

## Machine Learning An Algorithmic Perspective Second Edition Chapman Hallcrc Machine Learning Pattern Recognition

Familiarizes machine learning experts with imitation learning, statistical supervised learning theory, and reinforcement learning. It also roboticists and experts in applied artificial intelligence with a broader appreciation for the frameworks and tools available for imitation learning.

Just Enough R! An Interactive Approach to Machine Learning and Analytics presents just enough of the R language, machine learning algorithms, statistical methodology, and analytics for the reader to learn how to find interesting structure in data. The approach might be called "seeing then doing" as it first gives step-by-step explanations using simple, understandable examples of how the various machine learning algorithms work independent of any programming language. This is followed by detailed scripts written in R that apply the algorithms to solve nontrivial problems with real data. The script code is provided, allowing the reader to execute the scripts as they study the explanations given in the text. Features Gets you quickly using R as a problem-solving tool Uses RStudio's integrated development environment Shows how to interface R with SQLite Includes examples using R's Rattle graphical user interface Requires no prior knowledge of R, machine learning, or computer programming Offers over 50 scripts written in R, including several problem-solving templates that, with slight modification, can be used again and again Covers the most popular machine learning techniques, including ensemble-based methods and logistic regression Includes end-of-chapter exercises, many of which can be solved by modifying existing scripts Includes datasets from several areas, including business, health and medicine, and science About the Author Richard J. Roiger is a professor emeritus at Minnesota State University, Mankato, where he taught and performed research in the Computer and Information Science Department for over 30 years.

Algorithmic Culture: How Big Data and Artificial Intelligence are Transforming Everyday Life explores the complex ways in which algorithms and big data, or algorithmic culture, are simultaneously reshaping everyday culture while perpetuating inequality and intersectional discrimination. Contributors situate issues of humanity, identity, and culture in relation to free will, surveillance, capitalism, neoliberalism, consumerism, solipsism, and creativity, offering a critique of the myriad constraints enacted by algorithms. This book argues that consumers are undergoing an ontological overhaul due to the enhanced manipulability and increasingly mandatory nature of algorithms in the market, while also positing that algorithms may help navigate through chaos that is intrinsically present in the market democracy. Ultimately, Algorithmic Culture calls attention to the present-day cultural landscape as a whole as it has been reconfigured and re-

presented by algorithms.

Table of contents

Leading technology scholars examine how networks powered by algorithms are transforming humanity, posing deep questions about power, freedom, and fairness. This title is also available as Open Access on Cambridge Core.

A Proven, Hands-On Approach for Students without a Strong Statistical Foundation Since the best-selling first edition was published, there have been several prominent developments in the field of machine learning, including the increasing work on the statistical interpretations of machine learning algorithms. Unfortunately, computer science students without a strong statistical background often find it hard to get started in this area. Remediating this deficiency, *Machine Learning: An Algorithmic Perspective, Second Edition* helps students understand the algorithms of machine learning. It puts them on a path toward mastering the relevant mathematics and statistics as well as the necessary programming and experimentation. New to the Second Edition Two new chapters on deep belief networks and Gaussian processes Reorganization of the chapters to make a more natural flow of content Revision of the support vector machine material, including a simple implementation for experiments New material on random forests, the perceptron convergence theorem, accuracy methods, and conjugate gradient optimization for the multi-layer perceptron Additional discussions of the Kalman and particle filters Improved code, including better use of naming conventions in Python Suitable for both an introductory one-semester course and more advanced courses, the text strongly encourages students to practice with the code. Each chapter includes detailed examples along with further reading and problems. All of the code used to create the examples is available on the author's website.

Drawing from computer science, statistics, mathematics, and engineering, the multidisciplinary nature of machine learning is underscored by its applicability to areas ranging from finance to biology. This text provides the ideal blend of theory and practical, applicable knowledge.

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.

Machine learning, one of the top emerging sciences, has an extremely broad range of applications. However, many books on the subject provide only a theoretical approach, making it difficult for a newcomer to grasp the subject material. This book provides a more practical approach by explaining the concepts of machine learning algorithms and describing the areas of application for each algorithm, using simple

## Read Free Machine Learning An Algorithmic Perspective Second Edition Chapman Hallcrc Machine Learning Pattern Recognition

practical examples to demonstrate each algorithm and showing how different issues related to these algorithms are applied.

