OPTIMIZATION AND CONTROL

Richard Weber

Contents

Schedules

The first 6 lectures are devoted to **dynamic programming** in discrete-time and cover both finite and infinite-horizon problems; discounted-cost, positive, negative and average-cost programming; the time-homogeneous Markov case; stopping problems; value iteration and policy improvement.

The next 5 lectures are devoted to the **LQG model** (linear systems, quadratic costs) and cover the important ideas of controllability and observability; the Ricatti equation; imperfect observation, certainly equivalence and the Kalman filter.

The final 5 lectures are devoted to **continuous-time models** and include treatment of Pontryagin's maximum principle and the Hamiltonian; Markov decision processes on a countable state space and controlled diffusion processes.

Each of the 16 lectures is designed to be a somewhat self-contained unit, e.g., there will be one lecture on 'Negative Programming', one on 'Controllability', etc. Examples and applications are important in this course, so there are one or more worked examples in each lecture.

Examples sheets

There are three examples sheets, corresponding to the thirds of the course. There are two or three questions for each lecture, some theoretical and some of a problem nature. Each question is marked to indicate the lecture with which it is associated.

Lecture Notes and Handouts

There are printed lecture notes for the course and other occasional handouts. There are sheets summarising notation and what you are expected to know for the exams.

The notes include a list of keywords and I will be drawing your attention to these as we go along. If you have a good grasp of the meaning of each of these keywords, then you will be well on your way to understanding the important concepts of the course.

WWW pages

Notes for the course, and other information are on the web at
`http://www.statslab.cam.ac.uk/~rrw1/oc/index.html`.

Books

The following books are recommended.

D. P. Bertsekas, *Dynamic Programming*, Prentice Hall, 1987.
D. P. Bertsekas, *Dynamic Programming and Optimal Control*,
Volumes I and II, Prentice Hall, 1995.
L. M. Hocking, *Optimal Control: An introduction to the theory and applications*, Oxford 1991.
S. Ross, *Introduction to Stochastic Dynamic Programming*, Academic Press, 1983.
P. Whittle, *Optimization Over Time*. Volumes I and II, Wiley, 1982-83.

Ross's book is probably the easiest to read. However, it only covers Part I of the course. Whittle's book is good for Part II and Hocking's book is good for Part III. The recent book by Bertsekas is useful for all parts. Many other books address the topics of the course and a collection can be found in Sections 3B and 3D of the DPMMS library. Notation differs from book to book. My notation will be closest to that of Whittle's books and consistent throughout. For example, I will always denote a minimal cost function by $F(\cdot)$ (whereas, in the recommended books you will find F, V, ϕ, J and many others symbols used for this quantity.)