Syllabus
CAAM/DATA/STAT 31019
Optimization for Data Science - Autumn 2026

Description

Optimization problems arising in modern data science often fall outside the traditional assumptions of classical optimization theory: objectives may be nonconvex, nonsmooth, highly overparameterized, and accessible only through noisy gradient information. This course develops algorithms and theory for this regime, including stochastic and variance-reduced gradient methods, adaptive and preconditioned algorithms, mirror descent and non-Euclidean geometry, nonconvex complexity theory and benign landscape analysis, weakly convex optimization, automatic and implicit differentiation, and min-max formulations. A central theme is the interplay between optimization and statistical estimation: algorithmic choices can influence regularization, generalization, and implicit bias, while statistical considerations determine the target accuracy. Throughout, the course will balance theoretical analysis with computational implementation.

Coordinates

Time: MW, 3:00–4:20 p.m.
Location: Ryerson Physical Laboratory, Room 176
Website: mateodd25.github.io/opti-ds/
Assignments and announcements: Canvas.

Personnel

Instructor: Mateo Díaz
Email: mateodd@uchicago.edu
Office: Suzuki Center, 304
Office hours: M, 4:45–5:45 p.m.

Textbook

There is no required textbook. Lecture notes and selected readings will be posted on the course website and Canvas. The following are useful background references; additional papers will accompany the more specialized topics.

- D. Drusvyatskiy, *Convex Analysis and Nonsmooth Optimization* (2020), lecture notes.
- J. Nocedal and S. J. Wright, *Numerical Optimization*, second edition, Springer (2006), book information.
- S. Bubeck, *Convex Optimization: Algorithms and Complexity* (2015), available online.
- L. Bottou, F. E. Curtis, and J. Nocedal, *Optimization Methods for Large-Scale Machine Learning* (revised 2018), available online.

Topics and Tentative Sequence

The course will develop a selection of the following topics in depth.

1. **Foundations and the statistical setting.** Empirical and population objectives; optimization and estimation error; smoothness, convexity, conditioning, and optimality conditions; gradient descent and basic complexity bounds. Least squares and regularized regression as running examples.

2. **Stochastic optimization.** Stochastic gradient descent, minibatching, step-size schedules, and averaging; the roles of noise and target accuracy; finite-sum problems and variance reduction.

3. **Geometry, adaptivity, and preconditioning.** Mirror descent and Bregman divergences; non-Euclidean geometry; adaptive gradient methods and diagonal preconditioning; the role of curvature and the assumptions behind convergence guarantees.

4. **Smooth nonconvex optimization.** Complexity of finding approximate stationary points; saddle points and the distinction between local and global guarantees; benign landscapes and overparameterization in selected models.

5. **Nonsmooth and weakly convex optimization.** Subgradients, proximal maps, and composite objectives; weak convexity and the Moreau envelope; appropriate stationarity measures and stochastic methods.

6. **Differentiation through computation and optimization.** Automatic differentiation, Jacobian-vector and vector-Jacobian products; differentiating through iterations; implicit differentiation through optimality conditions and its regularity requirements. Automatic differentiation will also appear earlier in computational exercises.

7. **Min-max optimization.** Saddle-point formulations, gradient descent-ascent and its limitations, and extragradient methods; selected applications to robust estimation or learning.

8. **Optimization and statistical estimation.** Early stopping, explicit and implicit regularization, and algorithm-dependent solution selection; case studies connecting convergence, generalization, and implicit bias. These connections will also recur throughout the preceding topics.

Grading scheme

Course grades are based on homework, an in-class midterm, a final project, and participation. We will use an optimization-based grading scheme, invented by Ben Grimmer, that aims to choose the weights for the different components so as to maximize your grade subject to some reasonable constraints. For each student, let C_H, C_M, C_F, C_P be the homework, midterm, final, and participation scores, respectively, each on a scale from 0 to 100. Let H, M, F, P be the corresponding weights, measured in percentage points. The course grade, also on a scale from 0 to 100, is the optimal value of

$$
\begin{aligned}
\max_{H,M,F,P\in\mathbb{R}} \quad & \frac{C_H H + C_M M + C_F F + C_P P}{100} \\
\text{subject to} \quad & H + M + F + P = 100, \\
& 10 \leq H \leq 25, \\
& 30 \leq M, \\
& M \leq F, \\
& M + F \leq 80, \\
& 0 \leq P \leq 10.
\end{aligned}
$$

Thus, the weight of take-home assignments is at most 25%. The midterm and final project receive more weight, but their combined weight is at most 80%. Participation receives between 0 and 10%.

