

14D005

6 ECTS

Machine Learning

Overview and Objectives

The purpose of this course is to give a solid basis of statistical learning. Both theoretical foundations and algorithmic issues of supervised and unsupervised learning will be discussed in depth.

Prerequisites

Students are expected to have taken the course "Statistical Modeling and Inference".

Course Outline

- Classification. Basic model, Bayes-optimal classifiers.
- Nearest-neighbor rules: asymptotic probability of error, k-NN rules.
- Nonparametric classification: histogram and kernel rules.
- Empirical risk minimization: concentration inequalities (Chernoff bounds), overfitting.
- Uniform convergence and VC dimension, Rademacher averages.
- Model selection and complexity regularization
- Large margin classifiers: perceptron, high-dimensional linear classifiers.
- Convex surrogate losses. Boosting, support vector machines, kernel methods.
- Algorithms for learning: gradient descent and stochastic gradient descent.
- Unsupervised learning: principal component analysis, k-means clustering. Spectral clustering, the graph Laplacian, recursive clustering
- High-dimensional phenomena: random projections and the Johnson-Lindenstrauss lemma. Sparsity and compressed sensing.

Required Activities

Evaluation is based on homework assignments, a practical project, and a final exam.

14D005

6 ECTS

Machine Learning

Evaluation

Evaluation is based on homework assignments, a practical project, and a final exam.

Materials

John Hopcroft and Ravindran Kannan,
Foundations of Data Science.

G. James, D. Witten, T. Hastie, and R. Tibshirani.
An Introduction to Statistical Learning: with
Applications in R, Springer, 2013.

Hastie, Tibshirani, Friedman, The elements of
statistical learning, Springer, 2009.

Tom Mitchell: Machine Learning, McGraw Hill,
1997.

Duda, Hart, and Stork, Pattern Classification, 2nd
Edition, Wiley, 2000

C. M. Bishop, Pattern Recognition and Machine
Learning, Springer, 2007

D. Koller and N. Friedman, Probabilistic
Graphical Models: Principles and Techniques,
2009.

Kearns and Vazirani, Introduction to
Computational Learning Theory (MIT Press
1994)

J. Shawe-Taylor and N Crisianini, Kernel
Methods for Pattern Analysis, Cambridge
University Press, 2004

R. Sutton and Andrew Barto, Reinforcement
Learning: An Introduction, MIT Press, 1998

N. Cesa-Bianchi, and G. Lugosi, Prediction,
Learning, and Games. Cambridge University
Press, New York, 2006.

L. Devroye, L. Györfi, and G. Lugosi. A
Probabilistic Theory of Pattern Recognition.
Springer, New York, 1996.