

Encyclopedia Of Machine Learning And Data Mining

This authoritative reference work will provide readers with a complete overview of artificial intelligence (AI), including its historic development and current status; existing and projected AI applications; and present and potential future impact on the United States and the world. Some people believe that artificial intelligence (AI) will revolutionize modern life in ways that improve human existence. Others say that the promise of AI is overblown. Still others contend that AI applications could pose a grave threat to the economic security of millions of people by taking their jobs and otherwise rendering them "obsolete"—or, even worse, that AI could actually spell the end of the human race. This volume will help users understand the reasons AI development has both spirited defenders and alarmed critics; explain theories and innovations like Moore's Law, mindcloning, and Technological Singularity that drive AI research and debate; and give readers the information they need to make their own informed judgment about the promise and peril of this technology. All of this coverage is presented using language and terminology accessible to a lay audience. Introduction explaining the historical evolution of AI Chronology of important AI-related events Authoritative entries on leading pioneers, entrepreneurs, and thinkers; AI concepts and theories; AI's potential impact on different facets of society; and major movies and other cultural touchstones exploring AI technology

Feature engineering plays a vital role in big data analytics. Machine learning and data mining algorithms cannot work without data. Little can be achieved if there are few features to represent the underlying data objects, and the quality of results of those algorithms largely depends on the quality of the available features. Feature Engineering for Machine Learning and Data Analytics provides a comprehensive introduction to feature engineering, including feature generation, feature extraction, feature transformation, feature selection, and feature analysis and evaluation. The book presents key concepts, methods, examples, and applications, as well as chapters on feature engineering for major data types such as texts, images, sequences, time series, graphs, streaming data, software engineering data, Twitter data, and social media data. It also contains generic feature generation approaches, as well as methods for generating tried-and-tested, hand-crafted, domain-specific features. The first chapter defines the concepts of features and feature engineering, offers an overview of the book, and provides pointers to topics not covered in this book. The next six chapters are devoted to feature engineering, including feature generation for specific data types. The subsequent four chapters cover generic approaches for feature engineering, namely feature selection, feature transformation based feature engineering, deep learning based feature engineering, and pattern based feature generation and engineering. The last three chapters discuss feature engineering for social bot detection, software management, and Twitter-based applications respectively. This book can be used as a reference for data analysts, big data scientists, data preprocessing workers, project managers, project developers, prediction modelers, professors, researchers, graduate students, and upper level undergraduate students. It can also be used as the primary text for courses on feature engineering, or as a supplement for courses on machine learning, data mining, and big data analytics.

The fundamental mathematical tools needed to understand machine learning include linear algebra, analytic geometry, matrix decompositions, vector calculus, optimization, probability and statistics. These topics are traditionally taught in disparate courses, making it hard for data science or computer science students, or professionals, to efficiently learn the mathematics. This self-contained textbook bridges the gap between mathematical and machine learning texts, introducing the mathematical concepts with a minimum of prerequisites. It uses these concepts to derive four central machine learning methods: linear regression, principal component analysis, Gaussian mixture

models and support vector machines. For students and others with a mathematical background, these derivations provide a starting point to machine learning texts. For those learning the mathematics for the first time, the methods help build intuition and practical experience with applying mathematical concepts. Every chapter includes worked examples and exercises to test understanding. Programming tutorials are offered on the book's web site.

Artificial Intelligence in the Age of Neural Networks and Brain Computing demonstrates that existing disruptive implications and applications of AI is a development of the unique attributes of neural networks, mainly machine learning, distributed architectures, massive parallel processing, black-box inference, intrinsic nonlinearity and smart autonomous search engines. The book covers the major basic ideas of brain-like computing behind AI, provides a framework to deep learning, and launches novel and intriguing paradigms as future alternatives. The success of AI-based commercial products proposed by top industry leaders, such as Google, IBM, Microsoft, Intel and Amazon can be interpreted using this book. Developed from the 30th anniversary of the International Neural Network Society (INNS) and the 2017 International Joint Conference on Neural Networks (IJCNN) Authored by top experts, global field pioneers and researchers working on cutting-edge applications in signal processing, speech recognition, games, adaptive control and decision-making Edited by high-level academics and researchers in intelligent systems and neural networks

