Probabilistic Graphical Models Solutions

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 idea of modelling systems using graph theory has its origin in several scientific areas: in statistical physics (the study of large particle systems), in genetics (studying inheritable properties of natural species), and in interactions in contingency tables. The use of graphical models in statistics has increased considerably over recent years and the theory has been greatly developed and extended. This book provides the first comprehensive and authoritative account of the theory of graphical models and is written by a leading expert in the field. It contains the fundamental graph theory required and a thorough study of Markov properties associated with various type of graphs. The statistical theory of log-linear and graphical models for contingency tables, covariance selection models, and graphical models with mixed discrete-continous variables in developed detail. Special topics, such as the application of graphical models to probabilistic expert systems, are described briefly, and appendices give details of the multivarate normal distribution and of the theory of regular exponential families. The author has recently been awarded the RSS Guy Medal in Silver 1996 for his innovative contributions to statistical theory and practice, and especially for his work on graphical models.

As the power of Bayesian techniques has become more fully realized, the field of artificial intelligence has embraced Bayesian methodology and integrated it to the point where an introduction to Bayesian techniques is now a core course in many computer science programs. Unlike other books on the subject, Bayesian Artificial Intelligence keeps mathematical detail to a minimum and covers a broad range of topics. The authors integrate all of Bayesian net technology and learning Bayesian net technology and apply them both to knowledge engineering. They emphasize understanding and intuition but also provide the algorithms and technical background needed for applications. Software, exercises, and solutions are available on the authors' website.

The fundamental mathematical tools needed to understand machine learning include linear algebra, analytic geometry, matrix decompositions, vector calculus, optimization, probability and statistics. These topics are traditionally taught in disparate courses, making it hard for data science or computer science students, or professionals, to efficiently learn the mathematics. This self-contained textbook bridges the gap between mathematical and machine learning texts, introducing the mathematical concepts with a minimum of prerequisites. It uses these concepts to derive four central machine learning methods: linear regression, principal component analysis, Gaussian mixture models and support vector machines. For students and others with a mathematical background, these derivations provide a starting point to machine learning texts. For those learning the mathematics for the first time, the methods help build intuition and practical experience with applying mathematical concepts. Every chapter includes worked examples and exercises to test understanding. Programming tutorials are offered on the book's web site.

This fully updated new edition of a uniquely accessible textbook/reference provides a general introduction to probabilistic graphical models (PGMs) from an engineering perspective. It features new material on partially observable Markov decision processes, graphical models, and deep learning, as well as an even greater number of exercises. The book covers the fundamentals for each of the main classes of PGMs, including representation, inference and learning principles, and reviews real-world applications for each type of model. These applications are drawn from a broad range of disciplines, highlighting the many uses of Bayesian classifiers, hidden Markov models, Bayesian networks, dynamic and temporal Bayesian networks, Markov random fields, influence diagrams, and Markov decision processes. Topics and features: Presents a unified framework encompassing all of the main classes of PGMs Explores the fundamental aspects of representation, inference and learning for each technique Examines new material on partially observable Markov decision processes, and graphical models Includes a new chapter introducing deep neural networks and their relation with probabilistic graphical models Covers multidimensional Bayesian classifiers, relational graphical models, and causal models Provides substantial chapter-ending exercises, suggestions for further reading, and ideas for research or programming projects Describes classifiers such as Gaussian Naive Bayes, Circular Chain Classifiers, and Hierarchical Classifiers with Bayesian Networks Outlines the practical application of the different techniques Suggests possible course outlines for instructors This classroom-tested work is suitable as a textbook for an advanced undergraduate or a graduate course in probabilistic graphical models for students of computer science, engineering, and physics. Professionals wishing to apply probabilistic graphical models in their own field, or interested in the basis of these techniques, will also find the book to be an invaluable reference. Dr. Luis Enrique Sucar is a Senior Research Scientist at the National Institute for Astrophysics, Optics and Electronics (INAOE), Puebla, Mexico. He received the National Science Prize en 2016.

Sampling consists of selection, acquisition, and quantification of a part of the population. While selection and acquisition apply to physical sampling units of the population, quantification pertains only to the variable of interest, which is a particular characteristic of the sampling units. A sampling procedure is expected to provide a sample that is representative with respect to some specified criteria. Composite sampling, under idealized conditions, incurs no loss of information for estimating the population means. But an important limitation to the method has been the loss of information on individual sample values, such as, the extremely large value. In many of the situations where individual sample values are of interest or concern, composite sampling methods can be suitably modified to retrieve the information on individual sample values that may be lost due to compositing. This book presents statistical solutions to issues that arise in the context of applications of composite sampling.

