
Site To Download Learning Data Yaser S Abu Mostafa

Right here, we have countless ebook **Learning Data Yaser S Abu Mostafa** and collections to check out. We additionally meet the expense of variant types and with type of the books to browse. The okay book, fiction, history, novel, scientific research, as with ease as various other sorts of books are readily friendly here.

As this Learning Data Yaser S Abu Mostafa, it ends taking place being one of the favored book Learning Data Yaser S Abu Mostafa collections that we have. This is why you remain in the best website to look the incredible ebook to have.

9D2 - SUTTON AGUIRRE

Genetic algorithms are founded upon the principle of evolution, i.e., survival of the fittest. Hence evolution programming techniques, based on genetic algorithms, are applicable to many hard optimization problems, such as optimization of functions with linear and nonlinear constraints, the traveling salesman problem, and problems of scheduling, partitioning, and control. The importance of these techniques is still growing, since evolution programs are parallel in nature, and parallelism is one of the most promising directions in computer science. The book is self-contained and the only prerequisite is basic undergraduate mathematics. This third edition has been substantially revised and extended by three new chapters and by additional appendices containing working material to cover recent developments and a change in the perception of evolutionary computation.

This open access book explores machine learning and its impact on how we make sense of the world. It does so by bringing together two 'revolutions' in a surprising analogy: the revolution of

machine learning, which has placed computing on the path to artificial intelligence, and the revolution in thinking about the law that was spurred by Oliver Wendell Holmes Jr in the last two decades of the 19th century. Holmes reconceived law as prophecy based on experience, prefiguring the buzzwords of the machine learning age-prediction based on datasets. On the path to AI introduces readers to the key concepts of machine learning, discusses the potential applications and limitations of predictions generated by machines using data, and informs current debates amongst scholars, lawyers and policy makers on how it should be used and regulated wisely. Technologists will also find useful lessons learned from the last 120 years of legal grappling with accountability, explainability, and biased data.

November 28-December 1, 1994, Denver, Colorado NIPS is the longest running annual meeting devoted to Neural Information Processing Systems. Drawing on such disparate domains as neuroscience, cognitive science, computer science, statistics, mathematics, engineering, and theoretical physics, the papers collected in the proceedings of NIPS7 reflect the enduring scientific and

practical merit of a broad-based, inclusive approach to neural information processing. The primary focus remains the study of a wide variety of learning algorithms and architectures, for both supervised and unsupervised learning. The 139 contributions are divided into eight parts: Cognitive Science, Neuroscience, Learning Theory, Algorithms and Architectures, Implementations, Speech and Signal Processing, Visual Processing, and Applications. Topics of special interest include the analysis of recurrent nets, connections to HMMs and the EM procedure, and reinforcement-learning algorithms and the relation to dynamic programming. On the theoretical front, progress is reported in the theory of generalization, regularization, combining multiple models, and active learning. Neuroscientific studies range from the large-scale systems such as visual cortex to single-cell electrotonic structure, and work in cognitive scientific is closely tied to underlying neural constraints. There are also many novel applications such as tokamak plasma control, Glove-Talk, and hand tracking, and a variety of hardware implementations, with particular focus on analog VLSI.

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.

The annual conference on Neural Information Processing Systems (NIPS) is the flagship conference on neural computation. These proceedings contain all of the papers that were presented.

This text is a semester course in the basic mathematical and theoretical foundations of computer science. Students who make heavy use of computing should learn these foundations well, setting a base for a follow-on course in algorithms. A solid theoretical and algorithmic foundation in computer science sets the stage for developing good programs, programs that work, always and efficiently. Each chapter is a lecture that has been taught as such. Part I starts with basic logic, proofs and discrete mathematics, including: induction, recursion, summation, asymptotics and number theory. We then continue with graphs, counting and combinatorics, and wrap up the coverage of discrete mathematics with discrete probability. Part II presents the blockbuster application of discrete mathematics: the digital computer and a theory of computing. The goal is to understand what a computer can and cannot do. We start small, with automata, and end big with Turing Machines. Our approach is Socratic. The reader is encouraged to participate actively in the learning process by doing the quizzes and exercises that are liberally sprinkled through the text. The pace and level is appropriate for readers with one year of training in programming and calculus (college sophomores).

