ECE 6270: Convex Optimization: Theory, Algorithms, and Applications

Fall 2026 Syllabus

Course Description

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

Justin Romberg
Office hours and location: TBD, will be announced on Piazza the first week of class.

Teaching Assistants

We will almost certainly have 1 or 2 TAs for this course. The TAs will grade the homework and will be available for weekly office hours. When they are assigned, they will introduce themselves on Piazza and announce their plan for office hours.

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 Materials

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/`
- Beck: *First-Order Methods in Optimization*
  Available at `https://epubs.siam.org/doi/book/10.1137/1.9781611974997`,
  also available for download through the GT Library.
- Luenberger: *Optimization by Vector Space Methods*
  `http://amzn.to/2GZs0Cx`
- 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`

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.

Grading

- **Homeworks (30%):** There will be ≈ 8 homework assignments. See further details below.

- **Midterm quiz (30%):** Date will be announced the first two weeks of class.
- **Final exam (30%):** This is scheduled (by the registrar).
- **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%

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

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 homework** in the absence of prior approval.

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.**

Attendance Policy

Lecture attendance is mandatory and will count towards the participation portion of your grade. 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).

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. It is amazing what you can learn from a peer trying to understand the same material as you. I would also **strongly encourage** you to work as hard as you can on problems before searching the internet or feeding questions into ChatGPT, Gemini, Claude, etc. The point of assigning hard homework questions is for you to learn by working through them. You learn almost nothing by simply copying from a colleague, a solution that has been posted online, or a chatbot.

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.

Specific policies for this class include:

1. **Each student must write up and turn in their own homework solutions written in their own words**. As discussed above, collaboration on homework assignments is encouraged, but ultimately we expect that what you turn in reflects your understanding. **Cases where solutions appear to be identical or nearly identical will be immediately referred to the Office of Student Integrity.**
2. **Unauthorized use of any previous semester course materials, such as tests, quizzes, and homework, is prohibited in this course.** This obviously includes previous semester course material that you find online.
3. **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.**

For any questions involving these or any other Academic Honor Code issues, please consult me or www.honor.gatech.edu.

Student-Faculty Expectations Agreement

At Georgia Tech we believe that it is important to strive for an atmosphere of mutual respect, acknowledgement, 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.

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.

Significant effort is being made throughout the semester for the course materials distributed for this class to be Title II compliant. If you have comments or suggestions about how to improve the accessibility of the lecture notes, please contact the instructor.

Distance Learning

(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 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.

Online Resources

The course webpage is at:
`https://sites.gatech.edu/ece6270-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. The link to the course Piazza site will be posted on Canvas the first week of class. 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 make a private post — if you have a question you are probably not the only one, and other students may benefit from seeing the discussion.

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) Fenchel duality: a geometric approach to duality
   (e) algorithms: projected gradient descent
   (f) algorithms: log barrier (interior point) methods
   (g) algorithms: primal-dual alternating descent methods
5. Additional topics
   (a) distributed optimization
   (b) convex relaxations for non-convex problems
   (c) stochastic gradient descent
   (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).