Probabilistic Graphical Models Principles

26th International Conference, CP 2020, Louvain-la-Neuve, Belgium, September 7–11, 2020, Proceedings
Introduction to Graphical Modelling
Graph Representation Learning
Principles and Applications
Modeling and Reasoning with Bayesian Networks
An Introduction to Lifted Probabilistic Inference
Probabilistic Graphical Models
A Scalable Approach to Structure and Parameter Learning in Probabilistic Graphical Models
10th International Workshop, WILF 2013, Genoa, Italy, November 19-22, 2013, Proceedings
Fuzzy Logic and Applications
A Probabilistic Perspective
Mastering Probabilistic Graphical Models Using Python
Volume III: Interfaces and Applications of Artificial Intelligence
Statistical Relational Artificial Intelligence
Boosting
Principles and Practice of Structural Equation Modeling, Fourth Edition
Exact Algorithms, Second Edition
A New Way of Thinking in Financial Modelling
High-Dimensional Statistics
Foundations and Learning Algorithms
Probabilistic Graphical Models for Genetics, Genomics and Postgenomics
Languages, Semantics, Inference and Learning
Logic, Probability, and Computation
Constraint Processing
Bayesian Networks and Decision Graphs
Principles and Techniques
Graphical Models, Exponential Families, and Variational Inference
Foundations and Algorithms
Fundamentals, Techniques, and Applications
Introduction to Statistical Relational Learning
Foundations of Neural Computation
Machine Learning
Machine Learning
Practical Probabilistic Programming
Hybrid Random Fields
Deep Learning on Graphs
A Non-Asymptotic Viewpoint
Principles and Practice of Constraint Programming

Knowledge Graphs
Handbook of Graphical Models

LUCIANO KENDRA

26th International Conference, CP 2020, Louvain-la-Neuve, Belgium, September 7–11, 2020, Proceedings Springer
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.
Introduction to Graphical Modelling MIT Press
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.
Graph Representation Learning Morgan & Claypool Publishers
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.
Principles and Applications Springer Nature
Graph-structured data is ubiquitous throughout the natural and social sciences, from telecommunication networks to quantum chemistry. Building relational inductive biases into deep learning architectures is crucial for creating systems that can learn, reason, and generalize from this kind of data. Recent years have seen a surge in research on graph representation learning, including techniques for deep graph embeddings, generalizations of convolutional neural networks to graph-structured data, and neural message-passing approaches inspired by belief propagation. These advances in graph representation learning have led to new state-of-the-art results in numerous domains, including chemical synthesis, 3D vision, recommender systems, question answering, and social network analysis. This book provides a synthesis and overview of graph representation learning. It begins with a discussion of the goals of graph representation learning as well as key methodological foundations in graph theory and network analysis. Following this, the book introduces and reviews methods for learning node embeddings, including random-walk-based methods and applications to knowledge graphs. It then provides a technical synthesis and introduction to the highly successful graph neural network (GNN) formalism, which has become a dominant and fast-growing paradigm for deep learning with graph data. The book concludes with a synthesis of recent advancements in deep generative models for graphs—a nascent but quickly growing subset of graph representation learning.
Modeling and Reasoning with Bayesian Networks MIT Press
A practical introduction perfect for final-year undergraduate and graduate students without a solid background in linear algebra and calculus.
An Introduction to Lifted Probabilistic Inference Oxford University Press, USA
A coherent introductory text from a groundbreaking researcher, focusing on clarity and motivation to build intuition and understanding.
Probabilistic Graphical Models Packt Publishing Ltd
A comprehensive text on foundations and techniques of graph neural networks with applications in NLP, data mining, vision and healthcare.
A Scalable Approach to Structure and Parameter Learning in

Probabilistic Graphical Models Springer Nature
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.
10th International Workshop, WILF 2013, Genoa, Italy, November 19-22, 2013, Proceedings Guilford Publications
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.
Fuzzy Logic and Applications MIT Press
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 Probabilistic Perspective Academic Press
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.
Mastering Probabilistic Graphical Models Using Python CRC Press
Recent advances in the area of lifted inference, which exploits the structure inherent in relational probabilistic models. Statistical relational AI (StaRAI) studies the integration of reasoning under uncertainty with reasoning about individuals and relations. The representations used are often called relational probabilistic models. Lifted inference is about how to exploit the structure inherent in relational probabilistic models, either in the way they are expressed or by extracting structure from observations. This book covers recent significant advances in the area of lifted inference, providing a unifying introduction to this very active field. After providing necessary background on probabilistic graphical models, relational probabilistic models, and learning inside these models, the book turns to lifted inference, first covering exact inference and then approximate inference. In

addition, the book considers the theory of liftability and acting in relational domains, which allows the connection of learning and reasoning in relational domains.

Volume III: Interfaces and Applications of Artificial Intelligence Springer

This book brings together important topics of current research in probabilistic graphical modeling, learning from data and probabilistic inference. Coverage includes such topics as the characterization of conditional independence, the learning of graphical models with latent variables, and extensions to the influence diagram formalism as well as important application fields, such as the control of vehicles, bioinformatics and medicine.

Statistical Relational Artificial Intelligence Cambridge University 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.

Boosting Springer Science & Business Media

Probabilistic Graphical ModelsPrinciples and TechniquesMIT Press

Principles and Practice of Structural Equation Modeling, Fourth Edition MIT Press

Emphasizing concepts and rationale over mathematical minutiae, this is the most widely used, complete, and accessible structural equation modeling (SEM) text. Continuing the tradition of using real data examples from a variety of disciplines, the significantly revised fourth edition incorporates recent developments such as Pearl's graphing theory and the structural causal model (SCM), measurement invariance, and more. Readers gain a comprehensive understanding of all phases of SEM, from data collection and screening to the interpretation and reporting of the results. Learning is enhanced by exercises with answers, rules to remember, and topic boxes. The companion website supplies data, syntax, and output for the book's examples--now including files for Amos, EQS, LISREL, Mplus, Stata, and R (lavaan). New to This Edition *Extensively revised to cover important new topics: Pearl's graphing theory and the SCM, causal inference frameworks, conditional process modeling, path models for longitudinal data, item response theory, and more. *Chapters on best practices in all stages of SEM, measurement invariance in confirmatory factor analysis, and significance testing issues and bootstrapping. *Expanded coverage of psychometrics. *Additional computer tools: online files for all detailed examples, previously provided in EQS, LISREL, and Mplus, are now also given in Amos, Stata, and R (lavaan). *Reorganized to cover the specification, identification, and analysis of observed variable models separately from latent variable models. Pedagogical Features *Exercises with answers, plus end-of-chapter annotated lists of further reading. *Real examples of troublesome data, demonstrating how to handle typical problems in analyses. *Topic boxes on specialized issues, such as causes of nonpositive definite correlations. *Boxed rules to remember. *Website promoting a learn-by-doing approach, including syntax and data files for six widely used SEM computer tools.

Exact Algorithms, Second Edition MIT Press

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 New Way of Thinking in Financial Modelling Packt Pub Limited

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.

High-Dimensional Statistics MIT Press

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.

Foundations and Learning Algorithms Springer Science & Business Media

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

Related with Probabilistic Graphical Models Principles:

• Ssvf Program Guide 2022 : click here