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 Engi-

neering 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.

Table of contents

An introduction to a broad range of topics in deep learning, covering mathematical and conceptual background, deep learning techniques used in industry, and research perspectives. “Written by three experts in the field, *Deep Learning* is the only comprehensive book on the subject.” —Elon Musk, cochair of OpenAI; cofounder and CEO of Tesla and SpaceX Deep learning is a form of machine learning that enables computers to learn from experience and understand the world in terms of a hierarchy of concepts. Because the computer gathers knowledge from experience, there is no need for a human computer operator to formally specify all the knowledge that the computer needs. The hierarchy of concepts allows the computer to learn complicated concepts by building them out of simpler ones; a graph of these hierarchies would be many layers deep. This book introduces a broad range of topics in deep learning. The text offers mathematical and conceptual background, covering relevant concepts in linear algebra, probability theory and information theory, numerical computation, and machine learning. It describes deep learning techniques used by practitioners in industry, including deep feedforward networks,

regularization, optimization algorithms, convolutional networks, sequence modeling, and practical methodology; and it surveys such applications as natural language processing, speech recognition, computer vision, online recommendation systems, bioinformatics, and videogames. Finally, the book offers research perspectives, covering such theoretical topics as linear factor models, autoencoders, representation learning, structured probabilistic models, Monte Carlo methods, the partition function, approximate inference, and deep generative models. Deep Learning can be used by undergraduate or graduate students planning careers in either industry or research, and by software engineers who want to begin using deep learning in their products or platforms. A website offers supplementary material for both readers and instructors.

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 mod-

els, 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

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.

This textbook provides a comprehensive and current understanding of signal detection and estimation, including problems and solutions for each chapter. Signal detection plays an important role in fields such as radar, sonar, digital communications, image processing, and failure detection. The book explores both Gaussian detection and detection of Markov chains, presenting a unified treatment of coding and modulation topics. Addresses asymptotic of tests with the theory of large deviations, and robust detection. This text is appropriate for students of Electrical Engineering in graduate courses in Signal Detection and Estimation.

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

In this textbook, basic mathematical models used in Big Data Analytics are presented and application-oriented references to relevant practical issues are made. Necessary mathematical tools are examined and applied to current problems of data analysis, such as brand loyalty, portfolio selection, credit investigation, quality control, product clustering, asset pricing etc. - mainly in an economic context. In addition, we discuss interdisciplinary applications to biology, linguistics, sociology, electrical engineering, computer science and artificial intelligence. For the models, we make use of a wide range of mathematics - from basic disciplines of numerical linear algebra, statistics and optimization to more special-

ized game, graph and even complexity theories. By doing so, we cover all relevant techniques commonly used in Big Data Analytics. Each chapter starts with a concrete practical problem whose primary aim is to motivate the study of a particular Big Data Analytics technique. Next, mathematical results follow - including important definitions, auxiliary statements and conclusions arising. Case-studies help to deepen the acquired knowledge by applying it in an interdisciplinary context. Exercises serve to improve understanding of the underlying theory. Complete solutions for exercises can be consulted by the interested reader at the end of the textbook; for some which have to be solved numerically, we provide descriptions of algorithms in Python code as supplementary material. This textbook has been recommended and developed for university courses in Germany, Austria and Switzerland.

Over the past three decades or so, research on machine learning and data mining has led to a wide variety of algorithms that learn general functions from experience. As machine learning is maturing, it has begun to make the successful transition from academic research to various practical applications. Generic techniques such as decision trees and artificial neural networks, for example, are now being used in various commercial and industrial applications. Learning to Learn is an exciting new research direction within machine learning. Similar to traditional machine-learning algorithms, the methods described in Learning to Learn induce general functions from experience. However, the book investigates algorithms that can change the way they generalize, i.e., practice the task of learning itself, and improve on it. To illustrate the utility of learning to learn, it is worthwhile comparing machine learning with human learning. Humans encounter a continual stream

