Gaussian Processes For Machine Learning

The book deals mainly with three problems involving Gaussian stationary processes. The first problem consists of clarifying the conditions for mutual absolute continuity (equivalence) of probability distributions of a "random process segment" and of finding effective formulas for densities of the equiva lent distributions. Our second problem is to describe the classes of spectral measures corresponding in some sense to regular stationary processes (in par ticular, satisfying the well-known "strong mixing condition") as well as to describe the subclasses associated with "mixing rate". The third problem involves estimation of an unknown mean value of a random process, this random process being stationary except for its mean, i. e. , it is the problem of "distinguishing a signal from stationary noise". Furthermore, we give here auxiliary information (on distributions in Hilbert spaces, properties of sam ple functions, theorems on functions of a complex variable, etc.). Since 1958 many mathematicians have studied the problem of equivalence of various infinite-dimensional Gaussian distributions (detailed and sys tematic presentation of the basic results can be found, for instance, in [23]). In this book we have considered Gaussian stationary processes and arrived, we believe, at rather definite solutions. The second problem mentioned above is closely related with problems involving ergodic theory of Gaussian dynamic systems as well as prediction theory of stationary processes.

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.

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.

This text offers background in function theory, Hardy functions, and probability as preparation for surveys of Gaussian processes, strings and spectral functions, and strings and spaces of integral functions. It addresses the relationship between the past and the future of a real, one-dimensional, stationary Gaussian process. 1976 edition.

Surrogates

Bandit Algorithms

Some Theory for Kriging

A Probabilistic Perspective

Proceedings of the 2000 Conference

Artificial Neural Networks and Machine Learning - ICANN 2011

Machine Learning has become a key enabling technology for many engineering applications, investigating scientific questions and theoretical problems alike. To stimulate discussions and to disseminate new results, a summer school series was started in February 2002, the documentation of which is published as LNAI 2600. This book presents revised lectures of two subsequent summer schools held in 2003 in Canberra, Australia, and in Tübingen, Germany. The tutorial lectures included are devoted to statistical learning theory, unsupervised learning, Bayesian inference, and applications in pattern recognition; they provide in-depth overviews of exciting new developments and contain a large number of references. Graduate students, lecturers, researchers and professionals alike will find this book a useful resource in learning and teaching machine learning.

"This book introduces machine learning for readers with some background in basic linear algebra, statistics, probability, and programming. In a coherent statistical framework it covers a selection of supervised machine learning methods, from the most fundamental (k-NN, decision trees, linear and logistic regression) to more advanced methods (deep neural networks, support vector machines, Gaussian processes, random forests and boosting), plus commonly-used unsupervised methods (generative modeling, k-means, PCA, autoencoders and generative adversarial networks). Careful explanations and pseudo-code are presented for all methods. The authors maintain a focus on the fundamentals by drawing connections between methods and discussing general concepts such as loss functions, maximum likelihood, the bias-variance decomposition, ensemble averaging, kernels and the Bayesian approach along with generally useful tools such as regularization, cross validation, evaluation metrics and optimization methods. The final chapters offer practical advice for solving real-world supervised machine learning problems and on ethical aspects of modern machine learning"--

The latest edition of this classic is updated with new problem sets and material The Second Edition of this fundamental textbook maintains the book's tradition of clear, thought-provoking instruction. Readers are provided once again with an instructive mix of mathematics, physics, statistics, and information theory. All the essential topics in information theory are covered in detail, including entropy, data compression, channel capacity, rate distortion, network information theory, and hypothesis testing. The authors provide readers with a solid understanding of the underlying theory and applications. Problem sets and a telegraphic summary at the end of each chapter further assist readers. The historical notes that follow each chapter recap the main points. The Second Edition features: * Chapters reorganized to improve teaching * 200 new problems * New material on source coding, portfolio theory, and feedback capacity * Updated references Now current and enhanced, the Second Edition of Elements of Information Theory remains the ideal textbook for upper-level undergraduate and graduate courses in electrical engineering, statistics, and telecommunications.

Content Description. #Includes bibliographical references and index.

