UNIVERSITY OF NORTH CAROLINA
CHARLOTTE

Topics in Computer Science
Optimization for Machine Learning and Data Science

ITCS 6010/8010 - Fall 2023

Course Syllabus (Draft, as of March 12, 2023)

Instructor: Dr. Christian Kümmerle	**Course Number:** ITCS 6010/8010
Email: kuemmerle@uncc.edu	**Time:** Mondays & Wednesdays, 2:30 pm - 3:45 pm
	Location: Atkins 126

Summary: Optimization formulations and algorithms play a central role in modern data science, statistics and machine learning. This graduate-level course is an introduction to the basics of continuous optimization, with an emphasis on techniques that are most relevant to data science and machine learning and are amenable to large-scale implementations.
We also discuss basic convergence properties and computational aspects of relevant optimization techniques. The relationship and applications of optimization to common machine learning models such as (generalized) principal component analysis, sparse and logistic regression, regularized empirical risk minimization and the training of deep neural networks is another main focus of this course.
In addition to theory and algorithms, common optimization modeling software and computer solvers will be introduced to students.

Learning Objectives:

- Learn how to apply optimization algorithms to fit, train and evaluate machine learning models
- Understand the fundamental statistical and computational challenges in models motivating optimization
- Identify and understand the differences between linear, convex, non-convex, smooth and non-smooth optimization problems to select appropriate solution algorithms
- Develop a solid understanding of the basics of continuous optimization
- Understand the trade-offs and limitations of different numerical optimization methods
- Learn how to work with optimization modeling software and solvers
- Understand research papers at the intersection of optimization and data science/ machine learning (more focus on this in ITCS 8010)
- Gain practical experience implementing state-of-the-art customized optimization algorithms (more focus on this in ITCS 8010)

Prerequisites/ Expected Background:

- Calculus including multivariate differentiation,
- Linear Algebra
- One course in Probability and Statistics

- Some experience with Python
- Recommended: Introductory course in Machine Learning

Course Topics:

Data Models and Applications in Machine Learning

- Unsupervised Learning: Principal Component Analysis (w/ review of some linear algebra)
- Supervised Learning: Logistic Regression, Sparse Regression, Matrix Completion, Recommender Systems, Computer Vision
- Training of Deep Neural Networks
- Optimal Transport: Image Matching, Shape Registration

Algorithms and Theory of Continuous Optimization

- Basics of Convex Analysis & Optimization
- Multidimensional Calculus: Gradients, Optimality Conditions for Unconstrained Optimization
- Gradient Methods
- Autodifferentiation and Backpropagation
- Accelerated Methods
- Coordinate Descent Methods
- Stochastic Gradient Descent
- Subgradient Methods
- Proximal Gradient Methods
- Natural Gradient Methods
- Newton/Quasi-Newton Methods
- Frank-Wolfe
- Cubic Regularized Newton's Method
- Minimax Optimization

Projects: One important part of the course is a final project that is worked on by a *group* of 2-3 students (6000 level) or individually (8000 level). In the project, students apply and compare different optimization methods for a challenging real-world problem in machine learning/data science, prepare a final report, and present the results to their fellow students. For the 8000 level part, the study, implementation of recent scientific literature is part of the project.

Grading: Your final grade for the class will be given as a weighted average with the weights given as follows:

- Homework: 30%
- Midterm exam: 30%
- Final Project: 40%

Literature: The textbook that we follow most closely is the new textbook by Wright and Recht.

- Stephen J. Wright and Benjamin Recht (2022), *Optimization for Data Analysis*, Cambridge University Press.

The following book touches more on the relationship between structured data models and optimization, and its material will also be heavily used in this course.

- John Wright, Yi Ma (2022), *High-Dimensional Data Analysis with Low-Dimensional Models: Principles, Computation, and Applications*, Cambridge University Press.

Further references:

- Stephen Boyd and Lieven Vandenberghe (2009), *Convex Optimization*, Cambridge University Press.
- Marc P. Deisenroth and A. Aldo Faisal and Cheng Soon Ong (2020), *Mathematics for Machine Learning*, Cambridge University Press.
- Amir Beck (2017), *First-Order Methods in Optimization, SIAM.*
- Sébastien Bubeck (2016), *Convex Optimization: Algorithms and Complexity*, Foundations and Trends in Machine Learning.
- Elad Hazan (2019), *Lecture Notes: Optimization for Machine Learning*, Princeton University, arXiv:1909:03550.

Comparable courses at reference universities:

- Optimization for Machine Learning [Github], EPF Lausanne, Spring 2022, Prof. Martin Jaggi.
- Optimization for Data Science, ETH Zurich, Spring 2022, Prof. Niao He & Prof. Bernd Gärtner.
- Computational Principles for High-Dimensional Data Analysis, EECS 208, UC Berkeley, Fall 2021, Prof. Yi Ma
- Optimization for Machine Learning and Data Analysis, UNC Chapel Hill, Fall 2020, Prof. Quoc Tran-Dinh.
- Convex Optimization and Approximation: Optimization for Modern Data Analysis, EECS 227C/STAT 260, UC Berkeley, Spring 2019, Prof. Ben Recht.
- Optimization Methods for Signal & Image Processing and Machine Learning, EECS 559, University of Michigan, Fall 2021, Prof. Qing Qu.