Nowadays, the degree and scale of flood hazards has been massively increasing as a result of the changing climate, and large-scale floods jeopardize lives and properties, causing great economic losses, in the inundation-prone areas of the world. Early flood warning systems are promising countermeasures against flood hazards and losses. A collaborative assessment according to multiple disciplines, comprising hydrology, remote sensing, and meteorology, of the magnitude and impacts of flood hazards on inundation areas significantly contributes to model the integrity and precision of flood forecasting. Methodologically oriented countermeasures against flood hazards may involve the forecasting of reservoir inflows, river flows, tropical cyclone tracks, and flooding at different lead times and/or scales. Analyses of impacts, risks, uncertainty, resilience, and scenarios coupled with policy-oriented suggestions will give information for flood hazard mitigation. Emerging advances in computing technologies coupled with big-data mining have boosted data-driven applications, among which Machine Learning technology, with its flexibility and scalability in pattern extraction, has modernized not only scientific thinking but also predictive applications. This book explores recent Machine Learning advances on flood forecast and management in a timely manner and presents interdisciplinary approaches to modelling the complexity of flood hazards-related issues, with contributions to integrative solutions from a local, regional or global perspective.

This encyclopedia will be an essential resource for our times, reflecting the fact that we currently are living in an expanding data-driven world. Technological advancements and other related trends are contributing to the production of an astoundingly large and exponentially increasing collection of data and information, referred to in popular vernacular as “Big Data.” Social media and crowdsourcing platforms and various applications ? “apps” ? are producing reams of information ? from the instantaneous transactions and input of millions and millions of people around the globe. The Internet-of-Things (IoT), which is expected to comprise tens of billions of objects by the end of this decade, is actively sensing real-time intelligence on nearly every aspect of our lives and environment. The Global Positioning System (GPS) and other location-aware technologies are producing data that is specific down to particular latitude and longitude coordinates and seconds of the day. Large-scale instruments, such as the Large Hadron Collider (LHC), are collecting massive amounts of data on our planet and even distant corners of the visible universe. Digitization is being used to convert large collections of documents from print to digital format, giving rise to large

archives of unstructured data. Innovations in technology, in the areas of Cloud and molecular computing, Artificial Intelligence/Machine Learning, and Natural Language Processing (NLP), to name only a few, also are greatly expanding our capacity to store, manage, and process Big Data. In this context, the Encyclopedia of Big Data is being offered in recognition of a world that is rapidly moving from gigabytes to terabytes to petabytes and beyond. While indeed large data sets have long been around and in use in a variety of fields, the era of Big Data in which we now live departs from the past in a number of key respects and with this departure comes a fresh set of challenges and opportunities that cut across and affect multiple sectors and disciplines, and the public at large. With expanded analytical capacities at hand, Big Data is now being used for scientific inquiry and experimentation in nearly every (if not all) disciplines, from the social sciences to the humanities to the natural sciences, and more. Moreover, the use of Big Data has been well established beyond the Ivory Tower. In today's economy, businesses simply cannot be competitive without engaging Big Data in one way or another in support of operations, management, planning, or simply basic hiring decisions. In all levels of government, Big Data is being used to engage citizens and to guide policy making in pursuit of the interests of the public and society in general. Moreover, the changing nature of Big Data also raises new issues and concerns related to, for example, privacy, liability, security, access, and even the veracity of the data itself. Given the complex issues attending Big Data, there is a real need for a reference book that covers the subject from a multi-disciplinary, cross-sectoral, comprehensive, and international perspective. The Encyclopedia of Big Data will address this need and will be the first of such reference books to do so. Featuring some 500 entries, from "Access" to "Zillow," the Encyclopedia will serve as a fundamental resource for researchers and students, for decision makers and leaders, and for business analysts and purveyors. Developed for those in academia, industry, and government, and others with a general interest in Big Data, the encyclopedia will be aimed especially at those involved in its collection, analysis, and use. Ultimately, the Encyclopedia of Big Data will provide a common platform and language covering the breadth and depth of the topic for different segments, sectors, and disciplines.

This comprehensive encyclopedia, in A-Z format, provides easy access to relevant information for those seeking entry into any aspect within the broad field of Machine Learning. Most of the entries in this preeminent work include useful literature references.

Rapid technological advancement has given rise to new ethical dilemmas and security threats, while the development of appropriate ethical codes and security measures fail to keep pace, which makes the education of computer users and professionals crucial. The Encyclopedia of Information Ethics and Security is an original, comprehensive reference source on ethical and security issues relating to the latest technologies. Covering a wide range of themes, this valuable reference tool includes topics such as computer crime, information warfare, privacy, surveillance, intellectual property and education. This encyclopedia is a useful tool for students, academics, and professionals.