Adaptive and Natural Computing Algorithms

First International Workshop, Sheffield, UK, September 7-10, 2004. Revised Lectures

Managing Growing Pressures on Quantity and Quality

Applied Stochastic Differential Equations

Origin Story

From Motor Learning to Interaction Learning in Robots

The proceedings of the 2000 Neural Information Processing Systems (NIPS) Conference.The annual conference on Neural Information Processing Systems (NIPS) is the flagship conference on neural computation. The conference is interdisciplinary, with contributions in algorithms, learning theory, cognitive science, neuroscience, vision, speech and signal processing, reinforcement learning and control, implementations, and diverse applications. Only about 30 percent of the papers submitted are accepted for presentation at NIPS, so the quality is exceptionally high. These proceedings contain all of the papers that were presented at the 2000 conference.

Implement TensorFlow's offerings such as TensorBoard, TensorFlow.js, TensorFlow Probability, and TensorFlow Lite to build smart automation projects Key FeaturesUse machine learning and deep learning principles to build real-world projectsGet to grips with TensorFlow's impressive range of module offeringsImplement projects on GANs, reinforcement learning, and capsule networkBook Description TensorFlow has transformed the way machine learning is perceived. TensorFlow Machine Learning Projects teaches you how to exploit the benefits simplicity, efficiency, and flexibility of using TensorFlow in various real-world projects. With the help of this book, you'll not only learn how to build advanced projects using different datasets but also be able to tackle common challenges using a range of libraries from the TensorFlow ecosystem. To start with, you'll get to grips with using TensorFlow for machine learning projects; you'll explore a wide range of projects using TensorForest and TensorBoard for detecting exoplanets, TensorFlow.js for sentiment analysis, and TensorFlow Lite for digit classification. As you make your way through the book, you'll build projects in various real-world domains, incorporating natural language processing (NLP), the Gaussian process, autoencoders, recommender systems, and Bayesian neural networks, along with trending areas such as Generative Adversarial Networks (GANs), capsule networks, and reinforcement learning. You'll learn how to use the TensorFlow on Spark API and GPU-accelerated computing with TensorFlow to detect objects, followed by how to train and develop a recurrent neural network (RNN) model to generate book scripts. By the end of this book, you'll have gained the required expertise to build full-fledged machine learning projects at work. What you will learnUnderstand the TensorFlow ecosystem using various datasets and techniquesCreate recommendation systems for quality product recommendationsBuild projects using CNNs, NLP, and Bayesian neural networksPlay Pac-Man using deep reinforcement learningDeploy scalable TensorFlow-based machine learning systemsGenerate your own book script using RNNsWho this book is for TensorFlow Machine Learning Projects is for you if you are a data analyst, data scientist, machine learning professional, or deep learning enthusiast with basic knowledge of TensorFlow. This book is also for you if you want to build end-to-end projects in the machine learning domain using supervised, unsupervised, and reinforcement learning techniques

Bayesian nonparametrics comes of age with this landmark text synthesizing theory, methodology and computation.

This book constitutes the post-conference proceedings of the 5th International Conference on Machine Learning, Optimization, and Data Science, LOD 2019, held in Siena, Italy, in September 2019. The 54 full papers presented were carefully reviewed and selected from 158 submissions. The papers cover topics in the field of machine learning, artificial intelligence, reinforcement learning, computational optimization and data science presenting a substantial array of ideas, technologies, algorithms, methods and applications.

A First Course for Engineers and Scientists

Artificial Intelligence and Statistics

Gaussian Processes, Function Theory, and the Inverse Spectral Problem

Gaussian Random Processes

10th International Conference, ICANNGA 2011, Ljubljana, Slovenia, April 14-16, 2011, Proceedings, Part I

TensorFlow Machine Learning Projects

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 two volume set LNCS 6791 and LNCS 6792 constitutes the refereed proceedings of the 21th International Conference on Artificial Neural Networks, ICANN 2011, held in Espoo, Finland, in June 2011. The 106 revised full or poster papers presented were carefully reviewed and selected from numerous submissions. ICANN 2011 had two basic tracks: brain-inspired computing and machine learning research, with strong cross-disciplinary interactions and applications.

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 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.

Computer simulation experiments are essential to modern scientific discovery, whether that be in physics, chemistry, biology, epidemiology, ecology, engineering, etc. Surrogates are meta-models of computer simulations, used to solve mathematical models that are too intricate to be worked by hand. Gaussian process (GP) regression is a supremely flexible tool for the analysis of computer simulation experiments. This book presents an applied introduction to GP regression for modelling and optimization of computer simulation experiments. Features: · Emphasis on methods, applications, and reproducibility. · R code is integrated throughout for application of the methods. · Includes more than 200 full colour figures. · Includes many exercises to supplement understanding, with separate solutions available from the author. · Supported by a website with full code available to reproduce all methods and examples. The book is primarily designed as a textbook for postgraduate students studying GP regression from mathematics, statistics, computer science, and engineering. Given the breadth of examples, it could also be used by researchers from these fields, as well as from economics, life science, social science, etc.

