ISE 633 Large Scale Optimization for Machine Learning

Number of units: 03

Location and time: Monday/Wednesday 3:30-4:50pm, KAP 148 and also online at: https://usc.zoom.us/j/94648171113?pwd=bUZ5WkhkQWMwNGFYUS9FOFZaVFF3Zz09

Instructor: Meisam Razaviyayn
Email: razaviya@usc.edu
Office hours: 5-6PM Wednesdays.

Teaching Assistant: Sina Baharlouei
Email: baharlou@usc.edu
Office hours: 12-1PM Tuesdays.

Office Hours Zoom info:
https://usc.zoom.us/j/91010596889?pwd=NkovTkNmWk5acU9qWTA4NVlkbC9SZz09

Goal: The objective of the course is to introduce large scale optimization algorithms that arise in modern data science and machine learning applications.

Course Description: The course covers the theory and tools for large-scale optimization that arise in modern data science and machine learning applications. We will cover topics such as stochastic optimization, accelerated methods, parallelization, nonsmooth optimization, online optimization, variance reduction, differential privacy in optimization, min-max games, etc.

Textbook: There is no required textbook for the class. All course materials will be presented in class or will be available online as notes. The following textbooks cover parts of the course materials and you may find them useful:

- S. Boyd and L. Vandenberghe, Convex optimization, Cambridge university press, 2004.
  - The book is available for free here: http://web.stanford.edu/~boyd/cvxbook
- S. Shalev-Shwartz and S. Ben-David, Understanding Machine Learning: From Theory to Algorithms. Cambridge University Press, 2014.
  - The book is available for free here: http://www.cs.huji.ac.il/~shais/UnderstandingMachineLearning/
- A. Shapiro, D. Darinka, and A. Ruszczynski, Lectures on Stochastic Programming: Modeling and Theory, SIAM, 2009.
  - The book is available for free here: http://www2.isye.gatech.edu/~ashapiro/publications.html

Tentative Course Plan: (Course materials may change depending on the progress)

- Basics and preliminaries (Week 1)
  - Optimization overview, examples in machine learning, large-scale optimization, memory/time/CPU requirement
  - Mathematical Basics

- Unconstrained optimization: optimality conditions and algorithms (Weeks 2-3)
  - Necessary and sufficient optimality conditions
  - Convex versus non-convex and their examples
  - First-order methods (unconstrained)
  - Convergence analysis: Lower and upper-iteration complexity bounds
  - Momentum and accelerated methods

- Constrained optimization (Weeks 4-6)
  - Examples of constrained optimization in machine learning: fairness, safety, etc.
  - KKT optimality conditions and Lagrange multipliers
  - Projection-based algorithms, examples in machine learning

- Nonsmooth optimization (Weeks 7-8)
  - Examples of nonsmooth optimization in machine learning: Lasso, Netflix problem and Nuclear norm regularizers, nonsmooth activations in deep learning, etc.
  - Necessary and sufficient optimality conditions
  - Algorithms and their analysis: block coordinate descent, successive upper-bound minimization, alternating direction method of multipliers, subgradient descent, proximal gradient

- Empirical risk minimization and stochastic optimization (Weeks 9-11)
  - Stochastic/Online/Incremental optimization, SAA and SA
  - Bias/variance tradeoffs
  - Algorithms and their analysis: (robust) stochastic (sub-)gradient descent, stochastic variance reduced gradient descent
  - Parallel optimization: synchronous vs asynchronous

- Online convex optimization (Weeks 11-12)
  - Examples and regret definition
  - Algorithms and their analysis: follow-the-leader, follow-the-regularized-leader, online gradient descent

- Learning via min-max optimization (Weeks 13-14)
  - Examples: Generative Adversarial Networks, Adversarial Learning
  - Algorithms: Danskin's theorem, gradient descent-ascent algorithm, proximal-type methods

- Differential privacy and optimization (Weeks 14-15) – If we have time
  - Differential privacy definition and examples
  - Lower and upper-bounds on differentially private convex optimization

Course Requirement and Grading:

- During class time midterm on October 13 (30%)
- Final exam, Monday December 13, 2pm-4pm (35%)
- Homework assignments (20%)
- Participation (5%)
- Scribing (10%)

Homework assignments:

- All homework assignments are **due by 11:59pm** on the date indicated.
- Homework assignments must be submitted via Blackboard. **Only one pdf file** should be submitted for each homework assignment. You can submit latex pdf files, word converted pdfs, or scanned images which are converted to pdf format.
- Late homework submissions are not accepted **under any circumstances**. Start your homework assignments early.
- There will be 5-6 homework assignments. The two lowest scores will not be considered in your final grade.
- You are encouraged to discuss homework assignments with other students. However, each student is required to submit his/her own personal work.

Scribing: In order to gain experience with technical writing, each student is required to scribe notes for one lecture. These notes will be revised by the instructor and will be posted on the course website. The scribed notes should be written in a way that they are completely understandable to a student who may have missed the class. The **deadline for your scribing assignment is one week after the lecture** you are scribing.

Class Participation: Class participation is 5% of your entire grade. This grade will be based on the following two criteria:

- Being present in the lectures and actively participate in the discussions in class
- Answering the questions asked by other students in the class **Slack channel** and participate in the discussions on Slack

Scribing: In order to gain experience with technical writing, each student is required to scribe notes for two lectures. The scribed notes should be written *in a way that they are completely understandable to a student who may have missed the class*. The link to the **signup sheet** and the **overleaf document** is posted on Blackboard. For your scribing assignments, pay attention to the followings:

- The deadline for the assignment is 7 days after the lecture that you are scribing.
- The previous years' scribes is shared with you on Overleaf. Do not think of the previous years' scribes as perfect and complete. The topics are even changed in some parts of the course. Edit the document as you see fit. You should directly edit on the overleaf file. Your grade will be based on the improvement you made compared to last years' notes.

Where should you ask your questions?

- If you have a question regarding homework assignments or other parts of the course that you think other students, the TA, or the instructor can answer it, please ask the question in the course Slack channel. This will increase the interactions among all of us during this pandemic time.
- If you have questions regarding your homework assignment's grades, please email the TA.
- Email the instructor for other inquiries not listed above.

University policies:

- *Statement for Students with Disabilities.* Any student requesting academic accommodations based on a disability is required to register with Disability Services and Programs (DSP) each semester. A letter of verification for approved accommodations can be obtained from DSP. Please be sure the letter is delivered to your course instructor (or TA) as early in the semester as possible. DSP is located in STU 301 and is open from 8:30am to 5:00pm, Monday through Friday. Website and contact information for DSP: http://sait.usc.edu/academicsupport/centerprograms/dsp/home_index.html, (213) 740 – 0776n (Phone), (213) 740-6948 (TDD only), (213) 740-8216 (FAX), ability@usc.edu.

- *Statement on Academic Integrity.* USC seeks to maintain an optimal learning environment. General Principles of academic honesty include the concept of respect for the intellectual property of others, the expectation that individual work will be submitted unless otherwise allowed by an instructor, and the obligations both to protect ones own academic work from misuse by others as well as to avoid using another's work as ones own. All students are expected to understand and abide by these principles. SCampus, The Student Guidebook, contains the Student Conduct Code in Section 11.00, while the recommended sanctions are located in Appendix A: http://usc.edu/dept/publications/SCAMPUS/gov/. Students will be referred to the Office of Student Judicial Affairs and Community Standards for further review should there be any suspicion of academic dishonesty. The Review process can be found at: http://usc.edu/student-affaris/SJACS/. Information on intellectual property at USC is available at: http://usc.edu/academe/acsen/issues/ipr/index.html.

- *Emergency Preparedness/Course Continuity in a Crisis*. In case of emergency, when travel to campus is difficult, if not impossible, USC executive leadership will announce a digital way for instructors to teach students in their residence halls or homes using a combination of the Blackboard LMS (Learning Management System), teleconferencing, and other technologies. Instructors should be prepared to assign students a ``Plan B" project that can be completed ``at a distance". For additional information about maintaining your classes in an emergency, please access: http://cst.usc.edu/services/emergencyprep.html.