Introduction To Gaussian Processes

It was undoubtedly a necessary task to collect all the results on the concentration of measure during the past years in a monograph. The author did this very successfully and the book is an important contribution to the topic. It will surely influence further research in this area considerably. The book is very well written, and it was a great pleasure for the reviewer to read it. --Mathematical Reviews The observation of the concentration of measure phenomenon is inspired by isoperimetric inequalities. A familiar example is the way the uniform measure on the standard sphere S^n becomes concentrated around the equator as the dimension gets large. This property may be interpreted in terms of functions on the sphere with small oscillations, an idea going back to Levy. The phenomenon also occurs in probability, as a version of the law of large numbers, due to Emile Borel. This book offers the basic techniques and examples of the concentration of measure phenomenon. The concentration of measure phenomenon was put forward in the early seventies by V. Milman in the asymptotic geometry of Banach spaces. It is of powerful interest in applications in various areas, such as geometry, functional analysis and infinite-dimensional integration, discrete mathematics and complexity theory, and probability theory. Particular emphasis is on geometric, functional, and probabilistic tools to reach and describe measure concentration in a number of settings. The book presents concentration functions and inequalities, isoperimetric and functional examples, spectrum and topological applications, product measures, entropic and transportation methods, as well as aspects of M. Talagrand's deep investigation of concentration in product spaces and its application in discrete mathematics and probability theory, supremum of Gaussian and empirical processes, spin glass, random matrices, etc. Prerequisites are a basic background in measure theory, functional analysis, and probability theory.

Sensor data fusion is the process of combining error-prone, heterogeneous, incomplete, and ambiguous data to gather a higher level of situational awareness. In principle, all living creatures are fusing information from their complementary senses to coordinate their actions and to detect and localize danger. In sensor data fusion, this process is transferred to electronic systems, which rely on some "awareness" of what is happening in certain areas of interest. By means of probability theory and statistics, it is possible to model the relationship between the state space and the sensor data. The number of ingredients of the resulting Kalman filter is limited, but its applications are not.

Publisher's Note: Products purchased from Third Party sellers are not guaranteed by the publisher for quality, authenticity, or access to any online entitlements included with the product. A comprehensive introduction to the mathematical principles and algorithms in statistical signal processing and modern neural networks. This text is an expanded version of a graduate course on advanced signal processing at the Johns Hopkins University Whiting school program for professionals with students from electrical engineering, physics, computer and data science, and mathematics backgrounds. It covers the theory underlying applications in statistical signal processing including spectral estimation, linear prediction, adaptive filters, and optimal processing of uniform spatial arrays. Unique among books on the subject, it also includes a comprehensive introduction to modern neural networks with examples in time series and image classification. Coverage includes: Mathematical structures of signal spaces and matrix factorizations linear time-invariant systems and transforms Least squares filters

Random variables, estimation theory, and random processes Spectral estimation and autoregressive signal models linear prediction and adaptive filters Optimal processing of linear arrays Neural networks

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

Content Description. #Includes bibliographical references and index.

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. Isoperimetric, measure concentration and random process techniques appear at the basis of the modern understanding of Probability in Banach spaces. Based on these tools, the book presents a complete treatment of the main aspects of Probability in Banach spaces (integrability and limit theorems for vector valued random variables, boundedness and continuity of random processes) and of some of their links to Geometry of Banach spaces (via the type and cotype properties). Its purpose is to present some of the main aspects of this theory, from the foundations to the most important achievements. The main features of the investigation are the systematic use of isoperimetry and concentration of measure and abstract random process techniques (entropy and majorizing measures). Examples of these probabilistic tools and ideas to classical Banach space theory are further developed.
Gaussian Process Regression Analysis for Functional Data presents nonparametric statistical methods for functional regression analysis, specifically the methods based on a Gaussian process prior in a functional space. The authors focus on problems involving functional response variables and mixed covariates of functional and scalar variables. Covering the basics of Gaussian process regression, the first several chapters discuss functional data analysis, theoretical aspects based on the asymptotic properties of Gaussian process regression models, and new methodological developments for high dimensional data and variable selection. The remainder of the text explores advanced topics of functional regression analysis, including novel nonparametric statistical methods for curve prediction, curve clustering, functional ANOVA, and functional regression analysis of batch data, repeated curves, and non-Gaussian data. Many flexible models based on Gaussian processes provide

efficient ways of model learning, interpreting model structure, and carrying out inference, particularly when dealing with large dimensional functional data. This book shows how to use these Gaussian process regression models in the analysis of functional data. Some MATLAB® and C codes are available on the first author's website.

