

Probabilistic Graphical Models Principles And Techniques Adaptive Computation And Machine Learning Series

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.

INDUSTRIAL MOTOR CONTROL 7E is an integral part of any electrician training. Comprehensive and up to date, this book provides crucial information on basic relay control systems, programmable logic controllers, and solid state devices commonly found in an industrial setting. Written by a highly qualified and respected author, you will find easy-to-follow instructions and essential information on controlling industrial motors and commonly used devices in contemporary industry. INDUSTRIAL MOTOR CONTROL 7E successfully bridges the gap between industrial maintenance and instrumentation, giving you a fundamental understanding of the operation of variable frequency drives, solid state relays, and other applications that employ electronic devices. Important Notice: Media content referenced within the product description or the product text may not be available in the ebook version.

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.

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.

Graphical Models, Exponential Families, and Variational Inference

A New Way of Thinking in Financial Modelling

Probabilistic Graphical Models for Computer Vision.

Information Processing and Management of Uncertainty in Knowledge-Based Systems

Foundations and Learning Algorithms

Factor Graphs for Robot Perception

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.

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.

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

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.

Logic, Probability, and Computation

Mastering Probabilistic Graphical Models Using Python

Sparse Modeling

Handbook of Graphical Models

Graph Representation Learning

Exact Algorithms

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.

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

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.

Probabilistic Reasoning in Intelligent Systems

Quantified Representation of Uncertainty and Imprecision

Bayesian Reasoning and Machine Learning

Statistical Relational Artificial Intelligence

18th International Conference, IPMU 2020, Lisbon, Portugal, June 15-19, 2020, Proceedings, Part I

An Introduction to Lifted Probabilistic Inference

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 are 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.

The three-volume set LNCS 10860, 10861 + 10862 constitutes the proceedings of the 18th International Conference on Computational Science, ICCS 2018, held in Wuxi, China, in June 2018. The total of 155 full and 66 short papers presented in this book set was carefully reviewed and selected from 404 submissions. The papers were organized in topical sections named: Part I: ICCS Main Track Part II: Track of Advances in High-Performance Computational Earth Sciences:

Applications and Frameworks; Track of Agent-Based Simulations, Adaptive Algorithms and Solvers; Track of Applications of Matrix Methods in Artificial Intelligence and Machine Learning; Track of Architecture, Languages, Compilation and Hardware Support for Emerging Manycore Systems; Track of Biomedical and Bioinformatics Challenges for Computer Science; Track of Computational Finance and Business Intelligence; Track of Computational Optimization, Modelling and Simulation; Track of Data, Modeling, and Computation in IoT and Smart Systems; Track of Data-Driven Computational Sciences; Track of Mathematical-Methods-and-Algorithms for Extreme Scale; Track of Multiscale Modelling and Simulation Part III: Track of Simulations of Flow and Transport: Modeling, Algorithms and Computation; Track of Solving Problems with Uncertainties; Track of Teaching Computational Science; Poster Papers

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

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.

Learning Probabilistic Graphical Models in R

Probability for Machine Learning

Hybrid Random Fields

Constraint Processing

Computational Science - ICCS 2018

Discover How To Harness Uncertainty With Python

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 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. 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.

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.

Learning Bayesian Networks

Theory, Algorithms, and Applications

Machine Learning for Hackers

Case Studies and Algorithms to Get You Started

Untersuchung von dem Wesen des Geistes, oder des seltsamen Pietisten-Gespenstes, Welches heutigen Tages die Welt ä fftet ; Angestellet zur treuhertzigem ernstlichen Warnung aller frommen Christen, von einem Freunde der Piet ä t und Feinde der Pietisterey. Geschehen in demselben Jahr, da solche Warnung n ö thig war

Building Probabilistic Graphical Models with Python

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 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.

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. 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.

Principles and Applications

Probabilistic Graphical Models

A Probabilistic Perspective

Reasoning With Probabilistic and Deterministic Graphical Models

A Scalable Approach to Structure and Parameter Learning in Probabilistic Graphical Models

Networks of Plausible Inference

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 major concern of philosophers, logicians, artificial intelligence researchers and computer scientists, 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 Philosophical 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 Handbook series, relying on the DRUMS participants as the core of the authors for the Handbook together with external international experts.

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 and efficient 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 a prominent 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 the informativeness 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. In addition 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 graphical models are well suited for integrative biology and systems biology, hot topics guaranteed to be of lasting interest.

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.

Sparse models are particularly useful in scientific applications, such as biomarker discovery in genetic or neuroimaging data, where the interpretability of a predictive model is essential. Sparsity can also dramatically improve the cost efficiency of signal processing. Sparse Modeling: Theory, Algorithms, and Applications provides an introduction to the growing field of sparse modeling, including application examples, problem formulations that yield sparse solutions, algorithms for finding such solutions, and recent theoretical results on sparse recovery. The book gets you up to speed on the latest sparsity-related developments and will motivate you to continue learning about the field. The authors first present motivating examples and a high-level survey of key recent developments in sparse modeling. The book then describes optimization problems involving commonly used sparsity-enforcing tools, presents essential theoretical results, and discusses several state-of-the-art algorithms for finding sparse solutions. The authors go on to address a variety of sparse recovery problems that extend the basic formulation to more sophisticated forms of structured sparsity and to different loss functions. They also examine a particular class of sparse graphical models and cover dictionary learning and sparse matrix factorizations.

Deep Learning on Graphs

10th International Workshop, WILF 2013, Genoa, Italy, November 19-22, 2013, Proceedings

Model-Based Machine Learning

Learning in Graphical Models

Principles and Techniques

Machine Learning

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.

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.

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.

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.

Industrial Motor Control

Fuzzy Logic and Applications

18th International Conference, Wuxi, China, June 11–13, 2018, Proceedings, Part I

Elements of Causal Inference

Probabilistic Graphical Models for Genetics, Genomics and Postgenomics

Graphical Models

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.

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.

Probabilistic Reasoning in Intelligent Systems is a complete and accessible account of the theoretical foundations and computational methods that underlie plausible reasoning under uncertainty. The author provides a coherent explication of probability as a language for reasoning with partial belief and offers a unifying perspective on other AI approaches to uncertainty, such as the Dempster-Shafer formalism, truth maintenance systems, and nonmonotonic logic. The author distinguishes syntactic and semantic approaches to uncertainty—and offers techniques, based on belief networks, that provide a mechanism for making semantics-based systems operational. Specifically, network-propagation techniques serve as a mechanism for combining the theoretical coherence of probability theory with modern demands of reasoning-systems technology: modular declarative inputs, conceptually meaningful inferences, and parallel distributed computation. Application areas include diagnosis, forecasting, image interpretation, multi-sensor fusion, decision support systems, plan recognition, planning, speech recognition—in short, almost every task requiring that conclusions be drawn from uncertain clues and incomplete information. Probabilistic Reasoning in Intelligent Systems will be of special interest to scholars and researchers in AI, decision theory, statistics, logic, philosophy, cognitive psychology, and the management sciences. Professionals in the areas of knowledge-based systems, operations research, engineering, and statistics will find theoretical and computational tools of immediate practical use. The book can also be used as an excellent text for graduate-level courses in AI, operations research, or applied probability.

A comprehensive text on foundations and techniques of graph neural networks with applications in NLP, data mining, vision and healthcare.

Modeling and Reasoning with Bayesian Networks

Introduction to Graphical Modelling

Introduction to Statistical Relational Learning