Fall 2018
Foundations of Machine Learning
Course#: CSCI-GA.2566-001
Instructor: Mehryar Mohri
Graders/TAs: Dmitry Storcheus, Zhilei Wang, Ningshan Zhang.
Mailing List

Course Description

This course introduces the fundamental concepts and methods of machine learning, including the description and analysis of several modern algorithms, their theoretical basis, and the illustration of their applications. Many of the algorithms described have been successfully used in text and speech processing, bioinformatics, and other areas in real-world products and services. The main topics covered are:

- Probability tools, concentration inequalities
- PAC model
- Rademacher complexity, growth function, VC-dimension
- Perceptron, Winnow
- Support vector machines (SVMs)
- Kernel methods
- Boosting
- On-line learning
- Decision trees
- Density estimation, maximum entropy models
- Logistic regression, conditional maximum entropy models
- Regression problems and algorithms
- Ranking problems and algorithms
- Learning languages and automata
- Reinforcement learning, Markov decision processes (MDPs)

It is strongly recommended to those who can to also attend the Machine Learning Seminar. Those interested in further pursuing the study of machine learning could also attend the Advanced Machine Learning class.

Location and Time

Warren Weaver Hall Room 102,
251 Mercer Street.
Tuesdays 5:10 PM - 7:00 PM.

Prerequisite

Familiarity with basics in linear algebra, probability, and analysis of algorithms.

Projects and Assignments

There will be 3 to 4 assignments and a project. The final grade is essentially the average of the assignment and project grades. The standard high level of integrity is expected from all students, as with all math and computer science courses.

Lectures

- Lecture 00: Introduction to convex optimization.
- Lecture 01: Introduction to machine learning.
- Lecture 02: Learning guarantees for finite hypothesis sets.
- Lecture 03: Learning guarantees for infinite hypothesis sets.
- Lecture 04: Support vector machines (SVMs).
- Lecture 05: Kernel methods.
- Lecture 06: Boosting.
- Lecture 07: Maxent models.
- Lecture 08: On-line learning.
- Lecture 09: Ranking.
- Lecture 10: Multi-class classification.
- Lecture 11: Regression.
- Lecture 12: Reinforcement learning.
- Lecture 13: Learning languages.

Textbook

The following is the required textbook for the class. It covers all the material presented (and a lot more):

- Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of Machine Learning*. MIT Press, 2012.

Technical Papers

An extensive list of recommended papers for further reading is provided in the lecture slides.

Homework

- Homework 1 [solution].
- Homework 2 [solution].
- Homework 3 [solution].
- Project.

Previous years

- 2005
- 2006
- 2007
- 2008
- 2009
- 2010
- 2011
- 2012
- 2013
- 2014 (Spring)
- 2014 (Fall)
- 2015
- 2016
- 2017
- 2018