Probabilistic Graphical Models Principles And Techniques Solution

Probabilistic Graphical ModelsPrinciples and TechniquesMIT Press

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 book considers evolution at different scales: sequences, genes, gene families, organelles, genomes and species. The focus is on the mathematical and computational tools and concepts, which form an essential basis of evolutionary studies, indicate their limitations, and give them orientation. Recent years have witnessed rapid progress in the mathematics of evolution and phylogeny, with models and methods becoming more realistic, powerful, and complex. Aimed at graduates and researchers in phylogenetics, mathematicians, computer scientists and biologists, and including chapters by leading scientists: A. Bergeron, D. Bertrand, D. Bryant, R. Desper, O. Elemento, N. El-Mabrouk, N. Galtier, O. Gascuel, M. Hendy, S. Holmes, K. Huber, A. Meade, J. Mixtacki, B. Moret, E. Mossel, V. Moulton, M. Pagel, M.-A. Poursat, D. Sankoff, M. Steel, J. Stoye, J. Tang, L.-S. Wang, T. Warnow, Z. Yang, this book of contributed chapters explains the basis and covers the recent results in this highly topical area.

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

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

Constraint satisfaction is a simple but powerful tool. Constraints identify the impossible and reduce the realm of possibilities to effectively focus on the possible, allowing for a natural declarative formulation of what must be satisfied, without expressing how. The field of constraint reasoning has matured over the last three decades with contributions from a diverse community of researchers in artificial intelligence, databases and programming languages, operations research, management science, and applied mathematics. Today, constraint problems are used to model cognitive tasks in vision, language comprehension, default reasoning, diagnosis, scheduling, temporal and spatial reasoning. In Constraint Processing, Rina Dechter, synthesizes these contributions, along with her own significant work, to provide the first comprehensive examination of the theory that underlies constraint processing algorithms. Throughout, she focuses on fundamental tools and principles, emphasizing the representation and analysis of algorithms. ·Examines the basic practical aspects of each topic and then tackles more advanced issues, including current research challenges ·Builds the reader's understanding with definitions, examples, theory, algorithms and complexity analysis ·Synthesizes three decades of researchers work on constraint processing in AI, databases and programming languages, operations research, management science, and applied mathematics

Estimation of Distribution Algorithms: A New Tool for Evolutionary Computation is devoted to a new paradigm for evolutionary computation, named estimation of distribution algorithms (EDAs). This new class of algorithms generalizes genetic algorithms by replacing the crossover and mutation operators with learning and sampling from the probability distribution of the best individuals of the population at each iteration of the algorithm. Working in such a way, the relationships between the variables involved in the problem domain are explicitly and effectively captured and exploited. This text constitutes the first compilation and review of the techniques and applications of this new tool for performing evolutionary computation. Estimation of Distribution Algorithms: A New Tool for Evolutionary Computation is clearly divided into three parts. Part I is dedicated to the foundations of EDAs. In this part, after introducing some probabilistic graphical models - Bayesian and Gaussian networks - a review of existing EDA approaches is presented, as well as some new methods based on more flexible probabilistic graphical models. A mathematical modeling of discrete EDAs is also presented. Part II covers several applications of EDAs in some classical optimization problems: the travelling salesman problem, the job scheduling problem, and the knapsack problem. EDAs are also applied to the optimization of some well-known combinatorial and continuous functions. Part III presents the application of EDAs to solve some problems that arise in the machine learning field: feature subset selection, feature weighting in K-NN classifiers, rule induction, partial abductive inference in Bayesian networks, partitional clustering, and the search for optimal weights in artificial neural networks. Estimation of Distribution Algorithms: A New Tool for Evolutionary Computation is a useful and interesting tool for researchers working in the field of evolutionary computation and for engineers who face real-world optimization problems. This book may also be used by graduate students and researchers in computer science. `... I urge those who are interested in EDAs to study this well-crafted book today.' David E. Goldberg, University of Illinois Champaign-Urbana.

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.

