Bayesian Reasoning And Machine Learning David Barber

Machine learning is a relatively new branch of artificial intelligence. The field has undergone a significant period of growth in the 1990s, with many new areas of research and development being explored.

Bayesian Reasoning and Machine LearningCambridge University Press

This book constitutes the refereed proceedings of the 5th International Conference on Dependability in Sensor, Cloud, and Big Data Systems and Applications, DependSys, held in Guangzhou, China, in November 2019. The volume presents 39 full papers, which were carefully reviewed and selected from 112 submissions. The papers are organized in topical sections on ?dependability and security fundamentals and technologies; dependable and secure systems; dependable and secure applications; dependability and security measures and assessments; explainable artificial inteligence for cyberspace.

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 clear exposition begins with basic concepts and moves on to combination of events, dependent events and random variables, Bernoulli trials and the De Moivre-Laplace theorem, and more. Includes 150 problems, many with answers.

This tutorial text gives a unifying perspective on machine learning by covering both probabilistic and deterministic approaches -which are based on optimization techniques – together with the Bayesian inference approach, whose essence lies in the use of a hierarchy of probabilistic models. The book presents the major machine learning methods as they have been developed in different disciplines, such as statistics, statistical and adaptive signal processing and computer science. Focusing on the physical reasoning behind the mathematics, all the various methods and techniques are explained in depth, supported by examples and problems, giving an invaluable resource to the student and researcher for understanding and applying machine learning concepts. The book builds carefully from the basic classical methods to the most recent trends, with chapters written to be as self-contained as possible, making the text suitable for different courses: pattern recognition, statistical/adaptive signal processing, statistical/Bayesian learning, as well as short courses on sparse modeling, deep learning, and probabilistic graphical models. All major classical techniques: Mean/Least-Squares regression and filtering, Kalman filtering, stochastic approximation and online learning, Bayesian classification, decision trees, logistic regression and boosting methods. The latest trends: Sparsity, convex analysis and optimization, online distributed algorithms, learning in RKH spaces, Bayesian inference, graphical and hidden Markov models, particle filtering, deep learning, dictionary learning and latent variables modeling. Case studies - protein folding prediction, optical character recognition, text authorship identification, fMRI data analysis, change point detection, hyperspectral image unmixing, target localization, channel equalization and echo cancellation, show how the theory can be applied. MATLAB code for all the main algorithms are available on an accompanying website, enabling the reader to experiment with the code.

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

A thought-provoking and wide-ranging exploration of machine learning and the race to build computer intelligences as flexible as our own In the world's top research labs and universities, the race is on to invent the ultimate learning algorithm: one capable of discovering any knowledge from data, and doing anything we want, before we even ask. In The Master Algorithm, Pedro Domingos lifts the veil to give us a peek inside the learning machines that power Google, Amazon, and your smartphone. He assembles a blueprint for the future universal learner--the Master Algorithm--and discusses what it will mean for business, science, and society. If data-ism is today's philosophy, this book is its bible.

"A First Course in Machine Learning by Simon Rogers and Mark Girolami is the best introductory book for ML currently available. It combines rigor and precision with accessibility, starts from a detailed explanation of the basic foundations of Bayesian analysis in the simplest of settings, and goes all the way to the frontiers of the subject such as infinite mixture models, GPs, and MCMC." —Devdatt Dubhashi, Professor, Department of Computer Science and Engineering, Chalmers University, Sweden "This textbook manages to be easier to read than other comparable books in the subject while retaining all the rigorous treatment needed. The new chapters put it at the forefront of the field by covering topics that have become mainstream in machine learning over the last decade." —Daniel Barbara, George Mason University, Fairfax, Virginia, USA "The new edition of A First Course in Machine Learning by Rogers and Girolami is an excellent introduction to the use of statistical methods in machine learning. The book introduces concepts such as mathematical modeling, inference, and prediction, providing ‘just in time’ the essential background on linear algebra, calculus, and probability theory that the reader needs to understand these concepts." —Daniel Ortiz-Arroyo, Associate Professor, Aalborg University Esbjerg, Denmark "I was impressed by how closely the material aligns with the needs of an introductory course on machine learning, which is its greatest strength...Overall, this is a pragmatic and helpful book, which is well-aligned to the needs of an introductory course and one that I will be looking at for my own students in coming months." —David Clifton, University of Oxford, UK "The first edition of this book was already an excellent introductory text on machine learning for an advanced undergraduate or taught masters level course, or indeed for anybody who wants to learn about an interesting and important field of computer science. The additional chapters of advanced material on Gaussian process, MCMC and mixture modeling provide an ideal basis for practical projects, without disturbing the very clear and readable exposition of the basics contained in the first part of the book." —Gavin Cawley, Senior Lecturer, School of Computing Sciences, University of East Anglia, UK "This book could be used for junior/senior undergraduate students or first-year graduate students, as well as individuals who want to explore the field of machine learning...The book introduces not only the concepts but the underlying ideas on algorithm implementation from a critical thinking perspective." —Guangzhi Qu, Oakland University, Rochester, Michigan, USA

