerns. Aspects of predictive modelling using techniques such as classification, regression or clustering, and prediction model validation will be covered in the second section. The third section covers aspects of (mobile) clinical decision support systems, operational excellence and value-based healthcare. Fundamentals of Clinical Data Science is an essential resource for healthcare professionals and IT consultants intending to develop and refine their skills in personalized medicine, using solutions based on large datasets from electronic health records or telemonitoring programmes. The book's promise is "no math, no code" and will explain the topics in a style that is optimized for a healthcare audience.

Learning with Adults

A Reader

Springer Science & Business Media This anthology brings together some of the finest writers on different aspects of adult education and related areas to provide a complementary reader to the introductory text by Leona English and Peter Mayo Learning with Adults: A Critical Introduction. Areas tackled include Disability, Prisons, Third Age Universities, Lifelong Learning Policy, Learning Society, Poverty, LGBTQ, Sport, Women, Literacy, Transformative Learning, Community Arts, Aesthetics, Consumption, Migration, Libraries, Folk High Schools, Adult Education Policy, Subaltern Southern Social Movements, Social Creation, Community Radio, Social Film. Contexts focused on include Africa, Caribbean, Europe, Latin America, Asia (India), small island states. Over thirty authors involved including Zygmunt Bauman, Rosa Maria Torres, Oskar Negt, Antonia Darder, Jim Elmborg, D. W. Livingstone, Palle Rasmussen, Mae Shaw, Leona English, Asoke Bhattacharya, Cynthia L. Pemberton, Eileen Casey White, Daniel Schugurensky, Dip Kapoor, Peter Rule, John Myers, Joseph Giordmaina, Antonia De Vita, Alexis Kokkos, Marvin Formosa, Carmel Borg, Julia Preece, Patricia Cranton, Lyn Tett, Ali A. Abdi, Anna Maria Piussi, Behrang Foroughi, Taadi Ruth Modipa, Robert Hill, Edward Shiza, Kaela Jubas and Didacus Jules. ... Learning with Adults: A Reader constitutes the most valuable practical and theoretical reflection on adult education I have seen in a long time. Nelly P. Stromquist, Professor, International Education Policy, College of Education University of Maryland, College Park ... This book provides an opportunity at a very appropriate moment to discuss adult education issues during challenging times. Paula Guimarães, University of Lisbon ... Read and savour delights and surprises. Michael Welton, UBC and Athabasca University This book satisfies everything one could desire of a reader on the subject. Kenneth Wain, University of Malta

Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow Concepts, Tools, and Techniques to Build Intelligent Systems

O'Reilly Media Through a series of recent breakthroughs, deep learning has boosted the entire field of machine learning. Now, even programmers who know close to nothing about this technology can use simple, efficient tools to implement programs capable of learning from data. This practical book shows you how. By using concrete examples, minimal theory, and two production-ready Python frameworks—Scikit-Learn and TensorFlow—author Aurélien Géron helps you gain an intuitive understanding of the concepts and tools for building intelligent systems. You'll learn a range of techniques, starting with simple linear regression and progressing to deep neural networks. With exercises in each chapter to help you apply what you've learned, all you need is programming experience to get started. Explore the machine learning landscape, particularly neural nets Use Scikit-Learn to track an example machine-learning project end-to-end Explore several training models, including support vector machines, decision trees, random forests, and ensemble methods Use the TensorFlow library to build and train neural nets Dive into neural net architectures, including convolutional nets, recurrent nets, and deep reinforcement learning Learn techniques for training and scaling deep neural nets

