

Optimization Foundations for Machine Learning and Artificial Intelligence

Course Information

- **Instructor:** Katya Scheinberg (katya.scheinberg@isye.gatech.edu)
- **Course Prefix and Number:** ISYE 4803
- **Term:** Fall 2026

Course Description

Continuous optimization models and algorithms are at the heart of many problems of Data Science, Machine Learning, AI, signal processing, economics and other modern contexts. This course introduces foundations of optimization models and algorithms as applied in these settings. Topics include unconstrained optimization, gradient descent and its variants, Newton method, convergence of algorithms, parameter tuning, line-search, stochastic gradient descent and its variants, empirical risk minimization, constrained optimization, duality, convexity, Lagrange duality. We will use examples from linear regression, Lasso regression, logistic regression, support vector machines and neural networks.

See the table for the tentative plan for the course topics. It is subject to change.

Course Learning Outcomes

The objective of this course is for the student to learn how to "look under the hood" by understanding how and when optimization algorithms work, how to apply a correct algorithm for a given problem, how to test and tune the algorithms. Students will learn theoretical concepts such as convergence rate of an algorithm then learn to compare theory to empirical performance.

At the end of this course the students are expected to be able to

- apply gradient and stochastic descent algorithms in correct settings with correct parameters, recognize if the algorithm attains desirable solution.
- recognize and model linear optimization problems and some other constrained optimization problems,
- solve them using software packages and recognize when the output is correct.
- understand capabilities and limitations of the algorithms that are covered in this course.

Required Course Materials

No textbooks or materials are required.

Week	Topic
1	Introduction, optimization problems, constrained and unconstrained problems, local global optima Convex and nonconvex functions, level sets
2	Hessian, gradients, linear algebra Optimality conditions, least squares
3	Gradient descent, examples, Line search Convergence of gradient descent
4	Newton and Quasi-Newton methods Accelerated gradient descent, coordinate descent
5	Binary classification and Logistic regression Large scale data sets, stochastic gradient descent
6	Convergence of stochastic gradient descent Extensions and Neural Networks
7	Review Midterm
8	Linear and Integer programs, relaxations LP examples in signal processing
9	Linear programming, graphical method Duality, intro
10	Duality properties Complementary slackness
11	Lagrangian, KKT conditions Convex QP, duality
12	Portfolio Optimization Support Vector Machines and Lasso regression
13	Other convex problems and their optimality Methods for constrained problems
14	Other topics of interest Other topics of interest Review

Lecture slides: We make skeletal slides available on Canvas before the lecture. They can be used to take notes during the lecture. Full notes from the lecture will be posted after the lecture.

Additional resources: If you would like to read additional material related to this course, material from most lectures can be found in

- J. Nocedal and S. Wright *Numerical Optimization*; (2006).
<https://link.springer.com/book/10.1007/978-0-387-40065-5>
- Beck A. "Introduction to Nonlinear Optimization: Theory, Algorithms, and Applications with Python and MATLAB" Second Edition,
<https://epubs.siam.org/doi/book/10.1137/1.9781611977622> (2023).

Grading Policy

Course grade will be based on homework assignments, in-class quizzes and participation, midterm and final exam. The final exam will be cumulative, covering all topics studied during the semester.

Homework	10%	Weekly
In-class work	40%	Weekly
Midterm	25%	TBD
Final Exam	25%	TBD
Total	100%	

The course grade will be calculated using the percentages above. The natural cut-offs of 90, 80, 70 and 60 will be used for letter grades A, B, C and D. Letter grades can be increased using a curve.

Attendance Policy

Students are strongly encouraged to attend lectures and actively participate by asking for clarifications and answering questions.

We will use regular quizzes to help students check their knowledge and understanding. Some of the quizzes will be based on completed homework and will be graded for correctness. Other quizzes will be graded on completion (not correctness) and count toward participation grade.

80% attendance or higher will count as full attendance (so lectures skipped due to illness or other uncontrolled circumstances will not be penalized). In a student misses a quiz, the grade for that quiz is zero, thus the bottom 20% of the quizzes (by grade) will be dropped.

Homework Assignments Homework assignments will be posted on Canvas and are usually due at **9PM on Mondays or Wednesdays (check carefully)**.

- **Submitting homework:** The homework should be submitted by deadline posted with the homework and has to be individual work of each student. For the coding exercises, you may work with a partner, and you and your partner may turn in a single assignment.
- **Late homework:** Each student is allowed to submit up to two homework assignments up to two days late each. This is meant to cover regular issues, such as illness, sports competitions, interviews and other interruptions. After solutions have been released no late homework will be accepted.
- **Dropped homework grade:** Each student will have the lowest homework grade dropped from the total. If a student is missing one homework, this homework will have the grade 0 and will be dropped.
- **Regrade requests:** the regrade request are accepted at least 24 hours and at most 7 days after the grades are released. Any regrade request may require us to regrade your entire homework/exam, not just the grading mistake. Be aware that if the person who graded misunderstood your answer during the first grading, it was probably not clear. Explaining what you meant afterwards is not likely to earn you points.

Academic and Research Honesty/Integrity Statement

Georgia Tech aims to cultivate a community based on trust, academic integrity, and honor. Students are expected to act according to the highest ethical standards. Review the [Student Code of Conduct](#) and the [Academic Honor Code](#), especially [Appendix A: Graduate Addendum to the Academic Honor Code](#).

Students are expected to perform research in an ethical and responsible manner. All Doctoral and Master's Thesis students are required to take the [Responsible Conduct of Research training](#), and it is expected that students abide by the principles taught in that training while performing research for this thesis course.

Allegations of scientific or scholarly misconduct are handled in accordance with the procedures outlined by the [Policy for Responding to Allegations of Scientific or Other Scholarly Misconduct](#).

Core IMPACTS

Not applicable.

Accommodations for Students with Disabilities

If you are a student with learning needs that require special accommodation, [contact the Office of Disability Services](#) as soon as possible to make an appointment to discuss your special needs and to obtain an accommodations letter. Please also e-mail me as soon as possible in order to set up a time to discuss your learning needs.

Expectations of Advisors and Advisees

At Georgia Tech, we believe that it is important to strive for an atmosphere of mutual respect, acknowledgment, and responsibility between faculty members and the student body. [The Expectations of Advisors and Advisees](#) articulates some basic expectations that you can have of me and that I have of you. In the end, simple respect for knowledge, hard work, and cordial interactions will help build the environment we seek. Therefore, I encourage you to remain committed to the ideals of Georgia Tech while in this class.