Introduction To Machine Learning Ethem Alpaydin

Machine Learning in Non-Stationary EnvironmentsAn Introduction to Natural ComputationReliable Face Recognition MethodsIntroduction to Machine LearningLearning Kernel ClassifiersMathematics for Machine LearningNeural Networks for Pattern RecognitionApplied Machine LearningBayesian Reasoning and Machine LearningAI Crash CourseIntroduction to Machine LearningBioinformaticsMachine LearningDeep LearningAn Introduction to Machine LearningMachine Learning for Data StreamsAn Elementary Introduction to Statistical Learning TheoryLearning with KernelsThe Decision Maker's Handbook to Data ScienceLearning scikit-learn: Machine Learning in PythonMachine LearningFoundations of Machine LearningApplied Machine LearningMachine Learning For DummiesMachine Learning RefinedElements of Causal InferenceMachine LearningIntroduction to Machine LearningFoundations of Machine LearningData ScienceAn Introduction to Machine LearningMachine Learning, Second Edition: A Probabilistic PerspectiveDeep LearningMachine Learning for HackersMachine LearningIntroduction to Machine LearningIntroduction to Natural Language ProcessingIntroduction to Machine LearningA Concise Introduction to Machine LearningReinforcement Learning

Machine Learning in Non-Stationary Environments

An overview of the theory and application of kernel classification methods. Linear classifiers in kernel spaces have emerged as a major topic within the field of machine learning. The kernel technique takes the linear classifier—a limited, but well-established and comprehensively studied model—and extends its applicability to a wide range of nonlinear pattern-recognition tasks such as natural language processing, machine vision, and biological sequence analysis. This book provides the first comprehensive overview of both the theory and algorithms of kernel classifiers, including the most recent developments. It begins by describing the major algorithmic advances: kernel perceptron learning, kernel Fisher discriminants, support vector machines, relevance vector machines, Gaussian processes, and Bayes point machines. Then follows a detailed introduction to learning theory, including VC and PAC-Bayesian theory, data-dependent structural risk minimization, and compression bounds. Throughout, the book emphasizes the interaction between theory and algorithms: how learning algorithms work and why. The book includes many examples, complete pseudo code of the algorithms presented, and an extensive source code library.

An Introduction to Natural Computation

A concise introduction to the emerging field of data science, explaining its evolution, relation to machine learning, current uses, data infrastructure issues, and ethical challenges. The goal of data science is to improve decision making through the analysis of data. Today data science determines the ads we see online, the books and movies that are recommended to us online, which emails are filtered into our spam folders, and even how much we pay for health insurance. This volume in the MIT Press Essential Knowledge series offers a concise introduction to the emerging field of data science, explaining its evolution, current uses, data infrastructure issues, and ethical challenges. It has never been easier for organizations to gather, store, and process data. Use of data science is driven by the rise of big data and social media, the development of high-performance computing, and the emergence of such powerful methods for data analysis and modeling as deep learning. Data science encompasses a set of principles, problem definitions, algorithms, and processes for extracting non-obvious and useful patterns from large datasets. It is closely related to the fields of data mining and machine learning, but broader in scope. This book offers a brief history of the field, introduces fundamental data concepts, and describes the stages in a data science project. It considers data infrastructure and the challenges posed by integrating data from multiple sources, introduces the basics of machine learning, and discusses how to link machine learning expertise with real-world problems. The book also reviews ethical and legal issues, developments in data regulation, and computational approaches to preserving privacy. Finally, it considers the future impact of data science and offers principles for success in data science projects.