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.

Kosorok's brilliant text provides a self-contained introduction to empirical processes and semiparametric inference. These powerful research techniques are surprisingly useful for developing methods of statistical inference for complex models and in understanding the properties of such methods. This is an authoritative text that covers all the bases, and also a friendly and gradual introduction to the area. The book can be used as research reference and textbook.

In recent years the significance of Gaussian processes to communication networks has grown considerably. The inherent flexibility of the Gaussian traffic model enables the analysis, in a single mathematical framework, of systems with both long-range and short-range dependent input streams. Large Deviations for Gaussian Queues demonstrates how the Gaussian traffic model arises naturally, and how the analysis of the corresponding queuing model can be performed. The text provides a general introduction to Gaussian queues, and surveys recent research into the modelling of communications networks. Coverage includes: Discussion of the theoretical concepts and practical aspects related to Gaussian traffic models. Analysis of recent research asymptotic results for Gaussian queues, both in the large-buffer and many-

sources regime. An emphasis on rare-event analysis, relying on a variety of asymptotic techniques. Examination of single-node FIFO queuing systems, as well as queues operating under more complex scheduling disciplines, and queuing networks. A set of illustrative examples that directly relate to important practical problems in communication networking. A large collection of instructive exercises and accompanying solutions. Large Deviations for Gaussian Queues assumes minimal prior knowledge. It is ideally suited for postgraduate students in applied probability, operations research, computer science and electrical engineering. The book's self-contained style makes it perfect for practitioners in the communications networking industry and for researchers in related areas.
This book considers large and challenging multistage decision problems, which can be solved in principle by dynamic programming (DP), but their exact solution is computationally intractable. We discuss solution methods that rely on approximations to produce suboptimal policies with adequate performance. These methods are collectively known by several essentially equivalent names: reinforcement learning, approximate dynamic programming, neuro-dynamic programming. They have been at the forefront of research for the last 25 years, and they underlie, among others, the recent impressive successes of self-learning in the context of games such as chess and Go. Our subject has benefited greatly from the interplay of ideas from optimal control and from artificial intelligence, as it relates to reinforcement learning and simulation-based neural network methods. One of the aims of the book is to explore the common boundary between these two fields and to form a bridge that is accessible by workers with background in either field. Another aim is to organize coherently the broad mosaic of methods that have proved successful in practice while having a solid theoretical

and/or logical foundation. This may help researchers and practitioners to find their way through the maze of competing ideas that constitute the current state of the art. This book relates to several of our other books: Neuro-Dynamic Programming (Athena Scientific, 1996), Dynamic Programming and Optimal Control (4th edition, Athena Scientific, 2017), Abstract Dynamic Programming (2nd edition, Athena Scientific, 2018), and Nonlinear Programming (Athena Scientific, 2016). However, the mathematical style of this book is somewhat different. While we provide a rigorous, albeit short, mathematical account of the theory of finite and infinite horizon dynamic programming, and some fundamental approximation methods, we rely more on intuitive explanations and less on proof-based insights. Moreover, our mathematical requirements are quite modest: calculus, a minimal use of matrix-vector algebra, and elementary probability (mathematically complicated arguments involving laws of large numbers and stochastic convergence are bypassed in favor of intuitive explanations). The book illustrates the methodology with many examples and illustrations, and uses a gradual expository approach, which proceeds along four directions: (a) From exact DP to approximate DP: We first discuss exact DP algorithms, explain why they may be difficult to implement, and then use them as the basis for approximations. (b) From finite horizon to infinite horizon problems: We first discuss finite horizon exact and approximate DP methodologies, which are intuitive and mathematically simple, and then progress to infinite horizon problems. (c) From deterministic to stochastic models: We often discuss separately deterministic and stochastic problems, since deterministic problems are simpler and offer special advantages for some of our methods. (d) From model-based to model-free implementations: We first discuss model-based implementations, and then we identify schemes that can be appropriately modified to work with a simulator. The book is related and supplemented by the companion research monograph Rollout, Policy Iteration, and Distributed Reinforcement Learning (Athena Scientific, 2020), which focuses more closely on several topics related to rollout, approximate policy iteration, multiagent problems, discrete and Bayesian optimization, and distributed computation, which are either discussed in less detail or not covered at all in the present book. The author's website contains class notes, and a series of videolectures and slides from a 2021 course at ASU, which address a selection of topics from both books.

