ORIE 6365

Continuous Optimization: Algorithms and Complexity

Spring 2026

Lecture schedule: Mondays and Wednesdays, 10:10 AM - 11:25 AM
Location: 102 Upson Hall (Ithaca) and broadcasted to 397 Bloomberg Center (Cornell Tech)

Instructor: Nikita Doikov
Email: `nikita.doikov@cornell.edu`
Office: 218 Rhodes Hall
Office hours: Tuesdays 11:00 AM - 12:00 PM,
Thursdays 2:00 PM - 3:00 PM,
and by appointment.

Course description

This is a graduate-level course on the theory and algorithms of continuous optimization. It prepares students for research in optimization theory and for developing advanced methods for applications in operations research, machine learning, and related domains.

The main emphasis is on understanding different classes of optimization problems and their theoretical limitations. We will rigorously study convergence rates for first-order and second-order optimization methods on both convex and non-convex problems, with the aim of describing the range of their applicability. Some examples of applications will be drawn from machine learning, semidefinite programming, and large-scale graph optimization

Topics: convexity, smooth and non-smooth problems, duality. Ellipsoid and subgradient methods. Mirror descent and the geometry of optimization problems. Accelerated optimal methods, lower complexity bounds and resisting oracles. Composite problems. Stochastic and large-scale optimization, variance reduction techniques. Second-order algorithms: Newton, quasi-Newton, and interior-point methods.

Learning outcomes:

- Identify and classify continuous optimization problems based on their structure
- Design and implement efficient algorithms to solve different classes of optimization problems
- Analyze and compare convergence rates and complexity bounds of optimization methods

Prerequisites

Students should have a strong background in linear algebra and multivariable calculus. Knowledge of probability theory will be needed for certain topics. Basic proficiency in programming (e.g., Python) is necessary to complete practical assignments. Some previous exposure to introductory optimization (e.g., ORIE 6300) is highly recommended but not required.

Please talk with me if you have any questions regarding whether your background is sufficient to complete the course.

Grading Policy

Homeworks (around 3-4)	50%
Practical assignments (around 2)	30%
Final take-home exam	20%

These weights are approximate; the instructor reserves the right to change them later.

Schedule

This is a tentative schedule which is subject to change without notice, depending on available time.

- Introduction. Derivatives, optimality conditions, smooth functions. Convergence of gradient descent for non-convex smooth functions. Lower bounds for global optimization.
- Convexity, separation, and subgradients basics.
- Ellipsoid and subgradient algorithms and their complexity bounds.
- Smooth and stochastic problems: convergence rates for gradient methods.
- Lower complexity bounds. Optimal accelerated methods.
- Bregman divergence. Mirror descent algorithm. Smoothing.
- Newton's method. Local convergence for self-concordant functions. Regularization techniques and global convergence.
- Structural optimization and interior-point methods.

Textbook

There is no single required textbook. Each lecture will be supported by lecture notes and/or relevant references.

Academic Integrity

Each student is expected to abide by the Cornell University Code of Academic Integrity. While collaboration on homework is encouraged, final write-ups must be your own work.

Accommodations

Students with disabilities are encouraged to contact Student Disability Services to ensure necessary arrangements: `www.sds.cornell.edu`.