


The book was found

Gaussian Processes For Machine Learning (Adaptive Computation And Machine Learning Series)



Synopsis

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.

Book Information

Series: Adaptive Computation and Machine Learning series

Hardcover: 266 pages

Publisher: The MIT Press (November 23, 2005)

Language: English

ISBN-10: 026218253X

ISBN-13: 978-0262182539

Product Dimensions: 8 x 0.8 x 10 inches

Shipping Weight: 1.6 pounds (View shipping rates and policies)

Average Customer Review: 4.7 out of 5 stars [See all reviews](#) (3 customer reviews)

Best Sellers Rank: #108,242 in Books (See Top 100 in Books) #19 in [Books > Computers & Technology > Computer Science > AI & Machine Learning > Machine Theory](#) #95 in [Books > Computers & Technology > Computer Science > AI & Machine Learning > Intelligence & Semantics](#) #1457 in [Books > Textbooks > Computer Science](#)

Customer Reviews

A specific advantage of this book is that it is one of the few that dedicate a whole chapter on the

connection between Bayesian methods using Gaussian Processes and Reproducing Kernel Hilbert Spaces. Even if this connection is a posteriori pretty obvious, it is nice to have it broken down clearly into small understandable pieces. Otherwise, all the explanations concerning Gaussian Processes themselves for regression and classification are very clear and make this book a very worthwhile read. I would recommend also reading other books focusing more on Reproducing Kernel Hilbert Spaces in order to have a complete picture of these methods (e.g. "Learning with Kernels" by Scholkopf and Smola or for an even broader picture "Generalized Additive Models" by Hastie and Tibshirani). Finally, since GP and RKHS for classification are still evolving subjects, it is probably a good idea to keep reading more material on them after finishing this book.

Even though this is not a cookbook on Gaussian Processes, the explanations are clear and to the point. The book is highly technical but it also does a great job explaining how Gaussian Processes fit in the big picture regarding the last few decades in the Machine Learning field and how they are related in some ways to both SVM and Neural Networks. I'm still working my way through the book but so far I'm extremely pleased with it. As the first reviewer said, it's an evolving subject so keep looking for new material. It's a well-edited hardcover book and at this price it's a steal.

This is another great book on ML. Although title suggests that it is solely about GP, author manages to include a lot on general ML in such a small volume (but, yes it is mostly about GP). If you are already familiar with basics of ML, this book may help you understand some details. And, of course GP techniques produce really nice plots; even this fact alone is enough to try.

[Download to continue reading...](#)

Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning series)
Machine Learning: A Probabilistic Perspective (Adaptive Computation and Machine Learning series)
Foundations of Machine Learning (Adaptive Computation and Machine Learning series) Introduction to Machine Learning (Adaptive Computation and Machine Learning series) Bioinformatics: The Machine Learning Approach, Second Edition (Adaptive Computation and Machine Learning)
Unsupervised Machine Learning in Python: Master Data Science and Machine Learning with Cluster Analysis, Gaussian Mixture Models, and Principal Components Analysis Introduction to Statistical Relational Learning (Adaptive Computation and Machine Learning series) Reinforcement Learning: An Introduction (Adaptive Computation and Machine Learning series) Boosting: Foundations and Algorithms (Adaptive Computation and Machine Learning series) Probabilistic Graphical Models: Principles and Techniques (Adaptive Computation and Machine Learning series)

IntAR, Interventions Adaptive Reuse, Volume 03; Adaptive Reuse in Emerging Economies Deep Learning: Recurrent Neural Networks in Python: LSTM, GRU, and more RNN machine learning architectures in Python and Theano (Machine Learning in Python) Unsupervised Deep Learning in Python: Master Data Science and Machine Learning with Modern Neural Networks written in Python and Theano (Machine Learning in Python) Deep Learning in Python Prerequisites: Master Data Science and Machine Learning with Linear Regression and Logistic Regression in Python (Machine Learning in Python) Convolutional Neural Networks in Python: Master Data Science and Machine Learning with Modern Deep Learning in Python, Theano, and TensorFlow (Machine Learning in Python) Deep Learning in Python: Master Data Science and Machine Learning with Modern Neural Networks written in Python, Theano, and TensorFlow (Machine Learning in Python) Machine Learning with Spark - Tackle Big Data with Powerful Spark Machine Learning Algorithms Designing Adaptive and Personalized Learning Environments (Interdisciplinary Approaches to Educational Technology) Pharmacology - Text and Adaptive Learning Package, 8e A collection of Advanced Data Science and Machine Learning Interview Questions Solved in Python and Spark (II): Hands-on Big Data and Machine ... Programming Interview Questions) (Volume 7)

[Dmca](#)