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

updated syllabus, 1/30/2021[1]

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).

Both Professors Kohn and Stadler will be available for questions and discussion after each Mon/Wed class meeting. In addition, Professor Kohn will hold an office hour Thursday 9-10pm NY time; use the link you'll find in NYU Classes to join him (the meeting ID is different from the one we use for lectures, labs, and recitations).

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.

---

[1] *Main changes from the 12/31/2020 preliminary syllabus:* added office hour information, under "Instructors;" added a section "Additional books;" added dates of the Julia/JuMP tutorials to the section on "Computing;" adjusted both items (a) and (b) of the section "Synchronous vs asynchronous participation" (note especially the information in part (b) concerning our midterm); rewrote the section on "Course requirements and grading," though its essential content hasn't changed; and added a section "Some practical matters," with information that was also announced in an email sent via NYU Classes on Saturday Jan 30 (plus an additional item on "Quirks of the NYU Spring 2021 calendar").

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.

Additional books: This course will emphasize applications to economics and data science. The following books, while not suitable for use as textbooks for this class, are rich with applications to those areas:

(a) Rakesh Vohra, Advanced Mathematical Economics, Routledge, 2005; permalink: `https://bobcat.library.nyu.edu/permalink/f/1c17uag/nyu_aleph006302959`. This book is downloadable for free via Bobcat (use the permalink then choose online access). Its scope is broader than just optimization, and its level is more advanced than this class. But the first 115 pages are pretty accessible, and the topics they cover are within the scope of this class. The book is full of applications to economics – far more than we'll have time for.

(b) Gilbert Strang, Linear Algebra and Learning from Data, Wellesley Cambridge Press, 2019. Like Vohra, this book's level is more advanced than our class; moreover its focus isn't mainly on optimization (in fact, the emphasis is on tools from linear algebra). But the book is nevertheless worth mentioning, since it offers a broad perspective on how optimization, linear algebra, and related tools are used in data science. Unfortunately, it isn't available electronically, and only NYU Shanghai has a physical copy (on reserve). The book's website `https://math.mit.edu/~gs/learningfromdata/` offers links to several sellers, and Professor Strang's website `http://www-math.mit.edu/~gs/` has links to youtube and MIT OCW sites where you'll find lectures from a 2018 course based on a draft of the book.

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.

Information about downloading and installing Julia and JuMP will be distributed during the first week of classes. The first two recitations (Feb 5 and 12) will be devoted to Julia/JuMP tutorials, to help you get started using these computational tools. Each of these Julia/JuMP tutorials will be given twice: once in the usual Friday 3:30 recitation timeslot, and again at a time that's more convenient for people in China. *Students are strongly urged to attend the Julia/JuMP tutorials in person, to make sure they have access to our computational tools.*

Synchronous vs asynchronous participation: This class will be taught in an online-only format, and all lectures and recitations will be recorded using zoom. We strongly encourage you to participate synchronously, as we think this will help you get the most out of the class. However, we understand that our 3:30-4:45pm time slot 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 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 are permitted to choose which instance of each Lab they attend, synchronous participation in the Lab sessions is required. (Each lab session will have an associated worksheet or quiz; you'll be asked to turn in the worksheet or take the quiz at the end, to create a record of your participation.)

(b) There will be a midterm exam and a final exam. *We will attempt, early in the semester, to fix a day and time for the midterm exam that works for all students (regardless of time zone).* If that effort fails, however, our fallback is to hold the midterm in the Wed 3/17 lecture slot (3:30-4:45pm NY time).

If we succeed in finding a slot for the midterm that works for all students (regardless of time zone) then we'll try to do something similar for the final exam. If that can't be done, our fallback is to hold the final exam in the "official" slot that's set by the NYU Registrar. Each exam will be given just once, and synchronous participation will be required (regardless of your location or time zone).

Course requirements and grading: The course requirements are

- **Homework:** There will be approx 7 homework assignments, roughly one every two weeks. (HW1 will be an exception – it will be due at the end of week 2.) Each assignment will count equally toward your grade (regardless of the max possible points). Taken together, the homework will be worth 30% of your grade.
- **Exams:** The midterm exam and the final exam will each be worth 25% of your grade.
- **Numerically-oriented Lab sessions:** As discussed above (see item (a) under "Synchronous vs asynchronous participation"), synchronous participation is required for the Lab sessions (whose dates will be announced well in advance). This participation will account for 10% of your grade.

- **Quizzes:** There will be weekly quizzes. They will be short and you'll be able to do them quickly, provided you are up to date on the material. The quizzes will help make sure you don't forget about this class during the periods when no lectures are scheduled; in particular, they'll help you remember to stay up to date. They will usually be posted on Thursday (thus: after the week's lectures are finished) with a completion deadline just before the following week's first lecture (which is usually on Monday). Each quiz will count equally. Taken together, the quizzes will account for 10% of your grade.

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 – alternating between them in segments (each two or three lectures long).

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)

Some Practical Matters

Zoom links: Our lectures, recitations, and lab sessions will all be online-only. They will be recorded. You'll find links for each session (and links to the recordings of past sessions) in the the Zoom section of the NYU Classes site. All lectures, recitations, and lab sessions will use the same meeting ID.

Gradescope: We will use Gradescope for homework, quizzes and exams. You'll find Gradescope in the Lessons section of the NYU Classes site. We will inform you each time we post something on Gradescope.

Class materials: All class material (for example pdf versions of any handwritten notes created during lectures) will be available in the Resources section of the NYU Classes site. We will also maintain a table on the front page of the NYU Classes site, providing a summary of the materials available (with links).

Class communication forum: We will use Campuswire. This is an external tool, so you'll have to sign up if you haven't used it before (it is not integrated with NYU Classes). Those who joined the class by Jan 30, 2021 received an emailed invitation to join the class's Campuswire site. Those who join later (or who prefer not to use the emailed invitation) will find instructions for joining the Campuswire site on the front page of our NYU Classes site.

Campuswire will be our main communication forum. If you have questions, find a typo in an assignment, need help with a problem etc, please ask there. If you know an answer to another student's question, please help them; we will try to keep an eye on the forum but would like to encourage you to also interact and help with each other. Due to the virtual format of this class, maintaining interaction with each other is even more important than usual.

Quirks of the NYU Spring 2021 calendar: Our first class is Mon 2/1. There is no class Mon 2/15 (President's Day); instead, we have a class on Thurs 2/18 (a "Legislative Day," when only Monday classes meet). There is no recitation on Fri 3/19 and no lecture on Mon 4/19; these mark the beginning or end of NYU's two long-weekend "spring breaks." Our last class is Mon 5/10.

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