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 and self-contained introduction to Gaussian processes, which provide a principled, practical, probabilistic approach to learning in kernel machines. Gaussian processes (GPs) provide a principled, practical, probabilistic approach to learning in kernel machines. GPs have received increased attention in the machine-learning community over the past decade, and this book provides a long-needed systematic and unified treatment of theoretical and practical aspects of GPs in machine learning. The treatment is comprehensive and self-contained, targeted at researchers and students in machine learning and applied statistics. The book deals with the supervised-learning problem for both regression and classification, and includes detailed algorithms. A wide variety of covariance (kernel) functions are presented and their properties discussed. Model selection is discussed both from a Bayesian and a classical perspective. Many connections to other well-known techniques from machine learning and statistics are discussed, including support-vector machines, neural networks, splines, regularization networks, relevance vector machines and others. Theoretical issues including learning curves and the PAC-Bayesian framework are treated, and several approximation methods for learning with large datasets are discussed. The book contains illustrative examples and exercises, and code and datasets are available on the Web. Appendixes provide mathematical background and a discussion of Gaussian Markov processes.

Offering a fundamental basis in kernel-based learning theory, this book covers both statistical and algebraic principles. It provides over 30 major theorems for kernel-based supervised and unsupervised learning models. The first of the theorems establishes a condition, arguably necessary and sufficient, for the kernelization of learning models. In addition, several other theorems are devoted to proving mathematical equivalence between seemingly unrelated models. With over 25 closed-form and iterative algorithms, the book provides a step-by-step guide to algorithmic procedures and analysing which factors to consider in tackling a given problem, enabling readers to improve specifically designed learning algorithms, build models for new applications and develop efficient techniques suitable for green machine learning technologies. Numerous real-world examples and over 200 problems, several of which are Matlab-based simulation exercises, make this an essential resource for graduate students and professionals in computer science, electrical and biomedical engineering. Solutions to problems are provided online for instructors.

This text covers all the fundamentals and presents basic theoretical concepts and a wide range of techniques (algorithms) applicable to challenges in our day-to-day lives. The book recognizes that most of the ideas behind machine learning are simple and straightforward. It provides a platform for hands-on experience through self-study machine learning projects. Datasets for some benchmark applications have been explained to encourage the use of algorithms covered in this book. This is a comprehensive text book on machine learning for undergraduates in computer science and all engineering degree programs. Post graduates and research scholars will find it a useful initial exposure to the subject, before they go for highly theoretical depth in the specific areas of their research. For engineers, scientists, business managers and other practitioners, the book will help build the foundations of machine learning.

