STOR 712: Optimization for Machine Learning and Data Science

Time:	MW 11:15AM - 12:30PM
Place:	Hanes Hall 125
Instructor:	Michael O'Neill, STOR
email:	mikeoneill@unc.edu
office hours:	MW 2PM-3PM or by appointment
office:	Hanes Hall 323

Evaluation

- Final Project: 60%
  - *Milestone 1:* mid semester project proposal and presentation (20%)
  - *Milestone 2:* end semeter report and poster presentation (40%)
- Homework: 40%

Scope of the course

Syllabus: This course will provide a detailed and deep treatment for commonly used methods in continuous optimization, with applications in machine learning, statistics, data science, operations research, among others.

The main focus of this course is on continuous optimization algorithms, and it will also cover some core optimization theory as a foundation for the development of these algorithms. The discussions of algorithms will be accompanied with representative applications.

Topics

1. **Introduction:** Machine learning paradigm, empirical risk minimization, mathematical optimization, representative models and applications in machine learning, data science, operations research, signal/image processing, and others.
2. **Basic Theory:** Convex analysis and optimization theory, optimality conditions, smooth convex functions.
3. **First-Order Methods:** Gradient descent and accelerated variants.
4. **Stochastic Methods (Basic):** Stochastic gradient descent and accelerated variants.
5. **Stochastic Methods (Advanced):** Variance reduction and adaptive stochastic methods.
6. **First-Order Methods (Advanced):** Coordinate descent. Methods for constrained optimization: projected and conditional gradient. Methods for nonsmooth optimization: subgradient method and proximal gradient.
7. **Duality:** Lagrangian duality theory, optimality conditions, augmented lagrangian and ADMM.
8. **Second-order algorithms:** Newton and quasi-Newton methods and randomized variants.

9. **Algorithms for parallel and distributed computing:** Communication vs computation tradeoffs, synchronous and asynchronous methods, federated learning.

10. **Applications:** Neural networks/deep learning, logistic regression, support vector machines, lasso, low-rank matrix completion, least squares *(some of these applications may be merged into other topics where appropriate)*.

Text Book

"Optimization for Data Analysis" by Stephen J. Wright and Benjamin Recht.

Programming language

We will use Python for coding exercises in this course.

Homework

There will be roughly weekly homework assignments.

Final Project

Groups of 2-3 students will work on an open problem that is relevant to the course. The first component of the project will consist of a written (at most 3-page) project proposal that will be due approximately by the middle of the semester, followed by a presentation using slides. During the presentation (expected to last approximately 10-15 minutes) each group will receive feedback from the instructor and the rest of the class. The second component of the project will be an end-of-semester report, where the members of each group will present the progress of their project, their potential innovations, and their concluding remarks. All projects will be accompanied by a poster. All posters will be presented during a poster day (at a date to be determined). In past semesters, other faculty and students have been invited to attend the poster session.

Attendance

I strongly encourage you to attend all lectures and to actively participate. The material is much harder to learn on your own.

Additional References

1. "Optimization Methods for Large-Scale Machine Learning" by Léon Bottou, Frank E. Curtis, and Jorge Nocedal
   Source: `https://epubs.siam.org/doi/pdf/10.1137/16M1080173`

2. "Convex Optimization: Algorithms and Complexity" by Sébastien Bubeck
   Source: `https://arxiv.org/pdf/1405.4980.pdf`

3. "Understanding Machine Learning: From Theory to Algorithms" by Shai Ben-David and Shai Shalev-Shwartz
   Source: `https://www.cs.huji.ac.il/~shais/UnderstandingMachineLearning/index.html`

Other resources

All students should be aware of the following resources that are available at UNC-Chapel Hill.

- Honor Code: The Honor Code will be observed at all times in this course. The terms of the Honor Code are set out at `http://instrument.unc.edu`.
- Accessibility Resources and Service (ARS): The information about this service can be found at `https://ars.unc.edu`.
- Counseling and Psychological Services (CAPS): The information about this service can be found at `https://caps.unc.edu`.
- Title IX resources: Please visit `https://eoc.unc.edu/our-policies/state-and-federal-laws/title-ix-and-vawa/`.