Machine Learning and Optimization (MLOPT), Fall 2020
CSCI 4961/6961, RPI
Mondays and Thursdays, 10:10am–12pm, Online

Instructor: Alex Gittens
`gittea` at rpi.edu
Office Hours: Mon, Thur 12pm–1pm ET (via Submitty office hours queue) or by appointment

TA: Owen Xie
`xieo` at rpi.edu
Office Hours: Tue 9-10am ET, Thur 4-5pm ET (via Submitty office hours queue) or by appointment

Course synopsis. Modern (i.e. large-scale) machine learning and data science typically proceed by formulating the desired outcome as the solution to an optimization problem, then applying randomized algorithms to solve these problems efficiently. This class introduces the probability and optimization background necessary to understand these randomized algorithms, and surveys several popular randomized algorithms, placing the emphasis on those widely used in ML applications. The homeworks will involve hands-on applications and empirical characterizations of the behavior of these algorithms.

Prereqs. CSCI 2300 (Introduction to Algorithms), MATH 2010 (Multivariable Calculus and Matrix Algebra)

Texts. There is no official text for the course, but supplementary reading material will be added to the course website as we go along. I strongly recommend the following free online resources if you find you have difficulty with the linear algebra or probability used in the course, and for learning the Keras framework we use for the deep learning assignments:

Linear Algebra:	Introduction to Applied Linear Algebra: Vectors, Matrices, and Least Squares. Vandenberghe and Boyd.
Probability:	Introduction to Probability, Statistics, and Random Processes. Hossein Pishro-Nik.
	Jeff Erickson's notes on discrete probability. Jeff Erickson.
Keras/Deep Learning:	Deep Learning with Python. 2nd Edition. François Chollet.

Learning Outcomes (CSCI 4961) Upon successful completion of this course, each student:

- ✓ can formulate standard supervised and unsupervised machine learning tasks as optimization problems,
- ✓ can recall key definitions relating to convex functions, convex sets, and convex optimization
- ✓ can implement first-order and stochastic first-order solvers for convex optimization problems,
- ✓ can implement Newton's method and L-BFGS solvers for convex optimization problems,
- ✓ can identify the trade-offs inherent in using first-order vs second-order solvers for optimization problems arising in machine learning,
- ✓ can use randomization, including sketching, hashing, and dimensionality reduction, to efficiently solve optimization problems arising in machine learning,
- ✓ is familiar with the state-of-the-art classes of architectures used in deep learning, and their mapping onto standard problem domains, including computer vision and NLP
- ✓ has experience in applying machine learning techniques, including deep learning methods, to real-world problems

Additional Learning Outcomes (CSCI 6961) Upon successful completion of this course, each student:

- ✓ can identify convex optimization problems, their optimality conditions, and properties relevant to the selection and performance of optimization algorithms such as strong convexity and smoothness.
- ✓ can read the contemporary literature on randomized approaches to problems in large-scale machine

learning, evaluate the theoretical guarantees and empirical performance of novel methods, and effectively share their conclusions.

- ✓ can understand arguments guaranteeing the performance of stochastic optimization algorithms used in large-scale machine learning.

Grading.

Homeworks	Project	Weekly Participation
50%	**35%**	**15%**

There are no makeup participation scores or homeworks. Special circumstances will be handled case-by-case, if the student presents an institute letter requesting it and if the instructor deems the request reasonable.

Threshold	90%	85%	80%	75%	70%	65%	60%	$< 60\%$
Grade	A	B+	B	C+	C	D+	D	F

Collaboration and Academic Honesty. All assignments that are turned in for a grade must represent the student's own work. In particular:

- Discussion is allowed on homework but submitted work must be your own.
- **YOU ARE RESPONSIBLE FOR ENSURING THAT YOUR HOMEWORKS AND PARTICIPATIONS ARE NOT COPIED.**
- Any material used from **anywhere** other than the lectures or supplementary material, or your notes on those two sources must be properly cited.
- You must write and understand all solutions yourself.

In cases of academic dishonesty, the minimum penalty is a course grade of F, and other institute-mandated protocols may be invoked.