Boosting

Foundations and Algorithms

MIT Press An accessible introduction and essential reference for an approach to machine learning that creates highly accurate prediction rules by combining many weak and inaccurate ones. Boosting is an approach to machine learning based on the idea of creating a highly accurate predictor by combining many weak and inaccurate "rules of thumb." A remarkably rich theory has evolved around boosting, with connections to a range of topics, including statistics, game theory, convex optimization, and information geometry. Boosting algorithms have also enjoyed practical success in such fields as biology, vision, and speech processing. At various times in its history, boosting has been perceived as mysterious, controversial, even paradoxical. This book, written by the inventors of the method, brings together, organizes, simplifies, and substantially extends two decades of research on boosting, presenting both theory and applications in a way that is accessible to readers from diverse backgrounds while also providing an authoritative reference for advanced researchers. With its introductory treatment of all material and its inclusion of exercises in every chapter, the book is appropriate for course use as well. The book begins with a general introduction to machine learning algorithms and their analysis; then explores the core theory of boosting, especially its ability to generalize; examines some of the myriad other theoretical viewpoints that help to explain and understand boosting; provides practical extensions of boosting for more complex learning problems; and finally presents a number of advanced theoretical topics. Numerous applications and practical illustrations are offered throughout.

Federated Learning

Privacy and Incentive

Springer Nature This book provides a comprehensive and self-contained introduction to federated learning, ranging from the basic knowledge and theories to various key applications. Privacy and incentive issues are the focus of this book. It is timely as federated learning is becoming popular after the release of the General Data Protection Regulation (GDPR). Since federated learning aims to enable a machine model to be collaboratively trained without each party exposing private data to others. This setting adheres to regulatory requirements of data privacy protection such as GDPR. This book contains three main parts. Firstly, it introduces different privacy-preserving methods for protecting a federated learning model against different types of attacks such as data leakage and/or data poisoning. Secondly, the book presents incentive mechanisms which aim to encourage individuals to participate in the federated learning ecosystems. Last but not least, this book also describes how federated learning can be applied in industry and business to address data silo and privacy-preserving problems. The book is intended for readers from both the academia and the industry, who would like to learn about federated learning, practice its implementation, and apply it in their own business. Readers are expected to have some basic understanding of linear algebra, calculus, and neural network. Additionally, domain knowledge in FinTech and marketing would be helpful."

AI and Financial Markets

MDPI Artificial intelligence (AI) is regarded as the science and technology for producing an intelligent machine, particularly, an intelligent computer program. Machine learning is an approach to realizing AI comprising a collection of statistical algorithms, of which deep learning is one such example. Due to the rapid development of computer technology, AI has been actively explored for a variety of academic and practical purposes in the context of financial markets. This book focuses on the broad topic of "AI and Financial Markets", and includes novel research associated with this topic. The book includes contributions on the application of machine learning, agent-based artificial market simulation, and other related skills to the analysis of various aspects of financial markets.

Deep Learning in Computational Mechanics

An Introductory Course

Springer Nature This book provides a first course on deep learning in computational mechanics. The book starts with a short introduction to machine learning's fundamental concepts before neural networks are explained thoroughly. It then provides an overview of current topics in physics and engineering, setting the stage for the book's main topics: physics-informed neural networks and the deep energy method. The idea of the book is to provide the basic concepts in a mathematically sound manner and yet to stay as simple as possible. To achieve this goal, mostly one-dimensional examples are investigated, such as approximating functions by neural networks or the simulation of the temperature's evolution in a one-dimensional bar. Each chapter contains examples and exercises which are either solved analytically or in PyTorch, an open-source machine learning framework for python.

Interoperating Geographic Information Systems

Springer Science & Business Media Geographic information systems have developed rapidly in the past decade, and are now a major class of software, with applications that include infrastructure maintenance, resource management, agriculture, Earth science, and planning. But a lack of standards has led to a general inability for one GIS to interoperate with another. It is difficult for one GIS to share data with another, or for people trained on one system to adapt easily to the commands and user interface of another. Failure to interoperate is a problem at many levels, ranging from the purely technical to the semantic and the institutional. Interoperating Geographic Information Systems is about efforts to improve the ability of GISs to interoperate, and has been assembled through a collaboration between academic researchers and the software vendor community under the auspices of the US National Center for Geographic Information and Analysis and the Open GIS Consortium Inc. It includes chapters on the basic principles and the various conceptual frameworks that the research community has developed to think about the problem. Other chapters review a wide range of applications and the experiences of the authors in trying to achieve interoperability at a practical level. Interoperability opens