Over the past century, educational psychologists and researchers have posited many theories to explain how individuals learn, i.e. how they acquire, organize and deploy knowledge and skills. The 20th century can be considered the century of psychology on learning and related fields of interest (such as motivation, cognition, metacognition etc.) and it is fascinating to see the various mainstreams of learning, remembered and forgotten over the 20th century and note that basic assumptions of early theories survived several paradigm shifts of psychology and epistemology. Beyond folk psychology and its naïve theories of learning, psychological learning theories can be grouped into some basic categories, such as behaviorist learning theories, connectionist learning theories, cognitive learning theories, constructivist learning theories, and social learning theories. Learning theories are not limited to psychology and related fields of interest but rather we can find the topic of learning in various disciplines, such as philosophy and epistemology, education, information science, biology, and – as a

result of the emergence of computer technologies – especially also in the field of computer sciences and artificial intelligence. As a consequence, machine learning struck a chord in the 1980s and became an important field of the learning sciences in general. As the learning sciences became more specialized and complex, the various fields of interest were widely spread and separated from each other; as a consequence, even presently, there is no comprehensive overview of the sciences of learning or the central theoretical concepts and vocabulary on which researchers rely. The Encyclopedia of the Sciences of Learning provides an up-to-date, broad and authoritative coverage of the specific terms mostly used in the sciences of learning and its related fields, including relevant areas of instruction, pedagogy, cognitive sciences, and especially machine learning and knowledge engineering. This modern compendium will be an indispensable source of information for scientists, educators, engineers, and technical staff active in all fields of learning. More specifically, the Encyclopedia provides fast access to the most relevant theoretical terms provides up-to-date, broad and authoritative coverage of the most important theories within the various fields of the learning sciences and adjacent sciences and communication technologies; supplies clear and precise explanations of the theoretical terms, cross-references to related entries and up-to-date references to important research and publications. The Encyclopedia also contains biographical entries of individuals who have substantially contributed to the sciences of learning; the entries are written by a distinguished panel of researchers in the various fields of the learning sciences.

Summary Machine Learning in Action is unique book that blends the foundational theories of machine learning with the practical realities of building tools for everyday data analysis. You'll use the flexible Python programming language to build programs that implement algorithms for data classification, forecasting, recommendations, and higher-level features like summarization and simplification. About the Book A machine is said to learn when its performance improves with experience. Learning requires algorithms and programs that capture data and ferret out the interesting or useful patterns. Once the specialized domain of analysts and mathematicians, machine learning is becoming a skill needed by many. Machine Learning in Action is a clearly written tutorial for developers. It avoids academic language and takes you straight to the techniques you'll use in your day-to-day work. Many (Python) examples present the core algorithms of statistical data processing, data analysis, and data visualization in code you can reuse. You'll understand the concepts and how they fit in with tactical tasks like classification, forecasting, recommendations, and higher-level features like summarization and simplification. Readers need no prior experience with machine learning or statistical processing. Familiarity with Python is helpful. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book. What's Inside A no-nonsense introduction Examples showing common ML tasks Everyday data analysis Implementing classic algorithms like Apriori and Adaboos Table of Contents PART 1

CLASSIFICATION Machine learning basics Classifying with k-Nearest Neighbors Splitting datasets one feature at a time: decision trees Classifying with probability theory: naïve Bayes Logistic regression Support vector machines Improving classification with the AdaBoost meta algorithm PART 2 FORECASTING NUMERIC VALUES WITH REGRESSION Predicting numeric values: regression Tree-based regression PART 3 UNSUPERVISED LEARNING Grouping unlabeled items using k-means clustering Association analysis with the Apriori algorithm Efficiently finding frequent itemsets with FP-growth PART 4 ADDITIONAL TOOLS Using principal component analysis to simplify data Simplifying data with the singular value decomposition Big data and MapReduce