Reliable Face Recognition Methods

Theory, algorithms, and applications of machine learning techniques to overcome “covariate shift” non-stationarity. As the power of computing has grown over the past few decades, the field of machine learning has advanced rapidly in both theory and practice. Machine learning methods are usually based on the assumption that the data generation mechanism does not change over time. Yet real-world applications of machine learning, including image recognition, natural language processing, speech recognition, robot control, and bioinformatics, often violate this common assumption. Dealing with non-stationarity is one of modern machine learning's greatest challenges. This book focuses on a specific non-stationary environment known as covariate shift, in which the distributions of inputs (queries) change but the conditional distribution of outputs (answers) is unchanged, and presents machine learning theory, algorithms, and applications to overcome this variety of non-stationarity. After reviewing the state-of-the-art research in the field, the authors discuss topics that include learning under covariate shift, model selection, importance estimation, and active learning. They describe such real world applications of covariate shift adaption as brain-computer interface, speaker identification, and age prediction from facial images. With this book, they aim to encourage future research in machine learning, statistics, and engineering that strives to create truly autonomous learning machines able to learn under non-stationarity.

Introduction to Machine Learning

An accessible introduction to the artificial intelligence technology that enables computer vision, speech recognition, machine translation, and driverless cars. Deep learning is an artificial intelligence technology that enables computer vision, speech recognition in mobile phones, machine translation, AI games, driverless cars, and other applications. When we use consumer products from Google, Microsoft, Facebook, Apple, or Baidu, we are often interacting with a deep learning system. In this volume in the MIT Press Essential Knowledge series, computer scientist John Kelleher offers an accessible and concise but comprehensive introduction to the fundamental technology at the heart of the artificial intelligence revolution. Kelleher explains that deep learning enables data-driven decisions by identifying and extracting patterns from large datasets; its ability to learn from complex data makes deep learning ideally suited to take advantage of the rapid growth in big data and computational power. Kelleher also explains some of the basic concepts in deep learning, presents a history of advances in the field, and discusses the current state of the art. He describes the most important deep learning architectures, including autoencoders, recurrent neural networks, and long short-term networks, as well as such recent developments as Generative Adversarial Networks and capsule networks. He also provides a comprehensive (and comprehensible) introduction to the two fundamental algorithms in deep learning: gradient descent and backpropagation. Finally, Kelleher considers the future of deep learning—major trends, possible developments, and significant challenges.

Learning Kernel Classifiers

A thought-provoking look at statistical learning theory and its role in understanding human learning and inductive reasoning A joint endeavor from leading researchers in the fields of philosophy and electrical engineering, An Elementary Introduction to Statistical Learning Theory is a comprehensive and accessible primer on the rapidly evolving fields of statistical pattern recognition and statistical learning theory. Explaining these areas at a level and in a way that is not often found in other books on the topic, the authors present the basic theory behind contemporary machine learning and uniquely utilize its foundations as a framework for philosophical thinking about inductive inference. Promoting the fundamental goal of statistical learning, knowing what is achievable and what is not, this book demonstrates the value of a systematic methodology when used along with the needed techniques for evaluating the performance of a learning system. First, an introduction to machine learning is presented that includes brief discussions of applications such as image recognition, speech recognition, medical diagnostics, and statistical arbitrage. To enhance accessibility, two chapters on relevant aspects of probability theory are provided. Subsequent chapters feature coverage of topics such as the pattern recognition problem, optimal Bayes decision rule, the nearest neighbor rule, kernel rules, neural networks, support vector machines, and boosting. Appendices throughout the book explore the relationship between the discussed material and related topics from mathematics, philosophy, psychology, and statistics, drawing insightful connections between problems in these areas and statistical learning theory. All chapters conclude with a summary section, a set of practice questions, and a reference sections that supplies historical notes and additional resources for further study. An Elementary Introduction to Statistical

Learning Theory is an excellent book for courses on statistical learning theory, pattern recognition, and machine learning at the upper-undergraduate and graduate levels. It also serves as an introductory reference for researchers and practitioners in the fields of engineering, computer science, philosophy, and cognitive science that would like to further their knowledge of the topic.

