SE/EC/ME 724, Advanced Optimization Theory and Methods

Course Instructor: Alex Olshevsky, alexols@bu.edu.

Course summary: This course focuses on nonlinear optimization and introduces some recent developmens in the field. It complements SE/EC 524 which focuses on linear optimization.

Nonlinear optimization is about minimizing nonlinear functions subject to constraints. The functions are usually, but not always, convex. It finds applicability in a variety of fields ranging from just about every engineering discipline to economics, finance, and statistics. This course will focus on the fundamental theory of unconstrained and constrained optimization problems. We will discuss many different optimization methods with cross-cutting applicability and demonstrate a number of results pertaining to the success or failure of these techniques.

Please note that this will be a mathematically sophisticated class that will require you to be comfortable writing proofs. The course will involve homework assignments with challenging problems which will require creativity and inventiveness on your part.

Topics covered:

- Unconstrained optimization: optimality conditions, least squares, gradient methods, Newton's method. Congugate direction and quasi-Newton methods.
- Convex sets and functions. Semi-definite programming and modern convex optimization.
- KKT conditions, Lagrange multipliers and, duality. Constrained optimization methods.
- Optimization of a sum of functions: SVRG, SAGA, and related methods.
- The Alternating Direction Method of Multipliers.
- (Time permitting) Distributed optimization.

Prerequisites: The prerequisite is SE/EC 524. In my experience, facility and comfort with doing mathematics is the most important ingredient for success in this class.

Cautionary note: The ideal audience for this class is a PhD student whose research is either in optimization or closely connected to it. If you see yourself as purely a user of optimization without much interest in what is going on behind the scenes, you will likely not find this class very useful.

Indeed, a core component of the class is understanding the theory behind optimization methods. While you will have coding assignments which will ask you to apply the methods we discuss in class, most of your time will be spent on trying to understand the underlying theory.

To gauge your interest, ask yourself whether you are interested in spending a semester navigating a mathematical thicket to work out which methods converge geometrically as $O(\alpha^t)$ in terms of iteration t where $\alpha \in (0, 1)$, as compared to slower $O(1/t)$ or $O(1/t^2)$ convergence rates. If you do not find this prospect interesting, you will likely get bored quickly and perform poorly.

Grading: Homework (20%), Mid-term (50%), Final Project (30%). I anticipate assigning a homework every two weeks. The final project will involve a report and a presentation on a paper of your choice from a list provided by the instructor.

Textbook: *Introduction to Nonlinear Optimization: Theory, Algorithms, and Applications with MATLAB*, A. Beck, SIAM Publishing, 2014.

Academic Policy: BU takes academic integrity very seriously. Academic misconduct is conduct by which a student misrepresents his or her academic accomplishments, or impedes other students' opportunities of being judged fairly for their academic work. Knowingly allowing others to represent your work as their own is as serious an offense as submitting anothers work as your own. More information on BU's Academic Conduct Code, with examples, may be found at `http://www.bu.edu/academics/policies/academic-conduct-code`

Collaboration Policy: In this class you may use any textbooks or web sources when completing your homework and programming exercises. You may also use human collaborators from class subject to the following strictly enforced conditions:

- You must clearly acknowledge all your sources (including your collaborators) on the top of your homework.
- You must write all answers in your own words.
- You must write your own code.
- You must be able to fully explain your answers upon demand.
- Obviously, you may not collaborate with anyone on the midterm.

Failure to meet any of the above conditions would constitute plagiarism and will be considered cheating in this class. If you are not sure whether something is permitted by the course policy, ASK ME! (it's much more awkward to explain your actions after the fact to the college disciplinary committee). The penalty for academic misconduct at BU is severe.

Grade Scale: Final grades will be assigned according to the following scale:

A	93.0 – 100	C+	77.0 – 79.9
A–	90.0 – 92.9	C	73.0 – 76.9
B+	87.0 – 89.9	C–	70.0 – 72.9
B	83.0 – 86.9	D	60.0 – 69.9
B–	80.0 – 82.9	F	0 – 59.9

However, I may curve up the grades on individual assignments or the midterm.

Incomplete Grades: Incomplete grades will not be given to students who wish to improve their grade by taking the course in a subsequent semester. An incomplete grade may be given for medical reasons if a physician's note is provided. The purpose of an incomplete grade is to allow a student who has a legitimate interruption in the course to complete the remaining material in another semester.

Drop Date: Students are responsible for being aware of the drop dates for the current semester. Drop forms will not be back-dated.