enormous potential for new ways of using GIS and new mechanisms for exchanging data, and these are covered in chapters on information marketplaces, with special reference to geographic information. Institutional arrangements are also likely to be profoundly affected by the trend towards interoperable systems, and nowhere is the impact of interoperability more likely to cause fundamental change than in education, as educators address the needs of a new generation of GIS users with access to a new generation of tools. The book concludes with a series of chapters on education and institutional change. Interoperating Geographic Information Systems is suitable as a secondary text for graduate level courses in computer science, geography, spatial databases, and interoperability and as a reference for researchers and practitioners in industry, commerce and government.

Ad-Hoc, Mobile, and Wireless Networks

19th International Conference on Ad-Hoc Networks and Wireless, ADHOC-NOW 2020, Bari, Italy, October 19–21, 2020, Proceedings

Springer Nature This book constitutes the refereed proceedings of the 19th International Conference on Ad-Hoc, Mobile, and Wireless Networks, ADHOC-NOW 2020, held in Bari, Italy, in October 2020.* The 19 full and 4 short papers presented were carefully reviewed and selected from 39 submissions. The papers provide an in-depth and stimulating view on the new frontiers in the field of mobile, ad hoc and wireless computing. They are organized in the following topical sections: intelligent, programmable and delay- and disruption- tolerant networks; internet of drones and smart mobility; internet of things and internet of medical things; secure communication protocols and architectures; and wireless systems. *The conference was held virtually due to the COVID-19 pandemic.

Semantic Systems. The Power of AI and Knowledge Graphs

15th International Conference, SEMANTiCS 2019, Karlsruhe, Germany, September 9–12, 2019, Proceedings

Springer Nature This open access book constitutes the refereed proceedings of the 15th International Conference on Semantic Systems, SEMANTiCS 2019, held in Karlsruhe, Germany, in September 2019. The 20 full papers and 8 short papers presented in this volume were carefully reviewed and selected from 88 submissions. They cover topics such as: web semantics and linked (open) data; machine learning and deep learning techniques; semantic information management and knowledge integration; terminology, thesaurus and ontology management; data mining and knowledge discovery; semantics in blockchain and distributed ledger technologies.

Context in Computing

A Cross-Disciplinary Approach for Modeling the Real World

Springer This volume explores how context has been and can be used in computing to model human behaviors, actions and communications as well as to manage data and knowledge. It addresses context management and exploitation of context for sharing experience across domains. The book serves as a user-centric guide for readers wishing to develop context-based applications, as well as an intellectual reference on the concept of context. It provides a broad yet deep treatment of context in computing and related areas that depend heavily on computing. The coverage is broad because of its cross-disciplinary nature but treats topics at a sufficient depth to permit a reader to implement context in his/her computational endeavors. The volume addresses how context can be integrated in software and systems and how it can be used in a computing environment. Furthermore, the use of context to represent the human dimension, individually as well as collectively is explained. Contributions also include descriptions of how context has been represented in formal as well as non-formal, structured approaches. The last section describes several human behavior representation paradigms based on the concept of context as its central representational element. The depth and breadth of this content is certain to provide useful as well as intellectually enriching information to readers of diverse backgrounds who have an interest in or are intrigued by using context to assist in their representation of the real world.

