EECS 281A / STAT 241A
Statistical Learning Theory
Fall 2008

Practical information

- Lectures: Tues/Thurs from 14:00–15:30 Location: Leconte Hall 3
- Discussion section (optional): Wed. 17:00–18:00, 330 Evans Hall
- Course web page: `http://inst.eecs.berkeley.edu/~cs281a/fa08/`
  All announcements and homeworks will be posted at this site; please check it regularly.

- Instructor:

Name:	Martin Wainwright
Email:	wainwrig AT SYMBOL {stat,eecs} DOT berkeley DOT edu
Offices:	263 Cory Hall or 421 Evans Hall
Phone:	643-1978
Office hours:	Tuesday 11:00 – 12:00, 421 Evans Hall
	Thursday, 15:30–16:30, Location to be determined.

- Graduate student instructors:

Name:	Oleg Mayba
Email:	oleg AT SYMBOL stat DOT berkeley DOT edu
Office hours:	To be determined

Name:	Sahand Negahban
Email:	sahandn AT SYMBOL eecs DOT berkeley DOT edu
Office hours:	To be determined

- Course reader: *An Introduction to Probabilistic Graphical Models*, by M. Jordan. Available at Northside Copy Central, Hearst & Euclid. (Pick-up starting September 1).

Course outline

This course is a 3-unit course that provides an introduction to the area of probabilistic models based on graphs. This class of models provides a flexible and powerful framework for capturing statistical dependencies in complex, multivariate data. Key issues to be addressed include representation, efficient algorithms, and various aspects statistical inference with graphical models. The primary focus of the course is on theoretical and methodological aspects of graphical models and their associated algorithms. However, the concepts are relevant to a broad range of application areas, including statistical machine learning, signal processing, computer vision, natural language processing, neuroscience, communication theory, computational biology, econometrics etc.

- fundamentals of directed and undirected graphical models

- methods for exact inference
  - elimination; sum-product algorithm; max-product algorithm
  - multivariate Gaussians; Kalman filtering; Rauch-Tung Striebel smoothing
  - hidden Markov models; forward-backward algorithm
  - junction tree framework
- exponential families, generalized linear models and graphical models
- methods for approximate inference
  - sampling-based methods
  - variational methods
- estimation
  - basics: maximum likelihood, MAP, Bayesian
  - estimation based on discriminative criteria
  - mixture models; k-means; EM algorithm
  - model selection

Prerequisites

The prerequisites are previous coursework in linear algebra, multivariate calculus, basic probability and statistics (at the level of EECS 126). Some degree of mathematical maturity is also required. Coursework or background in graph theory, information theory, optimization theory and statistical physics is relevant, and could be helpful but is not required. Familiarity with a matrix-oriented programming language (e.g., MATLAB, R, Splus etc.) will be helpful.

Evaluation

Students will be evaluated based on a combination of regular homework assignments (60%), and a final course project (40%).

Homework: Although it is acceptable for students to discuss the homework assignments with one another, each student *must* write up his/her homework on an individual basis. Each student must indicate with whom (if anyone) they discussed the homework problems. Late homeworks will not be accepted.

Course project: The course project will involve independent work on a topic of the student's own choosing. Course projects will be presented in an informal poster session at the end of semester, and the work will be summarized in a write-up.

Academic policy: Please see the EECS department policy on academic dishonesty at: `http://www.eecs.berkeley.edu/Policies/acad.dis.shtml`.