

Syllabus for BAMS 518: Markov Decision Processes

Instructor:

Steven Shechter
Office: Henry Angus 477
Office hours: By appointment
E-mail: steven.shechter@sauder.ubc.ca
Office phone: 604-822-8340

Times/Location:

The course will be held on Mon/Weds, 10AM-12PM, in HA 232. The course starts on Mon, Feb 27 and the last lecture is Weds, March 28.

Course Objectives:

This course covers the theory and application of infinite-horizon MDPs (finite horizon MDPs are covered in BAMS 517). We will cover some of the key theoretical results and discuss examples from the literature that apply these concepts.

Topics covered will include:

- Discounted cost models
- Total (undiscounted) cost models
- Average cost models
- Proving structural properties
- Approximate dynamic programming and reinforcement learning

Prerequisites:

BAMS 517 (Decision Analysis) or equivalent, as well as coursework or exposure to real analysis, probability theory, Markov chains, linear algebra, linear programming, computer programming. This is a PhD level course, which covers theory, application, and implementation of MDPs. We will go over several theorems and proofs, and you will also be required to prove results in your assignments. Therefore, previous introduction to theorems and proofs is practically necessary. Students will also “get their hands dirty” with algorithmic implementations, so ability to code in some language (e.g, C, matlab) is also required.

Required Text:

Puterman, Martin L. Markov Decision Processes, J. Wiley and Sons, New York, NY, 2005.

Supplementary Texts on reserve at I.K. Barber Learning Centre:

The following text is also an excellent reference textbooks on dynamic programming:

Betsekas, DP. Dynamic programming and optimal control. (two volume set or volume 2 of the separated versions). This should be available on course reserve from the I.K. Barber learning centre.

WebCT:

Course announcements, assignments, class notes, and other items will be handled through WebCT. For general correspondence, e-mail me directly at the e-mail address above rather than through the WebCT e-mail system (I never check that e-mail).

Grading:

2 problem sets: 50% (25% each)
Term project: 50%.

Refer to <http://www.students.ubc.ca/calendar/index.cfm?tree=3,42,96,0> for UBC's grading scale.

All work (assignments and term project) is to be done individually. It is okay to discuss high-level concepts, but not to work out solutions together.

Term Project:

Details in separate document.