Deep Learning for NLP and Speech Recognition

Springer This textbook explains Deep Learning Architecture, with applications to various NLP Tasks, including Document Classification, Machine Translation, Language Modeling, and Speech Recognition. With the widespread adoption of deep learning, natural language processing (NLP), and speech applications in many areas (including Finance, Healthcare, and Government) there is a growing need for one comprehensive resource that maps deep learning techniques to NLP and speech and provides insights into using the tools and libraries for real-world applications. Deep Learning for NLP and Speech Recognition explains recent deep learning methods applicable to NLP and speech, provides state-of-the-art approaches, and offers real-world case studies with code to provide hands-on experience. Many books focus on deep learning theory or deep learning for NLP-specific tasks while others are cookbooks for tools and libraries, but the constant flux of new algorithms, tools, frameworks, and libraries in a rapidly evolving landscape means that there are few available texts that offer the material in this book. The book is organized into three parts, aligning to different groups of readers and their expertise. The three parts are: Machine Learning, NLP, and Speech Introduction The first part has three chapters that introduce readers to the fields of NLP, speech recognition, deep learning and machine learning with basic theory and hands-on case studies using Python-based tools and libraries. Deep Learning Basics The five chapters in the second part introduce deep learning and various topics that are crucial for speech and text processing, including word embeddings, convolutional neural networks, recurrent neural networks and speech recognition basics. Theory, practical tips, state-of-the-art methods, experimentations and analysis in using the methods discussed in theory on real-world tasks. Advanced Deep Learning Techniques for Text and Speech The third part has five chapters that discuss the latest and cutting-edge research in the areas of deep learning that intersect with NLP and speech. Topics including attention mechanisms, memory augmented networks, transfer learning, multi-task learning, domain adaptation, reinforcement learning, and end-to-end deep learning for speech recognition are covered using case studies.

Collage: A Process in Architectural Design

Springer Nature This book is about using “collage” among Iranian students in architecture studio, and in order to introduce the way these students use the technique to the English reader, we (Ali Yaser Jafari and Reihaneh Khorramrouei) have chosen this valuable book by AliAsghar Adibi to translate from Farsi to English. It provides a representative example of design through collage and culture. This book originally collected and published in three chapters: Collage history in different arts; Objectives and steps to make collage images; Two experienced examples.

Waste Management, Processing and Valorisation

Springer Nature This book highlights current efforts and research on waste management, processing and valorization, particularly in Asia-Africa countries. Chapters 1-2 highlight the overview of plastic waste management and the production of waste plastic oil (WPO). Chapters 3-5 discuss the landfill characterization and application of incineration and composting for waste processing. A new achievement in adsorbent production is highlighted in Chapters 6 and 7 while Chapters 10 and 11 focus on sewage characteristic and its utilization using microalgae. Enzyme production using waste is covered by Chapters 10-12. Chapter 13-14 dedicated to the advances in production of bioenergy. The book concludes with a discussion on life cycle analysis for solid waste management (Chapter 15).

Digital Libraries for Open Knowledge

24th International Conference on Theory and Practice of Digital Libraries, TPDL 2020, Lyon, France, August 25–27, 2020, Proceedings

Springer Nature This book constitutes the proceedings of the 24th International Conference on Theory and Practice of Digital Libraries, TPDL 2020, held in Lyon, France, in August 2020.* The 14 full papers and 4 short papers presented were carefully reviewed and selected from 53 submissions. TPDL 2020 attempts to facilitate establishing connections and convergences between diverse research communities such as Digital Humanities, Information Sciences and others that could benefit from ecosystems offered by digital libraries and repositories. The papers present a wide range of the following topics: knowledge graphs and linked data; quality assurance in digital libraries; ontology design; user requirements and behavior; research data management and discovery; and digital cultural heritage. * The conference was held virtually due to the COVID-19 pandemic.

Data Structures Using C++

Cengage Learning Now in its second edition, D.S. Malik brings his proven approach to C++ programming to the CS2 course. Clearly written with the student in mind, this text focuses on Data Structures and includes advanced topics in C++ such as Linked Lists and the Standard Template Library (STL). The text features abundant visual diagrams, examples, and extended Programming Examples, all of which serve to illuminate difficult concepts. Complete programming code and clear display of syntax, explanation, and example are used throughout the text, and each chapter concludes with a robust exercise set. Important Notice: Media content referenced within the product description or the product text may not be available in the ebook version.