of learning tasks. They do not just learn concepts or motor skills, they also learn bias, i.e., they learn how to generalize. As a result, humans are often able to generalize correctly from extremely few examples - often just a single example suffices to teach us a new thing. A deeper understanding of computer programs that improve their ability to learn can have a large practical impact on the field of machine learning and beyond. In recent years, the field has made significant progress towards a theory of learning to learn along with practical new algorithms, some of which led to impressive results in real-world applications. Learning to Learn provides a survey of some of the most exciting new research approaches, written by leading researchers in the field. Its objective is to investigate the utility and feasibility of computer programs that can learn how to learn, both from a practical and a theoretical point of view.

Learning and Intelligent Optimization (LION) is the combination of learning from data and optimization applied to solve complex and dynamic problems. The LION way is about increasing the automation level and connecting data directly to decisions and actions. More power is directly in the hands of decision makers in a self-service manner, without resorting to intermediate layers of data scientists. LION is a complex array of mechanisms, like the engine in an automobile, but the user (driver) does not need to know the inner workings of the engine in order to realize its tremendous benefits. LION's adoption will create a prairie fire of innovation which will reach most businesses in the next decades. Businesses, like plants in wildfire-prone ecosystems, will survive and prosper by adapting and embracing LION techniques, or they risk being transformed from giant trees to ashes by the spreading

competition.

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

129 6.2 Representation of hints. 131 6.3 Monotonicity hints ..
134 6.4 Theory 139 6.4.1 Capacity results 140 6.4.2 Deci-

sion boundaries 144 6.5 Conclusion 145 6.6 References..... ..
146 7 Analysis and Synthesis Tools for Robust SPRness 147 C.
Mosquera, J.R. Hernandez, F. Perez-Gonzalez 7.1 Introduc-
tion..... 147 7.2 SPR Analysis of Uncertain Systems. 153
7.2.1 The Poly topic Case . 155 7.2.2 The ZP-Ball Case 157
7.2.3 The Roots Space Case 159 7.3 Synthesis of LTI Filters
for Robust SPR Problems 161 7.3.1 Algebraic Design for Two
Plants 161 7.3.2 Algebraic Design for Three or More Plants
164 7.3.3 Approximate Design Methods. 165 7.4 Experimental re-
sults 167 7.5 Conclusions 168 7.6 References 169 8
Boundary Methods for Distribution Analysis 173 J.L. Sancho et al.
8.1 Introduction 173 8.1.1 Building a Classifier System .
175 8.2 Motivation 176 8.3 Boundary Methods as Fea-
ture-Set Evaluation 177 8.3.1 Results 179 8.3.2 Fea-
ture Set Evaluation using Boundary Methods: S- mary.
. 182 . . .

This is the first textbook on pattern recognition to present the Bayesian viewpoint. The book presents approximate inference algorithms that permit fast approximate answers in situations where exact answers are not feasible. It uses graphical models to describe probability distributions when no other books apply graphical models to machine learning. No previous knowledge of pattern recognition or machine learning concepts is assumed. Familiarity with multivariate calculus and basic linear algebra is required, and some experience in the use of probabilities would be helpful though not essential as the book includes a self-contained introduction to basic probability theory.

A pragmatic approach to Big Data by taking the reader on a journey between Big Data (what it is) and the Smart Data (what it is

for). Today's decision making can be reached via information (related to the data), knowledge (related to people and processes), and timing (the capacity to decide, act and react at the right time). The huge increase in volume of data traffic, and its format (unstructured data such as blogs, logs, and video) generated by the "digitalization" of our world modifies radically our relationship to the space (in motion) and time, dimension and by capillarity, the enterprise vision of performance monitoring and optimization. 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.

This completely revised second edition presents an introduction to statistical pattern recognition. Pattern recognition in general covers a wide range of problems: it is applied to engineering problems, such as character readers and wave form analysis as well as to brain modeling in biology and psychology. Statistical decision and estimation, which are the main subjects of this book, are regarded as fundamental to the study of pattern recognition. This book is appropriate as a text for introductory courses in pattern recognition and as a reference book for workers in the field. Each chapter contains computer projects as well as exercises.

