

Course

Markov Decision Processes: Theory and Applications

Date: 27-28 September, 9-10 November 2017
Time: Precise schedule to be determined (per block, start at 10:00 on the first day, end at 16.00 on the second day)
Location: Eindhoven University of Technology
Course leader: Prof. Lisa Maillart
Days: 4
ECTS: 1 (attendance) / 4 (attendance & assignment)
Course fee: Free for TRAIL/Beta/OML members; others please contact the TRAIL office
Registration: Via Research School Beta

Objectives

- Students will learn how to formulate finite and infinite horizon Markov decision processes (MDPs) under various objectives.
- Students will understand the concept of a structured policy and deriving sufficient conditions for its existence.
- Students will implement various MDP solution techniques (backward induction, value iteration, policy iteration, modified policy iteration, linear programming).
- Students will be exposed to the concepts of continuous time MDPs and partially observed MDPs (POMDPs).
- Students will be able to read and appreciate recent MDP publications in top journals.

Course Description

This course introduces the fundamentals of discrete sequential decision models when outcomes are uncertain. Topics covered include the formulation and analysis of stochastic dynamic programs under various criteria; executing solution algorithms; and applications in a variety of areas, e.g., inventory control, vehicle routing, maintenance planning, resource allocation, healthcare.

The first course meeting will focus on MDP formulation (Chapters 1-3) and finite horizon MDPs (Chapter 4), including backward induction and monotone optimal policies. The second course meeting will focus on infinite horizon MDPs with discounting (Chapter 6) and without discounting (parts of Chapter 7 and 8), and introduce continuous time MDPs (Chapter 11) and POMDPs (notes).

Assignments

Two sets of homework exercises (one set per block).

Course Materials

Puterman, M.L., Markov Decision Processes, John Wiley & Sons, New York, NY, 2006. We will cover parts of Chapters 1-8 and 11.

Prerequisites

- Basic probability and Markov chain theory: Conditional probability/expectation, Bayes rule, DTMCs, CTMCs, steady state behavior, uniformization (see e.g.: Introduction to Probability Models, S.M. Ross: Section 1.6, Chapters 3, 4 & 6)

- Stochastic ordering
(see e.g.: An Introduction to Stochastic Orders, Belzunce, Riquelme and Mulero: Sections 2.1 & 2.2)
- Basic linear programming theory: Primal/dual formulation
(see e.g.: Linear Programming, K.G. Murty: Chapter 1, Section 4.4)
- Basic programming skills (e.g., Matlab)

Reading

It will be helpful to have read Chapter 2 in advance of the first meeting. Advance reading for the second meeting will be announced at the end of the first meeting.

Notes

- The maximum number of participants is 20.
- It is possible to attend the first meeting only, but to attend the second meeting you must have attended the first meeting. Priority will be given to participants who are attending both meetings and will do the course with assignments.
- Participants are asked to subscribe by September 4 at the latest.

Hotel suggestions in Eindhoven

- Sandton Eindhoven Centre, Stratumsedijk 23D, 5611 NA Eindhoven, 040-212 1330
- Boutique Hotel Glow, Keizersgracht 13 A-B, 5611 GC Eindhoven, 040-782 0078
- Hotel La Reine, Wilhelminaplein 3, 5611 HE Eindhoven, 040-820 0311
- The Student Hotel Eindhoven, Stationsweg 1, 5611 AA Eindhoven, 040-231 9792
- Crown Inn, Markt 35, 5611 EC Eindhoven, 040-245 4545
- Queen Hotel, Markt 7, 5611 EB Eindhoven, 040-245 2480
- Holiday Inn Eindhoven, Veldmaarschalk Montgomerylaan 1, 5612 BA Eindhoven, 040- 235 8235
- Park Plaza Eindhoven, Geldropseweg 17, 5611 SC Eindhoven, 040-241 6500