Mathematics for Machine Learning

A comprehensive introduction to Support Vector Machines and related kernel methods. In the 1990s, a new type of learning algorithm was developed, based on results from statistical learning theory: the Support Vector Machine (SVM). This gave rise to a new class of theoretically elegant learning machines that use a central concept of SVMs---kernels--for a number of learning tasks. Kernel machines provide a modular framework that can be adapted to different tasks and domains by the choice of the kernel function and the base algorithm. They are replacing neural networks in a variety of fields, including engineering, information retrieval, and bioinformatics. Learning with Kernels provides an introduction to SVMs and related kernel methods. Although the book begins with the basics, it also includes the latest research. It provides all of the concepts necessary to enable a reader equipped with some basic mathematical knowledge to enter the world of machine learning using theoretically well-founded yet easy-to-use kernel algorithms and to understand and apply the powerful algorithms that have been developed over the last few years.

Neural Networks for Pattern Recognition

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

Applied Machine Learning

Your no-nonsense guide to making sense of machine learning Machine learning can be a mind-boggling concept for the masses, but those who are in the trenches of computer programming know just how invaluable it is. Without machine learning, fraud detection, web search results, real-time ads on web pages, credit scoring, automation, and email spam filtering wouldn't be possible, and this is only showcasing just a few of its capabilities. Written by two data science experts, Machine Learning For Dummies offers a much-needed entry point for anyone looking to use machine learning to accomplish practical tasks. Covering the entry-level topics needed to get you familiar with the basic concepts of machine learning, this guide quickly helps you make sense of the programming languages and tools you need to turn machine learning-based tasks into a reality. Whether you're maddened by the math behind machine learning, apprehensive about AI, perplexed by preprocessing data—or anything in between—this guide makes it easier to understand and implement machine learning seamlessly. Grasp how day-to-day activities are powered by machine learning Learn to 'speak' certain languages, such as Python and R, to teach machines to perform pattern-oriented tasks and data analysis Learn to code in R using R Studio Find out how to code in Python using Anaconda Dive into this complete beginner's guide so you are armed with all you need to know about machine learning!

Bayesian Reasoning and Machine Learning

This book provides a comprehensive introduction to the computational material that forms the underpinnings of the currently evolving set of brain models. It is now clear that the brain is unlikely to be understood without recourse to computational theories. The theme of An Introduction to Natural Computation is that ideas from diverse areas such as neuroscience, information theory, and optimization theory have recently been extended in ways that make them useful for describing the brains programs. This book provides a comprehensive introduction to the computational material that forms the underpinnings of the currently evolving set of brain models. It stresses the broad spectrum of learning models--ranging from neural network learning through reinforcement learning to genetic learning--and situates the various models in their appropriate neural context. To write about models of the brain before the brain is fully understood is a delicate matter. Very detailed models of the neural circuitry risk losing track of the task the brain is trying to solve. At the other extreme, models that represent cognitive constructs can be so abstract that they lose all relationship to neurobiology. An Introduction to Natural Computation takes the middle ground and stresses the computational task while staying near the neurobiology.

AI Crash Course

`Readers will emerge with a rigorous statistical grounding in the theory of how to construct and train neural networks in pattern recognition' New Scientist

Introduction to Machine Learning

Pierre Baldi and Soren Brunak present the key machine learning approaches and apply them to the computational problems encountered in the analysis of biological data. The book is aimed at two types of researchers and students. First are the biologists and biochemists who need to understand new data-driven algorithms, such as neural networks and hidden Markov models, in the context of biological sequences and their molecular structure and function. Second are those with a primary background in physics, mathematics, statistics, or computer science who need to know more about specific applications in molecular biology.

Bioinformatics

An intuitive approach to machine learning covering key concepts, real-world applications, and practical Python coding exercises.

Machine Learning

The emphasis of the book is on the question of Why – only if why an algorithm is successful is understood, can it be properly applied, and the results trusted. Algorithms are often taught side by side without showing the similarities and differences between them. This book addresses the commonalities, and aims to give a thorough and in-depth treatment and develop intuition, while remaining concise. This useful reference should be an essential on the bookshelves of anyone employing machine learning techniques.