Encyclopedia of Machine Learning and Data Mining Springer

The goal of machine learning is to program computers to use example data or past experience to solve a given problem. Many successful applications of machine learning exist already, including systems that analyze past sales data to predict customer behavior, optimize robot

behavior so that a task can be completed using minimum resources, and extract knowledge from bioinformatics data. Introduction to Machine Learning is a comprehensive textbook on the subject, covering a broad array of topics not usually included in introductory machine learning texts. Subjects include supervised learning; Bayesian decision theory; parametric, semi-parametric, and nonparametric methods; multivariate analysis; hidden Markov models; reinforcement learning; kernel machines; graphical models; Bayesian estimation; and statistical testing. Machine learning is rapidly becoming a skill that computer science students must master before graduation. The third edition of Introduction to Machine Learning reflects this shift, with added support for beginners, including selected solutions for exercises and additional example data sets (with code available online). Other substantial changes include discussions of outlier detection; ranking algorithms for perceptrons and support vector machines; matrix decomposition and spectral methods; distance estimation; new kernel algorithms; deep learning in multilayered perceptrons; and the nonparametric approach to Bayesian methods. All learning algorithms are explained so that students can easily move from the equations in the book to a computer program. The book can be used by both advanced undergraduates and graduate students. It will also be of interest to professionals who are concerned with the application of machine learning methods. Machine learning is an intimidating subject until you know the fundamentals. If you understand basic coding concepts, this introductory guide will help you gain a solid foundation in machine learning principles. Using the R programming language, you'll first start to learn with regression modelling and then move into more advanced topics such as neural networks and tree-based methods. Finally, you'll delve into the frontier of machine learning, using the caret package in R. Once you develop a familiarity with topics such as the difference between regression and classification models, you'll be able to solve an array of machine learning problems. Author Scott V. Burger provides several examples to help you build a working knowledge of machine learning. Explore machine learning models, algorithms, and data training Understand machine learning algorithms for supervised and unsupervised cases Examine statistical concepts for designing data for use in models Dive into linear regression models used in business and science Use single-layer and multilayer neural networks for calculating outcomes Look at how tree-based models work, including popular decision trees Get a comprehensive view of the machine learning ecosystem in R Explore the powerhouse of tools available in R's caret package

This book is your ultimate Machine Learning resource. Here you will find the most up-to-date information, facts, quotes and much more. In easy to read chapters, with extensive references and links to get you to know all there is to know about Machine Learning's whole picture right away. Get countless Machine Learning facts right at your fingertips with this essential resource. The Machine Learning Handbook is the single and largest Machine Learning reference book. This compendium of information is the authoritative source for all your entertainment, reference, and learning needs. It will be your go-to source for any Machine Learning questions. A mind-tickling encyclopedia on Machine Learning, a treat in its entirety and an oasis of learning about what you don't yet know...but are glad you found. The Machine Learning Handbook will answer all of your needs, and much more.

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.

As the age of Big Data emerges, it becomes necessary to take the five dimensions of Big Data- volume, variety, velocity, volatility, and veracity- and focus these dimensions towards one critical emphasis - value. The Encyclopedia of Business Analytics and Optimization confronts the challenges of information retrieval in the age of Big Data by exploring recent advances in the areas of knowledge management, data visualization, interdisciplinary communication, and others. Through its critical approach and practical application, this book will be a must-have reference for any professional, leader, analyst, or manager interested in making the most of the knowledge resources at their disposal.

Encyclopedia of Bioinformatics and Computational Biology: ABC of Bioinformatics combines elements of computer science, information technology, mathematics, statistics and biotechnology, providing the methodology and in silico solutions to mine biological data and processes. The book covers Theory, Topics and Applications, with a special focus on Integrative -omics and Systems Biology. The theoretical, methodological underpinnings of BCB, including phylogeny are covered, as are more current areas of focus, such as translational bioinformatics, cheminformatics, and environmental informatics. Finally, Applications provide guidance for commonly asked questions. This major reference work spans basic and cutting-edge methodologies authored by leaders in the field, providing an invaluable resource for students, scientists, professionals in research institutes, and a broad swath of researchers in biotechnology and the biomedical and pharmaceutical industries. Brings together information from computer science, information technology, mathematics, statistics and biotechnology Written and reviewed by leading experts in the field, providing a unique and authoritative resource Focuses on the main theoretical and methodological concepts before expanding on specific topics and applications Includes interactive images, multimedia tools and crosslinking to further resources and databases

"This book examines current, state-of-the-art research in the areas of data science, machine learning, data mining, optimization, artificial intelligence, statistics, and the interactions, linkages, and applications of knowledge-based business with information systems"--

"This book is a comprehensive and in-depth reference to the most recent developments in the field covering theoretical developments, techniques, technologies, among others"--Provided by publisher.

