

Series Foreword

The goal of building systems that can adapt to their environments and learn from their experience has attracted researchers from many fields, including computer science, engineering, mathematics, physics, neuroscience, and cognitive science. Out of this research has come a wide variety of learning techniques that have the potential to transform many scientific and industrial fields. Recently, several research communities have converged on a common set of issues surrounding supervised, unsupervised, and reinforcement learning problems. The MIT Press series on Adaptive Computation and Machine Learning seeks to unify the many diverse strands of machine learning research and to foster high quality research and innovative applications.

One of the most active directions in machine learning has been the development of practical Bayesian methods for challenging learning problems. *Gaussian Processes for Machine Learning* presents one of the most important Bayesian machine learning approaches based on a particularly effective method for placing a prior distribution over the space of functions. Carl Edward Rasmussen and Chris Williams are two of the pioneers in this area, and their book describes the mathematical foundations and practical application of Gaussian processes in regression and classification tasks. They also show how Gaussian processes can be interpreted as a Bayesian version of the well-known support vector machine methods. Students and researchers who study this book will be able to apply Gaussian process methods in creative ways to solve a wide range of problems in science and engineering.

Thomas Dietterich