How can machine learning--especially deep neural networks--make a real difference in your organization? This hands-on guide not only provides practical information, but helps you get started building efficient deep learning networks. The authors provide the fundamentals of deep learning--tuning, parallelization, vectorization, and building pipelines--that are valid for any library before introducing the open source DeepLearning4j (DL4J) library for developing production-class workflows. Through real-world examples, you'll learn methods and strategies for training deep network architectures and running deep learning workflows on Spark and Hadoop with DL4J.

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

Introduces machine learning and its algorithmic paradigms, explaining the principles behind automated learning approaches and the considerations underlying their usage.

A detailed and up-to-date introduction to machine learning, presented through the unifying lens of probabilistic modeling and Bayesian decision theory. This book offers a detailed and up-to-date introduction to machine learning (including deep learning) through the unifying lens of probabilistic modeling and Bayesian decision theory. The book covers mathematical background (including linear algebra and optimization), basic supervised learning (including linear and logistic regression and deep neural networks), as well as more advanced topics (including transfer learning and unsupervised learning). End-of-chapter exercises allow students to apply what they have learned, and an appendix covers notation. Probabilistic Machine Learning grew out of the author's 2012 book, Machine Learning: A Probabilistic Perspective. More than just a simple update, this is a completely new book that reflects the dramatic developments in the field since 2012, most notably deep learning. In addition, the new book is accompanied by online Python code, using libraries such as scikit-learn, JAX, PyTorch, and Tensorflow, which can be used to reproduce nearly all the figures; this code can be run inside a web browser using cloud-based notebooks, and provides a practical complement to the theoretical topics discussed in the book. This introductory text will be followed by a sequel that covers more advanced topics, taking the same probabilistic approach.

With an evolutionary advancement of Machine Learning (ML) algorithms, a rapid increase of data volumes and a significant improvement of computation powers, machine learning becomes hot in different applications. However, because of the nature of "black-box" in ML methods, ML still needs to be interpreted to link human and machine learning for transparency and user acceptance of delivered solutions. This edited book addresses such links from the perspectives of visualisation, explanation, trustworthiness and transparency. The book establishes the link between human and machine learning by exploring transparency in machine learning, visual explanation of ML processes, algorithmic explanation of ML models, human cognitive responses in ML-based decision making, human evaluation of machine learning and domain knowledge in transparent ML applications. This is the first book of its kind to systematically understand the current active research activities

## Read Free Machine Learning An Algorithmic Perspective Second Edition Chapman Hallcrc Machine Learning Pattern Recognition

and outcomes related to human and machine learning. The book will not only inspire researchers to passionately develop new algorithms incorporating human for human-centred ML algorithms, resulting in the overall advancement of ML, but also help ML practitioners proactively use ML outputs for informative and trustworthy decision making. This book is intended for researchers and practitioners involved with machine learning and its applications. The book will especially benefit researchers in areas like artificial intelligence, decision support systems and human-computer interaction.

Distributed Systems: An Algorithmic Approach, Second Edition provides a balanced and straightforward treatment of the underlying theory and practical applications of distributed computing. As in the previous version, the language is kept as unobscured as possible—clarity is given priority over mathematical formalism. This easily digestible text: Features significant updates that mirror the phenomenal growth of distributed systems Explores new topics related to peer-to-peer and social networks Includes fresh exercises, examples, and case studies Supplying a solid understanding of the key principles of distributed computing and their relationship to real-world applications, Distributed Systems: An Algorithmic Approach, Second Edition makes both an ideal textbook and a handy professional reference.

Featuring chapters by emerging and established scholars as well as by leading practitioners in the field, this Handbook both describes the state of algorithmic composition and also set the agenda for critical research on and analysis of algorithmic music.

New to the second edition of this advanced text are several chapters on regression, including neural networks and deep learning.