Bayesian Networks: With Examples in R, Second Edition introduces Bayesian networks using a hands-on approach. Simple yet meaningful examples illustrate each step of the modelling process and discuss side by side the underlying theory and its application using R code. The examples start from the simplest notions and gradually increase in complexity. In particular, this new edition contains significant new material on topics from modern machine-learning practice: dynamic networks, networks with heterogeneous variables, and model validation. The first three chapters explain the whole process of Bayesian network modelling, from structure learning to parameter learning to inference. These chapters cover discrete, Gaussian, and conditional Gaussian Bayesian networks. The following two chapters delve into dynamic networks (to model temporal data) and into networks including arbitrary random variables (using Stan). The book then gives a concise but rigorous treatment of the fundamentals of Bayesian networks and offers an introduction to causal Bayesian networks. It also presents an overview of R packages and other software implementing Bayesian networks. The final chapter evaluates two real-world examples: a landmark causal protein-signalling network published in Science and a probabilistic graphical model for predicting the composition of different body parts. Covering theoretical and practical aspects of Bayesian networks, this book provides you with an introductory overview of the field. It gives you a clear, practical understanding of the key points behind this modelling approach and, at the same time, it makes you familiar with the most relevant packages used to implement real-world analyses in R. The examples covered in the book span several application fields, data-driven models and expert systems, probabilistic and causal perspectives, thus giving you a starting point to work in a variety of scenarios. Online supplementary materials include the data sets and the code used in the book, which will all be made available from https://www.bnlearn.com/book-crc-2ed/

This book exemplifies the interplay between the general formal framework of graphical models and the exploration of new algorithm and architectures. The selections range from foundational papers of historical importance to results at the cutting edge of research. Graphical models use graphs to represent and manipulate joint probability distributions. They have their roots in artificial intelligence, statistics, and neural networks. The clean mathematical formalism of the graphical models framework makes it possible to understand a wide variety of network-based approaches to computation, and in particular to understand many neural network algorithms and architectures as instances of a broader probabilistic methodology. It also makes it possible to identify novel features of neural network algorithms and architectures and to extend them to more general graphical models.This book exemplifies the interplay between the general formal framework of graphical models and the exploration of new algorithms and architectures. The selections range from foundational papers of historical importance to results at the cutting edge of research. Contributors H. Attias, C. M. Bishop, B. J. Frey, Z. Ghahramani, D. Heckerman, G. E. Hinton, R. Hofmann, R. A. Jacobs, Michael I. Jordan, H. J. Kappen, A. Krogh, R. Neal, S. K. Riis, F. B. Rodríguez, L. K. Saul, Terrence J. Sejnowski, P. Smyth, M. E. Tipping, V. Tresp, Y. Weiss

The work reviewed in this book represents the synthesis of two important developments in modelling of complex stochastic phenomena. The book gives a thorough and rigorous mathematical treatment of the underlying ideas, structures, and algorithms.

This book provides a thorough introduction to the formal foundations and practical applications of Bayesian networks. It provides an extensive discussion of techniques for building Bayesian networks that model real-world situations, including techniques for synthesizing models from design, learning models from data, and debugging models using sensitivity analysis. It also treats exact and approximate inference algorithms at both theoretical and practical levels. The author assumes very little background on the covered subjects, supplying in-depth discussions for theoretically inclined readers and enough practical details to provide an algorithmic cookbook for the system developer.