Deep Learning

This graduate-level textbook introduces fundamental concepts and methods in machine learning. It describes several important modern algorithms, provides the theoretical underpinnings of these algorithms, and illustrates key aspects for their application. The authors aim to present novel theoretical tools and concepts while giving concise proofs even for relatively advanced topics. Foundations of Machine Learning fills the need for a general textbook that also offers theoretical details and an emphasis on proofs. Certain topics that are often treated with insufficient attention are discussed in more detail here; for example, entire chapters are devoted to regression, multi-class classification, and ranking. The first three chapters lay the theoretical foundation for what follows, but each remaining chapter is mostly self-contained. The appendix offers a concise probability review, a short introduction to convex optimization, tools for concentration bounds, and several basic properties of matrices and norms used in the book.The book is intended for graduate students and researchers in machine learning, statistics, and related areas; it can be used either as a textbook or as a reference text for a research seminar.

An Introduction to Machine Learning

A practical introduction perfect for final-year undergraduate and graduate students without a solid background in linear algebra and calculus.

Machine Learning for Data Streams

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

An Elementary Introduction to Statistical Learning Theory

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).

Learning with Kernels

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.

The Decision Maker's Handbook to Data Science

Machine learning methods are now an important tool for scientists, researchers, engineers and students in a wide range of areas. This book is written for people who want to adopt and use the main tools of machine learning, but aren’t necessarily going to want to be machine learning researchers. Intended for students in final year undergraduate or first year graduate computer science programs in machine learning, this textbook is a machine learning toolkit. Applied Machine Learning covers many topics for people who want to use machine learning processes to get things done, with a strong emphasis on using existing tools and packages, rather than writing one’s own code. A companion to the author's Probability and Statistics for Computer Science, this book picks up where the earlier book left off (but also supplies a summary of probability that the reader can use). Emphasizing the usefulness of standard machinery from applied statistics, this textbook gives an overview of the major applied areas in learning, including coverage of:• classification using standard machinery (naive bayes; nearest neighbor; SVM)• clustering and vector quantization (largely as in PSCS)• PCA (largely as in PSCS)• variants of PCA (NIPALS; latent semantic analysis; canonical correlation analysis)• linear regression (largely as in PSCS)• generalized linear models including logistic regression• model selection with Lasso, elasticnet• robustness and m-estimators• Markov chains and HMM’s (largely as in PSCS)• EM in fairly gory detail; long experience teaching this suggests one detailed example is required, which students hate; but once they’ve been through that, the next one is easy• simple graphical models (in the variational inference section)• classification with neural networks, with a particular emphasis onimage classification• autoencoding with neural networks• structure learning

Learning scikit-learn: Machine Learning in Python

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.

Machine Learning

A survey of computational methods for understanding, generating, and manipulating human language, which offers a synthesis of classical representations and algorithms with contemporary machine learning techniques. This textbook provides a technical perspective on natural language processing—methods for building computer software that understands, generates, and manipulates human language. It emphasizes contemporary data-driven approaches, focusing on techniques from supervised and unsupervised machine learning. The first section establishes a foundation in machine learning by building a set of tools that will be used throughout the book and applying them to word-based textual analysis. The second section introduces structured representations of language, including sequences, trees, and graphs. The third section explores different approaches to the representation and analysis of linguistic meaning, ranging from formal logic to neural word embeddings. The final section offers chapter-length treatments of three transformative applications of natural language processing: information extraction, machine translation, and text generation. End-of-chapter exercises include both paper-and-pencil analysis and software implementation. The text synthesizes and distills a broad and diverse research literature, linking contemporary machine learning techniques with the field's linguistic and computational foundations. It is suitable for use in advanced undergraduate and graduate-level courses and as a reference for software engineers and data scientists. Readers should have a background in computer programming and college-level mathematics. After mastering the material presented, students will have the technical skill to build and analyze novel natural language processing systems and to understand the latest research in the field.