This friendly and accessible guide to AI theory and programming in Python requires no maths or data science background. Key Features Roll up your sleeves and start programming AI models No math, data science, or machine

learning background required Packed with hands-on examples, illustrations, and clear step-by-step instructions 5 hands-on working projects put ideas into action and show step-by-step how to build intelligent software Book Description AI is changing the world – and with this book, anyone can start building intelligent software! Through his best-selling video courses, Hadelin de Ponteves has taught hundreds of thousands of people to write AI software. Now, for the first time, his hands-on, energetic approach is available as a book. Taking a graduated approach that starts with the basics before easing readers into more complicated formulas and notation, Hadelin helps you understand what you really need to build AI systems with reinforcement learning and deep learning. Five full working projects put the ideas into action, showing step-by-step how to build intelligent software using the best and easiest tools for AI programming: Google Colab Python TensorFlow Keras PyTorch AI Crash Course teaches everyone to build an AI to work in their applications. Once you've read this book, you're only limited by your imagination. What you will learn Master the key skills of deep learning, reinforcement learning, and deep reinforcement learning Understand Q-learning and deep Q-learning Learn from friendly, plain English explanations and practical activities Build fun projects, including a virtual-self-driving car Use AI to solve real-world business problems and win classic video games Build an intelligent, virtual robot warehouse worker Who this book is for If you want to add AI to your skillset, this book is for you. It doesn't require data science or machine learning knowledge. Just maths basics (high school level).

Table of contents

There are more than one billion documents on the Web, with the count continually rising at a pace of over one million new documents per day. As information increases, the motivation and interest in data warehousing and mining research and practice remains high in organizational interest. The Encyclopedia of Data Warehousing and Mining, Second Edition, offers thorough exposure to the issues of importance in the rapidly changing field of data warehousing and mining. This essential reference source informs decision makers, problem solvers, and data mining specialists in business, academia, government, and other settings with over 300 entries on theories, methodologies, functionalities, and applications.

Feature engineering is a crucial step in the machine-learning pipeline, yet this topic is rarely examined on its own. With this practical book, you'll learn techniques for extracting and transforming features—the numeric representations of raw data—into formats for machine-learning models. Each chapter guides you through a single data problem, such as how to represent text or image data. Together, these examples illustrate the main principles of feature engineering. Rather than simply teach these principles, authors Alice Zheng and Amanda Casari focus on practical application with exercises throughout the book. The closing chapter brings everything together by tackling a real-world, structured dataset with several feature-engineering techniques. Python packages including numpy, Pandas, Scikit-learn, and Matplotlib are used

in code examples. You'll examine: Feature engineering for numeric data: filtering, binning, scaling, log transforms, and power transforms Natural text techniques: bag-of-words, n-grams, and phrase detection Frequency-based filtering and feature scaling for eliminating uninformative features Encoding techniques of categorical variables, including feature hashing and bin-counting Model-based feature engineering with principal component analysis The concept of model stacking, using k-means as a featurization technique Image feature extraction with manual and deep-learning techniques Updated and revised second edition of the bestselling guide to exploring and mastering the most important algorithms for solving complex machine learning problems Key Features Updated to include new algorithms and techniques Code updated to Python 3.8 & TensorFlow 2.x New coverage of regression analysis, time series analysis, deep learning models, and cutting-edge applications Book Description Mastering Machine Learning Algorithms, Second Edition helps you harness the real power of machine learning algorithms in order to implement smarter ways of meeting today's overwhelming data needs. This newly updated and revised guide will help you master algorithms used widely in semi-supervised learning, reinforcement learning, supervised learning, and unsupervised learning domains. You will use all the modern libraries from the Python ecosystem – including NumPy and Keras – to extract features from varied complexities of data. Ranging from Bayesian models to the Markov chain Monte Carlo algorithm to Hidden Markov models, this machine learning book teaches you how to extract features from your dataset, perform complex dimensionality reduction, and train supervised and semi-supervised models by making use of Python-based libraries such as scikit-learn. You will also discover practical applications for complex techniques such as maximum likelihood estimation, Hebbian learning, and ensemble learning, and how to use TensorFlow 2.x to train effective deep neural networks. By the end of this book, you will be ready to implement and solve end-to-end machine learning problems and use case scenarios. What you will learn Understand the characteristics of a machine learning algorithm Implement algorithms from supervised, semi-supervised, unsupervised, and RL domains Learn how regression works in time-series analysis and risk prediction Create, model, and train complex probabilistic models Cluster high-dimensional data and evaluate model accuracy Discover how artificial neural networks work – train, optimize, and validate them Work with autoencoders, Hebbian networks, and GANs Who this book is for This book is for data science professionals who want to delve into complex ML algorithms to understand how various machine learning models can be built. Knowledge of Python programming is required.