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

Probabilistic Graphical ModelsPrinciples and TechniquesMIT Press

A general framework for constructing and using probabilistic models of complex systems that would enable a computer to use available information for making decisions. Most tasks require a person or an automated system to reason—to reach conclusions based on available information. The framework of probabilistic graphical models, presented in this book, provides a general approach for this task. The approach is model-based, allowing interpretable models to be constructed and then manipulated by reasoning algorithms. These models can also be learned automatically from data, allowing the approach to be used in cases where manually constructing a model is difficult or even impossible. Because uncertainty is an inescapable aspect of most real-world applications, the book focuses on probabilistic models, which make the uncertainty explicit and provide models that are more faithful to reality. Probabilistic Graphical Models discusses a variety of models, spanning Bayesian networks, undirected Markov networks, discrete and continuous models, and extensions to deal with dynamical systems and relational data. For each class of models, the text describes the three fundamental cornerstones: representation, inference, and learning, presenting both basic concepts and advanced techniques. Finally, the book considers the use of the proposed framework for causal reasoning and decision making under uncertainty. The main text in each chapter provides the detailed technical development of the key ideas. Most chapters also include boxes with additional material: skill boxes, which describe techniques; case study boxes, which discuss empirical cases related to the approach described in the text, including applications in computer vision, robotics, natural language understanding, and computational biology; and concept boxes, which present significant concepts drawn from the material in the chapter. Instructors (and readers) can group chapters in various combinations, from core topics to more technically advanced material, to suit their particular needs.

Appropriate - Many multivariate probabilistic models either use independent distributions or dependent Gaussian distributions. Yet, many real-world datasets contain count-valued or non-negative skewed data, e.g. bag-of-words text data and biological sequencing data. Thus, we develop novel probabilistic graphical models for use on count-valued and non-negative data including Poisson graphical models and multinomial graphical models. We develop one generalization that allows for triple-wise or k-wise graphical models going beyond the normal pairwise formulation. Furthermore, we also explore Gaussian-copula graphical models and derive closed-form solutions for the conditional distributions and marginal distributions (both before and after conditioning). Finally, we derive mixture and admixture, or topic model, generalizations of these graphical models to introduce more power and interpretability. Accessible - Previous multivariate models, especially related to text data, often have complex dependencies without a closed form and require complex inference algorithms that have limited theoretical justification. For example, hierarchical Bayesian models often require marginalizing over many latent variables. We show that our novel graphical models (even the k-wise interaction models) have simple and intuitive estimation procedures based on node-wise regressions that likely have similar theoretical guarantees as previous work in graphical models. For the copula-based graphical models, we show that simple approximations could still provide useful models; these copula models also come with closed-form conditional and marginal distributions, which make them amenable to exploratory inspection and manipulation. The parameters of these models are easy to interpret and thus may be accessible to a wide audience. Appealing - High-level visualization and interpretation of graphical models with even 100 variables has often been difficult even for a graphical model expert---despite visualization being one of the original motivators for graphical models. This difficulty is likely due to the lack of collaboration between graphical model experts and visualization experts. To begin bridging this gap, we develop a novel "what if?" interaction that manipulates and leverages the probabilistic power of graphical models. Our approach defines: the probabilistic mechanism via conditional probability; the query language to map text input to a conditional probability query; and the formal underlying probabilistic model. We then propose to visualize these query-specific probabilistic graphical models by combining the intuitiveness of force-directed layouts with the beauty and readability of word clouds, which pack many words into valuable screen space while ensuring words do not overlap via pixel-level collision detection. Although both the force-directed layout and the pixel-level packing problems are challenging in their own right, we approximate both simultaneously via adaptive simulated annealing starting from careful initialization. For visualizing mixture distributions, we also design a meaningful mapping from the properties of the mixture distribution to a color in the perceptually uniform CIELUV color space. Finally, we demonstrate our approach via illustrative visualizations of several real-world datasets.

Summary Practical Probabilistic Programming introduces the working programmer to probabilistic programming. In it, you'll learn how to use the PP paradigm to model application domains and then express those probabilistic models in code. Although PP can seem abstract, in this book you'll immediately work on practical examples, like using the Figaro language to build a spam filter and applying Bayesian and Markov networks, to diagnose computer system data problems and recover digital images. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology The data you accumulate about your customers, products, and website users can help you not only to interpret your past, it can also help you predict your future! Probabilistic programming uses code to draw probabilistic inferences from data. By applying specialized algorithms, your programs assign degrees of probability to conclusions. This means you can forecast future events like sales trends, computer system failures, experimental outcomes, and many other critical concerns. About the Book Practical Probabilistic Programming introduces the working programmer to probabilistic programming. In this book, you'll immediately work on practical examples like building a spam filter, diagnosing computer system data problems, and recovering digital images. You'll discover probabilistic inference, where algorithms help make extended predictions about issues like social media usage. Along the way, you'll learn to use functional-style programming for text analysis, object-oriented models to predict social phenomena like the spread of tweets, and open universe models to gauge real-life social media usage. The book also has chapters on how probabilistic models can help in decision making and modeling of dynamic systems. What's Inside Introduction to probabilistic modeling Writing probabilistic programs in Figaro Building Bayesian networks Predicting product lifecycles Decision-making algorithms About the Reader This book assumes no prior exposure to probabilistic programming. Knowledge of Scala is helpful. About the Author Avi Pfeffer is the principal developer of the Figaro language for probabilistic programming. Table of Contents PART 1 INTRODUCING PROBABILISTIC PROGRAMMING AND FIGARO Probabilistic programming in a nutshell A quick Figaro tutorial Creating a probabilistic programming application PART 2 WRITING PROBABILISTIC PROGRAMS Probabilistic models and probabilistic programs Modeling dependencies with Bayesian and Markov networks Using Scala and Figaro collections to build up models Object-oriented probabilistic modeling Modeling dynamic systems PART 3 INFERENCE The three rules of probabilistic inference Factored inference algorithms Sampling algorithms Solving other inference tasks Dynamic reasoning and parameter learning

