Ebook Directory
the best source of ebook

The book was found

Machine Learning: A Bayesian And Optimization Perspective (Net Developers)

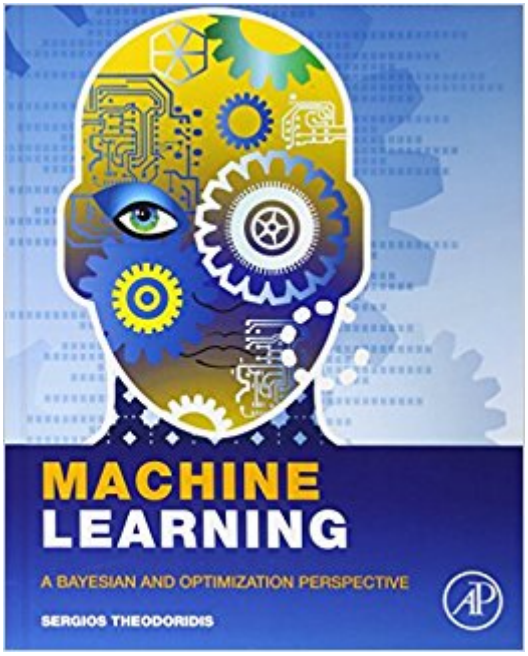

PDF
Adobe
DOWNLOAD EBOOK

Synopsis

This tutorial text gives a unifying perspective on machine learning by covering bothÃ Â probabilistic and deterministic approaches -which are based on optimization techniques Ã¢â ¬â œ together with the Bayesian inference approach, whose essence liesÃ Â 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 methodsÃ Â to Ã Â the most recent trends, with chapters written to be as self-contained as possible, making the text suitable for Ã Â 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.

Book Information

Series: Net Developers
Hardcover: 1062 pages
Publisher: Academic Press; 1 edition (April 10, 2015)
Language: English
ISBN-10: 0128015225
ISBN-13: 978-0128015223
Product Dimensions: 9.3 x 7.6 x 2 inches
Shipping Weight: 4.8 pounds (View shipping rates and policies)

Average Customer Review: 4.7 out of 5 stars 13 customer reviews

Best Sellers Rank: #61,903 in Books (See Top 100 in Books) #5 inÃ Â Books > Engineering & Transportation > Engineering > Telecommunications & Sensors > Signal Processing #25 inÃ Â Books > Textbooks > Computer Science > Artificial Intelligence #37 inÃ Â Books > Engineering & Transportation > Engineering > Mechanical > Machinery

Customer Reviews

"Overall, this text is well organized and full of details suitable for advanced graduate and postgraduate courses, as well as scholarsÃ¢â ¬Â¦" --Computing Reviews "Machine Learning: A Bayesian and Optimization Perspective", Academic Press, 2105, by Sergios Theodoridis is a wonderful book, up to date and rich in detail. It covers a broad selection of topics ranging from classical regression and classification techniques to more recent ones including sparse modeling, convex optimization, Bayesian learning, graphical models and neural networks, giving it a very modern feel and making it highly relevant in the deep learning era. While other widely used machine learning textbooks tend to sacrifice clarity for elegance, Professor Theodoridis provides you with enough detail and insights to understand the "fine print". This makes the book indispensable for the active machine learner." --Prof. Lars Kai Hansen, DTU Compute - Dept. Applied Mathematics and Computer Science Technical University of Denmark "Before the publication of Machine Learning: A Bayesian and Optimization Perspective, I had the opportunity to review one of the chapters in the book (on Monte Carlo methods). I have published actively in this area, and so I was curious how S. Theodoridis would write about it. I was utterly impressed. The chapter presented the material with an optimal mix of theoretical and practical contents in very clear manner and with information for a wide range of readers, from newcomers to more advanced readers. This raised my curiosity to read the rest of the book once it was published. I did it and my original impressions were further reinforced. S. Theodoridis has a great capability to disentangle the important from the unimportant and to make the most of the used space for writing. His text is rich with insights about the addressed topics that are not only helpful for novices but also for seasoned researchers. It goes without saying that my department adopted his book as a textbook in the course on machine learning." --Petar M. DjuriÃ„â ¡, Ph.D. SUNY Distinguished Professor Department of Electrical and Computer Engineering Stony Brook University, Stony Brook, USA. "As someone who has taught graduate courses in pattern recognition for over 35 years, I have always looked for a rigorous book that is current and appealing to students with widely varying backgrounds. The book on Machine Learning by Sergios Theodoridis has struck the perfect balance in explaining the key (traditional and new)