Foundations of Machine Learning

Data science is expanding across industries at a rapid pace, and the companies first to adopt best practices will gain a significant advantage. To reap the benefits, decision makers need to have a confident understanding of data science and its application in their organization. It is easy for novices to the subject to feel paralyzed by intimidating buzzwords, but what many don’t realize is that data science is in fact quite multidisciplinary—useful in the hands of business analysts, communications strategists, designers, and more. With the second edition of The Decision Maker’s Handbook to Data Science, you will learn how to think like a veteran data scientist and approach solutions to business problems in an entirely new way. Author Stylianos Kampakis provides you with the expertise and tools required to develop a solid data strategy that is continuously effective. Ethics and legal issues surrounding data collection and algorithmic bias are some common pitfalls that Kampakis helps you avoid, while guiding you on the path to build a thriving data science culture at your organization. This updated and revised second edition, includes plenty of case studies, tools for project assessment, and expanded content for hiring and managing data scientists Data science is a language that everyone at a modern company should understand across departments. Friction in communication arises most often when management does not connect with what a data scientist is doing or how impactful data collection and storage can be for their organization. The Decision Maker’s Handbook to Data Science bridges this gap and readies you for both the present and future of your workplace in this engaging, comprehensive guide. What You Will Learn Understand how data science can be used within your business. Recognize the differences between AI, machine learning, and statistics. Become skilled at thinking like a data scientist, without being one. Discover how to hire and manage data scientists. Comprehend how to build the right environment in order to make your organization data-driven. Who This Book Is For Startup founders, product managers, higher level managers, and any other non-technical decision makers who are thinking to implement data science in their organization and hire data scientists. A secondary audience includes people looking for a soft introduction into the subject of data science.

Applied Machine Learning

The book adopts a tutorial-based approach to introduce the user to Scikit-learn.If you are a programmer who wants to explore machine learning and data-based methods to build intelligent applications and enhance your programming skills, this the book for you. No previous experience with machine-learning algorithms is required.

Machine Learning For Dummies

This text covers all the fundamentals and presents basic theoretical concepts and a wide range of techniques (algorithms) applicable to challenges in our day-to-day lives. The book recognizes that most of the ideas behind machine learning are simple and straightforward. It provides a platform for hands-on experience through self-study machine learning projects. Datasets for some benchmark applications have been explained to encourage the use of algorithms covered in this book. This is a comprehensive text book on machine learning for undergraduates in computer science and all engineering degree programs. Post graduates and research scholars will find it a useful initial exposure to the subject, before they go for highly theoretical depth in the specific areas of their research. For engineers, scientists, business managers and other practitioners, the book will help build the foundations of machine learning.

Machine Learning Refined

A substantially revised fourth edition of a comprehensive textbook, including new coverage of recent advances in deep learning and neural networks. The goal of machine learning is to program computers to use example data or past experience to solve a given problem. Machine learning underlies such exciting new technologies as self-driving cars, speech recognition, and translation applications. This substantially revised fourth edition of a comprehensive, widely used machine learning textbook offers new coverage of recent advances in the field in both theory and practice, including developments in deep learning and neural networks. The book covers a broad array of topics not usually included in introductory machine learning texts, including supervised learning, Bayesian decision theory, parametric methods, semiparametric methods, nonparametric methods, multivariate analysis, hidden Markov models, reinforcement learning, kernel machines, graphical models, Bayesian estimation, and statistical testing. The fourth edition offers a new chapter on deep learning that discusses training, regularizing, and structuring deep neural networks such as convolutional and generative adversarial networks; new material in the chapter on reinforcement learning that covers the use of deep networks, the policy gradient methods, and deep reinforcement learning; new material in the chapter on multilayer perceptrons on autoencoders and the word2vec network; and discussion of a popular method of dimensionality reduction, t-SNE. New appendixes offer background material on linear algebra and optimization. End-of-chapter exercises help readers to apply concepts learned. Introduction to Machine Learning can be used in courses for advanced undergraduate and graduate students and as a reference for professionals.