Master Bayesian Inference through Practical Examples and Computation–Without Advanced Mathematical Analysis Bayesian methods of inference are deeply natural and extremely powerful. However, most discussions of Bayesian inference rely on intensely complex mathematical analyses and artificial examples, making it inaccessible to anyone without a strong mathematical background. Now, though, Cameron Davidson-Pilon introduces Bayesian inference from a computational perspective, bridging theory to practice–freeing you to get results using computing power. Bayesian Methods for Hackers illuminates Bayesian inference through probabilistic programming with the powerful PyMC language and the closely related Python tools NumPy, SciPy, and Matplotlib. Using this approach, you can reach effective solutions in small increments, without extensive mathematical intervention. Davidson-Pilon begins by introducing the concepts underlying Bayesian inference, comparing it with other techniques and guiding you through building and training your first Bayesian model. Next, he introduces PyMC through a series of detailed examples and intuitive explanations that have been refined after extensive user feedback. You'll learn how to use the Markov Chain Monte Carlo algorithm, choose appropriate sample sizes and priors, work with loss functions, and apply Bayesian inference in domains ranging from finance to marketing. Once you've mastered these techniques, you'll constantly turn to this guide for the working PyMC code you need to jumpstart future projects. Coverage includes • Learning the Bayesian "state of mind" and its practical implications • Understanding how computers perform Bayesian inference • Using the PyMC Python library to program Bayesian analyses • Building and debugging models with PyMC • Testing your model's "goodness of fit" • Opening the "black box" of the Markov Chain Monte Carlo algorithm to see how and why it works • Leveraging the power of the "Law of Large Numbers" • Mastering key concepts, such as clustering, convergence, autocorrelation, and thinning • Using loss functions to measure an estimate's weaknesses based on your goals and desired outcomes • Selecting appropriate priors and understanding how their influence changes with dataset size • Overcoming the "exploration versus exploitation" dilemma: deciding when "pretty good" is good enough • Using Bayesian inference to improve A/B testing • Solving data science problems when only small amounts of data are available Cameron Davidson-Pilon has worked in many areas of applied mathematics, from the evolutionary dynamics of genes and diseases to stochastic modeling of financial prices. His contributions to the open source community include lifelines, an implementation of survival analysis in Python. Educated at the University of Waterloo and at the Independent University of Moscow, he currently works with the online commerce leader Shopify.

Advanced statistical modeling and knowledge representation techniques for a newly emerging area of machine learning and probabilistic reasoning; includes introductory material, tutorials for different proposed approaches, and applications. Handling inherent uncertainty and exploiting compositional structure are fundamental to understanding and designing large-scale systems. Statistical relational learning builds on ideas from probability theory and statistics to address uncertainty while incorporating tools from logic, databases and programming languages to represent structure. In Introduction to Statistical Relational Learning, leading researchers in this emerging area of machine learning describe current formalisms, models, and algorithms that enable effective and robust reasoning about richly structured systems and data. The early chapters provide tutorials for material used in later chapters, offering introductions to representation, inference and learning in graphical models, and logic. The book then describes object-oriented approaches, including probabilistic relational models, relational Markov networks, and probabilistic entity-relationship models as well as logic-based formalisms including Bayesian logic programs, Markov logic, and stochastic logic programs. Later chapters discuss such topics as probabilistic models with unknown objects, relational dependency networks, reinforcement learning in relational domains, and information extraction. By presenting a variety of approaches, the book highlights commonalities and clarifies important differences among proposed approaches and, along the way, identifies important representational and algorithmic issues. Numerous applications are provided throughout.