The recent rapid growth in the variety and complexity of new machine learning architectures requires the development of improved methods for designing, analyzing, evaluating, and communicating machine learning technologies. Statistical Machine Learning: A Unified Framework provides students, engineers, and scientists with tools from mathematical statistics and nonlinear optimization theory to become experts in the field of machine learning. In particular, the material in this text directly supports the mathematical analysis and design of old, new, and not-yet-invented nonlinear high-dimensional machine learning algorithms. Features: Unified empirical risk minimization framework supports rigorous mathematical analyses of widely used supervised, unsupervised, and reinforcement machine learning algorithms Matrix calculus methods for supporting machine learning analysis and design applications Explicit conditions for ensuring convergence of adaptive, batch, minibatch, MCEM, and MCMC learning algorithms that minimize both unimodal and multimodal objective functions Explicit conditions for characterizing asymptotic properties of M-estimators and model selection criteria such as AIC and BIC in the presence of possible model misspecification This advanced text is suitable for graduate students or highly motivated undergraduate students in statistics, computer science, electrical engineering, and applied mathematics. The text is self-contained and only assumes knowledge of lower-division linear algebra and upper-division probability theory. Students, professional engineers, and multidisciplinary scientists possessing these minimal prerequisites will find this text challenging yet accessible. About the Author: Richard M. Golden (Ph.D., M.S.E.E., B.S.E.E.) is Professor of Cognitive Science and Participating Faculty Member in Electrical Engineering at the University

of Texas at Dallas. Dr. Golden has published articles and given talks at scientific conferences on a wide range of topics in the fields of both statistics and machine learning over the past three decades. His long-term research interests include identifying conditions for the convergence of deterministic and stochastic machine learning algorithms and investigating estimation and inference in the presence of possibly misspecified probability models.

Every other day we hear about new ways to put deep learning to good use: improved medical imaging, accurate credit card fraud detection, long range weather forecasting, and more. PyTorch puts these superpowers in your hands, providing a comfortable Python experience that gets you started quickly and then grows with you as you—and your deep learning skills—become more sophisticated. Deep Learning with PyTorch will make that journey engaging and fun.

Summary Every other day we hear about new ways to put deep learning to good use: improved medical imaging, accurate credit card fraud detection, long range weather forecasting, and more. PyTorch puts these superpowers in your hands, providing a comfortable Python experience that gets you started quickly and then grows with you as you—and your deep learning skills—become more sophisticated. Deep Learning with PyTorch will make that journey engaging and fun.

Foreword by Soumith Chintala, Cocreator of PyTorch. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications.

About the technology Although many deep learning tools use Python, the PyTorch library is truly Pythonic. Instantly familiar to anyone who knows PyData tools like NumPy and scikit-learn, PyTorch simplifies deep learning without sacrificing advanced features. It's excellent for building quick models, and it scales smoothly from laptop to enterprise. Because companies like Apple, Facebook, and JPMorgan Chase rely on PyTorch, it's a great skill to have as you expand your career options. It's easy to get started with PyTorch. It minimizes cognitive overhead without sacrificing the access to advanced features, meaning you can focus on what matters the most - building and training the latest and greatest deep learning models and contribute to making a dent in the world. PyTorch is also a snap to scale and extend, and it partners well with other Python tooling. PyTorch has been adopted by hundreds of deep learning practitioners and several first-class players like FAIR, OpenAI, FastAI and Purdue.

About the book Deep Learning with PyTorch teaches you to create neural networks and deep learning systems with PyTorch. This practical book quickly gets you to work building a real-world example from scratch: a tumor image classifier. Along the way, it covers best practices for the entire DL pipeline, including the PyTorch Tensor API, loading data in Python, monitoring training, and visualizing results. After covering the basics, the book will take you on a journey through larger projects. The centerpiece of the book is a neural network designed for cancer detection. You'll discover ways for training networks with limited inputs and start processing data to get some results. You'll sift through the unreliable initial results and focus on how to diagnose and fix the problems in your neural network. Finally, you'll look at ways to improve your results by

