Probabilistic Graphical Models Solutions

Familiarize yourself with probabilistic graphical models through real-world problems and illustrative code examples in R About This Book Predict and use a probabilistic graphical models (PGM) as an expert system Comprehend how your computer can learn Bayesian modeling to solve real-world problems Know how to prepare data and feed the models by using the appropriate algorithms from the appropriate R package Who This Book Is For This book is for anyone who has to deal with lots of data and draw conclusions from it, especially when the data is noisy or uncertain. Data scientists, machine learning enthusiasts, engineers, and those who curious about the latest advances in machine learning will find PGM interesting. What You Will Learn Understand the concepts of PGM and which type of PGM to use for which problem Tune the model's parameters and explore new models automatically Understand the basic principles of Bayesian models, from simple to advanced Transform the old linear regression model into a powerful probabilistic model Use standard industry models but with the power of PGM Understand the advanced models used throughout today's industry See how to compute posterior distribution with exact and approximate inference algorithms In Detail Probabilistic graphical models (PGM, also known as graphical models) are a marriage between probability theory and graph theory. Generally, PGMs use a graph-based representation. Two branches of graphical representations of distributions are commonly used, namely Bayesian networks and Markov networks. R has many packages to implement graphical models. We'll start by showing you how to transform a classical statistical model into a modern PGM and then look at how to do exact inference in graphical models. Proceeding, we'll introduce you to many modern R packages that will help you to perform inference on the models. We will then run a Bayesian linear regression and you'll see the advantage of going probabilistic when you want to do prediction. Next, you'll master using R packages and implementing its techniques. Finally, you'll be presented with machine learning applications that have a direct impact in many fields. Here, we'll cover clustering and the discovery of hidden information in big data, as well as two important methods, PCA and ICA, to reduce the size of big problems. Style and approach This book gives you a detailed and step-by-step explanation of each mathematical concept, which will help you build and analyze your own machine learning models and apply them to real-world problems. The mathematics is kept simple and each formula is explained thoroughly.

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

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.

Developed from celebrated Harvard statistics lectures, Introduction to Probability provides essential language and tools for understanding statistics, randomness, and uncertainty. The book explores a wide variety of applications and examples, ranging from coincidences and paradoxes to Google PageRank and Markov chain Monte Carlo (MCMC). Additional

Machine Learning and Probabilistic Graphical Models for Decision Support Systems

Graphical Models, Exponential Families, and Variational Inference

Introduction to Statistical Relational Learning

Machine Learning Quick Reference

Handbook of Graphical Models

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.

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 presents recent advancements in research, a review of new methods and techniques, and applications in decision support systems (DSS) with Machine Learning and Probabilistic Graphical Models, which are very effective techniques in gaining knowledge from Big Data and in interpreting decisions. It explores Bayesian network learning, Control Chart, Reinforcement Learning for multicriteria DSS, Anomaly Detection in Smart Manufacturing with Federated Learning, DSS in healthcare, DSS for supply chain management, etc. Researchers and practitioners alike will benefit from this book to enhance the understanding of machine learning, Probabilistic Graphical Models, and their uses in DSS in the context of decision making with uncertainty. The real-world case studies in various fields with guidance and recommendations for the practical applications of these studies are introduced in each chapter.

Probabilistic Graphical ModelsPrinciples and TechniquesMIT Press

Sampling Algorithms for Probabilistic Graphical Models with Determinism

Quick and essential machine learning hacks for training smart data models

7th European Workshop, PGM 2014, Utrecht, The Netherlands, September 17-19, 2014. Proceedings

Bayesian Reasoning and Machine Learning

Mastering Probabilistic Graphical Models Using Python

Your hands-on reference guide to developing, training, and optimizing your machine learning models Key FeaturesYour guide to learning efficient machine learning processes from scratchExplore expert techniques and hacks for a variety of machine learning conceptsWrite effective code in R, Python, Scala, and Spark to solve all your machine learning problemsBook Description Machine learning makes it possible to learn about the unknowns and gain hidden insights into your datasets by mastering many tools and techniques. This book guides you to do just that in a very compact manner. After giving a quick overview of what machine learning is all about, Machine Learning Quick Reference jumps right into its core algorithms and demonstrates how they can be applied to real-world scenarios. From model evaluation to optimizing their performance, this book will introduce you to the best practices in machine learning. Furthermore, you will also look at the more advanced aspects such as training neural networks and work

