

Markov Decision Processes, 7.5 hp

Course period:

March 18 - May 29, 2019

Last day for application:

March 18, 2019

Course leader / Address for applications:

Marina Axelson-Fisk / marinaa@chalmers.se

Course description (Advertisement for Ph.D. students):

Markov decision processes (MPDs) is a set of models used for decision making in stochastic environments. The goal is to find a policy, which gives the optimal action in each state of the environment. When the state space of the environment is too large or partially unknown, one way is to use reinforcement learning to solve the MPD. The course begins with a brief introduction to Markov processes and hidden Markov models. Then we go through theory and applications of some different MDP variants, as well as give a brief introduction to reinforcement learning and dynamic programming.

The course will start in mid-March and run twice a week (2 hours) until end of May (LP4). The schedule will be decided by participants at an introductory meeting.

Responsible department and other participation departments/organisations:

Mathematical Sciences

Teacher:

Marina Axelson-Fisk

Examiner:

Marina Axelson-Fisk

Markov Decision Processes, 7.5 hp

1. Confirmation

The syllabus was confirmed by the Head of the Department of XXX 200X-XX-XX, 200X-XX-XX.

Disciplinary domain: Science

Department in charge: Department of Mathematical Sciences

Main field of study: Mathematical Statistics

2. Position in the educational system

Elective course; third-cycle education

3. Entry requirements

Integration theory, Foundations of probability

4. Course content

Related topics that will be touched upon in more or less depth include:

- Markov processes, discrete and continuous
- Semi-Markov processes and Hidden Markov Models
- Operations research and decision theory (brief overview)
- Dynamic programming
- Reward functions, discount factors and optimal stopping
- Partial observability
- Reinforcement learning

5. Outcomes

At the end of the course, the students will have acquired knowledge about some of the main results and techniques of Markov decision processes in a number of common application situations.

6. Required reading

To be decided.

7. Assessment

Written exam at the end of the course.

A Ph.D. student who has failed a test twice has the right to change examiners, if it is possible. A written application should be sent to the Department.

In cases where a course has been discontinued or major changes have been made a Ph.D. should be guaranteed at least three examination occasions (including the ordinary examination occasion) during a time of at least one year from the last time the course was given.

8. Grading scale

The grading scale comprises Fail, (U), Pass (G)

9. Course Evaluation

The course evaluation is carried out together with the Ph.D. students at the end of the course, and is followed by an individual, anonymous survey. The results and possible changes in the course will be shared with the students who participated in the evaluation and to those who are beginning the course.

10. Language of instruction

The language of instruction is English.