Convex Optimization
10-725/36-725

Instructors:
Pradeep Ravikumar, pradeepr@cs.cmu.edu
Aarti Singh, aarti@cs.cmu.edu

Teaching Assistants:
Hao Gu, hg1@andrew.cmu.edu
Devendra Sachan, dsachan@andrew.cmu.edu
Yifeng Tao, yifengt@andrew.cmu.edu
Yichong Xu, yichongx@andrew.cmu.edu
Hongyang Zhang, hongyanz@andrew.cmu.edu

Lectures: Wean 7500, Mondays and Wednesdays, 1:30 - 2:50 PM

Office Hours:

Pradeep Ravikumar: GHC 8111, Mondays 3:00 – 4:00 PM
Aarti Singh: GHC 8207, Wednesdays 3:00 – 4:00 PM

Hao Gu: TBD
Devendra Sachan: TBD
Yifeng Tao: TBD
Yichong Xu: TBD
Hongyang Zhang: TBD

Course Description:
Nearly every problem in machine learning can be formulated as the optimization of some function, possibly under some set of constraints. This universal reduction may seem to suggest that such optimization tasks are intractable. Fortunately, many real world problems have special structure, such as convexity, smoothness, separability, etc., which allow us to formulate optimization problems that can often be solved efficiently. This course is designed to give a graduate-level student a thorough grounding in the formulation of optimization problems that exploit such structure, and in efficient solution methods for these problems. The main focus is on the formulation and solution of convex optimization problems, though we will discuss some recent advances in guarantees for non-convex optimization. These general concepts will also be illustrated through applications in machine learning and statistics. Students will work on an extensive optimization-based project throughout the semester.

Learning Objectives:
Upon completing the course, students should be able to approach an optimization problem, often derived from a statistics or machine learning context, and:

1. identify key properties such as convexity, smoothness, sparsity, etc., and/or possibly reformulate the problem so that it possesses such desirable properties;
2. select an algorithm for this optimization problem, with an understanding of the advantages and disadvantages of applying one method over another, given the problem and properties at hand;
3. implement this algorithm or use existing software to efficiently compute the solution.

Pre-requisites:
Students entering the class should have a pre-existing working knowledge of algorithms, though the class has been designed to allow students with a strong numerate background to catch up and fully participate. Though not required, having taken 10-701 or an equivalent machine learning or statistics class is strongly encouraged, since we will use applications in machine learning and statistics to demonstrate the concepts we cover in class.

Outline of material:

- Basics of Convex Analysis
- Unconstrained Optimization:
  - Descent Algorithms: Direction, Step-size selection
  - Line Search, Trust Region Methods
  - First Order, Second Order Methods
- Constrained Optimization:
  - Approaches for simple convex sets
  - Lagrange Multiplier Theory
  - Lagrange Methods: Barrier, Augmented Lagrangian, Interior Point Methods
  - Duality Theory
  - Dual Methods: Dual Ascent, Cutting Plane Methods, Decomposition Methods
- Advanced Topics:
  - Stochastic Optimization
  - Non-smooth Optimization
  - Non-convex Optimization

Tentative Course Schedule: Available on website.

Class Website:
http://www.cs.cmu.edu/~pradeepr/convexopt
http://www.cs.cmu.edu/~aarti/Class/10725_Fall17
The class schedule, logistics, and lecture materials will be posted there.

Discussion, Announcements:
We will use Piazza for announcements, as well as the discussion board for the class.

Textbooks:
Lectures are intended to be self-contained. For supplementary readings, in each lecture, we will have pointers to chapters from the following books:

- Convex Optimization, Stephen Boyd and Lieven Vandenberghe
  (available online for free: http://www.stanford.edu/~boyd/cvxbook/).
- Nonlinear Programming, Dimitri P. Bertsekas
- Numerical Optimization, Jorge Nocedal and Stephen Wright
- Introductory lectures on convex optimization: a basic course, Yurii Nesterov

Quizzes:
There will be a short, easy quiz due at midnight on the day of each lecture, consisting of multiple choice or true/false questions. The quizzes will be taken online, and the links will be given on the course website.

Homeworks:
There will be 5 homework assignments, approximately evenly spaced throughout the semester. The assignments will be posted on the course website, and on Piazza. We will use Gradescope for submitting, and grading assignments.

You will get a late day quota of 3 days, which you can distribute among the five homeworks as you wish. Homeworks submitted after your late day quota will not be accepted. We expect you to use the late day quota for conference deadlines and events of the like, so we cannot provide an additional extension for such cases. In the case of an emergency (sudden sickness, family problems, etc.), we can give you a reasonable extension. But we emphasize that this is reserved for true emergencies.

The homeworks are structured to give you experience in both written mathematical exercises and programming exercises. While it is completely acceptable for you to collaborate with other students in order to solve the problems, we assume that you will be taking full responsibility in terms of writing up your own solutions and implementing your own code. You must indicate on each homework the students with whom you collaborated.

Little tests:
There will be two little tests, scheduled to be about halfway through and at the end of the semester. Precise dates are on the course website. These will consist of multiple choice and true/false questions, as well as short-answer questions.

Class project:
There will be a class project. You can form groups of up to three students. Further details on the project will be available on the website.

Grading:

Homework	45%
Quizzes	5%
Little tests	25%
Project	25%

Accommodations for Students with Disabilities:
If you have a disability and have an accommodations letter from the Disability Resources office, we encourage you to discuss your accommodations and needs with the instructors as early in the semester as possible. We will work with you to ensure that accommodations are provided as appropriate. If you suspect that you may have a disability and would benefit from accommodations but are not yet registered with the Office of Disability Resources, we encourage you to contact them at access@andrew.cmu.edu.

Take care of yourself:
Take care of yourself. Do your best to maintain a healthy lifestyle this semester by eating well, exercising, avoiding drugs and alcohol, getting enough sleep and taking some time to relax. This will help you achieve your goals and cope with stress.

All of us benefit from support during times of struggle. There are many helpful resources available on campus and an important part of the college experience is learning how to ask for help. Asking for support sooner rather than later is almost always helpful.

If you or anyone you know experiences any academic stress, difficult life events, or feelings like anxiety or depression, we strongly encourage you to seek support. Counseling and Psychological Services (CaPS) is here to help: call 412-268-2922 and visit their website athttp://www.cmu.edu/counseling/. Consider reaching out to a friend, faculty or family member you trust for help getting connected to the support that can help.

If you or someone you know is feeling suicidal or in danger of self-harm, call someone immediately, day or night:

- CaPS: 412-268-2922
- Re:solve Crisis Network: 888-796-8226

If the situation is life threatening, call the police

- On campus: CMU Police: 412-268-2323
- Off campus: 911

If you have questions about this or your coursework, please let me know. Thank you, and have a great semester.