with different kinds of data, such as text, time-series, and sequential data. Advanced methods and techniques such as causal inference, deep Gaussian processes, and more are also covered. By the end of this book, you will be able to train fast, accurate machine learning models at your fingertips, which you can easily use as a point of reference. What you will learnGet a quick rundown of model selection, statistical modeling, and cross-validationChoose the best machine learning algorithm to solve your problemExplore kernel learning, neural networks, and time-series analysisTrain deep learning models and optimize them for maximum performanceBriefly cover Bayesian techniques and sentiment analysis in your NLP solutionImplement probabilistic graphical models and causal inferencesMeasure and optimize the performance of your machine learning modelsWho this book is for If you're a machine learning practitioner, data scientist, machine learning developer, or engineer, this book will serve as a reference point in building machine learning solutions. You will also find this book useful if you're an intermediate machine learning developer or data scientist looking for a quick, handy reference to all the concepts of machine learning. You'll need some exposure to machine learning to get the best out of this book.

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.*

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.

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

Topics in Probabilistic Graphical Models

Machine Learning

Reasoning With Probabilistic and Deterministic Graphical Models

with Applications in Systems Biology

12th International Symposium, ISMM 2015, Reykjavik, Iceland, May 27-29, 2015. Proceedings

Graphical models in their modern form have been around since the late 1970s and appear today in many areas of the sciences. Along with the ongoing developments of graphical models, a number of different graphical modeling software programs have been written over the years. In recent years many of these software developments have taken place within the R community, either in the form of new packages or by providing an R interface to existing software. This book attempts to give the reader a gentle introduction to graphical modeling using R and the main features of some of these packages. In addition, the book provides examples of how more advanced aspects of graphical modeling can be represented and handled within R. Topics covered in the seven chapters include graphical models for contingency tables, Gaussian and mixed graphical models, Bayesian networks and modeling high dimensional data.

Nowadays bioinformaticians and geneticists are faced with myriad high-throughput data usually presenting the characteristics of uncertainty, high dimensionality and large complexity. These data will only allow insights into this wealth of so-called 'omics' data if represented by flexible and scalable models, prior to any further analysis. At the interface between statistics and machine learning, probabilistic graphical models (PGMs) represent a powerful formalism to discover complex networks of relations. These models are also amenable to incorporating a priori biological information. Network reconstruction from gene expression data represents perhaps the most emblematic area of research where PGMs have been successfully applied. However these models have also created renewed interest in genetics in the broad sense, in particular regarding association genetics, causality discovery, prediction of outcomes, detection of copy number variations, and epigenetics. This book provides an overview of the applications of PGMs to genetics, genomics and postgenomics to meet this increased interest. A salient feature of bioinformatics, interdisciplinarity, reaches its limit when an intricate cooperation between domain specialists is requested. Currently, few people are specialists in the design of advanced methods using probabilistic graphical models for postgenomics or genetics. This book deciphers such models so that their perceived difficulty no longer hinders their use and focuses on fifteen illustrations showing the mechanisms behind the models. Probabilistic Graphical Models for Genetics, Genomics and Postgenomics covers six main themes: (1) Gene network inference (2) Causality discovery (3) Association genetics (4) Epigenetics (5) Detection of copy number variations (6) Prediction of outcomes from high-dimensional genomic data. Written by leading international experts, this is a collection of the most advanced work at the crossroads of probabilistic graphical models and genetics, genomics, and postgenomics. The self-contained chapters provide an enlightened account of the pros and cons of applying these powerful techniques.

This book contains the thoroughly refereed proceedings of the 12th International Symposium on Mathematical Morphology, ISMM 2015 held in Reykjavik, Iceland, in May 2015. The 62 revised full papers were carefully reviewed and selected from 72 submissions. The papers are organized in topical sections on evaluations and applications; hierarchies; color, multivalued and orientation fields; optimization, differential calculus and probabilities; topology and discrete geometry; and algorithms and implementation.

As the need for proficient power resources continues to grow, it is becoming increasingly important to implement new strategies and technologies in energy distribution to meet consumption needs. The employment of smart grid networks assists in the efficient allocation of energy resources. Smart Grid as a Solution for Renewable and Efficient Energy features emergent research and trends in energy consumption and management, as well as communication techniques utilized to monitor power transmission and usage. Emphasizing developments and challenges occurring in the field, this book is a critical resource for researchers and students concerned with signal processing, power demand management, energy storage procedures, and control techniques within smart grid networks.

Bayesian Networks

With Examples in R

A New Way of Thinking in Financial Modelling

A Probabilistic Perspective

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.

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.

