# Machine Learning

## A Probabilistic Perspective

Kevin P. Murphy

**New Textbook**
**Forthcoming**
**August 2012**

This textbook provides a comprehensive, up-to-date, and accessible presentation of the field of machine learning from a unified, coherent, probabilistic perspective. For example, traditional supervised learning (classification and regression) is viewed as fitting models of the form p(y|x), where y is the output or response to be predicted, and x are the observed input features; unsupervised learning is viewed in a similar way, namely fitting models of the form p(y), where there are multiple outputs y but no inputs. This bridges the gap between the two main branches of machine learning, and makes it easy to consider variations on a theme, such as structured output prediction, multi-task learning and semi-supervised learning.

A clear distinction is made between probability models and algorithms for fitting such models. However, algorithms for learning statistical models and algorithms for inference in probabilistic models are treated in a unified way. In particular, the popular regularization-based approach to learning is treated as a form of probabilistic inference where the posterior distribution is approximated by a single "best guess", i.e., MAP estimation (maximum likelihood estimation being a special case). Thus the book covers both Bayesian and non-Bayesian approaches to machine learning.

Despite the mathematical nature of the subject, the book is written in an accessible, informal style. All the major algorithms are described in simple pseudo-code (with accompanying Matlab code freely available on the web), and many examples of the methods applied to real-world data are included.

**About the Author**

Kevin P. Murphy is Associate Professor, Dept of Computer Science & Statistics, University of British Columbia, and a visiting Research Scientist at Google.

**Early comments about the book:**

"*An astonishing machine learning book: intuitive, full of examples, fun to read but still comprehensive, strong and deep! A great starting point for any university student — and a must have for anybody in the field.*" — Professor Jan Peters, Darmstadt University of Technology and Max-Planck Institute for Intelligent Systems

"Machine Learning: A Probabilistic Perspective *covers an impressive range of the state-of-the-art in statistical machine learning. It defines a clear and broadly accessible path that begins with the fundamentals of probability, and leads to a rich toolbox of statistical models and learning algorithms.*" — Professor Erik B. Sudderth, Brown University

"*Prof. Murphy excels at unravelling the complexities of machine learning methods while motivating the reader with a stream of illustrated examples and real world case studies. The accompanying software package includes source code for many of the figures, making it both easy and very tempting to dive in and explore these methods for yourself. A must-buy for anyone interested in machine learning or curious about how to extract useful knowledge from big data.*" — Dr John Winn, Microsoft Research (Cambridge)

More information on this title will be available in early 2012 •
Sign up for our new book announcements at http://mitpress.mit.edu

# Table of Content