Convex Optimization: Theory, Algorithms, and Applications
ECE 88xx Proposed Syllabus

August 26, 2014

Course Description

This course will cover theory, algorithms, and applications in modern convex optimization. Practical problems from signal and image processing, machine learning, communications, control, operations research, and other disciplines will be used to motivate general mathematical results and numerical methods.

Prerequisites

Students should have background knowledge in linear algebra and least-squares problems commensurate with the successful completion of ECE 6250. It is also expected that students are very comfortable using MATLAB, as it will be used extensively for the homework problem sets.

Grading

50% Homework, 8 (±1) assignments
25% Project (individuals or small groups depending on class size)
25% Final exam

Text

Boyd and Vandenberghe, *Convex Optimization*
A pdf of this book is freely available at
http://web.stanford.edu/~boyd/cvxbook/

Other useful texts:
Bertsekas and Nedic, *Convex Analysis and Optimization*
Nocedal and Wright, *Numerical Optimization*

We will also reference recent review papers on advanced optimization techniques towards the end of the course.

Extensive use will be made of the (freely available) CVX Matlab software package.

Topical outline

1. Introduction to optimization, basic geometric and algebraic concepts
2. Convexity
   (a) convex sets, closest points, separating hyperplanes
   (b) convex functions, subgradients, conditions at minima
   (c) convex optimization problems, basic properties
3. Unconstrained minimization of smooth functions
   (a) line search methods for 1D problems
   (b) gradient and steepest descent, convergence analysis
   (c) Newton's method, convergence analysis
   (d) implementation: small scale vs. large scale
4. Theory for Constrained optimization
   (a) Lagrange multipliers and the Lagrangian
   (b) duality
   (c) Karush-Kuhn-Tucker (KKT) conditions
5. Methods for Constrained optimization
   (a) barrier techniques
   (b) projected gradient descent
   (c) splitting methods, alternating direction method of multipliers
   (d) proximal algorithms
   (e) dual averaging and distributed optimization

Throughout the course, we will return multiple times to several important classes of convex programs, including least squares, linear programming, semi-definite programming, cone programming, and geometric programming.

The course lectures and homeworks will also be interspersed with relevant programming applications including

1. Signal and image processing
   Including: image deblurring and restoration; computed imaging (tomography, MRI); antenna beamforming; spectral factorization; maximum likelihood estimation; compressed sensing and sparse recovery; covariance estimation; phase retrieval; blind deconvolution.
2. Machine learning and statistics
   Including: linear classifiers and support vector machines; ridge regression and the LASSO; logistic regression; cluster analysis using MAXCUT (SDP relaxation); matrix completion and collaborative filtering.
3. Communications
   Including: channel estimation; maximum likelihood decoding (SDP relaxation).
4. Model predictive control