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. We will attempt to give page references to supplementary reading throughout the notes.

The books on this page are also listed along with their availability by the library on Leganto.

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

- Pattern Recognition and Machine Learning. Christopher Bishop.

  Beautifully-presented overview of probabilistic machine learning as it was at the time of publication. Depending on your background, it may not be a gentle introduction. We provide separate notes for this course so that we don't have to load everything with probabilities from the start, and so we can discuss some more recent developments.

  A free PDF is now available from the official site for the book.

- 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.

  Like Bishop, Murphy aims to give a unified probabilistic view from the beginning. It can be difficult to read small sections of this book in isolation.

  You can view the Murphy book on the University library web site. If the link doesn't work, search the library main page.

The following books may also be useful or of interest:

- Bayesian Reasoning and Machine Learning. David Barber.

  Parts of this course are covered by parts of this book. The book also contains a lot more material on probabilistic reasoning than we can cover in this course.

  A free PDF is available. However, the page numbers in the published and free editions don't match. We'll refer to section numbers instead of page numbers.

- 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. The second edition has substantially more material, including Gaussian processes.

- 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.