Clever Algorithms

Nature-inspired Programming Recipes

Jason Brownlee This book provides a handbook of algorithmic recipes from the fields of Metaheuristics, Biologically Inspired Computation and Computational Intelligence that have been described in a complete, consistent, and centralized manner. These standardized descriptions were carefully designed to be accessible, usable, and understandable. Most of the algorithms described in this book were originally inspired by biological and natural systems, such as the adaptive capabilities of genetic evolution and the acquired immune system, and the foraging behaviors of birds, bees, ants and bacteria. An encyclopedic algorithm reference, this book is intended for research scientists, engineers, students, and interested amateurs. Each algorithm description provides a working code example in the Ruby Programming Language.

Neural Machine Translation

Cambridge University Press Learn how to build machine translation systems with deep learning from the ground up, from basic concepts to cutting-edge research.

Reactive Oxygen Species and Oxidative Damage in Plants Under Stress

Springer This book provides detailed and comprehensive information on oxidative damage caused by stresses in plants with especial reference to the metabolism of reactive oxygen species (ROS). In plants, as in all aerobic organisms, ROS are common by-products formed by the inevitable leakage of electrons onto O₂ from the electron transport activities located in chloroplasts, mitochondria, peroxisomes and in plasma membranes or as a consequence of various metabolic pathways confined in different cellular loci. Environmental stresses such as heat, cold, drought, salinity, heavy-metal toxicity, ozone and ultraviolet radiation as well as pathogens/contagion attack lead to enhanced generation of ROS in plants due to disruption of cellular homeostasis. ROS play a dual role in plants; at low concentrations they act as signaling molecules that facilitate several responses in plant cells, including those promoted by biotic and abiotic agents. In divergence, at high levels they cause damage to cellular constituents triggering oxidative stress. In either case, small antioxidant molecules and enzymes modulate the action of these ambivalent species.

Dynamics of Community Formation

Developing Identity and Notions of Home

Springer This interdisciplinary work discusses the construction, maintenance, evolution, and destruction of home and community spaces, which are central to the development of social cohesion. By examining how people throughout the world form different communities to establish a sense of home, the volume surveys the formation of identity within the context of rapid development, global and domestic neoliberal and political governmental policies, and various societal pressures. The themes of cooperation, conflict, inclusion, exclusion, and balance require negotiation between different actors (e.g., the state, professional developers, social activists, and residents) as homes and communities develop.

Deep Learning-Based Approaches for Sentiment Analysis

Springer Nature This book covers deep-learning-based approaches for sentiment analysis, a relatively new, but fast-growing research area, which has significantly changed in the past few years. The book presents a collection of state-of-the-art approaches, focusing on the best-performing, cutting-edge solutions for the most common and difficult challenges faced in sentiment analysis research. Providing detailed explanations of the methodologies, the book is a valuable resource for researchers as well as newcomers to the field.

Advanced Data Analytics Using Python

With Machine Learning, Deep Learning and NLP Examples

Apres Gain a broad foundation of advanced data analytics concepts and discover the recent revolution in databases such as Neo4j, Elasticsearch, and MongoDB. This book discusses how to implement ETL techniques including topical crawling, which is applied in domains such as high-frequency algorithmic trading and goal-oriented dialog systems. You'll also see examples of machine learning concepts such as semi-supervised learning, deep learning, and NLP. Advanced Data Analytics Using Python also covers important traditional data analysis techniques such as time series and principal component analysis. After reading this book you will have experience of every technical aspect of an analytics project. You'll get to know the concepts using Python code, giving you samples to use in your own projects. What You Will Learn Work with data analysis techniques such as classification, clustering, regression, and forecasting Handle structured and unstructured data, ETL techniques, and different kinds of databases such as Neo4j, Elasticsearch, MongoDB, and MySQL Examine the different big data frameworks, including Hadoop and Spark Discover advanced machine learning concepts such as semi-supervised learning, deep learning, and NLP Who This Book Is For Data scientists and software developers interested in the field of data analytics.

