1 Course Description

This is a graduate-level course on the design and analysis of iterative algorithms for convex and non-convex optimization. Iterative methods based on gradient descent have 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 main topics covered are the design and analysis of gradient descent methods for convex problems; adaptive, stochastic and non-convex optimization; linear programming and duality; online learning; mirror descent. 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 and mirror descent;
- Discuss linear programming and duality;
- 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, linear algebra, and probability. It is expected that students have extensive experience with proof-based Mathematics and CS courses, and are proficient in writing rigorous proofs and algorithm analyses.

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
- Shalev-Shwartz, Ben-David. Understanding Machine Learning: From Theory to Algorithms
- Blum, Hopcroft, Kannan. Foundations of Data Science
- Hazan. Online Convex Optimization
- Shalev-Shwartz. Online Learning and Online Convex Optimization

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

The course grade will break down as follows:

- **Homeworks:** 50% (the lowest homework score will be dropped)
- **Midterm exam:** 20%
- **Final exam (cumulative):** 30%

Exams: There will be a midterm exam and a final exam. The midterm exam will be held in class, as indicated on the schedule. 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 every week on Friday at 10pm. Assignments will involve both analytic (math) problems and programming in Python. All assignments will be submitted electronically via Gradescope as a PDF for the math portion and as a Jupyter notebook for the programming portion. 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 the midterm exam week.

Each student receives two 24h extensions that can be used on any two homeworks. Note that submitting *any* portion of the homework late — either the math or programming portion — counts as the homework being submitted late. We will not accept homework submissions that are more than 24h late and we will not grant additional extensions. We will also automatically drop the lowest homework score when computing your homework grade. 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 and discussion. When students are at a borderline between grades, we will factor in attendance and participation before making a final determination.

6 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.

7 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 to Gradescope before submitting any work.*

8 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.

9 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