16:711:558 Convex Analysis and Optimization

Rutgers University, Fall 2011
Professor Jonathan Eckstein

Instructor: Professor Jonathan Eckstein
jeckstei@rci.rutgers.edu
(732) 445-3272 (RUTCOR), (848) 445-0510 (Business School)
http://rutcor.rutgers.edu/~jeckstei

Meeting time: RUTCOR 139, Wednesdays 1:40-4:40
Exceptions: Thanksgiving week, we will meet *Monday*, November 21, not November 23
The class scheduled December 7 might be moved to the evening of December 5

Text: *Convex Analysis and Optimization*
D. P. Bertsekas, with A. Nedić and A. E. Ozdaglar
Athena Scientific, 2003
http://www.athenasc.com/convexity.html

Additional reference material:

R. T. Rockafellar. *Convex Analysis*. Princeton University Press, 1970.
R. Vanderbei and E. Çinlar. *Real Analysis for Engineers*, free online book, 2000.
A. Ruszczyński. *Nonlinear Optimization*. Princeton University Press, 2006.
D. P. Bertsekas. *Nonlinear Programming*. Athena Scientific, 1995/1999.

Additional notes will be given out on specific topics.

Prerequisites

I will expect you to be familiar with the fundamentals of finite-dimensional real analysis, linear algebra, and multivariable differential calculus, as summarized in Section 1.1 of the textbook. I suggest you review this section of the text; if most of the results there are familiar to you, you should have sufficient mathematical background for the course.

Course Content

Convex analysis, the study of convexity and convex bodies, is a field of mathematical analysis that is extremely useful throughout the study of optimization theory and algorithms. This course will cover the basics of finite-dimensional convex analysis and how convex analysis applies to various kinds of optimization problems. Some of the concepts we will study, such as Lagrange multipliers and duality, are also central topics in nonlinear optimization courses; if you take or have taken a course in nonlinear optimization, there will be some unavoidable overlap, but I believe you will get a different, complementary perspective from this course. This course will not cover some other central nonlinear optimization topics, such as gradient, conjugate gradient, Newton, and barrier methods, along with convergence rate analysis – they are very important, but I will leave them to other courses.

The course will primarily follow the principle textbook, except for its treatment of conjugate functions and duality, which will instead be patterned after the classic approach of Rockafeller and the treatment in the online book by Vanderbei and Çinlar. I hope to cover the following topics:

1. Basic convexity concepts (2 weeks; Sections 1.2-1.5)
    a. Convex sets and functions
    b. Convex and affine hulls
    c. Relative interior, closure of functions, continuity, and semicontinuity
    d. Convex cones
    e. Recession cones
2. Separation and optimization (1 week; Chapter 2)
    a. Global and local optima
    b. Projection
    c. Recession and existence of optima
    d. Separating and supporting hyperplanes
3. Polarity (1 week; Chapter 3)
    a. Polar and dual cones and their representation
    b. The Minkowski-Weyl theorem
    c. The Farkas lemma and generalizations
4. Subgradients and locally generated cones (3 weeks; Chapter 4 and additional material)
    a. Subgradients and the subdifferential mapping
    b. Monotone point-to-set operators (additional notes)
    c. Elements of subdifferential calculus
    d. Approximate subgradients
    e. Subgradients and directional derivatives
    f. Normal cones and normal cone mappings
    g. Simple optimality conditions
    h. Deriving Lagrange multipliers and Slater's condition from subdifferential calculus
5. Conjugate functions and duality (3 weeks)
    a. Conjugate functions (Section 7.1 and additional notes)
    b. Simple Fenchel duality constructs (additional notes)
    c. Bifunction conjugate duality (notes from Vanderbei and Çinlar)
    d. Primal, dual, and primal-dual monotone operators associated with a convex optimization problem (additional notes)
6. Convex-analytic computational methods (4 weeks)
    a. Subgradient and Lagrangian relaxation methods (Sections 8.1-8.2)
    b. Extragradient methods (additional notes)
    c. Proximal methods – abstract, primal, dual, and primal-dual (additional notes)
    d. Cutting plane and bundle methods (Section 8.3, additional notes, and Ruszczyński text)

This is an ambitious plan, so a few topics might be dropped along the way.

Note that I will not explicitly cover infinite-dimensional spaces in this class. The key concepts and results are similar in such settings, but there tend to be tricky details that are distracting from the main flow of ideas.

In each class, I will state the key results and explain the most important or illuminating proofs. For some of the material, however, I may omit the proofs during class, and ask you to read the relevant portions of the textbook or supplementary notes.

Additional references

The Rockafellar book mentioned above is generally considered the “bible” of convex analysis. Despite being 40 years old, this book is still in print and widely available. It is an exceedingly useful comprehensive reference and very well written, and I encourage anybody working in optimization to purchase a copy. However, it is more of a monograph or reference work than a textbook, containing no explanatory figures and no exercises.

The Ruszczyński and Bertesekas textbooks on nonlinear optimization contain material related to the class and may also prove useful. Both of these texts contain appendices on convex analysis and contain many applications of the theory. I am tentatively planning to use the treatment of bundle methods from the Ruszczyński book, which is also usually the textbook in RUTCOR’s nonlinear optimization class, so you may also wish to purchase of that book, or at least be prepared to borrow a copy. There are also many other nonlinear optimization books that contain related material.

Assignments and grading

I will hand out a homework assignment every one or two weeks. The assignments will consist predominantly of mathematical proofs. Two of the assignments, one near the middle of the course, and one at the end, will be longer the rest and serve as take-home exams. Casual collaboration between students will be allowed on the regular assignments, but not on the take-home exam assignments. Your grade will be based on your performance on the assignments, probably 35% for each take-home, and 30% spread among the remaining assignments.

Office hours

I plan to be available 2:00-3:30PM on Tuesdays at RUTCOR, subject to occasional changes and cancellations. I can meet with students at other times by appointment, and I accept e-mail questions at all times.

Staying in touch

I may periodically e-mail the class through Rutgers’ RAMS system to make announcements, such as corrections to assignments. Please make sure that the e-mail address you have on file in Rutgers’ administrative information system is current, and that you check mail at that address regularly. I will also maintain a website of course material and announcements at http://eckstein.rutgers.edu/convex. Limited-distribution materials will be placed on a different site, probably Sakai (http://sakai.rutgers.edu).