Build 13 real-world projects with advanced numerical computations using the Python ecosystem

A Unified Framework

From Theory to Practice

An Introduction

Algorithms for Pattern Recognition

Gaussian Process Modeling, Design, and Optimization for the Applied Sciences

A comprehensive and self-contained introduction to Gaussian processes, which provide a principled, practical, probabilistic approach to learning in kernel machines. Gaussian processes (GPs) provide a principled, practical, probabilistic approach to learning in kernel machines. GPs have received increased attention in the machine-learning community over the past decade, and this book provides a long-needed systematic and unified treatment of theoretical and practical aspects of GPs in machine learning. The treatment is comprehensive and self-contained, targeted at researchers and students in machine learning and applied statistics. The book deals with the supervised-learning problem for both regression and classification, and includes detailed algorithms. A wide variety of covariance (kernel) functions are presented and their properties discussed. Model selection is discussed both from a Bayesian and a classical perspective. Many connections to other well-known techniques from machine learning and statistics are discussed, including support-vector machines, neural networks, splines, regularization networks, relevance vector machines and others. Theoretical issues including learning curves and the PAC-Bayesian framework are treated, and several approximation methods for learning with large datasets are discussed. The book contains illustrative examples and exercises, and code and datasets are available on the Web. Appendixes provide mathematical background and a discussion of Gaussian Markov processes.

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.

Getting the most out of neural networks and related data modelling techniques is the purpose of this book. The text, with the accompanying Netlab toolbox, provides all the necessary tools and knowledge. Throughout, the emphasis is on methods that are relevant to the practical application of neural networks to pattern analysis problems. All parts of the toolbox interact in a coherent way, and implementations and descriptions of standard statistical techniques are provided so that they can be used as benchmarks against which more sophisticated algorithms can be evaluated. Plenty of examples and demonstration programs illustrate the theory and help the reader understand the algorithms and how to apply them.

From an engineering standpoint, the increasing complexity of robotic systems and the increasing demand for more autonomously learning robots, has become essential. This book is largely based on the successful workshop "From motor to interaction learning in robots" held at the IEEE/RSJ International Conference on Intelligent Robot Systems. The major aim of the book is to give students interested the topics described above a chance to get started faster and researchers a helpful compandium.

Kernels for Vector-Valued Functions

Deterministic and Statistical Methods in Machine Learning

Bayesian Reasoning and Gaussian Processes for Machine Learning Applications

Semi-Supervised Learning

Efficient Reinforcement Learning Using Gaussian Processes

5th International Conference, LOD 2019, Siena, Italy, September 10–13, 2019, Proceedings

This monograph reviews different methods to design or learn valid kernel functions for multiple outputs, paying particular attention to the connection between probabilistic and regularization methods.

Gaussian Processes for Machine LearningMIT Press

Groundwater allocation determines who is able to use groundwater resources, how, when and where. It directly affects the value (economic, ecological, socio-cultural) that individuals and society obtain from groundwater, today and in the future.

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

Bayesian Time Series Models

Gaussian Processes

OECD Studies on Water Groundwater Allocation Managing Growing Pressures on Quantity and Quality

European Conference, Antwerp, Belgium, September 15-19, 2008, Proceedings, Part I

NETLAB

21st International Conference on Artificial Neural Networks, Espoo, Finland, June 14-17, 2011, Proceedings, Part I

A comprehensive review of an area of machine learning that deals with the use of unlabeled data in classification problems: state-of-the-art algorithms, a taxonomy of the field, applications, benchmark experiments, and directions for future research. In the field of machine learning, semi-supervised learning (SSL) occupies the middle ground, between supervised learning (in which all training examples are labeled) and unsupervised learning (in which no label data are given). Interest in SSL has increased in recent years, particularly because of application domains in which unlabeled data are plentiful, such as images, text, and bioinformatics. This first comprehensive overview of SSL presents state-of-the-art algorithms, a taxonomy of the field, selected applications, benchmark experiments, and perspectives on ongoing and future research.Semi-Supervised Learning first presents the key assumptions and ideas underlying the field: smoothness, cluster or low-density separation, manifold structure, and transduction. The core of the book is the presentation of SSL methods, organized according to algorithmic strategies. After an examination of generative models, the book describes algorithms that implement the low-density separation assumption, graph-based methods, and algorithms that perform two-step learning. The book then discusses SSL applications and offers guidelines for SSL practitioners by analyzing the results of extensive benchmark experiments. Finally, the book looks at interesting directions for SSL research. The book closes with a discussion of the relationship between semi-supervised learning and transduction.