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

Probability as an Alternative to Boolean LogicWhile logic is the mathematical foundation of rational reasoning and the fundamental principle of computing, it is restricted to problems where information is both complete and certain. However, many real-world problems, from financial investments to email filtering, are incomplete or uncertain in natur

Probabilistic Methods for Financial and Marketing Informatics aims to provide students with insights and a guide explaining how to apply probabilistic reasoning to business problems. Rather than dwelling on rigor, algorithms, and proofs of theorems, the authors concentrate on showing examples and using the software package Netica to represent and solve problems. The book contains unique coverage of probabilistic reasoning topics applied to business problems, including marketing, banking, operations management, and finance. It shares insights about when and why probabilistic methods can and cannot be used effectively. This book is recommended for all R&D professionals and students who are involved with industrial informatics, that is, applying the methodologies of computer science and engineering to business or industry information. This includes computer science and other professionals in the data management and data mining field whose interests are business and marketing information in general, and who want to apply AI and probabilistic methods to their problems in order to better predict how well a product or service will do in a particular market, for instance. Typical fields where this technology is used are in advertising, venture capital decision making, operational risk measurement in any industry, credit scoring, and investment science. Unique coverage of probabilistic reasoning topics applied to business problems, including marketing, banking, operations management, and finance Shares insights about when and why probabilistic methods can and cannot be used effectively Complete review of Bayesian networks and probabilistic methods for those IT professionals new to informatics.

Machine learning methods extract value from vast data sets quickly and with modest resources. They are established tools in a wide range of industrial applications, including search engines, DNA sequencing, stock market analysis, and robot locomotion, and their use is spreading rapidly. People who know the methods have their choice of rewarding jobs. This hands-on text opens these opportunities to computer science students with modest mathematical backgrounds. It is designed for final-year undergraduates and master's students with limited background in linear algebra and calculus. Comprehensive and coherent, it develops everything from basic reasoning to advanced techniques within the framework of graphical models. Students learn more than a menu of techniques, they develop analytical and problem-solving skills that equip them for the real world. Numerous examples and exercises, both computer based and theoretical, are included in every chapter. Resources for students and instructors, including a MATLAB toolbox, are available online.

The increasing complexity of our world demands new perspectives on the role of technology in decision making. Human decision making has its li- tations in terms of information-processing capacity. We need new technology to cope with the increasingly complex and information-rich nature of our modern society. This is particularly true for critical environments such as crisis management and tra?c management, where humans need to engage in close collaborations with arti?cial systems to observe and understand the situation and respond in a sensible way. We believe that close collaborations between humans and arti?cial systems will become essential and that the importance of research into Interactive Collaborative Information Systems (ICIS) is self-evident. Developments in information and communication technology have ra- cally changed our working environments. The vast amount of information available nowadays and the wirelessly networked nature of our modern so- ety open up new opportunities to handle di?cult decision-making situations such as computer-supported situation assessment and distributed decision making. To make good use of these new possibilities, we need to update our traditional views on the role and capabilities of information systems. The aim of the Interactive Collaborative Information Systems project is to develop techniques that support humans in complex information en- ronments and that facilitate distributed decision-making capabilities. ICIS emphasizes the importance of building actor-agent communities: close c- laborations between human and arti?cial actors that highlight their comp- mentary capabilities, and in which task distribution is ?exible and adaptive.

Bayesian methods for machine learning have been widely investigated, yielding principled methods for incorporating prior information into inference algorithms. This monograph provides the reader with an in-depth review of the role of Bayesian methods for the reinforcement learning (RL) paradigm. The major incentives for incorporating Bayesian reasoning in RL are that it provides an elegant approach to action-selection (exploration/exploitation) as a function of the uncertainty in learning, and it provides a machinery to incorporate prior knowledge into the algorithms. Bayesian Reinforcement Learning: A Survey first discusses models and methods for Bayesian inference in the simple single-step Bandit model. It then reviews the extensive recent literature on Bayesian methods for model-based RL, where prior information can be expressed on the parameters of the Markov model. It also presents Bayesian methods for model-free RL, where priors are expressed over the value function or policy class. Bayesian Reinforcement Learning: A Survey is a comprehensive reference for students and researchers with an interest in Bayesian RL algorithms and their theoretical and empirical properties.

