Convex optimization: Concepts, Algorithms and Analysis

CS292F Spring 2020

Instructor: Prof. Yu-Xiang Wang

Lectures: On Zoom. Monday and Wednesday 1:00 - 2:50

Piazza: `piazza.com/ucsb/winter2020/cs292f/home`
(We use Piazza for Q&A, discussion, and most annoucement!)

Gradescope: `https://www.gradescope.com/courses/109563`
(We will be using the gradescope only for homeworks and reading assignments.)

Gauchospace: The Gauchospace is active, but it is only reserved for communications that are restricted to people who are officially registered.

1 Overview

Nearly every machine learning problems can be written as an optimization problem with an objective function to optimize. A subset of these problems — those that come with a convex objective function and possibly a set of convex constraints — can be solved using standardized techniques in convex optimization. The associated properties, such as convexity, smoothness, sparsity, separability and so on are essential for understanding the computational properties of these problems and can be used to derive efficient iterative algorithms. The primary focus of the course will be on a class of highly practical algorithm that uses only black-box accesses to the gradient information of the objective. These are what we called — first order optimization algorithms. These algorithms are intimately connected to the field of online learning — a family of game-theoretic models in machine learning that assume very little about the data generating process and efficiently learn to predict on the fly. The course will explicitly discuss the connections between these methods and problem setups.

(New in 2020 Spring) While the bulk of the course will be about first order optimization, we will go beyond first-order methods and provide a functional understanding of the second-order methods as well.

The course is designed to help a graduate student to learn the theoretical foundation and practical algorithms of machine learning using tools from convex optimization. The techniques we learned are instrumental for understanding research papers in the field of machine learning and will be more generically applicable to problems outside machine learning that involve continuous optimization.

2 What you will learn?

1. You will be able to identify properties such as convexity, smoothness, sparsity and so on; reformulate problems into ones that have these properties;
2. You will learn how to analyze the convergence and computational complexity of first-order algorithms for each class of problems and choose appropriate algorithms for each problem of interest.

3. You will understand the meaning of regret and know the analytical tools for designing and analyzing regret-minimizing online learning algorithms.

4. You will gain some functional knowledge about second order methods such as Newton's method, interior point methods and understand their computational properties.

3 Prerequisites

As this is a graduate level course, there is no hard pre-requisite, but students entering the class are expected of the following.

- Students are expected to be able to follow rigorous mathematical arguments and perform derivation and proofs using notations, definitions and theorems from calculus, linear algebra, probability and statistics.
- Students are expected to have working knowledge of basic algorithms and data structures, as well as the necessary programming experience (Your choice of Python/numpy, Matlab, or Julia).
- Basic familiarity in machine learning problems are highly recommended as most of our examples will be coming from machine learning.

4 Textbooks

We will be using two textbooks.

1. "Convex optimization" by Boyd and Vandenberghe: [Link]
2. "Introduction to Online Convex Optimization" by Elad Hazan: [Link]
3. (Optional) "Convex Analysis" by Tyrell Rockefellar: [Link]
4. (Optional) "Prediction, Learning, and Games" by Cesa-Bianchi and Lugosi: [Link]

The other two optional reference books are excellent references on these topics as well. We will not be assigning them as reading materials.

On the other hands, the scope of this course is not limited to the textbooks, we will heavily refer to lecture notes from the Stanford Course EE364B (Boyd) and CMU Course 10-725 (Tibshirani), as well as relevant research papers. These will be assigned as reading materials.

5 Assignments and Grades

The grades will be based on homework assignments, the quality of your submitted reading notes and lecture attendance.

- 80% Homework assignments.
- 15% Reading notes (Due at the beginning of every lecture, from the second lecture onwards, submit through Gradescope.)
- 5% Participation
- Bonus 5% Volunteer to scribe the lectures.

What are "Reading notes"? We will be assigning reading materials for every lecture. You are required to read the book chapters and the associated papers *before each lecture* and write a summary note (at least one page long). You are required to submit the reading notes to Gradescope before the start of every lecture.

6 Logistics

Zoom lectures: Only registered students are given access to the real-time lectures over Zoom. You have been emailed instructions on how to connect to the lecture. If you are having trouble, please contact the instructor or the TA.

The live Zoom sessions will be recorded for students who may not be able to attend synchronous. By default, your microphone and camera will be muted when you join the session. If you do not want to be included in the recording, simply keep your camera and microphone off. You may ask questions in the chat window. Please refer to the "copyright" section for more details.

