EPFL

EE-556

Mathematics of data: from theory to computation

Cevher Volkan

Cursus	Sem.	Type
Computational science and Engineering	MA1, MA3	Opt.
Data Science	MA1, MA3	Opt.
Electrical Engineering		Obl.
Electrical and Electronical Engineering	MA1, MA3	Opt.
MNIS	MA3	Opt.
Managmt, tech et entr.	MA1, MA3	Opt.

Language of teaching	English
Credits	4
Session	Winter
Semester	Fall
Exam	Written
Workload	120h
Weeks	14
Hours	**4 weekly**
Courses	2 weekly
Exercises	2 weekly
Number of positions	

Summary

This course reviews recent advances in convex optimization and statistical analysis in the wake of Big Data. We provide an overview of the emerging convex formulations and their guarantees, describe scalable solution techniques, and illustrate the role of parallel and distributed computation.

Content

The course consists of the following topics
Lecture 1: Introduction to convex optimization and iterative methods.
Lecture 2: Review of basic probability theory. Maximum likelihood, M-estimators, and empirical risk minimization as a motivation for convex optimization.
Lecture 3: Fundamental concepts in convex analysis. Review of linear algebra.
Lecture 4: Unconstrained smooth minimization I: Concept of an iterative optimization algorithm. Convergence rates. Characterization of functions.
Lecture 5: Unconstrained smooth minimization II: Gradient and accelerated gradient methods.
Lecture 6: Unconstrained smooth minimization III: The quadratic case. The conjugate gradient method. Variable metric algorithms.
Lecture 7: Stochastic gradient methods.
Lecture 8: Non-convex optimization and deep learning: Non0cnovex optimization. Neural Networks. Adaptive gradient methods.
Lecture 9: Composite convex minimization I: Subgradient method. Proximal and accelerated proximal gradient methods.
Lecture 10: Composite convex minimization II: Proximal Newton-type methods. Stochastic proximal gradient methods.
Lecture 11: Constrained convex minimization I: The primal-dual approach. Smoothing approaches for non-smooth convex minimization.
Lecture 12: Constrained convex minimization II: The Frank-Wolfe method. The universal primal-dual gradient method. The alternating direction method of multipliers (ADMM).
Lecture 13: Constrained Convex optimization III.

Learning Prerequisites

Required courses

Previous coursework in calculus, linear algebra, and probability is required. Familiarity with optimization is useful.

Learning Outcomes

By the end of the course, the student must be able to:

- Choose an appropriate convex formulation for a data analytics problem at hand
- Estimate the underlying data size requirements for the correctness of its solution

- Implement an appropriate convex optimization algorithm based on the available computational platform
- Decide on a meaningful level of optimization accuracy for stopping the algorithm
- Characterize the time required for their algorithm to obtain a numerical solution with the chosen accuracy