

Engineering Tripos Part IIA, 3F8: Inference, 2017-18

Leader

[Dr Jose Miguel Hernandez-Lobato](#) [1]

Lecturer

Dr Jose Miguel Hernandez-Lobato

Lab Leader

Dr Jose Miguel Hernandez-Lobato

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, dynamic programming and Monte Carlo methods

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, clustering and sequence modelling problems.
- Implement simple optimisation methods (gradient and coordinate descent, stochastic gradient descent), dynamic programming (Kalman filter or Viterbi decoding) and simple Monte Carlo methods (importance sampling, rejection sampling, ancestral sampling).

Content

Introduction to inference (2L)

- ~~Introduction to inference~~ decision theory, estimation

Regression (2L)

- [REDACTED]

Classification (2L)

- [REDACTED]

Dimensionality Reduction (2L)

- [REDACTED]

Clustering (2L)

- [REDACTED]

Sequence models (3L)

- [REDACTED]

Basic Monte Carlo (2L)

- [REDACTED]

Further notes

Lecture allocations above are approximate.

Coursework

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

[Coursework Title]

Learning objectives:

-
-
-

Practical information:

- Sessions will take place in [Location], during week(s) [xxx].
- This activity [involves/doesn't involve] preliminary work ([estimated duration]).
-

Full Technical Report:

Students [will/won't] 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 [2], the MIT Press (2012).

David Barber Bayesian Reasoning and Machine Learning [3], Cambridge University Press (2012), available freely on the web.

Christopher M. Bishop Pattern Recognition and Machine Learning [4]. Springer (2006)

David J.C. MacKay Information Theory, Inference, and Learning Algorithms [5], Cambridge University Press (2003), available freely on the web.

Examination Guidelines

Please refer to [Form & conduct of the examinations](#) [6].

Last modified: 03/08/2017 15:42

Source URL (modified on 03-08-17): <https://teaching17-18.eng.cam.ac.uk/content/engineering-tripos-part-ii-a-3f8-inference-2017-18-0>

Links

[1] <mailto:jmh233@cam.ac.uk>

[2] <http://www.cs.ubc.ca/~murphyk/MLbook>

[3] <http://www.cs.ucl.ac.uk/staff/d.barber/brml>

[4] <http://research.microsoft.com/~cmbishop/PRML/index.htm>

[5] <http://www.inference.phy.cam.ac.uk/mackay/itila/>

[6] <https://teaching17-18.eng.cam.ac.uk/content/form-conduct-examinations>