Handbook on Systemic Risk

Cambridge University Press Experts in the field provide an introduction to the multifaceted aspects of this critically important topic.

Resampling Methods

A Practical Guide to Data Analysis

Springer Science & Business Media "...the author has packaged an excellent and modern set of topics around the development and use of quantitative models...the author has the capability to work at a more modest level. He does that very effectively in this 2nd Edition... If you need to learn about resampling, this book would be a good place to start." -- Technometrics This work is a practical, table-free introduction to data analysis using the bootstrap, cross-validation, and permutation tests; new to the second edition are several additional examples and a chapter dedicated to regression, data mining techniques, and their limitations. The book's many exercises, practical data sets, and use of free shareware make it an essential resource for students and teachers, as well as industrial statisticians, consultants, and research professionals.

Nanotechnology for Lithium-Ion Batteries

Springer Science & Business Media This book combines two areas of intense interest: nanotechnology, and energy conversion and storage devices. In particular, Li-ion batteries have enjoyed conspicuous success in many consumer electronic devices and their projected use in vehicles that will revolutionize the way we travel in the near future. For many applications, Li-ion batteries are the battery of choice. This book consolidates the scattered developments in all areas of research related to nanotechnology and lithium ion batteries.

Statistical Machine Learning

A Unified Framework

CRC Press 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.

A First Course in Machine Learning

CRC Press "A First Course in Machine Learning by Simon Rogers and Mark Girolami is the best introductory book for ML currently available. It combines rigor and precision with accessibility, starts from a detailed explanation of the basic foundations of Bayesian analysis in the simplest of settings, and goes all the way to the frontiers of the subject such as infinite mixture models, GPs, and MCMC." —Devdatt Dubhashi, Professor, Department of Computer Science and Engineering, Chalmers University, Sweden "This textbook manages to be easier to read than other comparable books in the subject while retaining all the rigorous treatment needed. The new chapters put it at the forefront of the field by covering topics that have become mainstream in machine learning over the last decade." —Daniel Barbara, George Mason University, Fairfax, Virginia, USA "The new edition of A First Course in Machine Learning by Rogers and Girolami is an excellent introduction to the use of statistical methods in machine learning. The book introduces concepts such as mathematical modeling, inference, and prediction, providing 'just in time' the essential background on linear algebra, calculus, and probability theory that the reader needs to understand these concepts." —Daniel Ortiz-Arroyo, Associate Professor, Aalborg University Esbjerg, Denmark "I was impressed by how closely the material aligns with the needs of an introductory course on machine learning, which is its greatest strength...Overall, this is a pragmatic and helpful book, which is well-aligned to the needs of an introductory course and one that I will be looking at for my own students in coming months." —David Clifton, University of Oxford, UK "The first edition of this book was already an excellent introductory text on machine learning for an advanced undergraduate or taught masters level course, or indeed for anybody who wants to learn about an interesting and important field of computer science. The additional chapters of advanced material on Gaussian process, MCMC and mixture modeling provide an ideal basis for practical projects, without disturbing the very clear and readable exposition of the basics contained in the first part of the book." —Gavin Cawley, Senior Lecturer, School of Computing Sciences, University of East Anglia, UK "This book could be used for junior/senior undergraduate students or first-year graduate students, as well as individuals who want to explore the field of machine learning...The book introduces not only the concepts but the underlying ideas on algorithm implementation from a critical thinking perspective." —Guangzhi Qu, Oakland University, Rochester, Michigan, USA