Molecular biologists are performing increasingly large and complicated experiments, but often have little background in data analy-

sis. The book is devoted to teaching the statistical and computational techniques molecular biologists need to analyze their data. It explains the big-picture concepts in data analysis using a wide variety of real-world molecular biological examples such as eQTLs, ortholog identification, motif finding, inference of population structure, protein fold prediction and many more. The book takes a pragmatic approach, focusing on techniques that are based on elegant mathematics yet are the simplest to explain to scientists with little background in computers and statistics.

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...

Deploy deep learning solutions in production with ease using TensorFlow. You'll also develop the mathematical understanding and intuition required to invent new deep learning architectures and solutions on your own. Pro Deep Learning with TensorFlow provides practical, hands-on expertise so you can learn deep learning from scratch and deploy meaningful deep learning solutions. This book will allow you to get up to speed quickly using TensorFlow and to optimize different deep learning architectures. All of the practical aspects of deep learning that are relevant in any industry are emphasized in this book. You will be able to use the prototypes demonstrated to build new deep learning applications.

The code presented in the book is available in the form of iPython notebooks and scripts which allow you to try out examples and extend them in interesting ways. You will be equipped with the mathematical foundation and scientific knowledge to pursue research in this field and give back to the community. What You'll Learn Understand full stack deep learning using TensorFlow and gain a solid mathematical foundation for deep learning Deploy complex deep learning solutions in production using TensorFlow Carry out research on deep learning and perform experiments using TensorFlow Who This Book Is For Data scientists and machine learning professionals, software developers, graduate students, and open source enthusiasts

This in-depth guide provides managers with a solid understanding of data and data trends, the opportunities that it can offer to businesses, and the dangers of these technologies. Written in an accessible style, Steven Finlay provides a contextual roadmap for developing solutions that deliver benefits to organizations.

Get complete instructions for manipulating, processing, cleaning, and crunching datasets in Python. Updated for Python 3.6, the second edition of this hands-on guide is packed with practical case studies that show you how to solve a broad set of data analysis problems effectively. You'll learn the latest versions of pandas, NumPy, IPython, and Jupyter in the process. Written by Wes McKinney, the creator of the Python pandas project, this book is a practical, modern introduction to data science tools in Python. It's ideal for analysts new to Python and for Python programmers new to data science and scientific computing. Data files and related material are available on GitHub. Use the IPython shell and Jupyter notebook for exploratory computing Learn basic and ad-

vanced features in NumPy (Numerical Python) Get started with data analysis tools in the pandas library Use flexible tools to load, clean, transform, merge, and reshape data Create informative visualizations with matplotlib Apply the pandas groupby facility to slice, dice, and summarize datasets Analyze and manipulate regular and irregular time series data Learn how to solve real-world data analysis problems with thorough, detailed examples

This book explores the intuitive appeal of neural networks and the genetic algorithm in finance. It demonstrates how neural networks used in combination with evolutionary computation outperform classical econometric methods for accuracy in forecasting, classification and dimensionality reduction. McNelis utilizes a variety of examples, from forecasting automobile production and corporate bond spread, to inflation and deflation processes in Hong Kong and Japan, to credit card default in Germany to bank failures in Texas, to cap-floor volatilities in New York and Hong Kong. * Offers a balanced, critical review of the neural network methods and genetic algorithms used in finance * Includes numerous examples and applications * Numerical illustrations use MATLAB code and the book is accompanied by a website

