

Course Outline : Data Analysis and Probabilistic Inference.

Lecturers: Duncan Gillies (dfg@doc.ic.ac.uk) and Marc Deisenroth (m.deisenroth@imperial.ac.uk)

Course Website (linked through CATE): <http://www.doc.ic.ac.uk/dfg/ProbabilisticInference/Bayesian.html>

Lectures

The first ten lectures are given by Duncan Gillies, lectures 11-18 will be given by Marc Deisenroth

Tutorials

The tutorials are problem sheets intended to help you understand the material of the lectures, and get useful practice for the exams. One or two of the tutorials are considerably harder than anything that will appear in an exam, so don't despair if you occasionally find them hard. There will be tutors to help you at the tutorial sessions.

Timetable

Tuesday 16.00-18.00 room 311

Thursday 09.00-11.00 room 340

Tuesday 16 Jan	Lecture 1	Lecture 2
Thursday 18 Jan	Lecture 3	Tutorial 1
Tuesday 23 Jan	Lecture 4	Tutorial 2
Thursday 25 Jan	Lecture 5	Lecture 6
Tuesday 30 Jan	Lecture 7	Tutorial 3
Thursday 1 Feb	Lecture 8	Tutorial 4
Tuesday 6 Feb	Lecture 9	Lecture 10
Thursday 8 Feb	Lecture 11	Tutorial 5
Tuesday 13 Feb	Lecture 12	Tutorial 6
Thursday 15 Feb	Lecture 13	Lecture 14
Tuesday 20 Feb	tba	tba
Thursday 22 Feb	tba	tba
Tuesday 27 Feb	tba	tba
Thursday 1 Mar	tba	tba
Tuesday 13 Mar	Revision session	Revision session

Coursework

The coursework is a practical exercise using python to model and manipulate data sets. It is divided into four parts, the first of which you can begin after lecture 2. You can work in groups of 2 or 3 for the first two parts (more details in the handout).

Booklist

The material for this course has been drawn from a wide range of sources and unfortunately there is no single book that can be recommended. The following books are available in the IC library and may prove useful in supporting the course.

1. Christopher Bishop: Pattern Recognition and Machine Learning. Springer Verlag 2006. This is a very well written and comprehensive book. The approach is a little different from the course, and it covers a lot more material, but it is still an excellent source of information.

2. Judea Pearl: Probabilistic Reasoning in Intelligent Systems. Morgan Kaufmann 1988 This book is regarded by many as the seminal work on Bayesian Networks. It is however not an easy read as Pearl has a habit of making the notation over complicated.
3. Richard Neapolitan: Probabilistic Reasoning in Expert Systems. John Wiley 1990 This is an excellent book, but unfortunately it is out of print. There are copies in the IC libraries. It covers much of the material of lectures 1 to 5 and 10 to 11.
4. Finn Jensen: Introduction to Bayesian Networks. This is a clearly written book. The approach is different to the one we take in lectures 1 to 12, and is more focused on join trees but you may still find it helpful.

- Lecture 1 Bayes' Theorem and Simple Bayesian Inference
- Lecture 2 Bayesian Decision Trees
- Lecture 3 Evidence and message passing
- Lecture 4 Probability propagation in singly connected networks
- Lecture 5 Building networks from data
- Lecture 6 Cause and Independence
- Lecture 7 Model Accuracy
- Lecture 8 Approximate Inference
- Lecture 9 Exact Inference
- Lecture 10 Probability propagation in Join Trees