Learning From Data A Short Course Yaser S Abu Mostafa

Yeah, reviewing a books **Learning From Data A Short Course Yaser S Abu Mostafa** could grow your close connections listings. This is just one of the solutions for you to be successful. As understood, capability does not recommend that you have extraordinary points.

Comprehending as skillfully as union even more than other will meet the expense of each success. neighboring to, the notice as skillfully as insight of this Learning From Data A Short Course Yaser S Abu Mostafa can be taken as capably as picked to act.

Connectionist Models of Cognition and Perception - John Andrew Bullinaria 2002

Connectionist Models of Cognition and Perception collects together refereed versions of twenty-three papers presented at the Seventh Neural Computation and Psychology Workshop (NCPW7). This workshop series is a well-established and unique forum that brings together researchers from such diverse disciplines as artificial intelligence, cognitive science, computer science, neurobiology, philosophy and psychology to discuss their latest work on connectionist modelling in psychology.The articles have the main theme of connectionist modelling of cognition and perception, and are organised into six sections, on: cell assemblies, representation, memory, perception, vision and language. This book is an invaluable resource for researchers interested in neural models of psychological phenomena.

Deep Learning in Computational Mechanics - Stefan Kollmannsberger 2021-08-05

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.

Portfolio Management under Stress - Riccardo Rebonato 2014-01-09

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

Pro Deep Learning with TensorFlow - Santanu Pattanayak 2017-12-06

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

Discrete Mathematics and Computing - Malik Magdon-Ismail 2019-12-14

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).

On the Path to AI - Thomas D. Grant 2020-01-01

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.

Introduction to Statistical Pattern Recognition - Keinosuke Fukunaga 2013-10-22

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.

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

Machine learning allows computers to learn and discern patterns without actually being programmed. When Statistical techniques and machine learning are combined together they are a powerful tool for analysing various kinds of data in many computer science/engineering areas including, image processing, speech processing, natural language processing, robot control, as well as in fundamental sciences such as biology, medicine, astronomy, physics, and materials. Introduction to Statistical Machine Learning provides a general introduction to machine learning that covers a wide range of topics concisely and will help you bridge the gap between theory and practice. Part I discusses the fundamental concepts of statistics and probability that are used in describing machine learning algorithms. Part II and Part III explain the two major approaches of machine learning techniques; generative methods and discriminative methods. While Part III provides an in-depth look at advanced topics that play essential roles in making machine learning algorithms more useful in practice. The accompanying MATLAB/Octave programs provide you with the necessary practical skills needed to accomplish a wide range of data analysis tasks. Provides the necessary background material to understand machine learning such as statistics, probability, linear algebra, and calculus. Complete coverage of the generative approach to statistical pattern recognition and the discriminative approach to statistical machine learning. Includes MATLAB/Octave programs so that readers can test the algorithms numerically and acquire both mathematical and practical skills in a wide range of data analysis tasks Discusses a wide range of applications in machine learning and statistics and provides examples drawn from image processing, speech processing, natural language processing, robot control, as well as biology, medicine, astronomy, physics, and materials.

Intelligent Methods in Signal Processing and Communications - Domingo DoCampo 2012-12-06

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 Decision 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 Introduction..... 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 results 167 7.5 Conclusions 168 7.6 References ..... . 169 8 Boundary Methods for Distribution Analysis 173 J.L. Sancho et aZ. 8.1 Introduction ..... . 173 8.1.1 Building a Classifier System . 175 8.2 Motivation ..... . 176 8.3 Boundary Methods as Feature-Set Evaluation 177 8.3.1 Results ..... . 179 8.3.2 Feature Set Evaluation using Boundary Methods: S- mary. . . . . 182 . . .

Big Data Analytics with Spark - Mohammed Guller 2015-12-29

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 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.

Python for Data Analysis - Wes McKinney 2017-09-25

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 advanced 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

Grokking Deep Learning - Andrew W. Trask 2019-01-23

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

Machine Learning - Kevin P. Murphy 2012-08-24

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

A First Course in Machine Learning - Simon Rogers 2016-10-14

"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

Programming Collective Intelligence - Toby Segaran 2007-08-16

Want to tap the power behind search rankings, product recommendations, social bookmarking, and online matchmaking? This fascinating book demonstrates how you can build Web 2.0 applications to mine the enormous amount of data created by people on the Internet. With the sophisticated algorithms in this book, you can write smart programs to access interesting datasets from other web sites, collect data from users of your

