ECE 6270: Convex Optimization: Theory, Algorithms, and Applications Spring 2026 Syllabus

January 12, 2026

Summary

This course will cover the fundamentals of convex optimization. We will talk about mathematical fundamentals, modeling (i.e., how to set up optimization problems in different applications), and algorithms.

Prerequisites

Students should be familiar with linear algebra (e.g., solving systems of equations, least squares, matrix factorizations including SVD), basic probability (e.g., you should be comfortable with multivariate probability densities), and have basic Python programming skills.

Instructor

Mark Davenport
Office hours: Mondays, 11am-12pm
Location: Van Leer W208

Teaching Assistants

Ruizhe Wang
Office hours: TBD
Location: TBD. We will try to have a mix of face-to-face and online office hours.

Grading

- **Homeworks (30%):** There will be ≈ 8 homework assignments. See further details below.
- **Midterm quiz (30%):** This is tentatively scheduled for Wednesday March 4.
- **Final exam (30%):** This is scheduled (by the registrar) for Wednesday May 6, 8:00a-10:50a.
- **Class participation (10%):** Class attendance and Piazza; see below.

Your final grade will be assigned as a letter grade according to the scale:

A: 90-100% B: 80-89% C: 70-79% D: 60-69% F: $\leq 59\%$

I may exercise the option to "curve" exam scores as necessary (by adjusting the grades higher, but not lower) if I determine that an exam was more difficult than intended.

Lecture

Lectures are Mondays and Wednesdays from 9:30am-10:45am and will be held in Allen Sustainable Education 110.

Lecture attendance is mandatory and will count towards the participation portion of your grade (unless you are registered for the Q section). This might strike some of you as a controversial policy, but I believe it is ultimately in your best interest to attend lecture and this will give you an incentive to do so.

Attendance will be sampled through simple survey questions that appear once or twice per lecture.

Homework

Homework will be assigned weekly (approximately). Homework will be submitted online through Canvas/Gradescope (exact instructions will provided when the time comes).

Each homework assignment will be graded out of 100 points. Over the course of the semester, **the maximum number of homework points that you can earn is** $(N-1) \cdot 100$**, where** N **is the number of assignments** – this serves a similar role to allowing you to drop one homework assignment, but should encourage you to still submit a partially completed one (and avoid panic if there is occasionally a problem that you do not finish in time.)

The homework assignments will be hard; many of them will require significant amounts of time and effort to complete. But this is really where most of the learning takes place. You will get out of the assignments what you put into them. Students who complete all of the assignments in full will be rewarded with a deep understanding of the role that linear algebra and optimization play in data science, machine learning, robotics, and controls (among other things). Effectively, homework is worth much more than 30% of your grade. In teaching many courses over the years,

the instructors **have never seen a case where a student does not put effort into the homework assignments but does well on the exams**.

Students are *strongly* encouraged to discuss homework problems with one another. However, **each student must write up and turn in their own solutions written in their own words**. **Cases where solutions appear to be identical or nearly identical will be immediately referred to the Office of Student Integrity.**

Unauthorized use of any previous semester course materials, such as tests, quizzes, and homework, is prohibited in this course. Furthermore, redistributing materials from this semester is also prohibited. For any questions involving these or any other Academic Honor Code issues, please consult me or www.honor.gatech.edu.

Distance learning section

(This section is only relevant to those in the 'Q' section.)

This class does have a small distance learning section. We will follow the standard guidelines for delayed deadlines for homeworks and exams. Distance exams will be proctored using in-person proctoring. In the event that this cannot be arranged, a backup option is using HonorLock; please read about how this works (and the technical requirements) here:
https://tinyurl.com/yceme2ak
https://tinyurl.com/4dj2rp3z
Basically, you need a non-Linux computer, a webcam, and a microphone. Reach out to GTPE if you have additional questions.

Online resources

The course webpage is at:
https://mdav.ece.gatech.edu/ece-6270-spring2026.
This page will provide general course information, copies of the lecture notes, homework assignments, relevant papers, and other resources. Homework solutions and some additional resources will be posted in Canvas as necessary.