This book contains all refereed papers that were accepted to the second edition of the « Digital Enterprise Design & Management » (DED&M 2014) international conference that took place in Paris (France) from February 4 to February 5, 2014. These proceedings cover the most recent trends in the emerging field of Digital Enterprise, both from an academic and a professional perspective. A special focus is put on digital uses, digital strategies, digital infrastructures and digital governance from an Enterprise Architecture

point of view. The DED&M 2014 conference is organized under the guidance of the Center of Excellence on Systems Architecture, Management, Economy and Strategy and benefits from the supports of both the Orange - Ecole Polytechnique - Télécom ParisTech "Innovation and Regulation" Chair and the Dassault Aviation - DCNS - DGA - Thales - Ecole Polytechnique - ENSTA ParisTech - Télécom ParisTech "Complex Systems Engineering" Chair. Big Data Analytics with Spark is a step-by-step guide for learning Spark, which is an open-source fast and general-purpose cluster computing framework for large-scale data analysis. You will learn how to use Spark for different types of big data analytics projects, including batch, interactive, graph, and stream data analysis as well as machine learning. In addition, this book will help you become a much sought-after Spark expert. Spark is one of the hottest Big Data technologies. The amount of data generated today by devices, applications and users is exploding. Therefore, there is a critical need for tools that can analyze large-scale data and unlock value from it. Spark is a powerful technology that meets that need. You can, for example, use Spark to perform low latency computations through the use of efficient caching and iterative algorithms; leverage the features of its shell for easy and interactive Data analysis; employ its fast batch processing and low latency features to process your real time data streams and so on. As a result, adoption of Spark is rapidly growing and is replacing Hadoop MapReduce as the technology of choice for big data analytics. This book provides an introduction to Spark and related big-data technologies. It covers Spark core and its add-on libraries, including Spark SQL, Spark Streaming, GraphX, and MLlib. Big Data Analytics with Spark is therefore written for busy pro-

professionals who prefer learning a new technology from a consolidated source instead of spending countless hours on the Internet trying to pick bits and pieces from different sources. The book also provides a chapter on Scala, the hottest functional programming language, and the program that underlies Spark. You'll learn the basics of functional programming in Scala, so that you can write Spark applications in it. What's more, Big Data Analytics with Spark provides an introduction to other big data technologies that are commonly used along with Spark, like Hive, Avro, Kafka and so on. So the book is self-sufficient; all the technologies that you need to know to use Spark are covered. The only thing that you are expected to know is programming in any language. There is a critical shortage of people with big data expertise, so companies are willing to pay top dollar for people with skills in areas like Spark and Scala. So reading this book and absorbing its principles will provide a boost—possibly a big boost—to your career.

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

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

Presenting a theory of the theoryless, a computer scientist provides a model of how effective behavior can be learned even in a world as complex as our own, shedding new light on human nature.

Artificial neural networks are nonlinear mapping systems whose structure is loosely based on principles observed in the nervous

systems of humans and animals. The basic idea is that massive systems of simple units linked together in appropriate ways can generate many complex and interesting behaviors. This book focuses on the subset of feedforward artificial neural networks called multilayer perceptrons (MLP). These are the mostly widely used neural networks, with applications as diverse as finance (forecasting), manufacturing (process control), and science (speech and image recognition). This book presents an extensive and practical overview of almost every aspect of MLP methodology, progressing from an initial discussion of what MLPs are and how they might be used to an in-depth examination of technical factors affecting performance. The book can be used as a tool kit by readers interested in applying networks to specific problems, yet it also presents theory and references outlining the last ten years of MLP research.

Machine learning (ML) is changing virtually every aspect of our lives. Today ML algorithms accomplish tasks that until recently only expert humans could perform. As it relates to finance, this is the most exciting time to adopt a disruptive technology that will transform how everyone invests for generations. Readers will learn how to structure Big data in a way that is amenable to ML algorithms; how to conduct research with ML algorithms on that data; how to use supercomputing methods; how to backtest your discoveries while avoiding false positives. The book addresses real-life problems faced by practitioners on a daily basis, and explains scientifically sound solutions using math, supported by code and examples. Readers become active users who can test the proposed solutions in their particular setting. Written by a recognized expert and portfolio manager, this book will equip invest-

ment professionals with the groundbreaking tools needed to succeed in modern finance.

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.

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/OC-

TAVE-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.

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.

This exciting and innovative Handbook provides readers with a comprehensive and globally relevant overview of the instruments, actors and design features of social protection systems, as well as their application and impacts in practice. It is the first book that centres around system building globally, a theme that has gained political importance yet has received relatively little attention in academia.