

Machine Learning with Discriminative Models

<http://acberg.com/ml>

CS 790-134 (Spring 2015)

T/Th: 12:30-1:45 in Sitterson 014

Professor: Alex Berg

<http://acberg.com>

aberg@cs.unc.edu

office hours: after class and by appointment in Sitterson 256.

Description

Machine learning with discriminative models studies approaches and algorithms that fit models to data. Such data-driven models are allowing computer systems to semi-automatically process the data deluge in almost every field of scientific inquiry and industrial application. The course includes both hands on implementation and building a theoretical foundation for machine learning. The focus will be on supervised learning, where labeled data is available. Broad topics will include generative vs. discriminative training, parametric vs. non-parametric learning, deep networks, support vector machines, bias vs. variance trade-offs, VC dimension, and large margins. Students will be encouraged to use data and problems from their own research for homework and projects in the course.

Text

We will use “The elements of statistical learning” by Hastie, Tibshirani, and Friedman available online here <http://statweb.stanford.edu/~tibs/ElemStatLearn/> as a textbook.

Target audience

The course is targeted toward graduate students in computer science and can be taken as a first course in machine learning or to supplement COMP 790-125 Machine Learning in Computational Biology.

Prerequisites

There are no formal prerequisites, but familiarity with linear algebra and

calculus will help for the optimization components of the course, and some familiarity with statistics will be helpful generally. There will be assignments that involve coding, these will probably be most straightforward in Matlab.

Course requirements and grading policy

The course will involve reading background material, participating in class (10%), homeworks involving implementation and problem sets (50%), and a midterm (20%), and a final project (20%). There will be an optional final for students who need to bring up their grades, it can be used to replace one of the class participation, final project, or midterm grades. Students are encouraged to use their own research projects as a source of problems and data for the homeworks and final project.

Key Dates

The mid-term will be held around Thursday March 12 in class. The final project will take place over the course of the last 5 weeks of the class.

Honor Code

All students are expected to do their own work. Discussions are encouraged, but (relevant) discussions with others must be cited on each assignment.