FJL3380 Theoretical Foundations of Machine Learning 6.0 credits

Teoretiska grunder för maskininlärning

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 FJL3380 valid from Autumn 2018

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

When completing the course, the student should be able to:

• Derive and apply the essential theoretical tools used in modern machine learning

• Concentration of measure in probability theory

• Stochastic optimization methods

• VC theory

• Describe the historical development of supervised and unsupervised learning algorithms

• Reflect on the advantages and drawbacks of deep learning

• Describe and explain the basic reinforcement learning algorithms and their modern versions

Course contents

A preliminary structure is given below:

Lecture 1. Introduction

Lecture 2. Probably Approximately Correct Framework and Empirical Risk Minimization

Lecture 3. Concentration inequalities

Lecture 4. The Vapnik-Chervonenkis (VC) Theory

Lecture 5. Linear Classification and Regression

Lecture 6. Regularization, Stability and Optimization

Lecture 7. Support Vector Machines and Kernel Methods

Lecture 8. Deep neural networks

Lecture 9. Clustering. Cluster validation and algorithms

Lecture 10. Reinforcement learning: model-free vs model-based algorithms

Lecture 11. Reinforcement learning: function approximation and deep RL

Disposition

Lectures on selected topics.

Course literature

Understanding Machine Leanring: From theory to algorithms, Shalev-Shwartz et al., lecture slides

Equipment

N/A

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.

Other requirements for final grade

Passing Grade based on 72h home exam and final project. The project consists in reading a few recent papers published at relevant conferences (NIPS, ICML) on a selected topic (e.g. on theoretical justification of deep learning), and to write a state-of-the-art report on the topic including historical developments, recent results, and open problems (5 pages double column minimum).

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.