Statistics meets Optimization

36-741: Randomized iterative methods

1 Basic course information

Instructor:	Yuting Wei; `ytwei AT cmu DOT edu`
Lecture:	MW 1:30-2:50pm; Porter Hall, A19 Aug 26 to Oct 14, 6 units
Office hours:	Yuting Wei, by appointment, Location: BH 229J
Announcement & materials:	will be posted on *Canvas*

Synopsis: This is a Ph.D. mini class series in statistics. It contains two separate parts with course number 36-741 and 36-742 respectively. We will cover a selection of topics in modern statistics and optimization, with a focus on the analysis of iterative sketching algorithms and approximate message passing algorithm applied to statistics problems. One of the main goals is to provide students with some background to understand the ongoing statistics/optimization literature and provide motivated students with some mathematical tools for working on related research.

Prerequisites: This course is intended for Ph.D. students with strong mathematical background. There are no formal prerequisites for these minis, but students are expected to have completed at least one intermediate statistics course (like 10-705 at CMU), and preferably an advanced statistics course and one course on optimization. Students are suggested to be familiar with

1. basics for statistics estimation (loss function, least square estimator, efficiency)
2. basic concentration inequalities (such as Markov, Chebyshev, Hoeffding)
3. basics for optimization (convexity, convergence, line search, gradient methods)

2 Course descriptions

36-741: Randomized sketching methods

Descriptions: In this mini, we will discuss some aspects of the interface between statistics and optimization. The goal of these lectures is to touch on various evolving areas at this interface. The objectives of optimization can be influenced by underlying statistical objectives in many ways, for example, the statistics precision caused by not having enough sample size is often of higher order than the machine precision; worst-case instance can be too conservative compared to the random ensembles; polynomial-time complexity may still be too large to be tractable. To further discuss these issues, we will start with a dimension reduction technique based on random projections and analyze how this technique helps us achieve faster optimization convergence without hurting statistical precisions.

Tentative list of topics, some might be assigned as reading materials:

- Dimension reduction via random projection
- Random projection for solving optimization problems
- Information theoretic lower bounds
- Sketched Newton's method
- Gradient methods for high-dimensional statistical recovery
- Low-rank matrix estimation: return to non-convexity
- Matrix completion and sparse PCA (with applications)

Textbook: There is no textbook for this class. Here are some references and additional papers will be uploaded as the class proceeds.

References:

- *High-dimensional statistics: A non-asymptotic viewpoint*, Cambridge University Press, 2019, by M. Wainwright.
- *Lectures on modern convex optimization: analysis, algorithms, and engineering applications*, volume 2. Siam, 2001, by A. Ben-Tal and A. Nemirovski.
- *Lecture notes on randomized linear algebra*, 2016, by M. W. Mahoney.

3 Graded components

There will be one long homework (40%) whose questions will be progressively released. Homework for mini 1 is due on Oct 2. No late homework will be accepted. A hard copy of your homework must be turned in. Please use Latex to typeset your homework. It is okay to collaborate on the homework but you must hand in your own copy.

Another component of the grades will be a course project (10% for proposal, 50% for final report) for each of the minis. This project can either be a literature review or include original research:

- Literature review. We will provide a list of related papers that are not covered in the lectures, and the literature review should involve in-depth summaries and exposition of one of these papers.
- Original research. It can be either theoretical or experimental, with the approval from the instructor. If you choose this option, you can do it either individually or in groups of two.

Three timestamps for the course project:

- Proposal (Sep 23). Submit a short report stating the papers you plan to survey or the research problems that you plan to work on. Describe why they are important or interesting, and provide some appropriate references. If you elect to do original research, you are encouraged to connect this project with your current research (but is still related to our course content). Please do not propose an overly ambitious project. You will receive feedback from the instructor.

- In-class presentation (if applicable).

- A written report (Oct 18). You are expected to submit a final project report, up to 5 pages long (not including references), summarizing your findings/contributions.