training with augmented data, make improvements to the model architecture, and perform other fine tuning. What's inside Training deep neural networks Implementing modules and loss functions Utilizing pretrained models from PyTorch Hub Exploring code samples in Jupyter Notebooks About the reader For Python programmers with an interest in machine learning. About the author Eli Stevens had roles from software engineer to CTO, and is currently working on machine learning in the self-driving-car industry. Luca Antiga is cofounder of an AI engineering company and an AI tech startup, as well as a former PyTorch contributor. Thomas Viehmann is a PyTorch core developer and machine learning trainer and consultant. consultant based in Munich, Germany and a PyTorch core developer. Table of Contents PART 1 - CORE PYTORCH 1 Introducing deep learning and the PyTorch Library 2 Pretrained networks 3 It starts with a tensor 4 Real-world data representation using tensors 5 The mechanics of learning 6 Using a neural network to fit the data 7 Telling birds from airplanes: Learning from images 8 Using convolutions to generalize PART 2 - LEARNING FROM IMAGES IN THE REAL WORLD: EARLY DETECTION OF LUNG CANCER 9 Using PyTorch to fight cancer 10 Combining data sources into a unified dataset 11 Training a classification model to detect suspected tumors 12 Improving training with metrics and augmentation 13 Using segmentation to find suspected nodules 14 End-to-end nodule analysis, and where to go next PART 3 - DEPLOYMENT 15 Deploying to production

Traditional books on machine learning can be divided into two groups- those aimed at advanced undergraduates or early postgraduates with reasonable mathematical knowledge and those that are primers on how to code algorithms. The field is ready for a text that not only demonstrates how to use the algorithms that make up machine learning methods, but Observing the environment and recognising patterns for the purpose of decision making is fundamental to human nature. This book deals with the scientific discipline that enables similar perception in machines through pattern recognition (PR), which has application in diverse technology areas. This book is an exposition of principal topics in PR using an algorithmic approach. It provides a thorough introduction to the concepts of PR and a systematic account of the major topics in PR besides reviewing the vast progress made in the field in recent times. It includes basic techniques of PR, neural networks, support vector machines and decision trees. While theoretical aspects have been given due coverage, the emphasis is more on the practical. The book is replete with examples and illustrations and includes chapter-end exercises. It is designed to meet the needs of senior undergraduate and postgraduate students of computer science and allied disciplines.

Over the course of a generation, algorithms have gone from mathematical abstractions to powerful mediators of daily life. Algorithms have made our lives more efficient, more entertaining, and, sometimes, better informed. At the same time, complex algorithms are increasingly violating the basic rights of individual citizens. Allegedly anonymized datasets

routinely leak our most sensitive personal information; statistical models for everything from mortgages to college admissions reflect racial and gender bias. Meanwhile, users manipulate algorithms to "game" search engines, spam filters, online reviewing services, and navigation apps. Understanding and improving the science behind the algorithms that run our lives is rapidly becoming one of the most pressing issues of this century. Traditional fixes, such as laws, regulations and watchdog groups, have proven woefully inadequate. Reporting from the cutting edge of scientific research, *The Ethical Algorithm* offers a new approach: a set of principled solutions based on the emerging and exciting science of socially aware algorithm design. Michael Kearns and Aaron Roth explain how we can better embed human principles into machine code - without halting the advance of data-driven scientific exploration. Weaving together innovative research with stories of citizens, scientists, and activists on the front lines, *The Ethical Algorithm* offers a compelling vision for a future, one in which we can better protect humans from the unintended impacts of algorithms while continuing to inspire wondrous advances in technology.

Distills key concepts from linear algebra, geometry, matrices, calculus, optimization, probability and statistics that are used in machine learning.

This book compiles leading research on the development of explainable and interpretable machine learning methods in the context of computer vision and machine learning. Research progress in computer vision and pattern recognition has led to a variety of modeling techniques with almost human-like performance. Although these models have obtained astounding results, they are limited in their explainability and interpretability: what is the rationale behind the decision made? what in the model structure explains its functioning? Hence, while good performance is a critical required characteristic for learning machines, explainability and interpretability capabilities are needed to take learning machines to the next step to include them in decision support systems involving human supervision. This book, written by leading international researchers, addresses key topics of explainability and interpretability, including the following:

- Evaluation and Generalization in Interpretable Machine Learning
- Explanation Methods in Deep Learning
- Learning Functional Causal Models with Generative Neural Networks
- Learning Interpretable Rules for Multi-Label Classification
- Structuring Neural Networks for More Explainable Predictions
- Generating Post Hoc Rationales of Deep Visual Classification Decisions
- Ensembling Visual Explanations
- Explainable Deep Driving by Visualizing Causal Attention
- Interdisciplinary Perspective on Algorithmic Job Candidate Search
- Multimodal Personality Trait Analysis for Explainable Modeling of Job Interview Decisions
- Inherent Explainability Pattern Theory-based Video Event Interpretations

Personal motivation. The dream of creating artificial devices that reach or outperform human intelligence is an old one. It is also one of the dreams of my youth, which have never left me. What makes this challenge so interesting? A solution would have enormous implications on

## Read Free Machine Learning An Algorithmic Perspective Second Edition Chapman Hallcrc Machine Learning Pattern Recognition

our society, and there are reasons to believe that the AI problem can be solved in my expected lifetime. So, it's worth sticking to it for a lifetime, even if it takes 30 years or so to reap the benefits. The AI problem. The science of artificial intelligence (AI) may be defined as the construction of intelligent systems and their analysis. A natural definition of a system is anything that has an input and an output stream. Intelligence is more complicated. It can have many faces like creativity, solving problems, pattern recognition, classification, learning, induction, deduction, building analogies, optimization, surviving in an environment, language processing, and knowledge. A formal definition incorporating every aspect of intelligence, however, seems difficult. Most, if not all known facets of intelligence can be formulated as goal driven or, more precisely, as maximizing some utility function. It is, therefore, sufficient to study goal-driven AI; e. g. the (biological) goal of animals and humans is to survive and spread. The goal of AI systems should be to be useful to humans.

The field of machine learning has matured to the point where many sophisticated learning approaches can be applied to practical applications. Thus it is of critical importance that researchers have the proper tools to evaluate learning approaches and understand the underlying issues. This book examines various aspects of the evaluation process with an emphasis on classification algorithms. The authors describe several techniques for classifier performance assessment, error estimation and resampling, obtaining statistical significance as well as selecting appropriate domains for evaluation. They also present a unified evaluation framework and highlight how different components of evaluation are both significantly interrelated and interdependent. The techniques presented in the book are illustrated using R and WEKA, facilitating better practical insight as well as implementation. Aimed at researchers in the theory and applications of machine learning, this book offers a solid basis for conducting performance evaluations of algorithms in practical settings.

Machine Learning An Algorithmic Perspective, Second Edition CRC Press

Designing machines that can read handwriting like human beings has been an ambitious goal for more than half a century, driving talented researchers to explore diverse approaches. Obstacles have often been encountered that at first appeared insurmountable but were indeed overcome before long. Yet some open issues remain to be solved. As an indispensable branch, Chinese handwriting recognition has been termed as one of the most difficult Pattern Recognition tasks. Chinese handwriting recognition poses its own unique challenges, such as huge variations in strokes, diversity of writing styles, and a large set of confusable categories. With ever-increasing training data, researchers have pursued elaborate algorithms to discern characters from different categories and compensate for the sample variations within the same category. As a result, Chinese handwriting recognition has evolved substantially and amazing achievements can be seen. This book introduces integral algorithms used in Chinese handwriting recognition and the applications of Chinese handwriting recognizers. The first part of the book covers both widespread canonical algorithms to a reliable recognizer and newly developed scalable methods in Chinese handwriting recognition. The recognition of Chinese handwritten text is presented systematically, including instructive guidelines for collecting samples, novel recognition paradigms, distributed discriminative learning of appearance models and distributed estimation of contextual models for large categories, in addition to celebrated methods, e.g. Gradient features, MQDF and HMMs. In the second part of this book, endeavors are made to create a friendlier human-machine interface through application of Chinese handwriting recognition. Four scenarios are exemplified: grid-assisted input, shortest moving input, handwritten micro-blog, and instant handwriting messenger. All the while, the book moves from basic to more complex approaches, also providing a list for further reading with literature comments.

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

## Read Free Machine Learning An Algorithmic Perspective Second Edition Chapman Hallcrc Machine Learning Pattern Recognition

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.

