Convex Optimization: Fall 2018

Machine Learning 10-725

Instructor: Ryan Tibshirani, ryantibs@cmu.edu

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

Overview and objectives

Nearly every problem in machine learning and computational statistics can be formulated in terms of the optimization of some function, possibly under some set of constraints. As we obviously cannot solve every problem in machine learning, this means that we cannot generically solve every optimization problem (at least not efficiently). Fortunately, many problems of interest in machine learning can be posed as optimization tasks that have special properties—such as convexity, smoothness, sparsity, separability, etc.—permitting standardized, efficient solution techniques.

This course is designed to give a graduate-level student a thorough grounding in these properties and their role in optimization, and a broad comprehension of algorithms tailored to exploit such properties. The focus will be on convex optimization problems (though we also may touch upon nonconvex optimization problems at some points). We will visit and revisit important applications in machine learning and statistics. Upon completing the course, students should be able to approach an optimization problem (often derived from a machine learning or statistics 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.

Outline of material

Here is an rough outline of the course topics.

Core topics: Split into 4 parts, as follows.

- Theory I: Basics of convex analysis

- Algorithms I: First-order methods
- Theory II: Duality and the KKT conditions
- Algorithms II: Second-order methods

Advanced topics: We will choose a handful of advanced topics to cover, depending on current research trends and class interests.

Logistics

Prerequisities: Students entering the class are expected to have a pre-existing working knowledge of basic algorithms and data structures; however, the class has been designed to allow students that do not have this background, but have strong analytical/mathematical skills, to catch up and fully participate. Students are expected to be comfortable with rigorous mathematical arguments. Though not formally required, having taken 10-701 or an equivalent machine learning or statistics class will be very helpful, since we will frequently use applications in machine learning and statistics to demonstrate the concepts we learn in class. Finally, several homework assignments will involve programmatic implementation of an optimization algorithm being studied in the course.

Class website: The class website is `http://www.stat.cmu.edu/~ryantibs/convexopt/`. The class schedule, lecture notes, homeworks, etc., will be posted there.

Office hours: The weekly schedule for office hours can be found on the course website.

Discussion group: We will use Piazza for class discussions.

Textbook: Lectures are intended to be self-contained. But a terrific reference for a lot of the material covered in class is *Convex Optimization* by Boyd and Vandenberghe, which is available online for free (`http://www.stanford.edu/~boyd/cvxbook/`). We will often draw examples or problems from this book. In addition, a thorough mathematical reference for much of the underlying convex analysis we cover is *Convex Analysis* by Rockafellar. Other relevant references will be listed at the end of each set of lecture slides.

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.

Homework: There will be 5 homework assignments, approximately evenly spaced throughout the semester. The assignments will be posted on the course website, and your homeworks will be submitted according to the instructions given there.

Late homework will not be accepted. In the case of a conference deadline or something of the like, you must give us at least **5 days notice** if you are requesting an extension, and we can give you +2 days. In the case of an emergency (sudden sickness,

family problems, etc.), no notice is needed, and 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. As we will often reuse problems from previous versions of the course, or problems covered by papers and webpages, you **must not to copy, refer to, or look at** previous solutions in preparing your answers.

Also, 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 consistent of multiple choice and true/false questions.

Class project: There will be a class project, with two milestones spaced about one-third and two-thirds of the way through the semester. Details on the course website.

Scribing: Students can sign up for scribing lectures, with no more than three students signed up for a particular lecture. Scribed notes are due **one week after the date of the corresponding lecture**. While scribing is not strictly required, it can bump up your grade if you are on the border (e.g., B+ to an A-). The sign up sheet can be found on the course website.

Evaluation: The grading breakdown is as follows.

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

Gradebook: We will use Canvas for recording and keeping track of grades throughout the class.

Accommodations for students with disabilities

If you have a disability and are registered with the Office of Disability Resources, please use their online system to notify us of your accommodations and discuss with us your needs 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, consider contacting 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. You are not alone. There are many helpful resources available on campus. Asking for support sooner rather than later is often 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 at http://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 us know.