One of the goals of artificial intelligence (AI) is creating autonomous agents that must make decisions based on uncertain and incomplete information. The goal is to design rational agents that must take the best action given the information available and their goals. Decision Theory Models for Applications in Artificial Intelligence: Concepts and Solutions provides an introduction to different types of decision theory techniques, including MDPs, POMDPs, Influence Diagrams, and Reinforcement Learning, and illustrates their application in artificial intelligence. This book provides insights into the advantages and challenges of using decision theory models for developing intelligent systems.

This is the first textbook on pattern recognition to present the Bayesian viewpoint. The book presents approximate inference algorithms that permit fast approximate answers in situations where exact answers are not feasible. It uses graphical models to describe probability distributions when no other books apply graphical models to machine learning. No previous knowledge of pattern recognition or machine learning concepts is assumed. Familiarity with multivariate calculus and basic linear algebra is required, and some experience in the use of probabilities would be helpful though not essential as the book includes a self-contained introduction to basic probability theory.

Probabilistic Graphical Models for Computer Vision introduces probabilistic graphical models (PGMs) for computer vision problems and teaches how to develop the PGM model from training data. This book discusses PGMs and their significance in the context of solving computer vision problems, giving the basic concepts, definitions and properties. It also provides a comprehensive introduction to well-established theories for different types of PGMs, including both directed and undirected PGMs, such as Bayesian Networks, Markov Networks and their variants. Discusses PGM theories and techniques with computer vision examples Focuses on well-established PGM theories that are accompanied by corresponding pseudocode for computer vision Includes an extensive list of references, online resources and a list of publicly available and commercial software Covers computer vision tasks, including feature extraction and image segmentation, object and facial recognition, human activity recognition, object tracking and 3D reconstruction

Trains researchers and graduate students in state-of-the-art statistical and machine learning methods to build models with real-world data.

A graphical model is a statistical model that is represented by a graph. The factorization properties underlying graphical models facilitate tractable computation with multivariate distributions, making the models a valuable tool with a plethora of applications. Furthermore, directed graphical models allow intuitive causal interpretations and have become a cornerstone for causal inference. While there exist a number of excellent books on graphical models, the field has grown so much that individual authors can hardly cover its entire scope. Moreover, the field is interdisciplinary by nature. Through chapters by leading researchers from different areas, this handbook provides a broad and accessible overview of the state of the art. Key features: * Contributions by leading researchers from a range of disciplines * Structured in five parts, covering foundations, computational aspects, statistical inference, causal inference, and applications * Balanced coverage of concepts, theory, methods, examples, and applications * Chapters can be read mostly independently, while cross-references highlight connections The handbook is targeted at a wide audience, including graduate students, applied researchers, and experts in graphical models.

This book presents an exciting new synthesis of directed and undirected, discrete and continuous graphical models. Combining elements of Bayesian networks and Markov random fields, the newly introduced hybrid random fields are an interesting approach to get the best of both these worlds, with an added promise of modularity and scalability. The authors have written an enjoyable book---rigorous in the treatment of the mathematical background, but also enlivened by interesting and original historical and philosophical perspectives. -- Manfred Jaeger, Aalborg Universitet The book not only marks an effective direction of investigation with significant experimental advances, but it is also---and perhaps primarily---a guide for the reader through an original trip in the space of probabilistic modeling. While digesting the book, one is enriched with a very open view of the field, with full of stimulating connections. [...] Everyone specifically interested in Bayesian networks and Markov random fields should not miss it. -- Marco Gori, Università degli Studi di Siena Graphical models are sometimes regarded---incorrectly---as an impractical approach to machine learning, assuming that they only work well for low-dimensional applications and discrete-valued domains. While guiding the reader through the major achievements of this research area in a technically detailed yet accessible way, the book is concerned with the presentation and thorough (mathematical and experimental) investigation of a novel paradigm for probabilistic graphical modeling, the hybrid random field. This model subsumes and extends both Bayesian networks and Markov random fields. Moreover, it comes with well-defined learning algorithms, both for discrete and continuous-valued domains, which fit the needs of real-world applications involving large-scale, high-dimensional data.

