1 Course Description

This is a graduate-level course on optimization algorithms. The course covers continuous and discrete optimization through the lens of convex optimization. Convex optimization has had a profound impact on many areas of computer science and engineering, including algorithms and theoretical computer science, machine learning, data mining, finance, and operations research. The first part of the course covers the theory of convex optimization and its applications. The second part covers discrete optimization techniques that build on the machinery covered in the first part.

The main objectives of the course are:

- Develop basic concepts from convex analysis and optimization;
- Discuss optimization problems arising in diverse domains, such as mathematics, machine learning, statistics, operations research, engineering;
- Study classes of convex optimization algorithms along with their complexity analysis;
- Discuss general methods for large-scale convex optimization, such as gradient descent, accelerated gradient descent, and mirror descent;
- Discuss linear programming and duality;
- Develop general optimization methods for discrete problems based on mathematical programming formulations and continuous optimization;
- Discuss structured optimization and develop the capability of designing customized algorithms by exploiting problem structure.

2 Prerequisites

The class will require familiarity with fundamental algorithmic problems and techniques at the level of CS330/CS530 or equivalent, and a working knowledge of (multivariate) calculus and linear algebra. CS591 E2: Convex Optimization Algorithms is recommended for undergraduates and beginning graduate students; more advanced students with a strong mathematical background can take the course without having taken the Convex Optimization Algorithms course.

3 Textbooks

The class is self-contained and it does not follow a particular textbook. The following references are a very good resource.

- Chong and Zak. An Introduction to Optimization
- Boyd and Vanderberghe. Convex Optimization
- Bubeck. Convex Optimization: Algorithms and Complexity
- Blum, Hopcroft, Kannan. Foundations of Data Science
- Shalev-Shwartz. Online Learning and Online Convex Optimization
- Bach. Learning with Submodular Functions: A Convex Optimization Perspective

4 Communications

We will use Piazza for class discussion and questions. The system is highly catered to getting you answers to your questions fast and efficiently from classmates and the course staff. Please do not email questions to the course staff, post your questions on Piazza instead. We also encourage you to post answers to student questions there (but obviously, not answers to problems on the current homework).

5 Grading and Attendance

The course grade will break down as follows:

- **Homeworks:** 45% (the lowest 2 homework scores will be dropped)
- **Midterm exam:** 20% (only 15% for PhD students)
- **Final exam:** 30% (cumulative) (only 25% for PhD students)
- **Attendance and participation in lecture and Piazza:** 5%

PhD students will be asked to help with grading. This extra work will account for 10% of their grade.

Exams: There will be a midterm exam and a final exam. The midterm exam will be held in class in the middle of the semester (see lecture schedule on the course website). The cumulative final will be held during the normal two-hour final exam slot.

Homework Assignments, Submission, and Late Policy: Assignments will typically be due on Thursdays at 11:59pm (midnight). All assignments will be submitted electronically via Gradescope as a PDF. Assignments will involve both analytic (math) problems and programming in Python. You can either type your homework using LaTex (we will provide a template for each homework and there are some pointers on LaTex below) or scan your work. Plan on assignments being due every week except midterm week.

We will not accept late submissions and we will not grant extensions. To offset this policy, when computing your homework grade, we will automatically drop the lowest 2 homework scores. However, we strongly recommend putting your best effort in every homework, as they provide the best preparation for the exams. As you likely already know, assignments requiring substantial creativity can take more time than you expect, so plan to finish a day early.

Regrading Procedure: Please submit your regrade request on Gradescope within one week. Note that when we regrade a problem, your score may go up or down.

Attendance: It is expected that you will attend lecture. When students are at a borderline between grades, we will factor in attendance and participation before making a final determination.

6 Topics

1. Mathematical background and introduction to optimization:
   - Review of linear algebra and multivariate calculus
   - Convex functions and sets, basics of convex optimization
   - Examples of discrete and continuous optimization problems
   - Gradient descent algorithm
2. Continuous optimization:
   - More advanced gradient descent algorithms: accelerated gradient descent, mirror descent (if time permits: stochastic gradient descent, Newton's method)

- Online optimization and learning
- Constrained convex optimization: linear programming, cutting plane and interior point methods
- Duality for linear programming and convex optimization

3. Discrete optimization:
   - Graph optimization: maximum flow and minimum cut, maximum cut
   - Submodular function minimization
   - Submodular function maximization
   - Approximation algorithms for combinatorial optimization problems
4. Applications in algorithms and theoretical computer science, machine learning, statistics, data mining

7 Academic Conduct

Academic standards and the code of academic conduct are taken very seriously by our university, by the College of Arts and Sciences, and by the Department of Computer Science. Course participants must adhere to the CAS Academic Conduct Code; please take the time to review this document if you are unfamiliar with its contents.

8 Collaboration and Honesty Policy

Please read carefully the collaboration and honesty policy for this course. Please sign and submit a signed copy of the policy with your first homework.

9 LaTeX

We strongly recommend that you prepare your homework solutions using LaTeX, and submit PDFs to Gradescope. LaTeX is a scientific document preparation system; most CS technical publications are prepared using this tool. Great editors exist on most platforms, such as TexShop for Mac and TeXstudio for several platforms. An alternative to setting up LaTex on your machine is to use Overleaf.

The not so short introduction to Latex is a good reference to get you started.

10 Python

In this course we will be using Python for some of the homework exercises. Python is available on the BU computer systems. If you want to use it on your own computer you will need to install it. We recommend that you download a Python distribution such as Anaconda. Alternatively, you can download Python here (use Python 3.x and not 2.x).

Some basic tutorials include:

- The Python Wiki Guide
- The official Python tutorial