own applications, and analyze and understand the data once you've found it. Programming Collective Intelligence takes you into the world of machine learning and statistics, and explains how to draw conclusions about user experience, marketing, personal tastes, and human behavior in general -- all from information that you and others collect every day. Each algorithm is described clearly and concisely with code that can immediately be used on your web site, blog, Wiki, or specialized application. This book explains: Collaborative filtering techniques that enable online retailers to recommend products or media Methods of clustering to detect groups of similar items in a large dataset Search engine features -- crawlers, indexers, query engines, and the PageRank algorithm Optimization algorithms that search millions of possible solutions to a problem and choose the best one Bayesian filtering, used in spam filters for classifying documents based on word types and other features Using decision trees not only to make predictions, but to model the way decisions are made Predicting numerical values rather than classifications to build price models Support vector machines to match people in online dating sites Non-negative matrix factorization to find the independent features in a dataset Evolving intelligence for problem solving -- how a computer develops its skill by improving its own code the more it plays a game Each chapter includes exercises for extending the algorithms to make them more powerful. Go beyond simple database-backed applications and put the wealth of Internet data to work for you. "Bravo! I cannot think of a better way for a developer to first learn these algorithms and methods, nor can I think of a better way for me (an old AI dog) to reinvigorate my knowledge of the details." -- Dan Russell, Google "Toby's book does a great job of breaking down the complex subject matter of machine-learning algorithms into practical, easy-to-understand examples that can be directly applied to analysis of social interaction across the Web today. If I had this book two years ago, it would have saved precious time going down some fruitless paths." -- Tim Wolters, CTO, Collective Intellect

Machine Learning Pocket Reference - Matt Harrison 2019-08-27

With detailed notes, tables, and examples, this handy reference will help you navigate the basics of structured machine learning. Author Matt Harrison delivers a valuable guide that you can use for additional support during training and as a convenient resource when you dive into your next machine learning project. Ideal for programmers, data scientists, and AI engineers, this book includes an overview of the machine learning process and walks you through classification with structured data. You'll also learn methods for clustering, predicting a continuous value (regression), and reducing dimensionality, among other topics. This pocket reference includes sections that cover: Classification, using the Titanic dataset Cleaning data and dealing with missing data Exploratory data analysis Common preprocessing steps using sample data Selecting features useful to the model Model selection Metrics and classification evaluation Regression examples using k-nearest neighbor, decision trees, boosting, and more Metrics for regression evaluation Clustering Dimensionality reduction Scikit-learn pipelines

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.

Learning from Data - Yaser S. Abu-Mostafa 2012-01-01

Fundamentals of Clinical Data Science - Pieter Kubben 2018-12-21

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.

Advances in Neural Information Processing Systems 9 - Michael C. Mozer 1997

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 neural networks and genetic algorithms, cognitive science, neuroscience and biology, computer science, AI, applied mathematics, physics, and many branches of engineering. Only about 30% of the papers submitted are accepted for presentation at NIPS, so the quality is exceptionally high. All of the papers presented appear in these proceedings.

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

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

Digital Enterprise Design & Management - Pierre-Jean Benghozi 2013-12-24

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 – T□ l□ com ParisTech "Innovation and Regulation" Chair and the Dassault Aviation – DCNS – DGA – Thales – Ecole Polytechnique – ENSTA ParisTech – T□ l□ com ParisTech "Complex Systems Engineering" Chair.

Self-Aware Computing Systems - Samuel Kounev 2017-01-23

This book provides formal and informal definitions and taxonomies for self-aware computing systems, and explains how self-aware computing relates to many existing subfields of computer science, especially software engineering. It describes architectures and algorithms for self-aware systems as well as the benefits and pitfalls of self-awareness, and reviews much of the latest relevant research across a wide array of disciplines, including open research challenges. The chapters of this book are organized into five parts: Introduction, System Architectures, Methods and Algorithms, Applications and Case Studies, and Outlook. Part I offers an introduction that defines self-aware computing systems from multiple perspectives, and establishes a formal definition, a taxonomy and a set of reference scenarios that help to unify the remaining chapters. Next, Part II explores architectures for self-aware computing systems, such as generic concepts and notations that allow a wide range of self-aware system architectures to be described and compared with both isolated and interacting systems. It also reviews the current state of reference architectures, architectural frameworks, and languages for self-aware systems. Part III focuses on methods and algorithms for self-aware computing systems by addressing issues pertaining to system design, like modeling, synthesis and verification. It also examines topics such as adaptation, benchmarks and metrics. Part IV then presents applications and case studies in various domains including cloud computing, data centers, cyber-physical systems, and the degree to which self-aware computing approaches have been adopted within those domains. Lastly, Part V surveys open challenges and future research directions for self-aware computing systems. It can be used as a handbook for professionals and researchers working in areas related to self-aware computing, and can also serve as an advanced textbook for lecturers and postgraduate students studying subjects like advanced software engineering, autonomic computing, self-adaptive systems, and data-center resource management. Each chapter is largely self-contained, and offers plenty of references for anyone wishing to pursue the topic more deeply.