Master probabilistic graphical models by learning through real-world problems and illustrative code examples in Python About This Book Gain in-depth knowledge of Probabilistic Graphical Models Model time-series problems using Dynamic Bayesian Networks A practical guide to help you apply PGMs to real-world problems Who This Book Is For If you are a researcher or a machine learning enthusiast, or are working in the data science field and have a basic idea of Bayesian Learning or Probabilistic Graphical Models, this book will help you to understand the details of Graphical Models and use it in your data science problems. This book will also help you select the appropriate model as well as the appropriate algorithm for your problem. What You Will Learn Get to know the basics of Probability theory and Graph Theory Work with Markov Networks Implement Bayesian Networks Exact Inference Techniques in Graphical Models such as the Variable Elimination Algorithm Understand approximate Inference Techniques in Graphical Models such as Message Passing Algorithms Sample algorithms in Graphical Models Grasp details of Naive Bayes with real-world examples Deploy PGMs using various libraries in Python Gain working details of Hidden Markov Models with real-world examples In Detail Probabilistic Graphical Models is a technique in machine learning that uses the concepts of graph theory to compactly represent and optimally predict values in our data problems. In real world problems, it's often difficult to select the appropriate graphical model as well as the appropriate inference algorithm, which can make a huge difference in computation time and accuracy. Thus, it is crucial to know the working details of these algorithms. This book starts with the basics of probability theory and graph theory, then goes on to discuss various models and inference algorithms. All the different types of models are discussed along with code examples to create and modify them, and also to run different inference algorithms on them. There is a complete chapter devoted to the most widely used networks Naive Bayes Model and Hidden Markov Models (HMMs). These models have been thoroughly discussed using real-world examples. Style and approach An easy-to-follow guide to help you understand Probabilistic Graphical Models using simple examples and numerous code examples, with an emphasis on more widely used models.

At the crossroads between statistics and machine learning, probabilistic graphical models provide a powerful formal framework to model complex data. For instance, Bayesian networks and Markov random fields are two of the most popular probabilistic graphical models. With the rapid advance of high-throughput technologies and their ever decreasing costs, a fast-growing volume of biological data of various types - the so-called "omics" - is in need of accurate andefficient methods for modeling, prior to further downstream analysis. As probabilistic graphical models are able to deal with high-dimensional data, it is foreseeable that such models will have aprominent role to play in advances in genome-wide data analyses. Currently, few people are specialists in the design of cutting-edge methods using probabilistic graphical models for genetics, genomics and postgenomics. This seriously hinders the diffusion of such methods. The prime aim of the book is therefore to bring the concepts underlying these advanced models within reach of scientists who are not specialists of these models, but with no concession on theinformativeness of the book. The target readers include researchers and engineers who have to design novel methods for postgenomics data analysis, as well as graduate students starting a Masters or a PhD. Inaddition to an introductory chapter on probabilistic graphical models, a thorough review chapter focusing on selected domains in genetics and fourteen chapters illustrate the design of such advanced approaches in various domains: gene network inference, inference of causal phenotype networks, association genetics, epigenetics, detection of copy number variations, and prediction of outcomes from high-dimensional genomic data. Notably, most examples also illustrate that probabilistic graphicalmodels are well suited for integrative biology and systems biology, hot topics guaranteed to be of lasting interest.

Disk contains: Tool for building Bayesian networks -- Library of examples -- Library of proposed solutions to some exercises.

For advanced students of network data science, this compact account covers both well-established methodology and the theory of models recently introduced in the graphical model literature. It focuses on the discrete case where all variables involved are categorical and, in this context, it achieves a unified presentation of classical and recent results.

In the past decade, a number of different research communities within the computational sciences have studied learning in networks, starting from a number of different points of view. There has been substantial progress in these different communities and surprising convergence has developed between the formalisms. The awareness of this convergence and the growing interest of researchers in understanding the essential unity of the subject underlies the current volume. Two research communities which have used graphical or network formalisms to particular advantage are the belief network community and the neural network community. Belief networks arose within computer science and statistics and were developed with an emphasis on prior knowledge and exact probabilistic calculations. Neural networks arose within electrical engineering, physics and neuroscience and have emphasised pattern recognition and systems modelling problems. This volume draws together researchers from these two communities and presents both kinds of networks as instances of a general unified graphical formalism. The book focuses on probabilistic methods for learning and inference in graphical models, algorithm analysis and design, theory and applications. Exact methods, sampling methods and variational methods are discussed in detail. Audience: A wide cross-section of computationally oriented researchers, including computer scientists, statisticians, electrical engineers, physicists and neuroscientists.

