



FDD3359 Reinforcement Learning 6.0 credits

Reinforcement Learning

This is a translation of the Swedish, legally binding, course syllabus.

If the course is discontinued, students may request to be examined during the following two academic years

Establishment

Course syllabus for FDD3359 valid from Spring 2019

Grading scale

P, F

Education cycle

Third cycle

Specific prerequisites

Language of instruction

The language of instruction is specified in the course offering information in the course catalogue.

Intended learning outcomes

The course gives an introduction to the field reinforcement learning. The aim is that students should be acquainted with different methods that are used for learning based on feedback.

On completion of the course, you should be able to:

- * identify basic concepts, terminology, theories, models and methods in reinforcement learning
- * develop and systematically test a number of basic methods in reinforcement learning
- * evaluate different learning algorithms experimentally and interpret and document results of experimental studies
- * choose appropriate method to process automatically various types of data as e.g. sensor data that are used in controlling algorithms
- * account for basic methods and limitations in reinforcement learning
- * build a toolbox of different algorithms and be able to apply these on real problems

in order to

- * be familiar with basic possibilities and limitations for reinforcement learning and thereby be able to assess which problems in e.g. robot movement regulation and automatic decision-making that can be solved with these technologies
- * be able to implement, analyse and evaluate simple systems based on reinforcement learning
- * have a broad knowledge to be able to read and profit by literature in the area.

Course contents

The following fields, among others, will be treated:

Reinforcement learning with known and unknown models, discrete and continuous dynamic systems, Markov process formalism, Bellman optimality principle, exact and approximate algorithms, proofs of convergence, action policies, MDPs, discounted MDPs, POMDPs, reinforcement learning with temporal difference, Monte Carlo, and Q-learning.

The course also includes components, where students should prepare a lecture as well as develop a laboratory session where other students participate.

Examination

- EXA1 - Examination, 6.0 credits, grading scale: P, F

Based on recommendation from KTH's coordinator for disabilities, the examiner will decide how to adapt an examination for students with documented disability.

The examiner may apply another examination format when re-examining individual students.

SEM1 (4 credits, P/F), PRO1 (2 credits, P/F)

Ethical approach

- All members of a group are responsible for the group's work.
- In any assessment, every student shall honestly disclose any help received and sources used.
- In an oral assessment, every student shall be able to present and answer questions about the entire assignment and solution.