10-702/36-702
Statistical Machine Learning

Syllabus, Spring 2017

http://www.stat.cmu.edu/~ryantibs/statml/
Lectures: Tuesday and Thursday 1:30 - 2:50 pm in HH B103

Statistical Machine Learning is a second graduate level course in **advanced machine learning**, assuming students have taken Machine Learning (10-715) and Intermediate Statistics (36-705). The term "statistical" in the title reflects the emphasis on statistical theory and methodology. The course combines methodology with theoretical foundations. Theorems are presented together with practical aspects of methodology and intuition to help students develop tools for selecting appropriate methods and approaches to problems in their own research. The course includes topics in statistical theory that are important for researchers in machine learning, including nonparametric theory, consistency, minimax estimation, and concentration of measure.

Instructors and Teaching Assistants

Instructors: Ryan Tibshirani and Larry Wasserman

Teaching Assistants: Jisu Kim, Eric Lei, Petar Stojanov, Yining Wang

The course website has their emails and office hours.

Prerequisites

You should have taken 10-701 or 10-715, and 36-705. If you did not take these courses, then you can only take the course with special permission from the instructors. **Ultimately, if you are given the OK, then it is your responsibility to do background reading to make sure you understand the concepts in those courses.**

We will assume that you are familiar with the following concepts:

1. Convergence in probability and convergence in distribution.
2. The central limit theorem and the law of large numbers.
3. Maximum likelihood; Fisher information.
4. Bayesian inference.
5. Regression.
6. Regularization; the bias-bariance tradeoff.
7. Bayes classifiers; linear classifiers; support vector machines.
8. Determinants, eigenvalues and eigenvectors.

Text

There is no text but course notes will be posted. Useful reference are:

1. Chris Bishop (2006). *Pattern Recognition and Machine Learning.*
2. Luc Devroye, László Györfi, Gábor Lugosi (1996). *A Probabilistic Theory of Pattern Recognition.*
3. László Györfi, Michael Kolher, Adam Krzyżak, Harro Walk (2002). *A Distribution Free Theory of Nonparametric Regression.*
4. Trevor Hastie, Robert Tibshirani, Jerome Friedman (2001). *The Elements of Statistical Learning*, Available at `http://www-stat.stanford.edu/~tibs/ElemStatLearn/`.
5. Larry Wasserman (2004). *All of Statistics: A Concise Course in Statistical Inference.*
6. Larry Wasserman (2005). *All of Nonparametric Statistics.*

Grading

There will be:

1. **Four assignments**. They are due **Fridays at 3:00 pm**, with the exact schedule on the course website.
2. **Midterm exam**. The date is **Tuesday March 7**.
3. **Project and course conference**. There will be a course project, with two milestones, a final report, and a class conference. Details are given on the course website.

The grading breakdown will be as follows:

50% Assignments
25% Midterm
25% Project

Policy on Collaboration

Collaboration on homework assignments with fellow students is encouraged. However, such collaboration should be clearly acknowledged, by listing the names of the students with whom you have had discussions concerning your solution. You may not, however, share written work or code after discussing a problem with others. The solutions should be written by you.

Topics and Schedule

The schedule of topics, along with all the lectures materials, are found on the course website.