Mastering Cloud Computing

Foundations and Applications Programming

Newnes Mastering Cloud Computing is designed for undergraduate students learning to develop cloud computing applications. Tomorrow's applications won't live on a single computer but will be deployed from and reside on a virtual server, accessible anywhere, any time. Tomorrow's application developers need to understand the requirements of building apps for these virtual systems, including concurrent programming, high-performance computing, and data-intensive systems. The book introduces the principles of distributed and parallel computing underlying cloud architectures and specifically focuses on virtualization, thread programming, task programming, and map-reduce programming. There are examples demonstrating all of these and more, with exercises and labs throughout. Explains how to make design choices and tradeoffs to consider when building applications to run in a virtual cloud environment Real-world case studies include scientific, business, and energy-efficiency considerations

Arabic Natural Language Processing

Morgan & Claypool Publishers This book provides system developers and researchers in natural language processing and computational linguistics with the necessary background information for working with the Arabic language. The goal is to introduce Arabic linguistic phenomena and review the state-of-the-art in Arabic processing. The book discusses Arabic script, phonology, orthography, morphology, syntax and semantics, with a final chapter on machine translation issues. The chapter sizes correspond more or less to what is linguistically distinctive about Arabic, with morphology getting the lion's share, followed by Arabic script. No previous knowledge of Arabic is needed. This book is designed for computer scientists and linguists alike. The focus of the book is on Modern Standard Arabic; however, notes on practical issues related to Arabic dialects and languages written in the Arabic script are presented in different chapters. Table of Contents: What is "Arabic"? / Arabic Script / Arabic Phonology and Orthography / Arabic Morphology / Computational Morphology Tasks / Arabic Syntax / A Note on Arabic Semantics / A Note on Arabic and Machine Translation

Advances in Neural Information Processing Systems 12

Proceedings of the 1999 Conference

MIT Press The annual conference on Neural Information Processing Systems (NIPS) is the flagship conference on neural computation. It draws preeminent academic researchers from around the world and is widely considered to be a showcase conference for new developments in network algorithms and architectures. The broad range of interdisciplinary research areas represented includes computer science, neuroscience, statistics, physics, cognitive science, and many branches of engineering, including signal processing and control theory. Only about 30 percent of the papers submitted are accepted for presentation at NIPS, so the quality is exceptionally high. These proceedings contain all of the papers that were presented.

Handbook of Research on Intrusion Detection Systems

IGI Global Businesses in today's world are adopting technology-enabled operating models that aim to improve growth, revenue, and identify emerging markets. However, most of these businesses are not suited to defend themselves from the cyber risks that come with these data-driven practices. To further prevent these threats, they need to have a complete understanding of modern network security solutions and the ability to manage, address, and respond to security breaches. The Handbook of Research on Intrusion Detection Systems provides emerging research exploring the theoretical and practical aspects of prominent and effective techniques used to detect and contain breaches within the fields of data science and cybersecurity. Featuring coverage on a broad range of topics such as botnet detection, cryptography, and access control models, this book is ideally designed for security analysts, scientists, researchers, programmers, developers, IT professionals, scholars, students, administrators, and faculty members seeking research on current advancements in network security technology.

Machine Learning

An Algorithmic Perspective

CRC Press 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 Translation and the Information Soup

Third Conference of the Association for Machine Translation in the Americas, AMTA'98, Langhorne, PA, USA, October 28-31, 1998 Proceedings

Springer Machine Translation and the Information Soup! Over the past forty years, machine translation has grown from a tantalizing dream to a respectable and stable scientific-linguistic enterprise, with users, commercial systems, university research, and government participation. But until very recently, MT has been performed as a relatively distinct operation, somewhat isolated from other text processing. Today, this situation is changing rapidly. The explosive growth of the Web has brought multilingual text into the reach of nearly everyone with a computer. We live in a soup of information, an increasingly multilingual bouillabaisse. And to partake of this soup, we can use MT systems together with more and more tools and language processing technologies|information retrieval engines, -tomated text summarizers, and multimodal and multilingual displays. Though some of them may still be rather experimental, and though they may not quite fit together well yet, it is clear that the future will offer text manipulation systems that contain all these functions, seamlessly interconnected in various ways.