Covering all the main approaches in state-of-the-art machine learning research, this will set a new standard as an introductory textbook. With the help of this book, you'll build smart algorithmic models using machine learning algorithms covering tasks such as time series forecasting, backtesting, trade predictions, and more using easy-to-follow examples. By the end, you'll be able to adopt algorithmic trading in your own business and implement intelligent investigative strategies.

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

A comprehensive overview of data mining from an algorithmic perspective, integrating related concepts from machine learning and statistics.

This book introduces machine learning methods in finance. It presents a unified treatment of machine learning and various statistical and computational disciplines in quantitative finance, such as financial econometrics and discrete time stochastic control, with an emphasis on how theory and hypothesis tests inform the choice of algorithm for financial data modeling and decision making. With the trend towards increasing computational resources and larger datasets, machine learning has grown into an important skillset for the finance industry. This book is written for advanced graduate students and academics in financial

econometrics, mathematical finance and applied statistics, in addition to quants and data scientists in the field of quantitative finance. *Machine Learning in Finance: From Theory to Practice* is divided into three parts, each part covering theory and applications. The first presents supervised learning for cross-sectional data from both a Bayesian and frequentist perspective. The more advanced material places a firm emphasis on neural networks, including deep learning, as well as Gaussian processes, with examples in investment management and derivative modeling. The second part presents supervised learning for time series data, arguably the most common data type used in finance with examples in trading, stochastic volatility and fixed income modeling. Finally, the third part presents reinforcement learning and its applications in trading, investment and wealth management. Python code examples are provided to support the readers' understanding of the methodologies and applications. The book also includes more than 80 mathematical and programming exercises, with worked solutions available to instructors. As a bridge to research in this emergent field, the final chapter presents the frontiers of machine learning in finance from a researcher's perspective, highlighting how many well-known concepts in statistical physics are likely to emerge as important methodologies for machine learning in finance.

This textbook introduces linear algebra and optimization in the context of machine learning. Examples and exercises are provided throughout this text book together with access to a solution's manual. This textbook targets graduate level students and professors in computer science, mathematics and data science. Advanced undergraduate students can also use this textbook. The chapters for this textbook are organized as follows: 1. Linear algebra and its applications: The chapters focus on the basics of linear algebra together with their common applications to singular value decomposition, matrix factorization, similarity matrices (kernel methods), and graph analysis. Numerous machine learning applications have been used as examples, such as spectral clustering, kernel-based classification, and outlier detection. The tight integration of linear algebra methods with examples from machine learning differentiates this book from generic volumes on linear algebra. The focus is clearly on the most relevant aspects of linear algebra for machine learning and to teach readers how to apply these concepts. 2. Optimization and its applications: Much of machine learning is posed as an optimization problem in which we try to maximize the accuracy of regression and classification models. The "parent problem" of optimization-centric machine learning is least-squares regression. Interestingly, this problem arises in both linear algebra and optimization, and is one of the key connecting problems of the two fields. Least-squares regression is also the starting point for support vector machines, logistic regression, and recommender systems. Furthermore, the methods for dimensionality reduction and matrix factorization also require the development of optimization methods. A general view of optimization in computational graphs is discussed together with its applications to back propagation in neural networks. A frequent challenge faced by beginners in machine learning is the extensive background required in linear algebra and optimization. One problem is that the existing linear algebra and optimization courses are not specific to machine learning; therefore, one would typically have to complete more course material than is necessary to pick up machine learning. Furthermore, certain types of ideas and tricks from optimization and linear algebra recur more frequently in machine learning than other application-centric settings.

Therefore, there is significant value in developing a view of linear algebra and optimization that is better suited to the specific perspective of machine learning.

Sea and freshwater ports are a key component of critical infrastructure and essential for maintaining global and domestic economies. In order to effectively secure a dynamic port facility operation, one must understand the business of maritime commerce. Following in the tradition of its bestselling predecessor, Port Security Management, Second Edition continues to supply readers with this understanding. This fully updated edition covers the latest in continuously changing legislation regarding federal mandates, securing vessels, cargo security, and granting employee credentials. Focusing on best practices, it details real-world solutions that law enforcement authorities and security management professionals can put to use immediately. Assuming little prior knowledge of the industry, the book examines port security in the context of global transportation systems. It supplies practitioners and educators with a framework for managing port security and details risk assessment and physical security best practices for securing ships and ports. The book explains how the various stakeholders, including port management, security, government, and private industry, can collaborate to develop safe and secure best practices while maintaining efficient operations. Addressing the legislative measures, regulatory issues, and logistical aspects of port security, the book includes coverage of cruise ships, cargo security, CT-PAT, and emergency operations. Complete with a new chapter on intelligence, this book is ideal for anyone with a vested interest in secure and prosperous port facilities who wants to truly understand how to best tackle the management of port security.