concepts in machine learning in a way that can be appreciated by undergraduate and graduate students as well as practicing engineers and scientists. The chapters have been written in a self-consistent way, which will help instructors to assemble different sections of the book to suit the background of students" --Rama Cellappa, Distinguished University Professor, Minta Martin Professor of Engineering, Chair, Department of Electrical and Computer Engineering, University of Maryland, USA.

This tutorial text gives a unifying perspective on machine learning by covering bothÃ Â probabilistic and deterministic approaches -which are based on optimization techniques Ã¢â ¬â œ together with the Bayesian inference approach, whose essence liesÃ Â 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 methodsÃ Â to Ã Â the most recent trends, with chapters written to be as self-contained as possible, making the text suitable for Ã Â 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. Key Features Include: An introductory chapter on related mathematical toolsAll major classical techniques:Ã Â Mean/Least-Squares regression and filtering, Kalman filtering, stochastic approximation and online learning, Bayesian classification, decision trees, logistic regression and boosting methodsA presentation of the physical reasoning, mathematical modeling and algorithmic implementation of each methodThe latest trends:Ã Â Sparsity, convex analysis and optimization,Ã Â online distributed algorithms, learning in RKH spaces, Bayesian inference, graphical and hidden Markov models,Ã Â particle filtering, deep learning, dictionary learning and latent modelingCase studies -Ã Â protein folding prediction,Ã Â 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 appliedMATLAB code for all the main algorithms are available on an accompanying website, enabling the reader to experiment with the code

The author put the machine learning and parameter estimation in systemic and unifying framework.

This is a great book for professional engineers who want to know the whole picture of the machine learning, the classic and new advanced ones. It answers a lot of my questions that I cannot get from other books. I really enjoy reading it.This book is focused more on the application level, not verbose on the theory. It is exact what professional engineer needs.

As a practitioner of Machine Learning, I am so amassed about Theodoridis' abilities to deliver fresh and precise content about the so fast evolving field of Machine Learning. This book is a must on the shelves of anybody calling herself or himself a data scientist. Sections like the ones about sparse data, Learning Kernels, Bayesian Non-Parametric Models, Probabilistic Graphical Models and Deep Learning make of this book a forefront reference on a field that is transforming the world.

It is a great book!!! It covers a wide range of subjects related to machine leaning not found in other books. It is well written and includes detailed reference list in each subject matter. The book should be useful for practitioners, graduate students and academics. I am glad I bought it.

I'm personally not a big fan of the hype around "machine learning" but this book is a good start if you haven't taken any statistics courses.

A great book to learn ML from.

fantastic book!

I'm still looking for a "perfect machine learning theory book": the one which is a pleasure to read and that covers most of concepts you see here and there all the time but always wanted to know how exactly they work: log-linear, maximum likelihood, MAP, least squares and MLS, expectation maximization, stochastic gradient descent, CRFs, mixtures of gaussian, and many others. I would like that the book explain to me why should I use this model or algorithm, why previous one would not be good? And I would like that the author take the time to carefully guide the reader throughout the theory, without leaving him alone with a bunch of matrix equations or integrals like if they were evident.I'm not a novice in the AI: I have a PhD (not in the theoretical Machine Learning though) and several years of practical experience with the algorithms. But most of the time I use the algorithms and models like blackboxes. My goal, however, is not only be able to use the algorithms and know where and how each algorithm can be used, but really understand the math that drives each

