446.639A Vector Space Optimization
Fall Semester 2012
Course Syllabus

Course Description

This course offers a unified treatment of optimization on finite-dimensional and infinite-dimensional vector spaces. The first half of the course is devoted to finite-dimensional optimization, and covers material that would be ordinarily treated in a "classical" optimization course: first- and second-order optimality conditions for both unconstrained and constrained problems, iterative descent algorithms for numerical optimization, and the special cases of linear and convex programming. The second half is devoted to the calculus of variations and optimal control, and covers material that would be ordinarily treated in a "classical" optimal control course: the Euler-Lagrange equations and the calculus of variations, the optimal control problem and the Maximum Principle, the Hamilton-Jacobi-Bellman equation and dynamic programming, iterative numerical algorithms for optimal control, and the special cases of the linear quadratic regulator and time-optimal control. This course does not address, e.g., optimization problems of a combinatorial or discrete nature, although the techniques developed in the course are in many cases relevant to these class of problems.

Given that entire courses are devoted to what we propose to cover in each half of the semester, one can legitimately ask what is the motivation for undertaking such an ambitious (perhaps even foolhardy) task, and what do we sacrifice in terms of depth by doubling the course's scope.

Regarding the motivation behind the course, this course is targeted to students who can't afford to take a year's worth of optimization courses, but need both a reliable intuition of the theory and a working familiarity with practical algorithms. Much of what we do in engineering involves optimization in some form, and it is important to develop a unified perspective on the various types of optimization problems that one encounters. Students who plan to work in the frontiers of optimal control and optimization obviously will need to augment their training with more advanced courses, and may even be better served by taking specialized courses in, e.g., linear programming, optimization theory, and optimal control theory rather than in this course.

As to what we sacrifice in depth, some technical aspects must necessarily be sacrificed, particularly with respect to mathematical proofs and convergence analysis. This may not be such a bad thing, as I have tried to trim the details in a way that makes the contents more friendly and accessible, and allows the reader to build a reasonably solid and rigorous intuition of the material while avoiding the proverbial "getting bogged down in technicalities." For example, proofs of global convergence of descent methods, and the Maximum Principle of optimal control, require considerable effort, and to those more interested in the practical aspects of optimization the payoff may not be worth the effort. At the same time, I believe a careful treatment of Lagrange multipliers and the Karush-Kuhn-Tucker conditions, and the derivation of the Euler-Lagrange equations in a more general setting, are important in developing a solid understanding of the material. In certain introductory courses this material receives rather cavalier treatment, and in the long run this is not helpful.

It is my hope that by the end of the course, the student, when faced with an optimization problem (either finite or infinite-dimensional), will have developed the necessary mathematical and physical intuition to determine if the problem is meaningful and solvable, and if so, to be able to select an appropriate optimization algorithm, to make sense of the results, and if necessary to customize the algorithm to exploit any special features of the problem. Analytic solutions are always preferable—the student should first try to find analytic solutions when they exist, be able to rigorously prove optimality, and also develop the ability to recognize any special features in the problem.

Instructor

The instructor for the course is (Frank) Chongwoo Park. His office is located in Building 301, Room 1515, tel. 880-7133, email *fcp@snu.ac.kr*. There will be no designated office hours for the course; rather, students are encouraged to contact the instructor at any time to discuss matters related to the course.

Classroom and Time

The class will meet on Mondays and Wednesdays from 9:30AM to 11:45 AM, in Building 301, Room 306.

Course Webpage

A course webpage will be maintained at the ETL site (exact address to be announced later). Assignments, solutions, and course announcements will be posted on the course webpage.

Prerequisites

The prerequisites for the course are an understanding of linear algebra and differential equations at an advanced undergraduate level, and proficiency in Matlab and a computer programming language, preferably C or C++. For the second half of the course, which will focus on optimal control theory, a previous introductory course in multivariable control or linear systems theory is helpful but not required.

Term Project

A term project involving the implementation of a specific optimization algorithm will be required. Term projects may be undertaken individually or in teams of up to two people. The educational objective of the term project is to expose the student to the myriad practical issues underlying optimization, which are best confronted by the actual implementation and testing of an algorithm in a practical setting.

In most cases it will be convenient to use some generally available optimization software, and the term project is designed to expose the student to the many choices available for an optimization package or algorithm. Some problems, however, may not be solvable with existing solvers, *e.g.*, convex problems that do not admit closed-form gradients and Hessians, singular optimal control problems, etc. Also, some problems may require more specialized or modified algorithms to handle, *e.g.*, high-dimensional problems with sparsity or other type of special structure.