Course Assessment: Homework

Approximately three assignments will be posted on Canvas, combining proofs, algorithm analysis, and computational experiments. Most assignments will include implementation and testing of an optimization method. I expect you to upload your solutions in PDF to Canvas. For problems requiring a computational implementation, I will ask you to include your code in the PDF as well as a link to a public github repository where you will need to upload the code as well.

Course Assessment: Midterm

There will be one in-class midterm. The exact date will be announced in class and by email. There will not be a final exam.

Course Assessment: Final project

The final project consists of a written report and an in-person presentation. You will work in groups of three to six students, with each group choosing a topic broadly related to optimization for data science. Contributions may be theoretical, methodological, or focused on an application.

A central goal is to develop your research judgment: the ability to recognize interesting questions that are worth pursuing. Your group should read relevant papers, identify an open question, and work toward answering it over the quarter. You are not expected to produce a finished paper or a complete solution. A strong project may instead formulate a compelling question, explain its significance, and present meaningful partial results.

The project should demonstrate substantive work beyond what an LLM can produce in a single response. In particular, one important evaluation criterion will be whether the instructor can one-shot the project. That is, whether an LLM can solve your proposed question and produce results on par or superior to those in your report using one prompt.

You are encouraged to meet with the instructor periodically to get feedback concerning your project, and progress. Detailed project instructions will be provided early in the course.

Course Assessment: Participation

The grading scheme permits a participation weight between 0 and 10 percent; full course marks are possible without participation credit. Full participation credit can be earned through engagement during or after lectures, office-hour discussions, or thoughtful questions.

Collaboration, Sources, and AI Tools

Students are expected to follow the University's Academic Honesty and Plagiarism policy. Give credit for ideas, text, code, and other materials drawn from external sources. The work submitted must accurately represent your own contribution and understanding.

It is undeniable that we are in the middle of an AI revolution and the use of AI will be a fundamental part of our work in the years to come. That being said, you are expected to be able to think and operate with and without AI. With that philosophy in mind we will have the following policies for AI use.

Midterm

The in-class midterm will be a handwritten, closed-book exam. You may not consult any sources, including notes, books, or electronic resources. Phones and laptops must be put away for the duration of the exam.

Homework

You may consult any sources, including AI tools, when working on homework assignments. However, you are strongly encouraged to think through the problems independently. AI is best used as a learning partner that helps clarify points of confusion, rather than as a solution oracle.

All submitted solutions must be written from scratch, in your own words, and without assistance from an LLM during the writing process. Copying AI-generated solutions is not permitted. You are responsible for understanding every part of your submitted work and may be asked to explain and defend it in a one-on-one oral examination with the instructor. If you cannot demonstrate an understanding of your solutions, you may receive a zero on the assignment.

Final project

We encourage you to use AI in this project: a central goal is to develop your ability to use these tools effectively. Your written report should be clear, well organized, easy to follow, and present original ideas or results. You are responsible for understanding and being able to explain every aspect of your project. Your understanding will be assessed during the final presentation.

Deadlines and Extensions

Each student receives two one-day, no-questions-asked extensions that may be used for homework assignments or the final project. The extensions must apply to separate submissions; they cannot be combined into a two-day extension. Also, extensions between different students do not compound, the final project can only be submitted one day late.

If illness, an emergency, or other circumstances affect your ability to participate or complete work, contact the instructor as soon as reasonably possible so that we can discuss your situation.

Accessibility

Students seeking disability-related accommodations should contact Student Disability Services (SDS) and follow its procedures for requesting and using accommodations. Once accommodations are approved, please notify the instructor through the SDS process and discuss implementation early enough for arrangements to be made. SDS can be reached at disabilities@uchicago.edu or 773-702-6000.

Personal Well-being

UChicago Student Wellness provides information about medical services, mental health support, and health-promotion resources. If personal or health circumstances are affecting your coursework, please reach out to the instructor or your program's Dean of Students office for help identifying appropriate support.

Classroom Climate

We will work together to create a classroom in which students can ask questions, test ideas, and learn from one another. Treat classmates and their contributions with respect. Intellectual disagreement is welcome; harassment, discrimination, and personal attacks are not. If you have concerns about the classroom environment or barriers to participating, please speak with the instructor or a teaching assistant. Raising a concern will not affect your grade.

Family Accommodations

You are welcome to bring a family member to class occasionally when caregiving responsibilities require it, such as when emergency child care is unavailable. Please be considerate of the classroom environment; you may step out and return as needed.