LINEAR AND NONLINEAR OPTIMIZATION, FALL 2024
MATH-UA 0253 and MA-UY 3204,

Lectures: Tu&Th 3:30 - 4:45pm, 19 W 4th St Room 102, starting 09/03/2024
Recitations Fr 3:30 - 4:45pm, 40 W 4th St (Tisch Hall) Room LC11, starting 09/13/2024

Instructor: Timo Schorlepp (timo.schorlepp@nyu.edu)
Recitation leader: Albert Zhang

Office hours: TBA

Description: Optimization is a major part of the toolbox of an applied mathematician, and more broadly of researchers in quantitative sciences including economics, physics, chemistry, data science, computational biology, climate science, and machine learning. This course provides an introduction to both linear and nonlinear optimization, with a balanced combination of theory, algorithms, and numerical implementation. Examples and applications will be discussed along the way. While no prior experience in programming is expected, the required coursework will include some programming in Julia.

Theoretical topics will include linear programming, convexity, duality, and minimax theorems. Algorithmic topics will include the simplex method for linear programming, interior point methods for linear programming, selected techniques for unconstrained optimization (gradient descent, Newton's method, quasi-Newton methods, conjugate gradient) and selected techniques for constrained optimization (penalty methods and barriers). Applications will be drawn from various areas, but will emphasize economics (two-person zero-sum games, matching and assignment problems) and data science (regression, convex-relaxation-based approaches to sparse inverse problems, tuning of neural networks).

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.

Textbook: We will mostly follow *D. Luenberger and Y. Ye, Linear and Nonlinear Programming, 5th edition, Springer, 2021*[1], which is available to members of the NYU community as a free PDF download from SpringerLink. Occasionally, additional material from other textbooks may be used, such as *D. Bertsimas and J. N. Tsitsiklis, Introduction to linear optimization*, *R. J. Vanderbei, Linear programming: foundations and extensions*, *J. Nocedal and S. J. Wright, Numerical optimization*, and *S. Boyd and L. Vandenberghe, Convex optimization.*

[1] https://link-springer-com.proxy.library.nyu.edu/book/10.1007/978-3-030-85450-8

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[2] and JuMP optimization package[3] 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. The first two recitations will be devoted to Julia/JuMP tutorials, to help you get started using these computational tools.

Course requirements and grading: The course requirements are

- **Homework:** There will be roughly one assignment every two weeks. Taken together, the homework will be worth 30% of your grade.
- **Exams:** There will be one midterm exam which will be worth 30% of your overall grade. The final exam will be worth 40% of your overall grade. Exam dates are TBA.
- **Recitation:** There will be a weekly recitation. Recitation may take the form of a review of the week's material, or individual practice problems that may be similar to exam problems. Bringing a laptop to recitation is strongly encouraged in the event that there is a computationally-oriented lab session.

Overall goals and a tentative semester plan: Our material has two related but distinct threads, namely linear programming (LP) and nonlinear programming (NLP). We will begin with an introduction to linear optimization, move on to unconstrained nonlinear optimization, and then return to more advanced topics in both areas in the end. Computing with Julia and JuMP will permit us to do examples, and will bring the methods to life.

Planned topics:

- LP: standard form, geometry of LPs, extreme points
- LP: basic feasible solutions
- LP: simplex method
- LP: duality
- NLP: first and second-order optimality conditions, convexity
- NLP: gradient descent, line search
- NLP: Newton's method for optimization
- NLP: quasi-Newton methods
- NLP: basics of constrained NLP
- LP/NLP: interior point methods

[2] https://julialang.org/
[3] https://jump.dev/JuMP.jl/stable/

Some Practical Matters

Class materials: All class material will be available via Brightspace.

Ed Discussion: We will use the Ed Discussion forum for all communication and announcements in this class. You are also encouraged to ask questions there and discuss homework problems. You will find a link to Ed Discussion on Brightspace.

Gradescope: We will use Gradescope for homework assignments. You will find a link to Gradescope on Brightspace.

Some Policies

Collaboration on homework: Students are welcome – and encouraged – to discuss the homework problems with others. However, for both numerical and pencil/paper type questions, each student must implement and present their own solutions. 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