The book examines in some depth two important classes of point processes, determinantal processes and ``Gaussian zeros'', i.e., zeros of random analytic functions with Gaussian coefficients. These processes share a property of ``point-repulsion'', where distinct points are less likely to fall close to each other than in processes, such as the Poisson process, that arise from independent sampling. Nevertheless, the treatment in the book emphasizes the use of independence: for random power series, the independence of coefficients is key; for determinantal processes, the number of points in a domain is a sum of independent indicators, and this yields a satisfying explanation of the central limit theorem (CLT) for this point count. Another unifying theme of the book is invariance of considered point processes under natural transformation groups. The book strives for balance between general theory and concrete examples. On the one hand, it presents a primer on modern techniques on the interface of probability and analysis. On the other hand, a wealth of determinantal processes of intrinsic interest are analyzed; these arise from random spanning trees and eigenvalues of random matrices, as well as from special power series with determinantal zeros. The material in the book formed the basis of a graduate course given at the IAS-Park City Summer School in

2007; the only background knowledge assumed can be acquired in first-year graduate courses in analysis and probability.

Gaussian processes can be viewed as a far-reaching infinite-dimensional extension of classical normal random variables. Their theory presents a powerful range of tools for probabilistic modelling in various academic and technical domains such as Statistics, Forecasting, Finance, Information Transmission, Machine Learning - to mention just a few. The objective of these Briefs is to present a quick and condensed treatment of the core theory that a reader must understand in order to make his own independent contributions. The primary intended readership are PhD/Masters students and researchers working in pure or applied mathematics. The first chapters introduce essentials of the classical theory of Gaussian processes and measures with the core notions of reproducing kernel, integral representation, isoperimetric property, large deviation principle. The brevity being a priority for teaching and learning purposes, certain technical details and proofs are omitted. The later chapters touch important recent issues not sufficiently reflected in the literature, such as small deviations, expansions, and quantization of processes. In university teaching, one can build a one-semester advanced course upon these Briefs.?

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.

The two-volume set LNAI 12084 and 12085 constitutes the thoroughly refereed proceedings of the 24th Pacific-Asia Conference on Knowledge Discovery and Data Mining, PAKDD 2020, which was due to be held in Singapore, in May 2020. The conference was held virtually due to the COVID-19 pandemic. The 135 full papers presented were carefully reviewed and selected from 628 submissions. The papers present new ideas, original research results, and practical development experiences from all KDD related areas, including data mining, data warehousing, machine learning, artificial intelligence, databases, statistics, knowledge engineering, visualization, decision-making systems, and the emerging applications. They are organized in the following topical sections: recommender systems; classification; clustering; mining social networks; representation learning and embedding; mining behavioral data; deep learning; feature extraction and selection; human, domain, organizational and social factors in data mining; mining sequential data; mining imbalanced data; association; privacy and security; supervised learning; novel algorithms; mining multi-media/multi-dimensional data; application; mining graph and network data; anomaly detection and analytics; mining spatial, temporal, unstructured and semi-structured data; sentiment analysis; statistical/graphical model; multi-source/distributed/parallel/cloud computing.
High-dimensional probability offers insight into the behavior of random vectors, random matrices, random subspaces, and objects used to quantify uncertainty in high dimensions. Drawing on ideas from probability, analysis, and geometry, it lends itself to applications in mathematics, statistics, theoretical computer science, signal processing, optimization, and more. It is the first to integrate theory, key tools, and modern applications of high-dimensional probability. Concentration inequalities form the core, and it covers both classical results such

as Hoeffding's and Chernoff's inequalities and modern developments such as the matrix Bernstein's inequality. It then introduces the powerful methods based on stochastic processes, including such tools as Slepian's, Sudakov's, and Dudley's inequalities, as well as generic chaining and bounds based on VC dimension. A broad range of illustrations is embedded throughout, including classical and modern results for covariance estimation, clustering, networks, semidefinite programming, coding, dimension reduction, matrix completion, machine learning, compressed sensing, and sparse regression.

