STOR 612
Foundations of Optimization
(**Fall 2023**)

Course Overview

This course is designed to introduce the foundations of optimization to graduate students at UNC-Chapel Hill, especially to students at the STOR department. It aims at conveying the most common theoretical elements and methods for graduate students who wish to pursue research and/or a career in different disciplines such as operations research, statistics, machine learning, computer science, engineering, applied mathematics, and data science. The course covers both classical and modern topics including basic to advanced theory and numerical methods. It starts from common mathematical models: linear programming and quadratic programming and moves to unconstrained convex minimization, nonconvex optimization, and then stochastic optimization. As usual, the background theory of convex optimization is convex sets, convex functions, and theory of polyhedra, which will be formally introduced in this course. Simplex, gradient descent, and Newton methods, which are the foundational algorithms in numerical optimization, will also be presented in this course. Common models such as LASSO, logistic regression, portfolio optimization, and empirical risk minimization will be used as representative examples during the course. Recent advanced methods such as accelerated gradient, proximal-based, and stochastic gradient algorithms with variance reduction and adaptive learning rates will be discussed. In addition to theory and algorithms, common optimization modeling software and computer solvers will also be introduced.

Time and Place

***Lectures*: Tuesdays and Thursdays 9:30AM - 10:45AM.**
***Place:* Hanes 130.**

***Office hours*: 1PM-2PM on Tuesdays and Thursdays (or by appointment)**.
Zoom link (please email me to request office hours via zoom): unc.zoom.us/my/mikeoneill

Instructor and TA

Instructor: Mike O'Neill (mikeoneill@unc.edu)

Office: 323 Hanes Hall, UNC-Chapel Hill

TA: Yang Luo (yangluo@unc.edu), Office Hours: TBA

Course content

	Mathematical Optimization Models
Lecture 1	Mathematical optimization - Examples and basic concepts: decision variables, objective, constraints, feasible solutions, feasible set, optimal value, and optimal solutions, etc.

Lecture 2	Linear programming (LP) - Representative examples (old and new): production planning, blending, network flows, transportation, optimal transport, least absolute deviations (LAD), Wasserstein barycenter, etc (if time permits).
Lecture 3	Forms of LPs and preprocessing (converting to standard form, solution construction, redundancy, etc)
	Mathematical Tools from Convex Analysis
Lecture 4	Convex sets, convex hulls, polyhedra, and convex cones: definitions, examples, and basic properties
Lecture 5	Convex functions: definitions, examples, and basic properties
Lecture 6	Smooth convex functions and strongly convex functions
Lecture 7	Geometry of polyhedra and solution structures of LPs: basic feasible solutions, optimal solutions, and existence of solutions
	Linear Programming (LP)
Lecture 8	Mathematical aspects of simplex methods: matrix form, tableau form, and examples
Lecture 9	Two-phase simplex methods and complexity
Lecture 10	Applications and software for LPs: modeling software and LP solvers
Lecture 11	Dual problem of LP and its construction, and weak duality
Lecture 12	Strong duality, complementarity slackness, and Farkas' lemma
Lecture 13	Sensitivity analysis and robust LPs (if time permits)
	Quadratic Programming (QP)
Lecture 14	Mathematical formulations of convex QPs and examples (e.g., SVM, linear optimal control, and portfolio optimization)
Lecture 15	Simple QPs (e.g., least-squares and QPs with equality constraints) and constrained QPs, KKT conditions, and duality
Lecture 16	Solution methods for QPs: projected gradient, and active-set methods
Lecture 17	Introduction to interior-point algorithms.
	General Nonlinear Optimization (NLP)
	(*With an emphasis on convex optimization, but also covering some nonconvex problems* (*e.g., nonnegative matrix factorization*) *and methods*)
Lecture 18	Mathematical models and composite problems, representative examples (Empirical risk minimization, LASSO, logistic regression, nonnegative matrix factorization, etc).
Lecture 19	Subdifferentials, normal cone, optimality condition, proximal operators, and projection
Lecture 20	The gradient descent method, its variants, and convergence analysis
Lecture 21	The accelerated gradient descent method and its variants
Lecture 22	Proximal/projection operator calculations; proximal gradient methods; and accelerated proximal gradient methods
Lecture 23	Other advanced optimization methods: mirror descent, coordinate descent, and conditional gradient methods (backup, if time permits)
Lecture 24	Implementation aspects: per-iteration complexity analysis, line-search, and specialization, etc (homework, tutorials, or recitations).
Lecture 25	Newton methods, convergence analysis, and implementation (if time permits)

Lecture 26	Quasi-Newton methods and BFGS updates (if time permits)
	Stochastic optimization
Lecture 27	Introduction to stochastic optimization and examples
Lecture 28	Stochastic gradient (SGD) methods, convergence, and complexity theory
Lecture 29	Stochastic gradient methods with variance reduction and adaptive LR (if time permits)

Course materials

Lecture notes

Lecture notes and slides will be provided to students on Canvas. They must be used internally in this course. Please do not distribute these materials without instructor's permission.