Intelligent and Fuzzy Techniques for Emerging Conditions and Digital Transformation

Proceedings of the INFUS 2021 Conference, held August 24-26, 2021. Volume 1

Springer Nature This book presents recent research in intelligent and fuzzy techniques. Emerging conditions such as pandemic, wars, natural disasters and various high technologies force people for significant changes in business and social life. The adoption of digital technologies to transform services or businesses, through replacing non-digital or manual processes with digital processes or replacing older digital technology with newer digital technologies through intelligent systems is the main scope of this book. It focuses on revealing the reflection of digital transformation in our business and social life under emerging conditions through intelligent and fuzzy systems. The latest intelligent and fuzzy methods and techniques on digital transformation are introduced by theory and applications. The intended readers are intelligent and fuzzy systems researchers, lecturers, M.Sc. and Ph.D. students studying digital transformation. Usage of ordinary fuzzy sets and their extensions, heuristics and metaheuristics from optimization to machine learning, from quality management to risk management makes the book an excellent source for researchers.

Neural Network Learning

Theoretical Foundations

Cambridge University Press This work explores probabilistic models of supervised learning problems and addresses the key statistical and computational questions. Chapters survey research on pattern classification with binary-output networks, including a discussion of the relevance of the Vapnik Chervonenkis dimension, and of estimates of the dimension for several neural network models. In addition, the authors develop a model of classification by real-output networks, and demonstrate the usefulness of classification...

Grokking Deep Learning

Simon and Schuster Summary Grokking Deep Learning teaches you to build deep learning neural networks from scratch! In his engaging style, seasoned deep learning expert Andrew Trask shows you the science under the hood, so you grok for yourself every detail of training neural networks. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Deep learning, a branch of artificial intelligence, teaches computers to learn by using neural networks, technology inspired by the human brain. Online text translation, self-driving cars, personalized product recommendations, and virtual voice assistants are just a few of the exciting modern advancements possible thanks to deep learning. About the Book Grokking Deep Learning teaches you to build deep learning neural networks from scratch! In his engaging style, seasoned deep learning expert Andrew Trask shows you the science under the hood, so you grok for yourself every detail of training neural networks. Using only Python and its math-supporting library, NumPy, you'll train your own neural networks to see and understand images, translate text into different languages, and even write like Shakespeare! When you're done, you'll be fully prepared to move on to mastering deep learning frameworks. What's inside The science behind deep learning Building and training your own neural networks Privacy concepts, including federated learning Tips for continuing your pursuit of deep learning About the Reader For readers with high school-level math and intermediate programming skills. About the Author Andrew Trask is a PhD student at Oxford University and a research scientist at DeepMind. Previously, Andrew was a researcher and analytics product manager at Digital Reasoning, where he trained the world's largest artificial neural network and helped guide the analytics roadmap for the Synthesys cognitive computing platform. Table of Contents Introducing deep learning: why you should learn it Fundamental concepts: how do machines learn? Introduction to neural prediction: forward propagation Introduction to neural learning: gradient descent Learning multiple weights at a time: generalizing gradient descent Building your first deep neural network: introduction to backpropagation How to picture neural networks: in your head and on paper Learning signal and ignoring noise: introduction to regularization and batching Modeling probabilities and nonlinearities: activation functions Neural learning about edges and corners: intro to convolutional neural networks Neural networks that understand language: king - man + woman == ? Neural networks that write like Shakespeare: recurrent layers for variable-length data Introducing automatic optimization: let's build a deep learning framework Learning to write like Shakespeare: long short-term memory Deep learning on unseen data: introducing federated learning Where to go from here: a brief guide