

EE290 – High-dimensional Data Analysis with Low-dimensional Models (Principles, Computation, and Applications)

This 3-unit graduate level course introduces basic geometric and statistical concepts and principles of low-dimensional models for high-dimensional signal and data analysis, spanning basic theory, efficient algorithms, and diverse applications. We will discuss recovery theory, based on high-dimensional geometry and non-asymptotic statistics, for sparse, low-rank, and low-dimensional models – including compressed sensing theory, matrix completion, robust principal component analysis, and dictionary learning etc. We will introduce principled methods for developing efficient optimization algorithms for recovering low-dimensional structures, with an emphasis on scalable and efficient first-order methods, for solving the associated convex and nonconvex problems. We will illustrate the theory and algorithms with numerous application examples, drawn from computer vision, image processing, audio processing, communications, scientific imaging, bioinformatics, information retrieval etc. The course will provide ample mathematical and programming exercises with supporting algorithms, codes, and data. A final course project will give students additional hands-on experience with an application area of their choosing. Throughout the course, we will discuss strong conceptual, algorithmic, and theoretical connections between low-dimensional models with other popular data-driven methods such as deep neural networks (DNNs), providing new perspectives to understand deep learning.

Administrative

Time and place: TuTh 3:30-5pm Online

Instructor: Professor Yi Ma

Tentative office hours: TBA.

Instructor email: yima@eecs.berkeley.edu.

Course webpage: We will use a piazza website to post lecture materials, homeworks, data, code examples, etc.

Prerequisites

Linear algebra and probability. Background in signal processing, optimization, and statistics may allow you to appreciate better certain aspects of the course material, but not necessary all at once. If you're curious about whether you would benefit from this course, contact the instructor for details. The course is open to senior undergraduates, with consent from the instructor.

Textbook

A new textbook by the instructor:

High-Dimensional Data Analysis with Low-Dimensional Models: Principles, Computation, and Applications, by John Wright and Yi Ma, to be published by Cambridge Press, 2020.

Students will be provided with drafts of this manuscript. We will also provide references to original research papers on the course website. Many of these papers contain additional results and elaborations that go far beyond what we cover in lecture and in the manuscript.

Computing

We will use the Matlab, Jupyter notebooks, and Google Colab environments for many of the in-lecture demos, examples and homeworks.

Grades

The course will be graded based on class participation (10%), homework (50%) and a course project (40%). For the course project, you can work on a topic of your choice – experimental, theoretical, or a combination of both. Be creative! Virtually any topic related to the course material is acceptable, provided the project is well-executed.

You may work alone, or in a team of two students. For teams of two, you will be expected to document who did what. Your deliverables will be a **project report** and a **short (15 min) talk** during the final exam slot for this class. If you did experimental work, you will also need to submit your **code**. You will be required to submit a brief (<1 page) **project proposal** by midterm, and to discuss your ideas with me before that date.

Tentative Syllabus (subject to changes)

- Course introduction, motivating examples
- Sparse solutions, ℓ^0 minimization, ℓ^0 uniqueness, NP-hardness
- ℓ^1 relaxation, ℓ^1 recovery under incoherence
- Recovery under RIP, random matrices
- Noise and inexact sparsity
- Rank minimization: motivating examples, nuclear norm relaxation
- Rank RIP (briefly), matrix completion
- Robust PCA and principal component pursuit
- General low-dim models, atomic norms, and geometry of random cones
- Convex optimization: first order methods, proximal gradient, acceleration
- Convex optimization: augmented Lagrangian, ADMM
- Nonconvex formulation: low-rank recovery, dictionary learning, blind deconvolution etc.
- Nonconvex optimization: from second to first order methods, randomized or regularized gradient descent
- Generalization: deep networks and low-dimensional structures.
- Applications: scientific imaging, face recognition, 3D reconstruction, photometric stereo, spectrum sensing, bioinformatics, etc.
- Course project presentations