Deck 0: Course Overview

Math 7870: Topics in Randomized Numerical Linear Algebra

Spring 2026

Akil Narayan

Math 7870 – Randomized NLA for DS/ML

This is a topics course (not part of the standard curriculum).

- Target audience: graduate students, perhaps advanced undergraduates
- Prerequisites: None, formally.
- What knowledge I will implicitly expect:
  - Significant comfort with linear algebra
  - Significant comfort with discrete probability
  - Some experience with numerical linear algebra
  - Some experience with continuous probability
- What background would be helpful
  - Probability theory (discrete is likely fine)
  - Analysis (‘real’ or ‘proof-based’)
  - (Convex) Optimization
  - Familiarity with large-scale applications in data science or machine learning

What is this course about?

This is a topics course on randomized methods for numerical linear algebra.

Thematically, this course revolves around some central themes:

- How do we tackle large-scale linear algebra-type problems using randomness?
- What kinds of algorithmic guarantees or performance expectations can we obtain?

These questions are somewhat intentionally vague – we will make them precise through examples throughout the semester.

It's not too hard to understand the basic algorithms and ideas that are the culmination of answering these questions.

Understanding the quantitative *why* for algorithm performance requires more depth, and this is what we'll spend much/most of the course focusing on.

If time permits: we'll delve into specific DS/ML applications, and perhaps into stochastic approximation and optimization.

Grading

Your grade is based entirely on two components:

- (50%) One project/'homework' assignment. Due date probably around/just after spring break
- (50%) One end-of-semester in-class presentation

These components are largely graded based on completion.

Some logistics

We'll always meet here (M LI 1735) over the semester. I won't record/post lectures videos.

Online portals:

- The "course website"
  `http://www.sci.utah.edu/~akil/math7870.html`
  This is the primary location I'll post info: lecture slides, the project assignment, and info on presentations.
- Canvas (accessible through your CIS portal)
  Nothing terribly important here: probably just the syllabus.

Office hours: TBA, let's discuss now.

Text/resources

There isn't really a single text or resource for this course. I'll pull material from several resources (which I will cite).

Randomized NLA is a somewhat recently-crystallized field – it was still the wild west around 7-8 years ago, but now there are trusted and pioneering sources.

Software is a completely different issue: the development of hardened computational routines to accomplish specific tasks is ongoing, but there is clear progress.

The project/homework assignment

I plan for this to be an exploratory computational project. Likely I'll provide 1-3 algorithmic exercises, and your submission will entail:

- Numerical investigation of a topic. You may use any programming language.
- A narrative description of the algorithm, and it's empirical performance.
- A discussion of the theory behind the method.
- I encourage you to work with others. (But submissions must be individual.)
- Submit online through Canvas/Gradescope (I'll expect a typed submission: we're only doing one of these over the whole semester....)

The end-of-semester presentation

Logistics:

- Each student will deliver an in-class presentation toward the end of the semester.
- Exact dates and length of presentation are yet to be determined: it will depend on the number of students in this course.
- Likely: 20-25 minutes in length, happening during the last 2 weeks of the course (mid/late April).
- During the middle of the semester, I'll arrange to have you sign up for a slot + topic.

Content:

- While your own independent coding + work is certainly fine, it's perfectly acceptable to present existing work (e.g., a paper).
- The topic is of your choosing. It can be an application near + dear to your heart. Alternatively, I will provide some papers, and you can choose one of them to read + investigate + present.
- Randomized linear algebra, or the potential for using randomized linear algebra, must be a core theme of whatever you decide to present.
- I do not want in-depth technical discussion (you won't have time for that). Rather, I'd like a birds-eye view of the application/problem statement, and how randomized linear algebra can be used to tackle it.

Slides and material

The material for this course is an amalgamation from a few sources.

Slides (like these) will be posted online.

Due to the nature of this course, it'll probably be fairly difficult for me to identify which parts of which source(s) any particular topic is taken from.

Instead, at the end of each set of slides I'll provide a fairly complete list of references for the whole course. You needn't read these references, but can look at them if you're interested in learning about a particular topic in more depth.

An overview of topics

Roughly speaking, this course will follow the anatomy:

- 2-3 weeks: NLA + probability review. General technical goals.
- 2 weeks: Prototypical quantitative concentration: Johnson-Lindenstrauss, high-dimensional probability
- 2 weeks: Sketching distributions, subspace embeddings
- 2 weeks: Low-rank approximation and eigenvalue estimation
- 2 weeks: Least squares
- 2 weeks: Stochastic approximation/optimization
- 2 weeks: In-class presentations

The use of AI/LLM's

You may generally interact with an AI/LLM in this course.

However, submission of direct AI-generated content for homework is expressly forbidden. You may use AI to help with understanding how to address problems, but you should understand the ideas yourself, and should compose and craft your own solutions.

"Why do any of this if AI can do it for me?"

The goal is (and has always been) to train you to reason, think, plan, and execute an investigation. If there's a "right answer", it's important to identify it, but it's arguably more important to develop independent skills that allow you to reach that answer.