Machine learning is a relatively new field, without a unanimous definition. In many ways, actuaries have been machine learners. In both pricing and reserving, but also more recently in capital modelling, actuaries have combined statistical methodology with a deep understanding of the problem at hand and how any solution may affect the company and its customers. One aspect that has, perhaps, not been so well

developed among actuaries is validation. Discussions among actuaries' "preferred methods" were often without solid scientific arguments, including validation of the case at hand. Through this collection, we aim to promote a good practice of machine learning in insurance, considering the following three key issues: a) who is the client, or sponsor, or otherwise interested real-life target of the study? b) The reason for working with a particular data set and a clarification of the available extra knowledge, that we also call prior knowledge, besides the data set alone. c) A mathematical statistical argument for the validation procedure.

This authoritative, expanded and updated second edition of Encyclopedia of Machine Learning and Data Mining provides easy access to core information for those seeking entry into any aspect within the broad field of Machine Learning and Data Mining. A paramount work, its 800 entries - about 150 of them newly updated or added - are filled with valuable literature references, providing the reader with a portal to more detailed information on any given topic. Topics for the Encyclopedia of Machine Learning and Data Mining include Learning and Logic, Data Mining, Applications, Text Mining, Statistical Learning, Reinforcement Learning, Pattern Mining, Graph Mining, Relational Mining, Evolutionary Computation, Information Theory, Behavior Cloning, and many others. Topics were selected by a distinguished international advisory board. Each peer-reviewed, highly-structured entry includes a definition, key words, an illustration, applications, a bibliography, and links to related literature. The entries are expository and tutorial, making this reference a practical resource for students, academics, or professionals who employ machine learning and data mining methods in their projects. Machine learning and data mining techniques have countless applications, including data science applications, and this reference is essential for anyone seeking quick access to vital information on the topic.

The Encyclopedia of Big Data Technologies provides researchers, educators, students and industry professionals with a comprehensive authority over the most relevant Big Data Technology concepts. With over 300 articles written by worldwide subject matter experts from both industry and academia, the encyclopedia covers topics such as big data storage systems, NoSQL database, cloud computing, distributed systems, data processing, data management, machine learning and social technologies, data science. Each peer-reviewed, highly structured entry provides the reader with basic terminology, subject overviews, key research results, application examples, future directions, cross references and a bibliography. The entries are expository and tutorial, making this reference a practical resource for students, academics, or professionals. In addition, the distinguished, international editorial board of the encyclopedia consists of well-respected scholars, each developing topics based upon their expertise.

A thought-provoking and wide-ranging exploration of machine learning and the race to build computer intelligences as flexible as our own In the world's top research labs and universities, the race is on to invent the ultimate learning algorithm: one capable of discovering any knowledge from data, and doing anything we want, before we even ask. In The Master Algorithm, Pedro Domingos lifts the veil to give us a peek inside the learning machines that power Google, Amazon, and your smartphone. He assembles a blueprint for the future universal learner--the Master Algorithm--and discusses what it will mean for business, science, and society. If data-ism is today's philosophy, this book is its bible.

"This set of books represents a detailed compendium of authoritative, research-based entries that define the contemporary state of knowledge on technology"--Provided by publisher.

A hands-on approach to tasks and techniques in data stream mining and real-time analytics, with examples in MOA, a popular freely available open-source software framework. Today many information sources—including sensor networks, financial markets, social networks,

and healthcare monitoring—are so-called data streams, arriving sequentially and at high speed. Analysis must take place in real time, with partial data and without the capacity to store the entire data set. This book presents algorithms and techniques used in data stream mining and real-time analytics. Taking a hands-on approach, the book demonstrates the techniques using MOA (Massive Online Analysis), a popular, freely available open-source software framework, allowing readers to try out the techniques after reading the explanations. The book first offers a brief introduction to the topic, covering big data mining, basic methodologies for mining data streams, and a simple example of MOA. More detailed discussions follow, with chapters on sketching techniques, change, classification, ensemble methods, regression, clustering, and frequent pattern mining. Most of these chapters include exercises, an MOA-based lab session, or both. Finally, the book discusses the MOA software, covering the MOA graphical user interface, the command line, use of its API, and the development of new methods within MOA. The book will be an essential reference for readers who want to use data stream mining as a tool, researchers in innovation or data stream mining, and programmers who want to create new algorithms for MOA.