them.Unfortunately, this is not the book that can help me with my goal. In the beginning of each chapter the author really tries to move slowly with a care to details, but very fast the math becomes the only language used on the page. If, in the middle of a section you didn’t understand how equation 12 follows from equation 11, your only option is to skip the remainder of the section and this is very frustrating.As an example, when presenting the "central limit theorem", the author writes "Consider N mutually independent random variables, each following its own distribution with mean values ... and variances ... Define a new random variable **as their sum**: ... Then the mean and variance of the new variable are given by...". Here, or before, no definition of a **sum of two random variables** was presented. But this is very important to understand, because later, for example, in the "Linear Regression" section of Chapter 3, the author writes "If we model the system as a linear combiner, the dependence relationship is written as: " (a linear combination of several random variables follows). What does this mean: a linear combination of ***random variables***? How is this related to the central limit theorem which says that by adding up several random variables, the resulting variable tends to have a gaussian distribution? Author, please don’t hurry up, it’s a book, not a NIPS paper!Furthermore, the whole section "3.10.1 LINEAR REGRESSION: THE NONWHITE GAUSSIAN NOISE CASE" on page 84 cannot be directly understood from the text because the author does not explain how the joint log-likelihood function L(theta) for the model of y dependent on theta, x and nu can be constructed. The equation 3.57 gives the final expression for L(theta) but no clues on how to build it if we only have a linear model for y. I spent the whole evening just to understand that to build the joint log-likelihood function one has to transform the y = theta*x + nu into the expression p(y=yn | x=xn, theta, nu) and in order to obtain one such expression for each yn, one has to write p(y=yn | x=xn, theta, nu) = sum_k p(xn*theta=k)p(nu=yn - xn*theta). Then, the joint log likelihood L(theta) can be obtained as ln p(y=y1 | x=x1, theta, nu) + ln p(y=y2 | x=x2, theta, nu) + ... + ln p(y=yn | x=xn, theta, nu).The internet is full of information on the subject of machine learning. Almost every subject is already explained by multiple sources. The problem with the information of the Web is that it is dispersed and often incomplete. If one decides to write a book on this subject, it has to be complete and self-contained. With this book, unfortunately, one still has to google, decrypt and guess things just too often to call the reading process a pleasure.

A much needed and long awaited handbook on modern (Statistical) Machine Learning, including the most current trends of Bayesian Non-parametrics and Deep Learning. Very well-written, it helps the students grasp the main concepts as well as the technical details of the method. The companion MATLAB source codes are also very useful for better understanding of the methods. Being an

instructor myself, I would rate it as exceptional.

Download to continue reading...

Machine Learning: A Bayesian and Optimization Perspective (Net Developers) Draw in Perspective: Step by Step, Learn Easily How to Draw in Perspective (Drawing in Perspective, Perspective Drawing, How to Draw 3D, Drawing 3D, Learn to Draw 3D, Learn to Draw in Perspective) Machine Learning: A Probabilistic Perspective (Adaptive Computation and Machine Learning series) Machine Learning: An Algorithmic Perspective, Second Edition (Chapman & Hall/Crc Machine Learning & Pattern Recognition) Step-by-Step Medical Coding 2016 Edition - Text and Workbook Package, 1e (.Net Developers) Guide to the Dissection of the Dog - E-Book (.Net Developers) Elsevier's Veterinary Assisting Textbook, 1e (.NET Developers Series) Introduction to Machine Learning (Adaptive Computation and Machine Learning series) Machine Learning: For Beginners: Definitive Guide for Neural Networks, Algorithms, Random Forests and Decision Trees Made Simple (Machine Learning, Book 1) Pro ASP.NET MVC 5 (Expert's Voice in ASP.Net) The Cross-Entropy Method: A Unified Approach to Combinatorial Optimization, Monte-Carlo Simulation and Machine Learning (Information Science and Statistics) Genetic Algorithms in Search, Optimization, and Machine Learning BREAD MACHINE COOKBOOK: 120 Most Delicious Bread Machine Recipes (bread, bread bible, bread makers, breakfast, bread machine cookbook, bread baking, bread making, healthy, healthy recipes) The Little Book on Digital Marketing SEO - Search Engine Optimization: Tips and tricks for keyword research in SEO or Search Engine Optimization Introduction to Linear Optimization (Athena Scientific Series in Optimization and Neural Computation, 6) Pyomo Ã¢â ¬â ¢ Optimization Modeling in Python (Springer Optimization and Its Applications) Engineering Design Optimization using Calculus Level Methods: A Casebook Approach: Math Modeling, Simulation, & Optimization Bayesian Decision Analysis: Principles and Practice Bayesian Filtering and Smoothing (Institute of Mathematical Statistics Textbooks) Doing Bayesian Data Analysis, Second Edition: A Tutorial with R, JAGS, and Stan

Contact Us

DMCA

Privacy

FAQ & Help