This tutorial text gives a unifying perspective on machine learning by covering both probabilistic and deterministic approaches -which are based on optimization techniques – together with the Bayesian inference approach, whose essence lies in the use of a hierarchy of probabilistic models. The book presents the major machine learning methods as they have been developed in different disciplines, such as statistics, statistical and adaptive signal processing and computer science. Focusing on the physical reasoning behind the mathematics, 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. The book builds carefully from the basic classical methods to the most recent trends, with chapters 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 short courses on sparse modeling, deep learning, and probabilistic graphical models. All major classical techniques: Mean/Least-Squares regression and filtering, Kalman filtering, stochastic approximation and online learning, Bayesian classification, decision trees, logistic regression and boosting methods. The latest trends: 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. Case studies - protein folding prediction, optical character recognition, text authorship identification, fMRI data analysis, change point detection, hyperspectral image unmixing, target localization, channel equalization and echo cancellation, show how the theory can be applied. MATLAB code for

all the main algorithms are available on an accompanying website, enabling the reader to experiment with the code.

Digitalization of societies has important ramifications for citizens and businesses. The digital landscape is rapidly changing, whereas at the same time there are growing concerns about how market access in the EU's digital market as well as fundamental rights can be sufficiently safeguarded in the shadow of 'big data' and algorithms. This timely and important book presents expert analyses of how digitalization raises questions of the future role for general principles of European Union (EU) law, including the foundational principles of the EU's fundamental economic freedoms and EU competition rules. Examining a number of issues revolving around the internal market, the digital single market, competition law, intellectual property, data protection, democracy and the rule of law, the contributors provide deeply informed insights of the challenges as to: effects of the technological disruption on the doctrine of general principles; the resilience of general principles in the digital society; the need to rely on new general principles in the digital society; the realization of the digital single market; the safeguarding of fundamental rights and the rule of law. The contributors are highly esteemed scholars from major European universities. A common theme is the need for judicial evolution of EU fundamental rights law in the digital era, alongside penetrating analyses of recent Court of Justice of the European Union case law on the impact of digitalization. Dealing as it does with an area of European law of particular complexity and rapidly growing importance, the anthology provides insights and knowledge about the ways in which digitalization is rapidly changing EU law. Are general principles of EU law as developed in an 'analogue world' sufficiently resilient to withstand the rapid and often disruptive developments taking place in the digital marketplace? The fresh look at the concept of 'general principles' taken by the various contributors helps to clarify the actual application in EU law in areas related to the ongoing digitalization of our society. It will be greatly appreciated by practitioners, policymakers and academics active in any of the many fields of law affected by the digital revolution.

AI framework intended to solve a problem of bias-variance tradeoff for supervised learning methods in real-life applications. The AI framework comprises of bootstrapping to create multiple training and testing data sets with various characteristics, design and analysis of statistical experiments to identify optimal feature subsets and optimal hyper-parameters for ML methods, data contamination to test for the robustness of the classifiers. Key Features: Using ML methods by itself doesn't ensure building classifiers that generalize well for new data Identifying optimal feature subsets and hyper-parameters of ML methods can be resolved using design and analysis of statistical experiments Using a bootstrapping approach to massive sampling of training and tests datasets with various data characteristics (e.g.: contaminated training sets) allows dealing with bias Developing of SAS-based table-driven environment allows managing all meta-data related to the proposed AI framework and creating interoperability with R libraries to accomplish variety of statistical and machine-learning tasks Computer programs in R and SAS that create AI framework are available on GitHub

[Copyright: 883e3d5d7b3ad92f96c742e4f70930c2](https://doi.org/10.1007/978-1-4939-9830-2)