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.

The most comprehensive book on the engineering aspects of building reliable AI systems. "If you intend to use machine learning to solve business problems at scale, I'm delighted you got your hands on this book." -Cassie Kozyrkov, Chief Decision Scientist at Google

"Foundational work about the reality of building machine learning models in production." -Karolis Urbonas, Head of Machine Learning and Science at Amazon

As computer and space technologies have been developed, geoscience information systems (GIS) and remote sensing (RS) technologies, which deal with the geospatial information, have been rapidly maturing. Moreover, over the last few decades, machine learning techniques including artificial neural network (ANN), deep learning, decision tree, and support vector machine (SVM) have been successfully applied to geospatial science and engineering research fields. The machine learning techniques have been widely applied to GIS and RS research fields and have recently produced valuable results in the areas of geoscience, environment, natural hazards, and natural resources. This book is a collection representing novel contributions detailing machine learning techniques as applied to geoscience information systems and remote sensing.

A very complete encyclopedia for Python and Machine Learning. First Volume

Getting numbers is easy; getting numbers you can trust is hard. This practical guide by experimentation leaders at Google, LinkedIn, and Microsoft will teach you how to accelerate innovation using trustworthy online controlled experiments, or A/B tests. Based on practical experiences at companies that each run more than 20,000 controlled experiments a year, the authors share examples, pitfalls, and advice for students and industry professionals getting started with experiments, plus deeper dives into advanced topics for practitioners who want to improve the way they make data-driven decisions. Learn how to

- Use the scientific method to evaluate hypotheses using controlled experiments
- Define key metrics and ideally an Overall Evaluation Criterion
- Test for trustworthiness of the results and alert experimenters to violated assumptions
- Build a scalable platform that lowers the marginal cost of experiments close to zero
- Avoid pitfalls like carryover

effects and Twyman's law • Understand how statistical issues play out in practice.

A modern, comprehensive 2010 overview providing an easy introduction for applied scientists who are not versed in mathematics.

This book is dedicated to intelligent systems of broad-spectrum application, such as personal and social biosafety or use of intelligent sensory micro-nanosystems such as "e-nose", "e-tongue" and "e-eye". In addition to that, effective acquiring information, knowledge management and improved knowledge transfer in any media, as well as modeling its information content using meta-and hyper heuristics and semantic reasoning all benefit from the systems covered in this book. Intelligent systems can also be applied in education and generating the intelligent distributed eLearning architecture, as well as in a large number of technical fields, such as industrial design, manufacturing and utilization, e.g., in precision agriculture, cartography, electric power distribution systems, intelligent building management systems, drilling operations etc. Furthermore, decision making using fuzzy logic models, computational recognition of comprehension uncertainty and the joint synthesis of goals and means of intelligent behavior biosystems, as well as diagnostic and human support in the healthcare environment have also been made easier.

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.

The Encyclopedia of Tourism Management and Marketing is, quite simply, the definitive reference work in the field. This is the largest tourism management and marketing ontology that has ever been put together and offers a holistic examination of this interdisciplinary field. Bringing together a wealth of expertise, a team of international authors from all parts of the world shed light on the current state of tourism research and practice around the globe and provide unique insights into the field. Carefully curated by leading tourism scholar Dimitrios Buhalis, the Encyclopedia is an invaluable resource for academics, students and practitioners and provides the ideal starting point for any research journey. The concise entries present an accessible and condensed overview of each topic and the selected references that follow each entry suggest directions for further detailed exploration. Key Features: - Over 1000 entries - Entries organized alphabetically for ease of navigation - Fully cross-referenced - Concise, structured entries by the world's foremost scholars in tourism - Selected references for further study -

Inclusive global authorship team. Contact us for a quote.

[Copyright: e46adf9705258766bf40a2cd9e63db95](#)