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

A step-by-step illustrated introduction to the astounding mathematics of symmetry This lavishly illustrated book provides a hands-on, step-by-step introduction to the intriguing mathematics of symmetry. Instead of breaking up patterns into blocks—a sort of potato-stamp method—Frank Farris offers a completely new waveform approach that enables you to create an endless variety of rosettes, friezes, and wallpaper patterns: dazzling art images where the beauty of nature meets the precision of mathematics. Featuring more than 100 stunning color illustrations and requiring only a modest background in math, Creating Symmetry begins by addressing the enigma of a simple curve, whose curious symmetry seems unexplained by its formula. Farris describes how complex numbers unlock the mystery, and how they lead to the next steps on an engaging path to constructing waveforms. He explains how to devise waveforms for each of the 17 possible wallpaper types, and then guides you through a host of other fascinating topics in symmetry, such as color-reversing patterns, three-color patterns, polyhedral symmetry, and hyperbolic symmetry. Along the way, Farris demonstrates how to marry waveforms with photographic images to construct beautiful symmetry patterns as he gradually familiarizes you with more advanced mathematics, including group theory, functional analysis, and partial differential equations. As you progress through the book, you'll learn how to create breathtaking art images of your own. Fun, accessible, and challenging, Creating Symmetry features numerous examples and exercises throughout, as well as engaging discussions of the history behind the mathematics presented in the book.

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.

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

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

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/

An intelligent agent interacting with the real world will encounter individual people, courses, test results, drugs prescriptions, chairs, boxes, etc., and needs to reason about properties of these individuals and relations among them as well as cope with uncertainty. Uncertainty has been studied in probability theory and graphical models, and relations have been studied in logic, in particular in the predicate calculus and its extensions. This book examines the foundations of combining logic and probability into what are called relational probabilistic models. It introduces representations, inference, and learning techniques for probability, logic, and their combinations. The book focuses on two representations in detail: Markov logic networks, a relational extension of undirected graphical models and weighted first-order predicate calculus formula, and Problog, a probabilistic extension of logic programs that can also be viewed as a Turing-complete relational extension of Bayesian networks.

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.

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.

We are happy to present the first volume of the Handbook of Defeasible Reasoning and Uncertainty Management Systems. Uncertainty pervades the real world and must therefore be addressed by every system that attempts to represent reality. The representation of uncertainty is a ma jor concern of philosophers, logicians, artificial intelligence researchers and com puter sciencists, psychologists, statisticians, economists and engineers. The present Handbook volumes provide frontline coverage of this area. This Handbook was produced in the style of previous handbook series like the Handbook of Philosoph ical Logic, the Handbook of Logic in Computer Science, the Handbook of Logic in Artificial Intelligence and Logic Programming, and can be seen as a companion to them in covering the wide applications of logic and reasoning. We hope it will answer the needs for adequate representations of uncertainty. This Handbook series grew out of the ESPRIT Basic Research Project DRUMS II, where the acronym is made out of the Handbook series title. This project was financially supported by the European Union and regroups 20 major European research teams working in the general domain of uncertainty. As a fringe benefit of the DRUMS project, the research community was able to create this Hand book series, relying on the DRUMS participants as the core of the authors for the Handbook together with external international experts.

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.

This book constitutes the refereed proceedings of the 15th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty, ECSQARU 2019, held in Belgrade, Serbia, in September 2019. The 41 full papers presented together with 3 abstracts of invited talks inn this volume were carefully reviewed and selected from 62 submissions. The papers are organized in topical sections named: Argumentation; Belief Functions; Conditional, Default and Analogical Reasoning; Learning and Decision Making; Precise and Imprecise Probabilities; and Uncertain Reasoning for Applications.

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.

A concise and self-contained introduction to causal inference, increasingly important in data science and machine learning. The mathematization of causality is a relatively recent development, and has become increasingly important in data science and machine learning. This book offers a self-contained and concise introduction to causal models and how to learn them from data. After explaining the need for causal models and discussing some of the principles underlying causal inference, the book teaches readers how to use causal models: how to compute intervention distributions, how to infer causal models from observational and interventional data, and how causal ideas could be exploited for classical machine learning problems. All of these topics are discussed first in terms of two variables and then in the more general multivariate case. The bivariate case turns out to be a particularly hard problem for causal learning because there are no conditional independences as used by classical methods for solving multivariate cases. The authors consider analyzing statistical asymmetries between cause and effect to be highly instructive, and they report on their decade of intensive research into this problem. The book is accessible to readers with a background in machine learning or statistics, and can be used in graduate courses or as a reference for researchers. The text includes code snippets that can be copied and pasted, exercises, and an appendix with a summary of the most important technical concepts.

Probability is the bedrock of machine learning. You cannot develop a deep understanding and application of machine learning without it. Cut through the equations, Greek letters, and confusion, and discover the topics in probability that you need to know. Using clear explanations, standard Python libraries, and step-by-step tutorial lessons, you will discover the importance of probability to machine learning, Bayesian probability, entropy, density estimation, maximum likelihood, and much more.