This accessible text/reference provides a general introduction to probabilistic graphical models (PGMs) from an engineering perspective. 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. Features: presents a unified framework encompassing all of the main classes of PGMs; describes the practical application of the different techniques; examines the latest developments in the field, covering multidimensional Bayesian classifiers, relational graphical models and causal models; provides exercises, suggestions for further reading, and ideas for research or programming projects at the end of each chapter.

Mixed constraint and probabilistic graphical models occur quite frequently in many real world applications. Examples include: genetic linkage analysis, functional/software verification, target tracking and activity modeling. Query answering and in particular probabilistic inference on such graphical models is computationally hard often requiring exponential time in the worst case. Therefore in practice sampling algorithms are widely used for providing an approximate answer. In presence of deterministic dependencies or hard constraints, however, sampling has to overcome some principal challenges. In particular, importance sampling type schemes suffer from what is known as the rejection problem in that samples having zero weight may be generated with probability arbitrarily close to one yielding useless results. On the other hand, Markov Chain Monte Carlo techniques do not converge at all often yielding highly inaccurate estimates. In this thesis, we address these problems in a two fold manner. First, we utilize research done in constraint satisfaction and satisfiability communities for processing constraints to reduce or eliminate rejection. Second, mindful of the time overhead in sample generation due to determinism, we both make and utilize advances in statistical estimation theory to make the "most" out of the generated samples. Utilizing constraint satisfaction and satisfiability research, we propose two classes of sampling algorithms - one based on consistency enforcement and the other based on systematic search. The consistency enforcement class of algorithms work by shrinking the domains of random variables, by strengthening constraints, or by creating new ones, so that some or all zeros in the problem space can be removed. This improves convergence because of dimensionality reduction and also reduces rejection because many zero weight samples will not be generated. Our systematic search based techniques called SampleSearch manage the rejection problem by interleaving sampling with backtracking search. In this scheme, when a sample is supposed to be rejected, the algorithm continues instead with systematic backtracking search until a strictly positive-weight sample is generated. The strength of this scheme is that any state-of-the-art constraint satisfaction or propositional satisfiability search algorithm can be used with minor modifications. Through large scale experimental evaluation, we show that SampleSearch outperforms all state-of-the-art schemes when a significant amount of determinism is present in the graphical model. Subsequently, we combine SampleSearch with known statistical techniques such as Sampling Importance Resampling and Metropolis Hastings yielding efficient algorithms for sampling solutions from a uniform distribution over the solutions of a Boolean satisfiability formula. Unlike state-of-the-art algorithms, our SampleSearch-based algorithms guarantee convergence in the limit. As to statistical estimation, we make two distinct contributions. First, we propose several new statistical inequalities extending the one-sample Markov inequality to multiple samples which can be used in conjunction with SampleSearch to probabilistically lower bound likelihood tasks over mixed networks. Second, we present a novel framework called "AND/OR importance sampling" which generalizes the process of computing sample mean by exploiting AND/OR search spaces for graphical models. Specifically we provide a spectrum of AND/OR sample means which are defined on the same set of samples but derive different estimates trading variance with time. At one end is the AND/OR sample tree mean which has smaller variance than the conventional OR sample tree mean and has the same time complexity. At the other end is the AND/OR graph sample mean which has even lower variance but has higher time and space complexity. We demonstrate empirically that AND/OR sample means are far closer to the exact answer than the conventional OR sample mean.

Introduction to Probability

Probabilistic Graphical Models

Bayesian Networks and Influence Diagrams: A Guide to Construction and Analysis

Probabilistic Graphical Models for Computer Vision

Bayesian Networks in R

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
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.
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.

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.

Factor Graphs for Robot Perception

Decision Theory Models for Applications in Artificial Intelligence: Concepts and Solutions

Pattern Recognition and Machine Learning

Bayesian Networks and Decision Graphs
Bayesian Data Analysis, Third Edition

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.*

*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 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.

Bayesian Networks in R with Applications in Systems Biology is unique as it introduces the reader to the essential concepts in Bayesian network modeling and inference in conjunction with examples in the open-source statistical environment R. The level of sophistication is also gradually increased across the chapters with exercises and solutions for enhanced understanding for hands-on experimentation of the theory and concepts. The application focuses on systems biology with emphasis on modeling pathways and signaling mechanisms from high-throughput molecular data. Bayesian networks have proven to be especially useful abstractions in this regard. Their usefulness is especially exemplified by their ability to discover new associations in addition to validating known ones across the molecules of interest. It is also expected that the prevalence of publicly available high-throughput biological data sets may encourage the audience to explore investigating novel paradigms using the approaches presented in the book.

A Bayesian and Optimization Perspective

Concepts and Solutions

Learning Probabilistic Graphical Models in R
An Introduction
Learning in Graphical Models

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.