Zoom office hours: Office hours will be conducted in a similar manner to the lectures. Details will be announced via email.

Attendance policy: The attendance to lectures is required. It is part of the course evaluation to attend the lectures. Send PM on Piazza if you will have to miss lectures due to other personal businesses.

Late homework policy: Each student will have exactly **four late days** without penalty in total (for instance, you may use 3 for HW1 and then your are left with 1 for HW4). If you are out of your late-days, the homeworks will count for at most 50% of their value. There will be no exception to this rule.

7 Policy on Academic Integrity

Please read this section carefully.

The university, the department, and this instructor all take the issue of academic integrity **very** seriously. A university requires an atmosphere of mutual trust and respect. While collaboration is an integral part of many scholarly activities, it is not always appropriate in a course, and it is never appropriate unless due credit is given to all participants in the collaboration. This goes for both ideas and programming or other work.

Here are some examples:

- Allowed: Discussion of lecture and textbook materials
- Allowed: Discussion of how to approach assignments, what techniques to consider, what textbook or lecture material is relevant
- Allowed: Collaboration on homeworks. You need to declare your collaborators, and each student still needs to write their own report / code independently.
- Allowed: Refer to materials
- Not allowed: Turning in someone else's work as your own, even with that person's permission.
- Not allowed: Allowing someone else to turn in your work as his or her own.

- Not allowed: Turning in work without proper acknowledgment of the sources of the content (including ideas) contained within the work.

For some views on academic integrity at UCSB see the Academic Integrity page of the Office of Judicial Affairs.

Summary: Academic integrity is absolutely required - dishonesty (cheating, plagiarism, etc.) benefits no one and hurts everyone. Violations of these honor codes on academic integrity will be reported to the Office of Student Conduct. If you find yourself in such a situation, please contact the instructor or the TA, and the penalty will be lighter. If you are not sure whether or not something is appropriate, please ask the instructor or TA.

8 Code of conduct

The University of California, Santa Barbara has a general code of conduct for all students published here: `http://www.sa.ucsb.edu/docs/default-source/student-conduct/conductofcode2017.pdf?sfvrsn=d3c07d4f_2`

In addition to the UCSB's code of conduct, the Computer Science Department holds students, staff, and faculty to the following standards:

- Treat all members of the academic community (students, staff, and faculty) with respect regardless of their experiences and background, including (but not limited to) their cultural backgrounds, socioeconomic status, disabilities, age, religion, sexual orientation, neuro(a)typicality, and gender identity.
- Physical or mental harm, sexual harassment, aggression, and derogatory language is not acceptable in any form.
- Respect the personal property of others and University resources. Unauthorized access, use, vandalism, or theft of equipment, computer servers, labs / offices / classrooms, etc. is prohibited.
- The exchange and challenge of ideas are done in a thoughtful, respectful and constructive manner.
- Disruption of departmental activities such as special events, talks, lectures, and meetings is not acceptable.
- Any violation of the given standards should be reported to the Computer Science chair and/or the CS diversity committee (`diversity@cs.ucsb.edu`). Consequences may include a formal warning, suspension, or expulsion from the University.

9 Students with Disabilities

If you are a student with a disability and would like to discuss special academic accommodations, please contact the instructor. In addition, students with temporary or permanent disabilities are referred to the Disabled Students Program (DSP) at UCSB. DSP will arrange for special services when appropriate (e.g., facilitation of access, note takers, readers, sign language interpreters). Please note that it is the student's responsibility to communicate his or her special needs to the instructor, along with a letter of verification from DSP.

10 Copyright of course materials

My lectures and course materials, including PowerPoint presentations, tests, outlines, and similar materials, are protected by U.S. copyright law and by University policy. I am the exclusive owner of the copyright in those materials I create. You may take notes and make copies of course materials for your own use. You may also share those materials with another student who is enrolled in or auditing this course. You may not reproduce, distribute or display (post/upload) lecture notes or recordings or course materials in any other way — whether or not a fee is charged — without my express prior written consent. You also may not allow others to do so.

If you do so, you may be subject to student conduct proceedings under the UC Santa Barbara Student Code of Conduct.

Similarly, you own the copyright in your original papers and exam essays. If I am interested in posting your answers or papers on the course web site, I will ask for your written permission.

The live Zoom session will be recorded for students who may not be able to attend at this time. By default, your microphone and camera will be muted when you join the session. If you do not want to be included in the recording, simply keep your camera and microphone off. You may ask questions in the chat window.