



ID2223 Scalable Machine Learning and Deep Learning 7.5 credits

Skalbar maskininlärning och djupinlärning

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 ID2223 valid from Autumn 2016

Grading scale

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

Education cycle

Second cycle

Main field of study

Computer Science and Engineering

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 studies fundamentals of distributed machine learning algorithms and the fundamentals of deep learning. We will cover the basics of machine learning and introduce techniques and systems that enable machine learning algorithms to be efficiently parallelized. The course complements courses in machine learning and distributed systems, with a focus on both the topic of Deep Learning as well as the intersection between distributed systems and machine learning. The course prepares the students for master projects, and Ph.D. studies in the area of Data Science and distributed computing.

The main objective of this course is to provide the students with a solid foundation for understanding large-scale machine learning algorithms, in particular, Deep Learning, and their application areas.

On successful completion of the course, the student will:

- be able to re-implement a classical machine learning algorithm as a scalable machine learning algorithm
- be able to design and train a layered neural network system
- apply a trained layered neural network system to make useful predictions or classifications in an application area
- be able to elaborate the performance tradeoffs when parallelizing machine learning algorithms as well as the limitations in different network environments
- be able to identify appropriate distributed machine learning algorithms to efficiently solve classification and pattern recognition problems.

Course contents

Topics:

- Machine Learning Principles
- Using Scalable Data Analytics Frameworks to parallelize machine learning algorithms
- Distributed Linear Regression
- Distributed Logistic Regression
- Distributed Principal Component Analysis
- Linear Algebra, Probability Theory and Numerical Computation
- Feedforward Deep Networks
- Regularization in Deep Learning
- Optimization for Training Deep Models
- Convolutional Networks
- Sequence Modelling: Recurrent and Recursive Nets
- Applications of Deep Learning

Course literature

Material from the the course is derived from the recent research publications as well as the following textbook:

Deep Learning, Yoshua Bengio, Ian Goodfellow and Aaron Courville, MIT Press (in preparation).

Examination

- LAB1 - Programming Assignments, 3.0 credits, grading scale: P, F
- TEN1 - Examination, 4.5 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.

Written examination. Laboratory tasks.

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.