FID3020 Advanced Course in Large Scale Machine Learning and Deep Learning 7.5 credits

Avancerad kurs i 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 FID3020 valid from Autumn 2017

Grading scale

Education cycle

Third cycle

Specific prerequisites

Enrolled as a doctoral student.

Language of instruction

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

Intended learning outcomes

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.
* be able to discuss, analyze, present, and critically review the very latest research advancements in the areas of Large Scale Machine Learning and Deep Learning and make connections to knowledge in related fields.

Course contents

This course is a graduate course that will cover both the basics and recent research in the area of Large Scale Machine Learning and Deep Learning. The course topics are:

Machine Learning Principles
Using Scalable Data Analytics Frameworks to parallelize machine learning algorithms
Distributed Linear Regression
Distributed Logistic Regression
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
Generative Adverserial Networks
Deep Reinforcement Learning
Applications of Deep Learning

Disposition

The course has both coursework and examination parts.

Equipment

Nothing.

Examination

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.

P/F

Other requirements for final grade

The course will be assessed with a Pass/Fail grade, based on attaining a passing grade in both the coursework and the exam.

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.