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.

Table of contents

Bayesian Networks: An Introduction provides a self-contained introduction to the theory and applications of Bayesian networks, a topic of interest and importance for statisticians, computer scientists and those involved in modelling complex data sets. The material has been extensively tested in classroom teaching and assumes a basic knowledge of probability, statistics and mathematics. All notions are carefully explained and feature exercises throughout. Features include: An introduction to Dirichlet Distribution, Exponential Families and their applications. A detailed description of learning algorithms and Conditional Gaussian Distributions using Junction Tree methods. A discussion of Pearl's intervention calculus, with an introduction to the notion of see and do conditioning. All concepts are clearly defined and illustrated with examples and exercises. Solutions are provided online. This book will prove a valuable resource for postgraduate students of statistics, computer engineering, mathematics, data mining, artificial intelligence, and biology. Researchers and users of comparable modelling or statistical techniques such as neural networks will also find this book of interest.

With the recent and enormous increase in the amount of available data sets of all kinds, applying effective and efficient techniques for analyzing and extracting information from that data has become a crucial task. Intelligent Data Analysis for Real-Life Applications: Theory and Practice investigates the application of Intelligent Data Analysis (IDA) to these data sets through the design and development of algorithms and techniques to extract knowledge from databases. This pivotal reference explores practical applications of IDA, and it is essential for academic and research libraries as well as students, researchers, and educators in data analysis, application development, and database management.

Bayesian Networks and Influence Diagrams: A Guide to Construction and Analysis, Second Edition, provides a comprehensive guide for practitioners who wish to understand, construct, and analyze intelligent systems for decision support based on probabilistic networks. This new edition contains six new sections, in addition to fully-updated examples, tables, figures, and a revised appendix. Intended primarily for practitioners, this book does not require sophisticated mathematical skills or deep understanding of the underlying theory and methods nor does it discuss alternative technologies for reasoning under uncertainty. The theory and methods presented are illustrated through more than 140 examples, and exercises are included for the reader to check his or her level of understanding. The techniques and methods presented for knowledge elicitation, model construction and verification, modeling techniques and tricks, learning models from data, and analyses of models have all been developed and refined on the basis of numerous courses that the authors have held for practitioners worldwide.

This monograph is about discrete energy minimization for discrete graphical models. It considers graphical models, or, more precisely, maximum a posteriori inference for graphical models, purely as a combinatorial optimization problem.

This is a brand new edition of an essential work on Bayesian networks and decision graphs. It is an introduction to probabilistic graphical models including Bayesian networks and influence diagrams. The reader is guided through the two types of frameworks with examples and exercises, which also give instruction on how to build these models. Structured in two parts, the first section focuses on probabilistic graphical models, while the second part deals with decision graphs, and in addition to the frameworks described in the previous edition, it also introduces Markov decision process and partially ordered decision problems.

The core of this paper is a general set of variational principles for the problems of computing marginal probabilities and modes, applicable to multivariate statistical models in the exponential family.

Electrical activity in the myocardium coordinates the contraction of the heart, and its knowledge could lead to a better understanding, diagnosis, and treatment of cardiac diseases. This electrical activity generates an electromagnetic field that propagates outside the heart and reaches the human torso surface, where it can be easily measured. Classical electrocardiography aims to interpret the 12-lead electrocardiogram (ECG) to determine cardiac activity and support the diagnosis of cardiac pathologies such as arrhythmias, altered activations, and ischemia. More recently, a higher number of leads is used to reconstruct a more detailed quantitative description of the electrical activity in the heart by solving the so-called inverse problem of electrocardiography. This technique is known as ECG imaging. Today, clinical applications of ECG imaging are showing promising results in guiding a variety of electrophysiological interventions such as catheter ablation of atrial fibrillation and ventricular tachycardia. However, in order to promote the adoption of ECG imaging in the routine clinical practice, further research is required regarding more accurate mathematical methods, further scientific validation under different preclinical scenarios and a more extensive clinical validation

Reviews the use of factor graphs for the modeling and solving of large-scale inference problems in robotics. Factor graphs are introduced as an economical representation within which to formulate the different inference problems, setting the stage for the subsequent sections on practical methods to solve them.