Engineering Tripos Part IIA, 3F8: Inference, 2023-24

Leader

Prof. Richard E. Turner [1]

Lecturer

Prof. Miguel Hernandez-Lobato, Prof. Richard E. Turner

Lab Leader

Prof Miguel Hernandez-Lobato [2]

Timing and Structure

Lent Term.

Prerequisites

3F3 Statistical Signal Processing

Aims

The aims of the course are to:

- Provide a thorough introduction into the topic of statistical inference including maximum-likelihood and Bayesian approaches
- Introduce inference algorithms for regression, classification, clustering and sequence modelling
- Introduce basic concepts in optimisation and dynamic programming

Objectives

As specific objectives, by the end of the course students should be able to:

- Understand the use of maximum-likelihood and Bayesian inference and the strengths and weaknesses of both approaches.
- Implement methods to solve simple regression, classification, dimensionality reduction, clustering and sequence modelling problems.
- Implement simple optimisation methods (gradient and coordinate descent, stochastic gradient descent) and dynamic programming (Kalman filter or forward algorithm).

Content

Introduction to inference (2L)

- [redacted] decision theory, [redacted] estimation

Regression (3L)

- [redacted]

Classification (2L)

- [redacted]

Dimensionality Reduction (2L)

- [redacted]

Clustering (3L)

- [redacted]

Sequence models (3L)

- [redacted] dynamic programming

Very Basic Monte Carlo (introduced through the lectures above)

- [redacted]

Further notes

Lecture allocations above are approximate.

Coursework

Title: **Logistic Regression for Binary Classification**

To implement an algorithm for performing classification, called logistic regression, using gradient descent optimisation.

Learning objectives:

- understand the logistic regression model through visualising predictions
- how to apply maximum likelihood and MAP fitting using optimisation
- how to implement gradient ascent
- understand how feature expansions can turn linear methods into non-linear methods

Practical information:

- Sessions will take place in the DPO, during week(s) [TBD].
- This activity involves a small amount of preliminary work [estimated duration 1hr].

Full Technical Report:

Students will have the option to submit a Full Technical Report.

Booklists

There is no required textbook. However, the material covered is treated excellent recent text books:

Kevin P. Murphy Machine Learning: a Probabilistic Perspective [3], the MIT Press (2012).

David Barber Bayesian Reasoning and Machine Learning [4], Cambridge University Press (2012), available freely on the web.

Christopher M. Bishop Pattern Recognition and Machine Learning [5]. Springer (2006)

David J.C. MacKay Information Theory, Inference, and Learning Algorithms [6], Cambridge University Press (2003), available freely on the web.

Examination Guidelines

Please refer to Form & conduct of the examinations [7].

Last modified: 05/06/2023 20:04

Links
[1] mailto:ret26@cam.ac.uk
[2] mailto:jmh233@cam.ac.uk
[3] http://www.cs.ubc.ca/~murphyk/MLbook
[4] http://www.cs.ucl.ac.uk/staff/d.barber/brml
[5] http://research.microsoft.com/~cmbishop/PRML/index.htm
[6] http://www.inference.phy.cam.ac.uk/mackay/itila/
[7] https://teaching.eng.cam.ac.uk/content/form-conduct-examinations