ORIE 6750 Optimal Learning
Fall 2009

Often, we must decide which information to collect, and which to leave behind. Examples of such situations include global optimization, where we must choose at which points to evaluate an objective function, and dealing with parameter uncertainty, where we must decide which parameters in a problem should be investigated more fully. In making decisions about which information to collect, we implicitly weigh the cost of collecting information against its benefit, which is the ability to make better decisions at a future point in time. If these information collection decisions are made so that we obtain the maximal net benefit possible, then our learning may be said to be optimal.

This course will consider a wide variety of such information collection problems, and will cover topics in decision theory, sequential design of experiments, Bayesian statistics, and dynamic programming.

The objectives of this course are

- To provide a general framework in which to understand and formulate information collection problems.
- To provide a survey of problems, methodologies, and theoretical results within optimal learning.
- To give an idea of selected research topics in optimal learning.

Lecturer

Peter Frazier, Email pf98@cornell.edu

Assessment

There will be approximately 5 homeworks and a project. Homeworks are worth a total of 50% with the lowest grade dropped under the condition that you make a conscientious attempt at all homeworks. The project (done individually) is worth 50%.

(Update Sep 10) For each lecture from 5 (Sep. 10) onward, one person taking the class for a grade will be asked to be the scribe for the day. The scribe will take notes and write them up in LaTeX. Each set of notes will be graded and considered as a homework, and will be due 1 week from the day of the lecture. Please send both the pdf and the LaTeX file, along with any images. The pdf will be made available on Blackboard for the class. It is likely that each person will be asked to scribe twice, or possibly three times.

Prerequisites

- Familiarity with basic probability and stochastic processes at, or close to, the level of ORIE 6500.
- Some programming experience. Students will be required to write small computer programs for the homeworks, and complete a larger program for the computer assignment. No assistance in programming will be given. You should be familiar with, and have access to, a high-level programming language. Matlab or R is recommended, but students may use other languages such as Java, Python, or C if they choose.

- Although there will be some overlap with other ORIE courses in statistics such as ORIE 6780 (Bayesian Statistics and Data Analysis) and ORIE 6720 (Sequential Methods in Statistics) and knowledge of this material may be helpful to students, it is NOT a prerequisite, and all required material will be covered in this course.

Readings
There is no textbook for the course, but readings will be posted on the course website. Readings will include book chapters and technical papers.

Blackboard
Please sign up for the class webpage..

Course Content
We will cover the following topics

- Overview of Bayesian statistics: Bayes rule, conjugate priors, maximum a posteriori estimates.
- Overview of dynamic programming, as applied to problems in sequential Bayesian statistics.
- Sequential hypothesis testing.
- Ranking and selection
- Bayesian global optimization
- Multi-armed bandits

As time permits, we may also cover the following topics
- Active learning
- Changepoint detection
- Partially observable Markov decision processes
- Non-Bayesian formulations of information collection problems