This book is devoted to a systematic analysis of asymptotic behavior of distributions of various typical functionals of Gaussian random variables and fields. The text begins with an extended introduction, which explains fundamental ideas and sketches the basic methods fully presented later in the book. Good approximate formulas and sharp estimates of the remainders are obtained for a large class of Gaussian and similar processes. The author devotes special attention to the development of asymptotic analysis methods, emphasizing the method of comparison, the double-sum method and the method of moments. The author has added an extended introduction and has significantly revised the text for this translation, particularly the material on the double-sum method.

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 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.

Now in its third edition, this classic book is widely considered the leading text on Bayesian methods, lauded for its accessible, practical approach to analyzing data and solving research problems. Bayesian Data Analysis, Third Edition continues to take an applied approach to analysis using up-to-date Bayesian methods. The authors—all leaders in the statistics community—introduce basic concepts from a data-analytic perspective before presenting advanced methods. Throughout the text, numerous worked examples drawn from real applications and research emphasize the use of Bayesian inference in practice. New to the Third Edition Four new chapters on nonparametric modeling Coverage of weakly informative priors and boundary-avoiding priors Updated discussion of cross-validation and predictive information criteria Improved convergence monitoring and effective sample size calculations for iterative simulation Presentations of Hamiltonian Monte Carlo, variational Bayes, and expectation propagation New and revised software code The book can be used in three different ways. For undergraduate students, it introduces Bayesian inference starting from first principles. For graduate students, the text presents effective current approaches to Bayesian modeling and computation in statistics and related fields. For researchers, it provides an assortment of Bayesian methods in applied statistics. Additional materials, including data sets used in the examples, solutions to selected exercises, and software instructions, are available on the book's web page.

This book serves as a standard reference, making this area accessible not only to researchers in probability and statistics, but also to graduate students and practitioners. The book assumes only a first-year graduate course in probability. Each chapter begins with a brief overview and concludes with a wide range of exercises at varying levels of difficulty. The authors supply detailed hints for the more challenging problems, and cover many advances made in recent years.

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 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.

This book was first published in 2006. Written by two of the foremost researchers in the field, this book studies the local times of Markov processes by employing isomorphism theorems that relate them to certain associated Gaussian processes. It builds to this material through self-contained but harmonized 'mini-courses' on the relevant ingredients, which assume only knowledge of measure-theoretic probability. The streamlined selection of topics creates an easy entrance for students and experts in related fields. The book starts by developing the fundamentals of Markov process theory and then of Gaussian process theory, including sample path properties. It then proceeds to more advanced results, bringing the reader to the heart of contemporary research. It presents the remarkable isomorphism theorems of Dynkin and Eisenbaum and then shows how they can be applied to obtain new properties of Markov processes by using well-established techniques in Gaussian process theory. This original, readable book will appeal to both researchers and advanced graduate students.

The fundamental question of characterizing continuity and boundedness of

Gaussian processes goes back to Kolmogorov. After contributions by R. Dudley and X. Fernique, it was solved by the author. This book provides an overview of "generic chaining", a completely natural variation on the ideas of Kolmogorov. It takes the reader from the first principles to the edge of current knowledge and to the open problems that remain in this domain.

In recent years neural computing has emerged as a practical technology, with successful applications in many fields. The majority of these applications are concerned with problems in pattern recognition, and make use of feedforward network architectures such as the multilayer perceptron and the radial basis function network. Also, it has become widely acknowledged that successful applications of neural computing require a principled, rather than ad hoc, approach. (From the preface to "Neural Networks for Pattern Recognition" by C.M. Bishop, Oxford Univ Press 1995.) This NATO volume, based on a 1997 workshop, presents a coordinated series of tutorial articles covering recent developments in the field of neural computing. It is ideally suited to graduate students and researchers.

