Large Scale Matrix Computation and Machine Learning
CSCI 6971/4971
Mondays and Thursdays, 2pm-3:50pm, Lally 102

Alex Gittens
Office Hours: Mon, Thurs 4pm-5pm or by appointment (use the link on my website)
Lally 316
gittea@rpi.edu
`http://www.cs.rpi.edu/~gittea/`

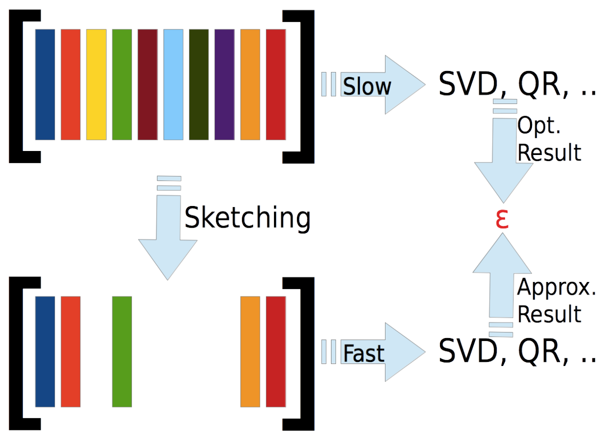

Course Description: Modern machine learning routinely deals with millions of points in high-dimensional spaces. Classical linear algebra and optimization algorithms can be prohibitively costly in such applications, as they aim at machine precision and/or scale super-linearly in the size of the input data. Randomization can be used to bring the costs of machine learning algorithms closer to linear in the size of the input data; this is done by sacrificing, in a principled manner, computational accuracy for increased speed. This course surveys modern randomized algorithms and their applications to machine learning, with the goal of providing a solid foundation for the use of randomization in large-scale machine learning. Topics covered will include time-accuracy tradeoffs, stochastic first-order and second-order methods, applications of low-rank approximation, approximate kernel learning, distributed optimization, and hyperparameter optimization.

Course Website: `http://www.cs.rpi.edu/~gittea/teaching/spring2019/csci6971-and-4971.html` Note that this document supercedes any statements made on the website.
Course Text(s): None required. As necessary, links to papers and other reading materials will be provided.

Course Objectives: At the end of the semester, students will understand how randomization allows us to efficiently tackle large-scale machine learning problems. To that end, they will be familiar with stochastic first- and second-order optimization techniques, the randomized numerical linear algebra primitives used in optimization, and large-scale kernel methods based on low-rank matrix approximations. They will be familiar with the analytical tools used in guaranteeing the performance of these algorithms, and be able to identify the tradeoffs involved in using randomization for machine learning.

Grading Criteria:

Homeworks	50%
In-class Pop Quizzes	20%
Project	30%

Students are expected to have writing supplies on hand in each class to complete the in-class pop quizzes. If you are an athlete or for some other reason are not able to attend each class, make alternative arrangements with the instructor in the first two weeks of the course.

Letter grades will be computed from the semester average. Maximum lower bound cutoffs for A, B, C and D grades are 90%, 80%, 70%, and 60%, respectively. These bounds may be moved lower at the instructors discretion.

Homework Policy

All assignments must be typeset and are due, via email, at the start of class (the first 15 minutes). Late homeworks will be penalized and accepted at the discretion of the instructor. There will be no makeup homeworks.

Project Policy

A presentation project will be assigned during the second half of the semester. It will entail reading a paper, summarizing the guarantees established in the paper, implementing the main algorithm and appropriately empirically evaluating it, then delivering a 15 minute presentation before the class.

Academic Integrity

The Rensselaer Handbook of Student Rights and Responsibilities and The Graduate Student Supplement define various forms of Academic Dishonesty and you should make yourself familiar with these. In this course, all assignments that are turned in for a grade must represent the students own work. In cases where help was received, or teamwork was allowed, a notation on the assignment should indicate your collaboration.

Submission of any assignment that is in violation of this policy will result in an administrative grade of F for the course.

If you have any question concerning this policy before submitting an assignment, please ask for clarification.