If you’re an experienced programmer interested in crunching data, this book will get you started with machine learning—a toolkit of algorithms that enables computers to train themselves to automate useful tasks. Authors Drew Conway and John Myles White help you understand machine learning and statistics tools through a series of hands-on case studies, instead of a traditional math-heavy presentation. Each chapter focuses on a specific problem in machine learning, such as classification, prediction, optimization, and recommendation. Using the R programming language, you’ll learn how to analyze sample datasets and write simple machine learning algorithms. Machine Learning for Hackers is ideal for programmers from any background, including business, government, and academic research. Develop a naïve Bayesian classifier to determine if an email is spam, based only on its text Use linear regression to predict the number of page views for the top 1,000 websites Learn optimization techniques by attempting to break a simple letter cipher Compare and contrast U.S. Senators statistically, based on their voting records Build a “whom to follow” recommendation system from Twitter 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.

This three volume set (CCIS 1237-1239) constitutes the proceedings of the 18th International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems, IPMU 2020, in June 2020. The conference was scheduled to take place in Lisbon, Portugal, at University of Lisbon, but due to COVID-19 pandemic it was held virtually. The 173 papers were carefully reviewed and selected from 213 submissions. The papers are organized in topical sections: homage to Enrique Ruspini; invited talks; foundations and mathematics; decision making, preferences and votes; optimization and uncertainty; games; real world applications; knowledge processing and creation; machine learning I; machine learning II; XAI; image processing; temporal data processing; text analysis and processing; fuzzy interval analysis; theoretical and applied aspects of imprecise probabilities; similarities in artificial intelligence; belief function theory and its applications; aggregation: theory and practice; aggregation: pre-aggregation functions and other generalizations of monotonicity; aggregation: aggregation of different data structures; fuzzy methods in data mining and knowledge discovery; computational intelligence for logistics and transportation problems; fuzzy implication functions; soft methods in statistics and data analysis; image understanding and explainable AI; fuzzy and generalized quantifier theory; mathematical methods towards dealing with uncertainty in applied sciences; statistical image processing and analysis, with applications in neuroimaging; interval uncertainty; discrete models and computational intelligence; current techniques to model, process and describe time series; mathematical fuzzy logic and graded reasoning models; formal concept analysis, rough sets, general operators and related topics; computational intelligence methods in information modelling, representation and processing.

This open access book constitutes the proceedings of the 30th European Symposium on Programming, ESOP 2021, which was held during March 27 until April 1, 2021, as part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2021. The conference was planned to take place in Luxembourg and changed to an online format due to the COVID-19 pandemic. The 24 papers included in this volume were carefully reviewed and selected from 79 submissions. They deal with fundamental issues in the specification, design, analysis, and implementation of programming languages and systems.

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.

This book serves as a textbook or reference for anyone with an interest in probabilistic modeling in the fields of computer science, computer engineering, and electrical engineering. This text is also a resource for courses on expert systems, machine learning, and artificial intelligence. Beginning with a basic theoretical introduction, the author then provides a discussion of inference, methods of learning, and applications based on Bayesian networks and beyond.

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.

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.

Many of the concepts and terminology surrounding modern causal inference can be quite intimidating to the novice. Judea Pearl presents a book ideal for beginners in statistics, providing a comprehensive introduction to the field of causality. Examples from classical statistics are presented throughout to demonstrate the need for causality in resolving decision-making dilemmas posed by data. Causal methods are also compared to traditional statistical methods, whilst questions are provided at the end of each section to aid student learning.

This book constitutes the proceedings of the 10th International Workshop on Fuzzy Logic and Applications, WILF 2013, held in Genoa, Italy, in November 2013. After a rigorous peer-review selection process, ultimately 19 regular papers were selected for inclusion in this volume from 29 submissions. In addition the book contains 3 keynote talks and 2 tutorials. The papers are organized in topical sections named: fuzzy machine learning and interpretability; theory and applications.

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

A useful introduction to this topic for both students and researchers, with an emphasis on applications and practicalities rather than on a formal development. It is based on the popular software package for graphical modelling, MIM, freely available for downloading from the Internet. Following a description of some of the basic ideas of graphical modelling, subsequent chapters describe particular families of models, including log-linear models, Gaussian models, and models for mixed discrete and continuous variables. Further chapters cover hypothesis testing and model selection. Chapters 7 and 8 are new to this second edition and describe the use of directed, chain, and other graphs, complete with a summary of recent work on causal inference.

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 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.