Deep Learning for NLP and Speech Recognition - Uday Kamath 2019-06-10

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.

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.

Gravitational N-Body Simulations - Sverre J. Aarseth 2003-10-23

This book discusses in detail all the relevant numerical methods for the classical N-body problem. It demonstrates how to develop clear and elegant algorithms for models of gravitational systems, and explains the fundamental mathematical tools needed to describe the dynamics of a large number of mutually attractive particles. Particular attention is given to the techniques needed to model astrophysical phenomena such as close encounters and the dynamics of black hole binaries. The author reviews relevant work in the field and covers applications to the problems of planetary formation and star cluster dynamics, both of Pleiades type and globular clusters. Self-contained and pedagogical, this book is suitable for graduate students and researchers in theoretical physics, astronomy and cosmology.

Advances in Neural Information Processing Systems 7 - Gerald Tesauro 1995

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.

Information Theory, Inference and Learning Algorithms

- David J. C. MacKay 2003-09-25
Information theory and inference, taught together in this exciting textbook, lie at the heart of many important areas of modern technology - communication, signal processing, data mining, machine learning, pattern recognition, computational neuroscience, bioinformatics and cryptography. The book introduces theory in tandem with applications. Information theory is taught alongside practical communication systems such as arithmetic coding for data compression and sparse-graph codes for error-correction. Inference techniques, including message-passing algorithms, Monte Carlo methods and variational approximations, are developed alongside applications to clustering, convolutional codes, independent component analysis, and neural networks. Uniquely, the book covers state-of-the-art error-correcting codes, including low-density-parity-check codes, turbo codes, and digital fountain codes - the twenty-first-century standards for satellite communications, disk drives, and data broadcast. Richly illustrated, filled with worked examples and over 400 exercises, some with detailed solutions, the book is ideal for self-learning, and for undergraduate or graduate courses. It also provides an unparalleled entry point for professionals in areas as diverse as computational biology, financial engineering and machine learning.

The Minimum Description Length Principle - Peter D. Gr nwald 2007
This introduction to the MDL Principle provides a reference accessible to graduate students and researchers in statistics, pattern classification, machine learning, and data mining, to philosophers interested in the foundations of statistics, and to researchers in other applied sciences that involve model selection.

Deep Learning - Ian Goodfellow 2016-11-10
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.

Probability & Statistics with R for Engineers and Scientists - Michael Akritas 2018-03-21
This title is part of the Pearson Modern Classics series. Pearson Modern Classics are acclaimed titles at a value price. Please visit www.pearsonhighered.com/math-classics-series for a complete list of titles. This text grew out of the author's notes for a course that he has taught for many years to a diverse group of undergraduates. The early introduction to the major concepts engages students immediately, which helps them see the big picture, and sets an appropriate tone for the course. In subsequent chapters, these topics are revisited, developed, and formalized, but the early introduction helps students build a true understanding of the concepts. The text utilizes the statistical software R, which is both widely used and freely available (thanks to the Free Software Foundation). However, in contrast with other books for the intended audience, this book by Akritas emphasizes not only the interpretation of software output, but also the generation of this output. Applications are diverse and relevant, and come from a variety of fields.

Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow - Aur lien G ron 2019-09-05
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 Aur lien G ron 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

The Lion Way - Roberto Battiti 2014-02-21

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.

Data Structures Using C++ - D. S. Malik 2009-07-31

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.

AI and Machine Learning for Coders - Laurence Moroney 2020-10-01

If you're looking to make a career move from programmer to AI specialist, this is the ideal place to start. Based on Laurence Moroney's extremely successful AI courses, this introductory book provides a hands-on, code-first approach to help you build confidence while you learn key topics. You'll understand how to implement the most common scenarios in machine learning, such as computer vision, natural language processing (NLP), and sequence modeling for web, mobile, cloud, and embedded runtimes. Most books on machine learning begin with a daunting amount of advanced math. This guide is built on practical lessons that let you work directly with the code. You'll learn: How to build models with TensorFlow using skills that employers desire The basics of machine learning by working with code samples How to implement computer vision, including feature detection in images How to use NLP to tokenize and sequence words and sentences Methods for embedding models in Android and iOS How to serve models over the web and in the cloud with TensorFlow Serving

Probably Approximately Correct - Leslie Valiant 2013-06-04

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.

From Big Data to Smart Data - Fernando Iafrate 2015-03-30

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.

Advances in Neural Information Processing Systems 10 - Michael I. Jordan 1998

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.

Learning to Learn - Sebastian Thrun 2012-12-06

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.

Machine Learning - Stephen Marsland 2011-03-23

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