Machine Learning

Credits: 7.5 credits

Course organizer and lecturer:
Department of Mathematics and Mathematical Statistics
Armin Eftekhari, e-mail: armin.eftekhari@gmail.com

Prerequisites:
Students are expected to have a basic knowledge of probability theory, linear algebra, and calculus. A previous course on advanced probability would be beneficial.

Objective:
This course will serve as an introduction to the modern theory of statistical learning, and a rigorous but selective survey of various learning machines and algorithms.

Content:
The first part of the course covers the key concepts from the theory of statistical learning, such as PAC learning, uniform convergence, and Rademakher complexity. The second part of the course provides a rigorous introduction to various learning machines, ranging from linear predictors to neural networks.

Examination:
The examination consists of a written exam and a course project.

Literature: The main textbook will be the book by Shalev-Shwartz and Ben-David, available digitally free of charge.

Shalev-Shwartz, S., Ben-David, S., *Understanding Machine Learning: From Theory to Algorithms*, Cambridge University Press, 2014.
Hastie, T., Tibshirani, R., Friedman, J., *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Springer, 2013.
Mohri, M., Rostamizadeh, A., Talwalkar, A., *Foundations of Machine Learning*, MIT Press, 2012.