Probabilistic Machine Learning

An Introduction

Kevin P. Murphy

A comprehensive undergraduate-level introduction integrating classical machine learning with deep learning

Kevin Murphy's landmark work on probabilistic machine learning and Bayesian decision theory has been updated for the deep learning era. The first of two volumes, this book makes machine learning accessible for advanced undergraduate courses.

Highlights:

- Background material on linear algebra, optimization, probability, and statistics
- A focus on supervised learning using various model types (linear, nonlinear/ deep, nonparametric) with some coverage of unsupervised learning
- Python code to reproduce all figures in the book, using scikit-learn, JAX, Tensorflow, and Pytorch
- End-of-chapter exercises offer practical complement to the theoretical topics

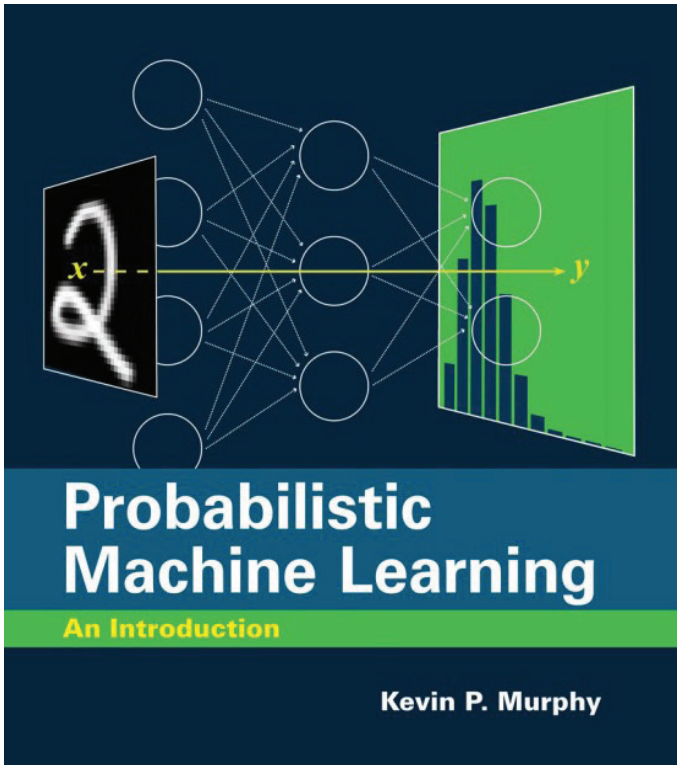

Order your digital exam copy:

Visit mitpress.mit.edu/PML
and click Request Exam/Desk Copy

For questions or to share your thoughts:
mitpress_textbooks@mit.edu

Probabilistic Machine Learning: An Introduction
By Kevin P. Murphy
March 1, 2022
Hardcover | $110.00 | 864 pp. | 8 in x 9 in
444 illustrations | ISBN: 9780262046824

Adopters of *Machine Learning: A Probabilistic Perspective* transitioning to Murphy's new book will find eight new chapters, a revamped organization to best reflect high level relationships between topics, many new references, and completely new software.

BRIEF CONTENTS, *Probabilistic Machine Learning: An Introduction*

COMING SOON: a deep dive into probabilistic modeling and inference

Probabilistic Machine Learning

Advanced Topics

FORTHCOMING 2023

An advanced book for researchers and graduate students working in machine learning and statistics that reflects the influence of deep learning

Features contributions from top researchers and scientists from places such as Google, Deepmind, Amazon, Purdue, NYU, and the University of Washington covering topics including deep generative modeling, graphical models, Bayesian inference, reinforcement learning, and causality.

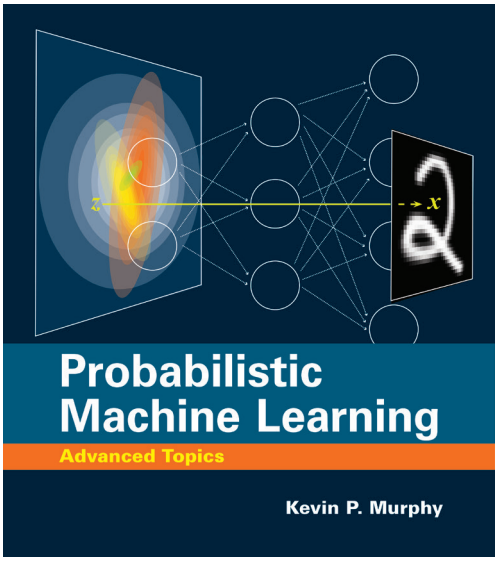

The MIT Press