LINEAR AND NONLINEAR OPTIMIZATION, SPRING 2021

MATH-UA 253 and MA-UY 3204, *online-only*

Lectures MW 3:30-4:45, Recitations F 3:30-4:45

Preliminary syllabus, 12/31/2020

Instructors: This class will be taught by Professor Robert V. Kohn (kohn@cims.nyu.edu) and Professor Georg Stadler (stadler@cims.nyu.edu). Our Teaching Assistant is Shanyin Tong (shanyin.tong@nyu.edu). Our office hours will be announced at the beginning of the semester.

Description: Optimization is a major part of the toolbox of the applied mathematician, and more broadly of researchers in quantitative sciences including economics, data science, machine learning, and quantitative social sciences. This course provides an application-oriented introduction to linear programming and nonlinear optimization, with a balanced combination of theory, algorithms, and numerical implementation. While no prior experience in programming is expected, the required coursework will include numerical implementations, including some programming; students will be introduced to appropriate computational tools, with which they will gain experience as they do the numerical assignments. Theoretical topics will include linear programming, convexity, duality, minimax theorems, and dynamic programming. Algorithmic topics will include the simplex method for linear programming, selected techniques for unconstrained optimization (eg gradient descent, stochastic gradient descent, Newton's method, and quasi-Newton methods) and selected techniques for constrained optimization (eg using penalty methods and barriers). Applications will be drawn from many areas, but will emphasize economics (eg two-person zero-sum games, matching and assignment problems), data science (eg regression, convex-relaxation-based approaches to sparse inverse problems, tuning of neural networks, prediction with expert advice) and operations research (eg shortest paths in networks and optimization of network flows).

Prerequisites: multivariable calculus and linear algebra. The multivariable calculus prerequisite is satisfied by MATH-UA 123 (Calculus III) or MATH-UA 129 (Honors Calculus III) or MATH-UA 213 (Math for Economics III) with a grade of C or better. The linear algebra prerequisite is satisfied by MATH-UA 140 (Linear Algebra) or MATH-UA 148 (Honors Linear Algebra) with a grade of C or better. The prerequisites can also be met by equivalent coursework elsewhere.

Textbooks: We will use the following two books as texts. Each is available to members of the NYU community as a free pdf download from SpringerLink, and a paperback copy can be purchased at a very low cost from the book's SpringerLink page; for access, find the book in Bobcat (or use the permalink given below) then click on SpringerLink.

(1) Robert Vanderbei, Linear Programming – Foundations and Extensions, 5th edition, Springer, 2020; permalink:
`https://bobcat.library.nyu.edu/permalink/f/1c17uag/nyu_aleph007821893`

(2) David Luenberger and Yinyu Ye, Linear and Nonlinear Programming, 4th edition, Springer-Verlag, 2015; permalink:
https://bobcat.library.nyu.edu/permalink/f/1c17uag/nyu_aleph006561000

We will, of course, cover only selected parts of each book. For a few topics that aren't covered well by these books, additional notes or readings will be made available.

Computing: While no prior experience in programming is expected, the required coursework will include numerical implementations, including some programming. Students will be introduced to the open-source Julia programming language and JuMP optimization package early in the semester. Julia is rather similar to MATLAB and Python, so it will be advantageous if you are familiar with one of these languages. Note that using the Julia software package and basic programming will be an essential and obligatory part of this course.

Synchronous vs asynchronous participation: This class will be taught in an online-only format. To get the most out of it, students are strongly encouraged to participate synchronously. However, we understand that our 3:30-4:45pm timeslot is not very convenient for some students (for example those in China); therefore synchronous participation at every class is a recommendation, not a requirement. Synchronous participation will, however, be required as follows:

(a) There will be some numerically-oriented "Lab Sessions" over the course of the semester (probably about 5). Each Lab Session will be offered twice: once in a regularly-scheduled lecture slot, and once in a different slot that's more convenient for students in distant time zones. While students will be permitted to choose which instance of each Lab they attend, synchronous participation in the Lab Sessions is required.

(b) There will be a midterm exam and a final exam. The midterm will use one full lecture slot (3:30-4:45pm New York time). The scheduling of the final exam is up to the NYU Registrar. Synchronous participation will be required for both the midterm and the final, regardless of your location or timezone.

Course requirements and grading: There will be regular homework assignments (roughly one every two weeks); short quizzes (roughly once/week, done in a way that does not require synchronous participation); numerically-oriented Lab Sessions (in which synchronous participation is required); and two exams (a midterm and a final, both requiring synchronous participation). The semester grade will depend on these elements as follows: 30% homework, 10% quizzes, 10% lab participation, 25% midterm, 25% final.

Overall goals and a tentative semester plan: Our material has two related but distinct threads, namely linear programming (for which we'll mainly use Vanderbei's book) and nonlinear optimization (for which we'll mainly use the book of Luenberger & Ye). These threads will be developed in parallel over the course of the semester.

As stated in the course description, the class will emphasize connections with and applications to economics and data science. These are sometimes conceptual (such as the use of linear programming to solve two-person zero-sum games) and sometimes quite practical (such as the use of stochastic gradient descent to tune neural networks – in effect, minimizing a nonlinear and nonconvex function of many variables). As we discuss various optimization techniques, we will focus on the technique's essential character, power, and limitations; computing with Julia and JuMP will permit us to do examples, bringing the methods to life and applying them.

This class is being offered for the first time in Spring 2021, so the semester plan is still a work in progress. Tentatively, we expect to cover the following topics (in approximately this order):

- Introduction to linear programming; geometry of the simplex method (Vanderbei, chapters 1-3)
- Introduction to nonlinear optimization; gradient descent and stochastic gradient descent (Luenberger & Ye, chapter 8)
- Duality in linear programming (Vanderbei, 5.1-5.5) and its application to game theory (Vanderbei, chapter 11)
- Some applications of linear programming to data science (Vanderbei, chapter 12)
- Newton's method, and quasi-Newton methods (drawing from Luenberger & Ye 10.1-10.4)
- Linear programming applied to network problems, and to matching and assignment problems (Vanderbei chapters 14 & 15, with additional sources for matching and assignment problems)
- Interior point methods for linear programming, and barrier & penalty methods for nonlinear optimization with constraints (drawing from Luenberger & Ye chapters 5 and 13)
- Dynamic programming – shortest paths in networks; resource planning; prediction with expert advice (not in our textbooks – additional sources will be provided)

An updated (and more detailed) plan will be provided at the start of the semester.

Some Policies

Collaboration on homework: Students are welcome – and even encouraged – to discuss the homework problems with others. However, for both numerical and pencil/paper type questions, each student must implement and present his or her own solutions (this is an important part of the learning process). Direct copying of another student's solution is *not* permitted – both because it amounts to cheating, and because it defeats the entire purpose of the homework (which is to gain familiarity with new concepts and techniques).

Academic integrity: Plagiarism and cheating will not be tolerated. NYU's College of Arts and Sciences has policies in this area, and they will be followed. See http://cas.nyu.edu/academic-integrity.html

Makeup exams: The date of the midterm will be announced by the first week of classes; our final exam slot will be set by the NYU Registrar. Makeup exams will be given only for legitimate reasons such as religious holidays, conflicts with university-sponsored activities (eg athletics), or documented illness. If the reason for requesting a makeup is known in advance, permission to take a makeup *must* be requested *before* the exam date. Requests based on personal convenience will *not* be granted.