A comprehensive overview of Monte Carlo simulation that explores the latest topics, techniques, and real-world applications More and more of today’s numerical problems found in engineering and finance are solved through Monte Carlo methods. The heightened popularity of these methods and their continuing development makes it important for researchers to have a comprehensive understanding of the Monte Carlo approach. Handbook of Monte Carlo Methods provides the theory, algorithms, and applications that helps provide a thorough understanding of the emerging dynamics of this rapidly-growing field. The authors begin with a discussion of fundamentals such as how to generate random numbers on a computer. Subsequent chapters discuss key Monte Carlo topics and methods, including: Random variable and stochastic process generation Markov chain Monte Carlo, featuring key algorithms such as the Metropolis-Hastings method, the Gibbs sampler, and hit-and-run Discrete-event simulation Techniques for the statistical analysis of simulation data including the delta method, steady-state estimation, and kernel density estimation Variance reduction, including importance sampling, latin hypercube sampling, and conditional Monte Carlo Estimation of derivatives and sensitivity analysis Advanced topics including cross-entropy, rare events, kernel density estimation, quasi Monte Carlo, particle systems, and randomized optimization The presented theoretical concepts are illustrated with worked examples that use MATLAB®, a related Web site houses the MATLAB® code, allowing readers to work hands-on with the material and also features the author's own lecture notes on Monte Carlo methods. Detailed appendices provide background material on probability theory, stochastic processes, and mathematical statistics as well as the key optimization concepts and techniques that are relevant to Monte Carlo simulation. Handbook of Monte Carlo Methods is an excellent reference for applied statisticians and practitioners working in the fields of engineering and finance who use or would like to learn how to use Monte Carlo in their research. It is also a suitable supplement for courses on Monte Carlo methods and computational statistics at the upper-undergraduate and graduate levels.

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

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

Table of contents

"Vast amounts of data present amajor challenge to all thoseworking in computer science, and its many related fields, who need to process and extract value from such data. Machine learning technology is already used to help with this task in a wide range of industrial applications, including search engines, DNA sequencing, stock market analysis and robot locomotion. As its usage becomes more widespread, no student should be without the skills taught in this book. Designed for final-year undergraduate and graduate students, this gentle introduction is ideally suited to readers without a solid background in linear algebra and calculus. It covers everything from basic reasoning to advanced techniques in machine learning, and rucially enables students to construct their own models for real-world problems by teaching them what lies behind the methods. Numerous examples and exercises are included in the text. Comprehensive resources for students and instructors are available online"--

This is an introduction to Bayesian statistics and decision theory, including advanced topics such as Monte Carlo methods. This new edition contains several revised chapters and a new chapter on model choice.

The first unified treatment of time series modelling techniques spanning machine learning, statistics, engineering and computer science.

Fun guide to learning Bayesian statistics and probability through unusual and illustrative examples. Probability and statistics are increasingly important in a huge range of professions. But many people use data in ways they don't even understand, meaning they aren't getting the most from it. Bayesian Statistics the Fun Way will change that. This book will give you a complete understanding of Bayesian statistics through simple explanations and un-boring examples. Find out the probability of UFOs landing in your garden, how likely Han Solo is to survive a flight through an asteroid shower, how to win an argument about conspiracy theories, and whether a burglary really was a burglary, to name a few examples. By using these off-the-beaten-track examples, the author actually makes learning statistics fun. And you'll learn real skills, like how to: - How to measure your own level of uncertainty in a conclusion or belief - Calculate Bayes theorem and understand what it's useful for - Find the posterior, likelihood, and prior to check the accuracy of your conclusions - Calculate distributions to see the range of your data - Compare hypotheses and draw reliable conclusions from them Next time you find yourself with a sheaf of survey results and no idea what to do with them, turn to Bayesian Statistics the Fun Way to get the most value from your data.

A Turing Award-winning computer scientist and statistician shows how understanding causality has revolutionized science and will revolutionize artificial intelligence "Correlation is not causation." This mantra, chanted by scientists for more than a century, has led to a virtual prohibition on causal talk. Today, that taboo is dead. The causal revolution, instigated by Judea Pearl and his colleagues, has cut through a century of confusion and established causality -- the study of cause and effect -- on a firm scientific basis. His work explains how we can know easy things, like whether it was rain or a sprinkler that made a sidewalk wet; and how to answer hard questions, like whether a drug cured an illness. Pearl's work enables us to know not just whether one thing causes another: it lets us explore the world that is and the worlds that could have been. It shows us the essence of human thought and key to artificial intelligence. Anyone who wants to understand either needs The Book of Why.