"Robotic Mapping and Exploration" is an important contribution in the area of simultaneous localization and mapping (SLAM) for autonomous robots, which has been receiving a great deal of attention by the research community in the

latest few years. The contents are focused on the autonomous mapping learning problem. Solutions include uncertainty-driven exploration, active loop closing, coordination of multiple robots, learning and incorporating background knowledge, and dealing with dynamic environments. Results are accompanied by a rich set of experiments, revealing a promising outlook toward the application to a wide range of mobile robots and field settings, such as search and rescue, transportation tasks, or automated vacuum cleaning.
This monograph opens up new horizons for engineers and researchers in academia and in industry dealing with or interested in new developments in the field of system identification and control. It emphasizes guidelines for working solutions and practical advice for their implementation rather than the theoretical background of Gaussian process (GP) models. The book demonstrates the potential of this recent development in probabilistic machine-learning methods and gives the reader an intuitive understanding of the topic. The current state of the art is treated along with possible future directions for research. Systems control design relies on mathematical models and these may be developed from measurement data. This process of system identification, when based on GP models, can play an integral part of control design in data-based control and its description as such is an essential aspect of the text. The background of GP

regression is introduced first with system identification and incorporation of prior knowledge then leading into full-blown control. The book is illustrated by extensive use of examples, line drawings, and graphical presentation of computer-simulation results and plant measurements. The research results presented are applied in real-life case studies drawn from successful applications including: a gas–liquid separator control; urban-traffic signal modelling and reconstruction; and prediction of atmospheric ozone concentration. A MATLAB® toolbox, for identification and simulation of dynamic GP models is provided for download.

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.

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.
For almost fifty years, Richard M. Dudley has been extremely influential in the development of several areas of Probability. His work on Gaussian processes led to the understanding of the basic fact that their sample boundedness and continuity should be characterized in terms of proper measures of complexity of their parameter spaces equipped with the intrinsic covariance metric. His sufficient condition for sample continuity in terms of metric entropy is widely used and was proved by X. Fernique to be necessary for stationary Gaussian processes, whereas its more subtle versions (majorizing measures) were proved by M. Talagrand to be necessary in general. Together with V. N. Vapnik and A. Y. Cervonenkis, R. M. Dudley is a founder of the modern theory of empirical

processes in general spaces. His work on uniform central limit theorems (under bracketing entropy conditions and for Vapnik-Cervonenkis classes), greatly extends classical results that go back to A. N. Kolmogorov and M. D. Donsker, and became the starting point of a new line of research, continued in the work of Dudley and others, that developed empirical processes into one of the major tools in mathematical statistics and statistical learning theory. As a consequence of Dudley's early work on weak convergence of probability measures on non-separable metric spaces, the Skorohod topology on the space of regulated right-continuous functions can be replaced, in the study of weak convergence of the empirical distribution function, by the supremum norm. In a further recent step Dudley replaces this norm by the stronger p-variation norms, which then allows replacing compact differentiability of many statistical functionals by Fréchet differentiability in the delta method. Richard M. Dudley has also made important contributions to mathematical statistics, the theory of weak convergence, relativistic Markov processes, differentiability of nonlinear operators and several other areas of mathematics. Professor Dudley has been the adviser to thirty PhD's and is a Professor of Mathematics at the Massachusetts Institute of Technology.

Branching Brownian motion (BBM) is a classical object in probability theory with deep connections to partial differential equations. This book highlights the connection to classical extreme value theory and to the theory of mean-field spin glasses in statistical mechanics. Starting with a concise review of classical extreme value statistics and a basic introduction to mean-field spin glasses, the author then focuses on branching Brownian motion. Here, the classical results of Bramson on the asymptotics of solutions of the F-KPP equation are reviewed in detail and applied to the recent construction of the extremal process of BBM. The extension of these results to branching Brownian motion with variable speed are then explained. As a self-contained exposition that is accessible to graduate students with some background in probability theory, this book makes a good introduction for anyone interested in accessing this exciting field of mathematics.

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.

Computer simulation experiments are essential to modern scientific discovery, whether that be in physics, chemistry, biology, epidemiology, ecology, engineering, etc. Surrogates are metamodels 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.

Simulation of materials at the atomistic level is an important tool in studying microscopic structures and processes. The atomic interactions necessary for the simulations are correctly described by Quantum Mechanics, but the size of systems and the length of processes that can be modelled are still limited. The framework of Gaussian Approximation Potentials that is developed in this thesis allows us to generate interatomic potentials automatically, based on quantum mechanical data. The resulting potentials offer several orders of magnitude faster computations, while maintaining quantum mechanical accuracy. The method has already been successfully applied for semiconductors and metals.

Stochastic Analysis for Gaussian Random Processes and Fields: With Applications presents Hilbert space methods to study deep analytic properties connecting probabilistic notions. In particular, it studies Gaussian random fields using reproducing kernel Hilbert spaces (RKHSs).The book begins with preliminary results on covariance and associated RKHS