

Books useful for MLPR

You don't have to purchase any textbooks for this course. However, many students find it useful to see the material from a different perspective. I will attempt to give page references to supplementary reading throughout the notes.

Where possible, I will prioritize giving references to the following books:

- **Machine Learning: A Probabilistic Perspective.** Kevin P Murphy.
This is a book that would remain useful into PhD-level studies and research. It covers far more material than required for this course. It gives succinct treatments of topics, key references to the literature, and example code.
Many past MLPR students have found it difficult to read small sections of this book in isolation however. That is why I am providing more detailed introductory notes for parts of the course.
You can view this book online via the University library web site. Hopefully, this link will work. Otherwise search from the library main page.
- **Bayesian Reasoning and Machine Learning.** David Barber.
Parts of this course are covered by parts of this book. The book also contains a lot of material on probabilistic reasoning that we cover in PMR.
A free PDF is available. However, the page numbers in the published and free editions don't match. I'll try to refer to section numbers instead of page numbers.

The following books may also be useful or of interest:

- **Pattern Recognition and Machine Learning.** Christopher Bishop.
A few years ago this textbook was the main recommended text for the course. Depending on your background, it may not be a gentle introduction.
- **A First Course in Machine Learning.** Simon Rogers and Mark Girolami.
An excellent introductory course. A slower pace than any of the above books, but doesn't cover everything we need.
(I own the first edition, but a second edition has come out, which seems to cover more of MLPR than before.)
- **Information Theory, Inference and Learning Algorithms.** David MacKay
Some parts of this wide-ranging text are useful for this course. The book is highly recommended for those with broad interests in learning algorithms. The chapters vary significantly in the mathematical level assumed throughout the book. A free PDF is available.
- **The Elements of Statistical Learning: Data Mining, Inference, and Prediction.** Trevor Hastie, Robert Tibshirani, and Jerome Friedman.
Machine learning is a broad field, and this book comes from a different part of the community to the above books. It emphasises different methods and approaches, and is well worth a look. A free PDF is available.