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

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.

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

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

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/

Learn how to apply test-driven development (TDD) to machine-learning algorithms—and catch mistakes that could sink your analysis. In this practical guide, author Matthew Kirk takes you through the principles of TDD and machine learning, and shows you how to apply TDD to several machine-learning algorithms, including Naive Bayesian classifiers and Neural Networks. Machine-learning algorithms often have tests baked in, but they can't account for human errors in coding. Rather than blindly rely on machine-learning results as many researchers have, you can mitigate the risk of errors with TDD and write clean, stable machine-learning code. If you're familiar with Ruby 2.1, you're ready to start. Apply TDD to write and run tests before you start coding Learn the best uses and tradeoffs of eight machine learning algorithms Use real-world examples to test each algorithm through engaging, hands-on exercises Understand the similarities between TDD and the scientific method for validating solutions Be aware of the risks of machine learning, such as underfitting and overfitting data Explore techniques for improving your machine-learning models or data extraction

This book, presented in three volumes, examines environmental disciplines in relation to major players in contemporary science: Big Data, artificial intelligence and cloud computing. Today, there is a real sense of urgency regarding the evolution of computer technology, the ever-increasing volume of data, threats to our climate and the sustainable development of our planet. As such, we need to reduce technology just as much as we need to bridge the global socio-economic gap between the North and South; between universal free access to data (open data) and free software (open source). In this book, we pay particular attention to certain environmental subjects, in order to enrich our understanding of cloud computing. These subjects are: erosion; urban air pollution and atmospheric pollution in Southeast Asia; melting permafrost (causing the accelerated release of soil organic carbon in the atmosphere); alert systems of environmental hazards (such as forest fires, prospective modeling of socio-spatial practices and land use); and web fountains of geographical data. Finally, this book asks the question: in order to find a pattern in the data, how do we move from a traditional computing model-based world to pure mathematical research? After thorough examination of this topic, we conclude that this goal is both transdisciplinary and achievable.

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.

This integrated collection covers a range of parallelization platforms, concurrent programming frameworks and machine learning settings, with case studies.

Explore the principles and practicalities of quantum computing Key Features Discover how quantum computing works and delve into the math behind it with this quantum computing textbook Learn how it may become the most important new computer technology of the century Explore the inner workings of quantum computing technology to quickly process complex cloud data and solve problems Book Description Quantum computing is making us change the way we think about computers. Quantum bits, a.k.a. qubits, can make it possible to solve problems that would otherwise be intractable with current computing technology. Dancing with Qubits is a quantum computing textbook that starts with an overview of why quantum computing is so different from classical computing and describes several industry use cases where it can have a major impact. From there it moves on to a fuller description of classical computing and the mathematical underpinnings necessary to understand such concepts as superposition, entanglement, and interference. Next up is circuits and algorithms, both basic and more sophisticated. It then nicely moves on to provide a survey of the physics and engineering ideas behind how quantum computing hardware is built. Finally, the book looks to the future and gives you guidance on understanding how further developments will affect you. Really understanding quantum computing requires a lot of math, and this book doesn't shy away from the necessary math concepts you'll need. Each topic is introduced and explained thoroughly, in clear English with helpful examples. What you will learn See how quantum computing works, delve into the math behind it, what makes it different, and why it is so powerful with this quantum computing textbook Discover the complex, mind-bending mechanics that underpin quantum systems Understand the necessary concepts behind classical and quantum computing Refresh and extend your grasp of essential mathematics, computing, and quantum theory Explore the main applications of quantum computing to the fields of scientific computing, AI, and elsewhere Examine a detailed overview of qubits, quantum circuits, and quantum algorithm Who this book is for Dancing with Qubits is a quantum computing textbook for those who want to deeply explore the inner workings of quantum computing. This entails some sophisticated mathematical exposition and is therefore best suited for those with a healthy interest in mathematics, physics, engineering, and computer science.

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.