



# Machine Learning: A Probabilistic Perspective

*Kevin P. Murphy*

[Download now](#)

[Read Online ➔](#)

# Machine Learning: A Probabilistic Perspective

*Kevin P. Murphy*

## Machine Learning: A Probabilistic Perspective Kevin P. Murphy

Today's Web-enabled deluge of electronic data calls for automated methods of data analysis. Machine learning provides these, developing methods that can automatically detect patterns in data and then use the uncovered patterns to predict future data. This textbook offers a comprehensive and self-contained introduction to the field of machine learning, a unified, probabilistic approach. The coverage combines breadth and depth, offering necessary background material on such topics as probability, optimization, and linear algebra as well as discussion of recent developments in the field, including conditional random fields, L1 regularization, and deep learning. The book is written in an informal, accessible style, complete with pseudo-code for the most important algorithms. All topics are copiously illustrated with color images and worked examples drawn from such application domains as biology, text processing, computer vision, and robotics. Rather than providing a cookbook of different heuristic methods, the book stresses a principled model-based approach, often using the language of graphical models to specify models in a concise and intuitive way. Almost all the models described have been implemented in a MATLAB software package--PMTK (probabilistic modeling toolkit)--that is freely available online. The book is suitable for upper-level undergraduates with an introductory-level college math background and beginning graduate students.

## Machine Learning: A Probabilistic Perspective Details

Date : Published August 24th 2012 by Mit Press

ISBN : 9780262018029

Author : Kevin P. Murphy

Format : Hardcover 1104 pages

Genre : Science, Computer Science, Technical, Artificial Intelligence, Mathematics



[Download Machine Learning: A Probabilistic Perspective ...pdf](#)



[Read Online Machine Learning: A Probabilistic Perspective ...pdf](#)

**Download and Read Free Online Machine Learning: A Probabilistic Perspective Kevin P. Murphy**

---

# From Reader Review Machine Learning: A Probabilistic Perspective for online ebook

## Kai\_S says

This book is amazing. I really enjoy reading it. Kevin Murphy is a great teacher and excellent researcher. You can get lots of insights that absent from practical books or blogs.

There are many typos in the first 3 printings. The 4th (and later) is much better. What I bought (11/24/2017) is the 6th printing (the same as the 4th).

---

## Trung Nguyen says

This can become a very good reference book for machine learning. A good complementary to Pattern Recognition and Machine Learning by Bishop.

---

## BCS says

This substantial book is a deep and detailed introduction to the field of machine learning, using probabilistic methods. It is aimed at a graduate-level readership and assumes a mathematical background that includes calculus, statistics and linear algebra.

The book opens with a brief survey of the kinds of problems to which machine learning can be applied, and sketches the types of methods that can be used to model these problems.

After a short introduction to probability, the remaining 27 chapters, over around a thousand pages, expand in depth on these and related topics.

Due to the nature of the material, much of the content is mathematical equations and proofs. Numerical examples are scarce. Rather, the authors rely upon the reader's mathematical intuition supported by a variety of graphical illustrations.

Where relevant, overviews of real-world applications of various techniques are provided, which help to make the rather abstract subject matter more concrete. In addition, MATLAB and GNU Octave code, which implements the algorithms given in the book, is freely available as a download from the book's website.

The author's style is easy but authoritative, with useful explanations and asides. That said, the book is intended as a course textbook and requires the reader to pay close attention. In the absence of an instructor,

I frequently found myself reaching for other resources in order to support the explanations given in the book, though this is probably due to my not understanding the requisite background topics sufficiently well.

Exercises are included at the end of each chapter, however while solutions are apparently available to instructors, they are not given in the book.

This comprehensive book should be of great interest to learners and practitioners in the field of machine learning.

Reviewed by Patrick Hill CEng MBCS CITP

---

### **Sami says**

Content of the book is fantastic (five stars), albeit slightly out of date in 2016. However, the first printing is so full of typos (zero stars) that it is difficult to understand how the version ever got printed. Clearly nobody read through it before printing approval. I would not recommend the first edition to anyone unless they are experts with the ability to verify and if necessary rewrite every single equation.

---

### **Aburnap says**

Hard pressed to say anyone has actually "read" this whole book--it reads like a smattering of all popular machine learning algorithms. I would not recommend it for an introduction to machine learning, not due to the technical prowess required (as it is actually much lighter on math than other similar books), but moreso due to the method and depth in which the author introduces the material.

That being said, this is perhaps the best modern "reference" text on machine learning methods. If you are already familiar with where many the methods exist in the overall landscape of machine learning, this book is absolutely fantastic.

This book describes itself as being Bayesian, but it is clearly less so than many other texts (e.g., Bishop's PRML or the frequentist slant of Hastie's ESL). Instead, most algorithms are motivated mostly by what is convention/en vogue in the machine learning community. In particular, just about every algorithm is posed as a convex relaxation to the actual posterior distribution, such that our modern optimization algorithms can compute MAP solutions given any reasonably large data set. True Bayesian inference methods are more of an afterthought, "evidenced" by their treatment only in much later chapters of MCMC/sampling and variational approaches.

Buy this book! Buy some page markers for it! But also buy an accompanying text for more principled approaches to the fundamentals.

---

### **Charles Siegel says**

Solid, but it needed better notation. The notation got very cumbersome by the end and obscured a lot of the intuition behind what was going on.

---

### **Aiham Taleb says**

Well, although this book is not made for reading purposes (in the common usage of the word reading). But I

found it really interesting. It contains every single thing that is related with Machine Learning, every algorithm that is used, every modern approach that is developed. I liked how Murphy ordered the book's topics.

Surely it is not recommended for everyone, but at least recommended for those who want to understand deeply Machine Learning in a very comprehensive way.

---