We plan to make exclusive use of Piazza to make announcements and answer questions:
https://piazza.com/gatech/spring2026/ece6270 .
Piazza is a great platform for you to work with your fellow students to discuss problems, form study/project groups, etc. **Please direct any questions you might have to Piazza.** Unless your question is personal in nature, please do not email me with questions — if you have a question you are probably not the only one, and other students may benefit from seeing the discussion. **Since your posts on Piazza are public, please do not simply post your solutions asking for confirmation that they are correct.**

Text

There is no required text. Extensive course notes will be provided that cover all of the required material in full. These will be posted as they become available at the course website. Textbooks that you might find useful (and from which a large portion of the notes were sourced) include:

- Boyd and Vanderberghe: *Convex Optimization*
  Available at http://amzn.to/2RBbH30,
  also available as a free pdf at http://web.stanford.edu/~boyd/cvxbook/
- Bertsekas, Nedic, and Ozdaglar: *Convex Analysis and Optimization*
  http://amzn.to/2C6cxek
- Nocedal and Wright: *Numerical Optimization*
  http://amzn.to/2VEpmp0
- Ben-Tal and Nemirovski: *Lectures on Modern Convex Optimization*
  http://amzn.to/2RDoKRx
- Luenberger: *Optimization by Vector Space Methods*
  http://amzn.to/2GZs0Cx

I may also provide some additional resources (e.g., papers and excerpts from other books) on the course website and/or using canvas as appropriate. We will also use the CVX Python package,
https://cvxopt.org
which makes it easy to prototype many of the optimization programs we will see this semester.

Course Objectives

Upon successful completion of this course, students should:

1. Be able to recognize and differentiate between common classes of optimization problems.
2. Have an understanding of how duality can be exploited to develop alternative approaches to solving an optimization problem.
3. Be able to implement and analyze the convergence properties of common iterative optimization algorithms.
4. Be able to translate practical engineering problems into optimization problems (modeling).

Course Expectations and Guidelines

Academic integrity

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. For information on Georgia Tech's Academic Honor Code, please visit www.catalog.gatech.edu/policies/honor-code. Any student suspected of cheating or plagiarizing on a quiz, exam, or assignment will be reported to the Office of Student Integrity, who will investigate the incident and identify the appropriate penalty for violations.

Redistributing materials from this course and/or using external sites for assistance (e.g., contributing to test banks, CourseHero, Chegg, or similar sites) is prohibited.

Collaboration and group work

Students are *strongly* encouraged to discuss homework problems with one another. However, **each student must write up and turn in their own solutions written in their own words**. **Cases where solutions appear to be identical or nearly identical will be immediately referred to the Office of Student Integrity.**

Absences, late assignments, and missed exams

Active participation in the class discussions is expected. Please attend class unless you have a compelling reason not to do so.)

We would like to be able to discuss the homework assignments in class the day after they are due, and thus **we cannot accept late homeworks** in the absence of prior approval. In the event that an excused absence prevents you from submitting an assignment, your homework grade will be calculated on a pro-rated basis. **Exams will be completed during specified time frames. If you expect to miss an exam, please contact me as soon as you realize this so we can make alternative arrangements.** We may consider options to take the exam at an alternate time or instead may adjust the grading allocation to place more emphasis on other exams, depending on the circumstances.

Accommodations for students with disabilities

If you are a student with learning needs that require special accommodation, contact the Office of Disability Services at (404)894-2563 or disabilityservices.gatech.edu, as soon as possible, to make an appointment to discuss your special needs and to obtain an accommodations letter. Please also e-mail us as soon as possible in order to set up a time to discuss your learning needs.

Student-Faculty expectations agreement

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. In the end, simple respect for knowledge, hard work, and cordial interactions will help build the environment we seek. Therefore, we encourage you to remain committed to the ideals of Georgia Tech while in this class. See www.catalog.gatech.edu/rules/22 for an articulation of some basic expectation that you can have of us and that we have of you.

Outline

The outline below should be treated as an approximation; it is subject to (hopefully small) changes.

1. Introduction to optimization, examples of convex optimization problems
2. Convexity
   (a) convex sets
   (b) convex functions
3. Unconstrained optimization
   (a) optimality conditions
   (b) line search methods for 1D problems
   (c) gradient descent, convergence analysis
   (d) accelerated first order methods
   (e) Newton's method, Quasi-Newton methods
   (f) non-smooth optimization: subgradients and proximal algorithms
4. Constrained optimization
   (a) geometric optimality conditions
   (b) Karush-Kuhn-Tucker (KKT) conditions
   (c) Lagrange duality, saddle points
   (d) algorithms: projected gradient descent
   (e) algorithms: log barrier (interior point) methods
   (f) algorithms: primal-dual alternating descent methods
5. Additional topics
   (a) stochastic gradient descent
   (b) distributed optimization
   (c) convex relaxations for non-convex problems
   (d) game theory and minimax strategies

Throughout the course, we will be using different applications to motivate the theory. These will cover some well-known (and not so well-known) problems in signal and image processing, communications, control, machine learning, and statistical estimation (among other things).