Reference Books

[B_1]. R. T. Rockafellar: Convex Analysis, 1970, Princeton Univ. Press. This book is now available online for free and can be downloaded from https://www.convexoptimization.com/TOOLS/AnalysisRock.pdf.

[B_2]. D. Bertsimas and J. Tsitsiklis: Introduction to Linear Optimization, Third Edition, Athena Scientic, Belmont, Massachusetts, 1997.

[B_3]. S. Boyd and L. Vandenberghe: Convex Optimization, 2006, Cambridge Univ. Press. This book is available for free at http://stanford.edu/~boyd/cvxbook/.

[B_4]. A. Beck: First-order methods in optimization, volume 25, SIAM, 2017.

[B_5]. Y. Nesterov: Introductory lectures on Convex Optimization, 2004 or 2018. The lectures can be found at http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.693.855&rep=rep1&type=pdf.

[B_6]. D. Bertsekas, Convex Optimization Theory/Algorithms, Athena Scientific, 2009.

[B_7]. J. Nocedal and S. Wright, Numerical Optimization, Springer-Verlag, 2006.

[B_8]. G. Lan, First-order and Stochastic Optimization Methods for Machine Learning, Springer-Nature, 2020.

Prerequisites

This course requires basic background on numerical linear algebra and some parts of multivariable calculus such as basic topology (e.g., open, closed sets, limits, continuity, derivatives, etc.), gradients, and Hessian of multivariable functions. Students are recommended to take MATH547/MATH577 or equivalence (e.g., STOR 415) before this course. If student does not have such background, then it is highly recommended to drop the course.

Class attendance and course website

Students are expected to attend all classes. Students are responsible for assignments or policies that are announced in class or in material handed out in class, whether or not students are in class. Students are also responsible for any material distributed electronically by email or via the course webpage. A course website is available at https://canvas.unc.edu/. Once student logs in with his/her ONYEN and password, go to the course site entitled STOR612.001.FA23. We will use it for posting lecture notes and grades, and for other purposes. Assignments will be turned in via Gradescope, which will be linked to Canvas assignments. Please assign pages when submitting a Gradescope assignment to ensure your work is graded accurately and efficiently. Please visit: https://curricula.unc.edu/curriculum-proposals/cim/syllabus/#details-0-6 for further information on the UNC Attendance Policy.

Homework and Projects

Homework (HW) or projects will be assigned in most weeks. It will be posted on the course site in the "Assignments" page of Canvas. It will be due on dates stated in the assignment information posted in Canvas; it will be graded, and the grades will be returned to you, usually within one week. Students are asked to turn in their work via Gradescope, which will be linked to the Canvas assignments. In general, late homework/project will receive no credit. Occasionally, reasonable exceptions may be made (e.g., travel for conferences or sickness), with the instructor's specific approval in each case. Verbatim copying of homework/project is absolutely forbidden and constitutes a violation of the Honor Code. Students who believe their grade on an assignment or examination is in error can request adjustment of the grade during a period of three weeks after the due date of the item in question. The three-week period may be shortened for the last one or two assignments of the semester. Any questions regarding homework/project grades should first be taken up with the TA; if these questions cannot be resolved with the TA, then feel free to discuss them with the instructor. Students are responsible for checking grade book regularly (e.g., weekly) and interacting with the instructor to prevent any mistake on grading for both homework/project assignments and exams during the semester.

Exams

There will be one in-class examination (75 minutes), scheduled in one class during the midterm week, and one final examination (three hours). The exact time of the midterm will be announced directly via Canvas. The time and place for the final exam are scheduled by the Registrar office, and can be found at https://registrar.unc.edu/academic-calendar/. This time will also be mentioned in class. Any questions regarding exam grades should be taken up with the instructor. There is no make-up exam for the midterm. If student has an official reason such as sickness, please contact the instructor to discuss possible solutions. This arrangement should be done as early as possible before the exam if student already has a plan, or after the exam, otherwise.

Course grade

A student's course grade will be based on the final course average, in computing which the graded work will be weighted as follows: regular homework/project assignments: 35%; in-class examination: 25%; final exam: 40%. No homework/project assignment is dropped. All questions about course registration and waitlists should be directed to Ms. Christine Keat (crikeat@email.unc.edu, Hanes 321, 919-962-2307). Again, each student is responsible for verifying his or her recorded scores (homework, projects, and exams) during the semester.

The Honor Code will be observed at all times in this course. The terms of the Honor Code are set out at http://instrument.unc.edu. Please carefully check it.

Community Standards and Classroom Policies

Please visit https://carolinatogether.unc.edu/community-standards/ for further information about Covid-19 and UNC's policies.

Other Resources

All students should be aware of the following resources that are available at UNC-Chapel Hill.

- **Accessibility Resources and Service (ARS):** The information of this service can be found at https://ars.unc.edu.
- **Counseling and Psychological Services (CAPS):** The information about this service can be found at https://caps.unc.edu.
- **Title IX resources:** Please visit https://eoc.unc.edu/our-policies/state-and-federal-laws/title-ix-and-vawa/.