Elements of Causal Inference

An introductory text in machine learning that gives a unified treatment of methods based on statistics, pattern recognition, neural networks, artificial intelligence, signal processing, control, and data mining.

Machine Learning

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.

Introduction to Machine Learning

A new edition of an introductory text in machine learning that gives a unified treatment of machine learning problems and solutions. 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. The second edition of 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. In order to present a unified treatment of machine learning problems and solutions, it discusses many methods from different fields, including statistics, pattern recognition, neural networks, artificial intelligence, signal processing, control, and data mining. All learning algorithms are explained so that the student can easily move from the equations in the book to a computer program. The text covers such topics as supervised learning, Bayesian decision theory, parametric methods, multivariate methods, multilayer perceptrons, local models, hidden Markov models, assessing and comparing classification algorithms, and reinforcement learning. New to the second edition are chapters on kernel machines, graphical models, and Bayesian estimation; expanded coverage of statistical tests in a chapter on design and analysis of machine learning experiments; case studies available on the Web (with downloadable results for instructors); and many additional exercises. All chapters have been revised and updated. Introduction to Machine Learning can be used by advanced undergraduates and graduate students who have completed courses in computer programming, probability, calculus, and linear algebra. It will also be of interest to engineers in the field who are concerned with the application of machine learning methods.

Foundations of Machine Learning

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.

Data Science

An Introduction to Machine Learning

Just like electricity, Machine Learning will revolutionize our life in many ways – some of which are not even conceivable today. This book provides a thorough conceptual understanding of Machine Learning techniques and algorithms. Many of the mathematical concepts are explained in an intuitive manner. The book starts with an overview of machine learning and the underlying Mathematical and Statistical concepts before moving onto machine learning topics. It gradually builds up the depth, covering many of the present day machine learning algorithms, ending in Deep Learning and Reinforcement Learning algorithms. The book also covers some of the popular Machine Learning applications. The material in this book is agnostic to any specific programming language or hardware so that readers can try these concepts on whichever platforms they are already familiar with. Offers a comprehensive introduction to Machine Learning, while not assuming any prior knowledge of the topic; Provides a complete overview of available techniques and algorithms in conceptual terms, covering various application domains of machine learning; Not tied to any specific software language or hardware implementation.

Machine Learning, Second Edition: A Probabilistic Perspective

The mathematization of causality is a relatively recent development, and has become increasingly important in data science and machine learning. This book offers a self-contained and concise introduction to causal models and how to learn them from data. After explaining the need for causal models and discussing some of the principles underlying causal inference, the book teaches readers how to use causal models: how to compute intervention distributions, how to infer causal models from observational and interventional data, and how causal ideas could be exploited for classical machine learning problems. All of these topics are discussed first in terms of two variables and then in the more general multivariate case. The bivariate case turns out to be a particularly hard problem for causal learning because there are no conditional independences as used by classical methods for solving multivariate cases. The authors consider analyzing statistical asymmetries between cause and effect to be highly instructive, and they report on their decade of intensive research into this problem. The book is accessible to readers with a background in machine learning or statistics, and can be used in graduate courses or as a reference for researchers. The text includes code snippets that can be copied and pasted, exercises, and an appendix with a summary of the most important technical concepts.

Deep Learning

Machine Learning: An Applied Mathematics Introduction covers the essential mathematics behind all of the following topics - K Nearest Neighbours; K Means Clustering; Naïve Bayes Classifier; Regression Methods; Support Vector Machines; Self-Organizing Maps; Decision Trees; Neural Networks; Reinforcement Learning