Graphical models (e.g., Bayesian and constraint networks, influence diagrams, and Markov decision processes) have become a central paradigm for knowledge representation and reasoning in both artificial intelligence and computer science in general. These models are used to perform many reasoning tasks, such as scheduling, planning and learning, diagnosis and prediction, design, hardware and software verification, and bioinformatics. These problems can be stated as the formal tasks of constraint satisfaction and satisfiability, combinatorial optimization, and probabilistic inference. It is well

known that the tasks are computationally hard, but research during the past three decades has yielded a variety of principles and techniques that significantly advanced the state of the art. This book provides comprehensive coverage of the primary exact algorithms for reasoning with such models. The main feature exploited by the algorithms is the model's graph. We present inference-based, message-passing schemes (e.g., variable-elimination) and search-based, conditioning schemes (e.g., cycle-cutset conditioning and AND/OR search). Each class possesses distinguished characteristics and in particular has different time vs. space behavior. We emphasize the dependence of both schemes on few graph parameters such as the treewidth, cycle-cutset, and (the pseudo-tree) height. The new edition includes the notion of influence diagrams, which focus on sequential decision making under uncertainty. We believe the principles outlined in the book would serve well in moving forward to approximation and anytime-based schemes. The target audience of this book is researchers and students in the artificial intelligence and machine learning area, and beyond.

This book constitutes the refereed proceedings of the 7th International Workshop on Probabilistic Graphical Models, PGM 2014, held in Utrecht, The Netherlands, in September 2014. The 38 revised full papers presented in this book were carefully reviewed and selected from 44 submissions. The papers cover all aspects of graphical models for probabilistic reasoning, decision making, and learning.

One major branch of enhancing the performance of evolutionary algorithms is the exploitation of linkage learning. This monograph aims to capture the recent progress of linkage learning, by compiling a series of focused technical chapters to keep abreast of the developments and trends in the area of linkage. In evolutionary algorithms, linkage models the relation between decision variables with the genetic linkage observed in biological systems, and linkage learning connects computational optimization methodologies and natural evolution mechanisms. Exploitation of linkage learning can enable us to design better evolutionary algorithms as well as to potentially gain insight into biological systems. Linkage learning has the potential to become one of the dominant aspects of evolutionary algorithms; research in this area can potentially yield promising results in addressing the scalability issues.

Exact Algorithms

Probabilistic Machine Learning

Mathematical Morphology and Its Applications to Signal and Image Processing

Exact Algorithms, Second Edition

Graphical Models with R

Complex networks involve interplay of several entities with each other that results in dependence in these entities. This dependence has often a modular structure that should be exploited for efficient inference and decision making. We present a survey of inference, learning, and decision making in such complex networks depicted as probabilistic graphical models. In particular, we focus on four particular problems in graphical models: (i) efficient computation of marginal and conditional probabilities; (ii) efficient parameter estimation; (iii) efficient structural learning; and (iv) decision making. We present the models in which exact solution is tractable and elucidate on the associated algorithms. We also present approximate solution techniques for general graphical models.

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.

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.

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.

Principles and Applications

Probabilistic Graphical Models for Genetics, Genomics, and Postgenomics

Smart Grid as a Solution for Renewable and Efficient Energy

Exploitation of Linkage Learning in Evolutionary Algorithms

Probabilistic Graphical Models for Computer Vision.

Machine Learning, a vital and core area of artificial intelligence (AI), is propelling the AI field ever further and making it one of the most compelling areas of computer science research. This textbook offers a comprehensive and unbiased introduction to almost all aspects of machine learning, from the fundamentals to advanced topics. It consists of 16 chapters divided into three parts: Part 1 (Chapters 1-3) introduces the fundamentals of machine learning, including terminology, basic principles, evaluation, and linear models; Part 2 (Chapters 4-10) presents classic and commonly used machine learning methods, such as decision trees, neural networks, support vector machines, Bayesian classifiers, ensemble methods, clustering, dimension reduction and metric learning; Part 3 (Chapters 11-16) introduces some advanced topics, covering feature selection and sparse learning, computational learning theory, semi-supervised learning, probabilistic graphical models, rule learning, and reinforcement learning. Each chapter includes exercises and further reading, so that readers can explore areas of interest. The book can be used as an undergraduate or postgraduate textbook for computer science, computer engineering, electrical engineering, data science, and related majors. It is also a useful reference resource for

researchers and practitioners of machine learning.
Principles and Techniques
Appropriate, Accessible and Appealing Probabilistic Graphical Models
Reasoning with Probabilistic and Deterministic Graphical Models