A summary of past work and a description of new approaches to thinking about kriging, commonly used in the prediction of a random field based on observations at some set of locations in mining, hydrology, atmospheric sciences, and geography.

A comprehensive and rigorous introduction for graduate students and researchers, with applications in sequential decision-making problems.

Artificial "neural networks" are widely used as flexible models for classification and regression applications, but questions remain about how the power of these models can be safely exploited when training data is limited. This book demonstrates how Bayesian methods allow complex neural network models to be used without fear of the "overfitting" that can occur with traditional training methods. Insight into the nature of these complex Bayesian models is provided by a theoretical investigation of the priors over functions that underlie them. A practical implementation of Bayesian neural network learning using Markov chain Monte Carlo methods is also described, and software for it is freely available over the Internet. Presupposing only basic knowledge of probability and statistics, this book should be of interest to researchers in statistics, engineering, and artificial intelligence.

A Big History of Everything

A Review

Introduction to Machine Learning

Advances in Neural Information Processing Systems 13

Elements of Information Theory

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 two-volume set LNCS 6593 and 6594 constitutes the refereed proceedings of the 10th International Conference on Adaptive and Natural Computing Algorithms, ICANNGA 2010, held in Ljubljana, Slovenia, in April 2010. The 83 revised full papers presented were carefully reviewed and selected from a total of 144 submissions. The first volume includes 42 papers and a plenary lecture and is organized in topical sections on neural networks and evolutionary computation.

Aimed at students and researchers in mathematics, communications engineering, and economics, this book describes the probabilistic structure of a Gaussian process in terms of its canonical representation (or its innovation process). Multiple Markov properties of a Gaussian process and equivalence problems of Gaussian processes are clearly presented. The authors' approach is unique, involving causality in time evolution and information-theoretic aspects. Because the book is self-contained and only requires background in the fundamentals of probability theory and measure theory, it would be suitable as a textbook at the senior undergraduate or graduate level.

A statistical view of uncertainty in expert systems. Knowledge, decision making, and uncertainty. Conceptual clustering and its relation to numerical taxonomy. Learning rates in supervised and unsupervised intelligent systems. Pinpoint good hypotheses with heuristics. Artificial intelligence approaches in statistics. REX review. Representing statistical computations: toward a deeper understanding. Student phase 1: a report on work in progress. Representing statistical knowledge for expert data analysis systems. Environments for supporting statistical strategy. Use of psychometric tools for knowledge acquisition: a case study. The analysis phase in development of knowledge based systems. Implementation and study of statistical strategy. Patterns in statisticalstrategy. A DIY guide to statistical strategy. An alphabet for statistician's expert systems.

Learning Kernel Classifiers

Machine Learning and Knowledge Discovery in Databases

Machine Learning

Statistical Machine Learning

Bayesian Learning for Neural Networks

Fundamentals of Nonparametric Bayesian Inference

This book constitutes the refereed proceedings of the joint conference on Machine Learning and Knowledge Discovery in Databases: ECML PKDD 2008, held in Antwerp, Belgium, in September 2008. The 100 papers presented in two volumes, together with 5 invited talks, were carefully reviewed and selected from 521 submissions. In addition to the regular papers the volume contains 14 abstracts of papers appearing in full version in the Machine Learning Journal and the Knowledge Discovery and Databases Journal of Springer. The conference intends to provide an international forum for the discussion of the latest high quality research results in all areas related to machine learning and knowledge discovery in databases. The topics addressed are application of machine learning and data mining methods to real-world problems, particularly exploratory research that describes novel learning and mining tasks and applications requiring non-standard techniques.

