

**Abstract:**

This is a course aimed at introducing various dimensionality reduction techniques, such as sampling and sketching, to speed up commonly occurring optimization problems, in particular those occurring in numerical linear algebra.

Today, there is a tremendous interest in processing “Big Data” due to the ubiquity of large data sets being generated in production systems. This has led to an interest in designing algorithms and systems for problems with very large inputs. The problems range from matrix problems and numerical linear algebra, optimization problems and graph problems.

This course will highlight the recent advances in algorithms for numerical linear algebra and optimization that have come from the technique of linear sketching, whereby given a matrix, one first compresses it to a much smaller matrix by pre-multiplying it by a (usually) random matrix with certain properties. Much of the expensive computation can then be performed on the smaller matrix, thereby accelerating the solution for the original problem. This technique has led to the fastest known algorithms for fundamental problems in this area. We will consider least squares as well as  $\ell_1$ -regression problems, low rank approximation, and many variants of these problems, such as those in distributed environments. We will also discuss connections of these methods with and using graph sparsifiers. Some of this work is partially covered in the speakers monograph: “Sketching as a Tool for Numerical Linear Algebra”, Foundations and Trends in Theoretical Computer Science, vol 10, issue 1-2, pp. 1-157, 2014, publisher: NOW Publishing.

This is a theoretical course. It is assumed that participants will have background in Linear Algebra and Probability and are familiar with mathematical proofs.