A New Course Proposal:
EECS 208: Computational Principles for High-dimensional Data Analysis

Administrative

Instructor in Charge: Professor Yi Ma
Instructor email: `yima@eecs.berkeley.edu`.
Other willing candidate instructors: Professor Shankar Sastry, Jiantao Jiao, and Michael Lustig.

Course Description

This 4 unit graduate course EECS208 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.

The course includes 3 hours lectures (by the Instructor) and 1 hour discussion session (by a GSI) per week. Homework includes both written exercises and programming exercises. A final course project includes a midterm proposal and final presentation and report.

Prerequisites

Undergraduate linear algebra (Math 110), statistics (Stat 134), and probability (EE126). Background in signal processing (EE123), optimization (EE 227), machine learning (CS189/289), and computer vision (CS 280) may allow you to appreciate better certain aspects of the course material, but not necessary all at once. The course is open to senior undergraduates, with consent from the instructor. If you're curious about whether you would benefit from this course, contact the instructor for details.

Textbook

A new textbook by the instructor:

> *High-Dimensional Data Analysis with Low-Dimensional Models: Principles, Computation, and Applications*, 726 pages, by John Wright and Yi Ma, Cambridge University Press, 2021.

A complete pre-production version is available online: `https://book-wright-ma.github.io`. This book will be the primary reading material for the students. Students will be guided to conduct literature survey to subjects related to their projects. We will also organize and maintain additional resources, references, and material on the book website, including similar course materials offered in other institutions.

Website and Computing Tools

We normally will use a piazza website to post lecture materials, homeworks, data, code examples, etc. We will use the Matlab, Jupyter notebooks, and Google Colab environments for many of the in-lecture demos, examples and homeworks.

Grading Policy

The course will be graded based on class participation (10%), homework (50%) and a course project (40%).

Homeworks: There will be about 5 bi-weekly homeworks for the first 10 weeks. Each homework contains both a written part and a programming part. If the course is offered with exams (instead of course project), a couple of more homeworks can be assigned for the later weeks.

Project: For the course project, you can work on a topic of your choice – experimental, theoretical, or a combination of both. Be creative! The topic should be closely related to the course material, and each team need to submit a midterm **project proposal** and give a 10min presentation. The proposal and presentation should show adequate literature survey on related topics, provide good motivations and even preliminary results to support the ideas. 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 **short (15 min) talk** during the final exam week and a 8-page conference paper style **project report**. If you did experimental work, you will also need to submit your **code**.

Tentative Syllabus

This syllabus is suggestive: it follows mainly the outline of the textbook. Nevertheless, in the past offerings, we found this outline is somewhat ambitious. The textbook has suggested several options for the instructor to select a coherent subset from the listed topics for a one-semester course.

- **Week 1**: Course introduction, motivating examples (Chapter 1 of the textbook).
- **Week 2**: Sparse solutions, ℓ^0 minimization, ℓ^0 uniqueness, NP-hardness (Chapter 2).
- **Week 3**: ℓ^1 relaxation, ℓ^1 recovery under incoherence (Chapter 3)
- **Week 4**: Recovery under RIP, random matrices, Noise and inexact sparsity (Chapter 3)
- **Week 5**: Rank minimization: motivating examples, nuclear norm relaxation (Chapter 4)
- **Week 6**: Rank RIP (briefly), matrix completion (Chapter 4)
- **Week 7**: Robust PCA and principal component pursuit (Chapter 5)
- **Week 8**: Nonconvex formulation: low-rank recovery, dictionary learning, blind deconvolution etc. (Chapter 7)
- **Week 8 (alternative)**: General low-dim models, atomic norms, and geometry of random cones (Chapter 6)
- **Midterm exam or project proposal presentation**
- **Week 10**: Convex optimization: first order methods, proximal gradient, acceleration (Chapter 8)
- **Week 11** Convex optimization: augmented Lagrangian, ADMM (Chapter 8)

- **Week 12**: Nonconvex optimization: from second to first order methods, randomized or regularized gradient descent (Chapter 9)
- **Week 13**: Selected topics in applications: scientific imaging, spectrum sensing, face recognition, 3D reconstruction, photometric stereo etc. (Chapters 10 – 14)
- **Week 14**: Selected topics in nonlinear low-dim models: transform invariant low-dimensional structures or deep networks (Chapter 15 or 16)
- **Final exam or course project presentation and report**