Around the tenth week of the course, students will be asked to submit term project proposals that describe the application setting and the ensuing optimization problem, and the numerical optimization algorithm to be used or developed. Precise guidelines for the term project proposal and final report will be distributed later. The algorithms must be implemented in C, C++, or Matlab.

Grading

Grading for the course will be based on a combination of problem sets (approximately 4-5 assignments), two exams, and a term project. The approximate grading formula will be as follows:

- Exam I (covering finite-dimensional optimization): 30%;
- Exam II (covering infinite-dimensional optimization): 30%;
- Term project and presentation: 30%;
- Homework and class participation: 10%.

Grading for the term projects will be done on an individual basis; for team projects the individual contribution of each team member must be made explicit.

You are permitted, even encouraged, to work together with fellow students on problem sets. However, the solutions must be written independently by each student. Since there is no TA for the course, students will be assigned homework grading duties on a rotating basis, including the preparation of solutions for distribution to the class.

Text

The primary text on finite-dimensional optimization will be *Linear and Nonlinear Programming* by D.G. Luenberger and Y. Ye. We will not follow the sequence of topics as presented in the text, as I believe it gives too much prominence to linear programming. We shall instead begin with the general nonlinear optimization problem, and work our way back to linear programming; the detailed sequence of topics is listed below. A helpful reference is *Convex Optimization* by S. Boyd and L. Vandenweghe; free electronic drafts of the latter book should be available online at Steve Boyd's website (http://www.stanford.edu/ boyd/cvxbook). For infinite-dimensional optimization, our primary reference will be *Optimal Control Theory* by Donald Kirk. Some useful references include *Applied Optimal Control* by A.E. Bryson and Y.C. Ho, *Optimum Systems Control* by A.P. Sage, *Optimization by Vector Space Methods* by D.G. Luenberger, and an as yet unpublished set of lecture notes by Daniel Liberzon.

Outline of Topics

We will attempt to closely adhere to the following chronological list of topics; omission or inclusion of certain topics may be possible depending on time constraints and other external factors:

- **Part I: Finite-dimensional optimization**—examples and applications of finite-dimensional optimization problems (*Luenberger* Ch. 1)
- Unconstrained problems: first- and second-order necessary and sufficient conditions for optimality (*Luenberger* Ch. 7.1-7.3)
- Problems with equality constraints: Lagrange multipliers and first- and second-order optimality conditions (*Luenberger* Ch. 11.1-11.5)
- Problems with inequality constraints: the Karush-Kuhn-Tucker optimality conditions (*Luenberger* Ch. 11.8)
- Convex optimization: convex sets and convex functions, properties of convex optimization problems (*Luenberger* Ch. 7.4-7.5)
- Duality (*Luenberger* Ch. 14.1-14.2)
- Line search algorithms (*Luenberger* Ch. 8.1-8.5)

- Descent methods for unconstrained problems: the steepest descent method and its convergence properties (*Luenberger* Ch. 8.6-8.7)
- Newton's method for unconstrained minimization (*Luenberger* Ch. 8.8)
- Quasi-Newton methods for unconstrained minimization: the DFP and BFGS algorithms (*Luenberger* Ch. 10.1-10.4)
- Nonlinear constrained minimization: the gradient projection algorithm (*Luenberger* Ch. 12.4-12.5)
- Nonlinear constrained minimization: penalty and barrier methods (*Luenberger* Ch. 13.1-13.4)
- Nonlinear constrained optimization: the sequential quadratic programming algorithm (SQP) and its variants (*Luenberger* Ch. 15.1-15.5)
- Linear programming: basic results, the simplex method (*Luenberger* Ch. 2, 3.1-3.4, 4.1-4.2)
- **Part II: Infinite-dimensional optimization**—examples of infinite-dimensional optimization problems (*Kirk* Ch. 1-2)
- Calculus of variations, Euler-Lagrange equations, second-order conditions (*Kirk* Ch. 4.1-4.3)
- Variational problems with integral and other constraints (*Kirk* Ch. 4.4-4.5)
- Optimal control: relation to the calculus of variations, first- and second-order necessary conditions for optimality, conditions for fixed/free-final time, fixed/free final state problems. (*Kirk* Ch. 5.1)
- The Maximum Principle, existence of optimal controls (*Kirk* Ch. 5.3)
- Time-optimal control: the double integrator problem (*Kirk* Ch. 5.4)
- Dynamic programming and the Hamilton-Jacobi-Bellman equation (*Kirk* Ch. 3)
- The linear quadratic regular (LQR) problem (*Kirk* Ch. 5.2, 3.12)
- Direct methods for numerical optimal control (Supplemental notes)
- Indirect methods for numerical optimal control (Supplemental notes)