EL2810 Machine Learning Theory 7.5 credits

Maskininlärningsteori

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

On 2020-10-13, the Head of the EECS School has decided to establish this official course syllabus to apply from spring semester 2021, registration number J-2020-1817.

Grading scale

A, B, C, D, E, FX, F

Education cycle

Second cycle

Main field of study

Computer Science and Engineering,Electrical Engineering

Specific prerequisites

Language of instruction

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

Intended learning outcomes

After passing the course, the student shall be able to

- derive and apply the basic theoretical tools that are used in modern machine learning
- describe known performance guarantees for important machine learning algorithms.

Course contents

Subject 1. Introduction
Main types of learning: supervised learning, unsupervised learning and reinforcement learning, and their mathematical formalisation (input and label spaces, hypothesis classes, loss function).
Subject 2. PAC framework and empirical risk minimization.
The concept of Probably Approximately Correct (PAC) learnability. Oracle inequalities and bias-variance trade-off Empirical risk minimization principle. Overfitting and the No-Free-Lunch Theorem. Uniform convergence.
Subject 3. Concentration inequalities.
Markov, Chebyshev and Chernoff bounds. Sub-gaussian random variables. Hoeffdings lemma and inequality. Bounded difference (McDiarmid) inequality.
Subject 4. Vapnik-Chervonenkis (VC) Theory
PAC learnability for finite hypothesis classes. Shattering and VC dimension. Sauer-Shelahs lemma. Rademacher complexity. Fundamental Theorem of PAC learning
Subject 5. Linear classification and regression
Linear predictors. Linear classification. Perceptron algorithms Application of VC theory to multilayer neural networks. Logistic and linear regression.
Subject 6. Regularisation, stability and optimisation
Regularized risk minimization Algorithmic stability and its application to generalization bounds for regularized risk minimization. Algorithms for convex learning: gradient descent, sub-gradient descent and stochastic gradient descent.
Subject 7. Support vector machines and kernel methods
Introduction to SVM with hard and soft margins. Performance bounds of hard and soft-margin SVM. Learning algorithms for SVM. Kernel methods; linear separability using embeddings Kernel trick and the representer theorem; admissible kernels
Subject 8. Deep neural networks
Neural networks and representation theorems. Training neural nets using backpropagation. Dropout as a regularization technique. Recent results about the lost surface and local minima of neural networks. Recent theoretical developments justifying deep learning.
Subject 9. Clustering. Cluster validation and algortihms.
Performance metrics for clusters. State-of-the-art clustering algortihms. Cluster evaluation. K-means and its performance guarantees. The EM-algorithm and its performance for Gaussian mixtures. Spectral clustering, random matrix theory and concentration.
Subject 10. Active learning, online optimization and sequential decisio making
Introduction to bandit problems and reinforcement learning. Exploration-exploitation trade-off. Fundamental limits via the change-of-measure arguments. Examples of algorithms and their guarantees. Best policy identification vs regret minimization.

Examination

- HEM1 - Homework, 1.0 credits, grading scale: P, F
- HEM2 - Homework, 1.0 credits, grading scale: P, F
- LAB1 - Laboratory assignment, 1.0 credits, grading scale: P, F
- LAB2 - Laboratory assignment, 1.0 credits, grading scale: P, F
- TEN1 - Written exam, 3.5 credits, grading scale: A, B, C, D, E, FX, 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.

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.