This New York Times bestseller "elegantly weaves evidence and insights . . . into a single, accessible historical narrative" (Bill Gates) and presents a captivating history of the universe -- from the Big Bang to dinosaurs to mass globalization and beyond. Most historians study the smallest slivers of time, emphasizing specific dates, individuals, and documents. But what would it look like to study the whole of history, from the big bang through the present day -- and even into the remote future? How would looking at the full span of time change the way we perceive the universe, the earth, and our very existence? These were the questions David Christian set out to answer when he created the field of "Big History," the most exciting new approach to understanding where we have been, where we are, and where we are going. In Origin Story, Christian takes readers on a wild ride through the entire 13.8 billion years we've come to know as "history." By focusing on defining events (thresholds), major trends, and profound questions about our origins, Christian exposes the hidden threads that tie everything together -- from the creation of the planet to the advent of agriculture, nuclear war, and beyond. With stunning insights into the origin of the universe, the beginning of life, the emergence of humans, and what the future might bring, Origin Story boldly reframes our place in the cosmos.

This volume contains selected tutorial and young scientist school papers of the 5th RAAI Summer School on Artificial Intelligence, held in July 2019 at Institute of Physics and Technology (MIPT) campus in Dolgoprudny, a suburb of Moscow, Russia. The 11 chapters in this volume present papers focusing on various important aspects of Multiagent systems; Behavior planning; Natural language processing; Modeling of reasoning; and Machine learning and data analysis.

Stochastic differential equations are differential equations whose solutions are stochastic processes. They exhibit appealing mathematical properties that are useful in modeling uncertainties and noisy phenomena in many disciplines. This book is motivated by applications of stochastic differential equations in target tracking and medical technology and, in particular, their use in methodologies such as filtering, smoothing, parameter estimation, and machine learning. It builds an intuitive hands-on understanding of what stochastic differential equations are all about, but also covers the essentials of It calculus, the central theorems in the field, and such approximation schemes as stochastic Runge-Kutta. Greater emphasis is given to solution methods than to analysis of theoretical properties of the equations. The book's practical approach assumes only prior understanding of ordinary differential equations. The numerous worked examples and end-of-chapter exercises include application-driven derivations and computational assignments. MATLAB/Octave source code is available for download, promoting hands-on work with the methods.

Probabilistic Machine Learning

Theory and Algorithms

ML Summer Schools 2003, Canberra, Australia, February 2-14, 2003, Tübingen, Germany, August 4-16, 2003, Revised Lectures

Gaussian Processes for Machine Learning

Advanced Lectures on Machine Learning

Artificial Intelligence

This book consitutes the refereed proceedings of the First International Workshop on Machine Learning held in Sheffield, UK, in September 2004. The 19 revised full papers presented were carefully reviewed and selected for inclusion in the book. They address all current issues in the rapidly maturing field of machine learning that aims to provide practical methods for data discovery, categorisation and modelling. The particular focus of the workshop was advanced research methods in machine learning and statistical signal processing.

"Gaussian processes (GPs) provide a principled, practical, probabilistic approach to learning in kernel machines. GPs have received increased attention in the machine-learning community over the past decade, and this book provides a long-needed systematic and unified treatment of theoretical and practical aspects of GPs in machine learning. The treatment is comprehensive and self-contained, targeted at researchers and students in machine learning and applied statistics."--Page 4 de la couverture

This book introduces Bayesian reasoning and Gaussian processes into machine learning applications. Bayesian methods are applied in many areas, such as game development, decision making, and drug discovery. It is very effective for machine learning algorithms in handling missing data and extracting information from small datasets. Bayesian Reasoning and Gaussian Processes for Machine Learning Applications uses a statistical background to understand continuous distributions and how learning can be viewed from a probabilistic framework. The chapters progress into such machine learning topics as belief network and Bayesian reinforcement learning, which is followed by Gaussian process introduction, classification, regression, covariance, and performance analysis of Gaussian processes with other models. FEATURES Contains recent advancements in machine learning Highlights applications of machine learning algorithms Offers both quantitative and qualitative research Includes numerous case studies This book is aimed at graduates, researchers, and professionals in the field of data science and machine learning.

The first unified treatment of time series modelling techniques spanning machine learning, statistics, engineering and computer science.

Interpolation of Spatial Data

Machine Learning in Finance

Machine Learning, Optimization, and Data Science

A Bayesian and Optimization Perspective

5th RAAI Summer School, Dolgoprudny, Russia, July 4 – 7, 2019, Tutorial Lectures

Graphical Models for Machine Learning and Digital Communication