Gaussian Processes for Machine Learning

Carl Edward Rasmussen and Christopher K. I. Williams

January, 2006

Abstract

Gaussian processes (GPs) provide a principled, practical, probabilistic approach to learning in kernel machines. GPs have received growing attention in the machine learning community over the past decade. The 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 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. Code and datasets can be obtained on the web. Appendices provide mathematical background and a discussion of Gaussian Markov processes.

Resources

The book Rasmussen and Williams "Gaussian Processes for Machine Learning" is published by the MIT Press, 272 pages, 2006, ISBN 0-262-18253-X or 978-0-262-18253-9. There is an associated web page at GaussianProcess.org/gpml.

The book is also avaiable on-line, either as chapters from the list of contents page at GaussianProcess.org/gpml/chapters, or in its entirety as a single pdf file.

The GPML Toolbox provides Gaussian Process software for octave and matlab, and is available from GaussianProcess.org/gpml/code and described the brief paper Gaussian Processes for Machine Learning (GPML) Toolbox.