Machine Learning for Hackers

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

Machine Learning: A Bayesian and Optimization Perspective, 2nd edition, gives a unified perspective on machine learning by covering both pillars of supervised learning, namely regression and classification. The book starts with the basics, including mean square, least squares and maximum likelihood methods, ridge regression, Bayesian decision theory classification, logistic regression, and decision trees. It then progresses to more recent techniques, covering sparse modelling methods, learning in reproducing kernel Hilbert spaces and support vector machines, Bayesian inference with a focus on the EM algorithm and its approximate inference variational versions, Monte Carlo methods, probabilistic graphical models focusing on Bayesian networks, hidden Markov models and particle filtering. Dimensionality reduction and latent variables modelling are also considered in depth. This palette of techniques concludes with an extended chapter on neural networks and deep learning architectures. The book also covers the fundamentals of statistical parameter estimation, Wiener and Kalman filtering, convexity and convex optimization, including a chapter on stochastic approximation and the gradient descent family of algorithms, presenting related online learning techniques as well as concepts and algorithmic versions for distributed optimization. Focusing on the physical reasoning behind the mathematics, without sacrificing rigor, all the various methods and techniques are explained in depth, supported by examples and problems, giving an invaluable resource to the student and researcher for understanding and applying machine learning concepts. Most of the chapters include typical case studies and computer exercises, both in MATLAB and Python. The chapters are written to be as self-contained as possible, making the text suitable for different courses: pattern recognition, statistical/adaptive signal processing, statistical/Bayesian learning, as well as courses on sparse modeling, deep learning, and probabilistic graphical models. New to this edition: Complete re-write of the chapter on Neural Networks and Deep Learning to reflect the latest advances since the 1st edition. The chapter, starting from the basic perceptron and feed-forward neural networks concepts, now presents an in depth treatment of deep networks, including recent optimization algorithms, batch normalization, regularization techniques such as the dropout method, convolutional neural networks, recurrent neural networks, attention mechanisms, adversarial examples and training, capsule networks and generative architectures, such as restricted Boltzman machines (RBMs), variational autoencoders and generative adversarial networks (GANs). Expanded treatment of Bayesian learning to include nonparametric Bayesian methods, with a focus on the Chinese restaurant and the Indian buffet processes. Presents the physical reasoning, mathematical modeling and algorithmic implementation of each method Updates on the latest trends, including sparsity, convex analysis and optimization, online distributed algorithms, learning in RKH spaces, Bayesian inference, graphical and hidden Markov models, particle filtering, deep learning, dictionary learning and latent variables modeling Provides case studies on a variety of topics, including protein folding prediction, optical character recognition, text authorship identification, fMRI data analysis, change point detection, hyperspectral image unmixing, target localization, and more

Introduction to Machine Learning

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.

Introduction to Natural Language Processing

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

Introduction to Machine Learning

This book seeks to comprehensively address the face recognition problem while gaining new insights from complementary fields of endeavor. These include neurosciences, statistics, signal and image processing, computer vision, machine learning and data mining. The book examines the evolution of research surrounding the field to date, explores new directions, and offers specific guidance on the most promising venues for future research and development. The book’s focused approach and its clarity of presentation make this an excellent reference work.

A Concise Introduction to Machine Learning

A substantially revised third edition of a comprehensive textbook that covers a broad range of topics not often included in introductory texts. 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.

Reinforcement Learning

This textbook presents fundamental machine learning concepts in an easy to understand manner by providing practical advice, using straightforward examples, and offering engaging discussions of relevant applications. The main topics include Bayesian classifiers, nearest-neighbor classifiers, linear and polynomial classifiers, decision trees, neural networks, and support vector machines. Later chapters show how to combine these simple tools by way of “boosting,” how to exploit them in more complicated domains, and how to deal with diverse advanced practical issues. One chapter is dedicated to the popular genetic algorithms. This revised edition contains three entirely new chapters on critical topics regarding the pragmatic application of machine learning in industry. The chapters examine multi-label domains, unsupervised learning and its use in deep learning, and logical approaches to induction. Numerous chapters have been expanded, and the presentation of the material has been enhanced. The book contains many new exercises, numerous solved examples, thought-provoking experiments, and computer assignments for independent work.