

CS 229r: Spectral Graph Theory in Computer Science Fall 2020

Syllabus

Course website: <https://salil.seas.harvard.edu/node/661604>

Lecture times: Tuesday & Thursday, 1:30-2:45pm Eastern time via Zoom

Live shopping week meeting: Thursday 8/20, 1:30-2:45pm Eastern time via Zoom (will be recorded). The Zoom link is available from [Canvas](#). Please email the course staff (below) if you are not able to access it.

Shopping week office hours: see <http://salil.seas.harvard.edu/> for Salil's office hours on 8/20 and 8/21.

Course Staff

Instructor: Salil Vadhan (he/him/his; they/them/theirs)
<http://salil.seas.harvard.edu/>, salil_vadhan@harvard.edu

Teaching Fellow: Chi-Ning Chou (he/him/his; they/them/theirs)
<https://cnchou.github.io/>, chiningchou@g.harvard.edu

Faculty Coordinator: Allison Choat (she/her/hers; they/them/theirs)
achoat@seas.harvard.edu

Whenever possible, please post questions for the course staff on Piazza (privately if needed) rather than emailing us, so that we can all see the question and responses.

Overview

Spectral graph theory is about how eigenvalues, eigenvectors, and other linear-algebraic quantities give us useful information about a graph, for example about how well-connected it is, how well we can cluster or color the nodes, and how quickly random walks converge to a limiting distribution. Spectral graph theory has turned out to be extremely useful in theoretical computer science, with applications ranging from solving linear systems, converting randomized

algorithms to deterministic algorithms, sampling via Markov Chain Monte Carlo, counting, web search, and maximum flow. In this course, we will study both the mathematics and the algorithmic applications of spectral graph theory, including some results from the past couple of years. Along the way, we will cover a number of concepts and tools that are quite broadly useful, such as mixing of Markov chains, expander graphs, algorithmic linear algebra, and the multiplicative weights method for optimization.

Learning Outcomes

By the end of the course, we hope that you will all be able to:

- Comfortably use linear-algebraic tools in designing algorithms and in studying graph-theoretic problems,
- Extract both the high-level ideas and low-level details when reading a text and identify interesting questions that are not answered,
- Explain and collaboratively work through an advanced theory subject with your peers,
- Understand the state of the art in algorithmic spectral graph theory as needed to engage in research in this area and/or apply it to other areas.
- Formulate and carry out an interesting, short-term independent project, and present the work in both written and oral form.

Prerequisites

CS 124 and linear algebra at the level of AM 22a or Math 21a.

Format and Expectations

The main class meetings (TuTh 1:30-2:45pm Eastern time via Zoom) In addition, Chi-Ning will hold a section every other week, and both Salil and Chi-Ning will hold online office hours at multiple times of the day to accommodate a variety of time zones. Furthermore, we will hold several online social gatherings (e.g. using [gather](#)) to create a community within the class and facilitate finding partners for problem sets and projects.

The main components of the course are as follows:

- **Main class meetings (TuTh 1:30-2:45pm Eastern):** These will be interactive lectures, with the exposition and discussion guided by student questions and comments, interspersed with occasional problem-solving and brainstorming in breakout groups. Students are expected to attend these meetings, and keep their cameras on during class.
- **Section (every other week, time TBD):** A section will be held every other week, at a time TBD to accommodate as many students as possible. The sections will include guided work on practice problems, review of difficult concepts, and supplementary material. Attendance is not required.
- **Office hours (every week, at multiple times):** Salil and Chi-Ning will hold office hours at multiple times of the day to accommodate students in a variety of time zones.
- **Reading and commenting:** to allow our class meetings to be more interactive, short

advance reading will be assigned for most class meetings, and you will be expected to comment on those readings using the social e-reader [Perusall](#). Salil will often record a short video to help guide you through the readings. Your commenting will help us focus our class meetings on the material that you all find most confusing or most interesting, and will give you experience in working through advanced material with your peers.

- **Problem sets:** There will be problem sets due approximately every other week, typically Fridays at 5:00pm Eastern time. Some of the problems may require deep thought, so you are strongly encouraged to begin them early and to brainstorm in groups of 2-3 students. (See Collaboration Policy below.)
- **Scribe notes:** students will be expected to type scribe notes for one or two class meetings during the semester. Depending on enrollment, this may be done solo or in pairs.
- **Final project:** You will do a final project on a topic of your choosing. Projects can be done individually or in pairs, with groups of three allowed for ambitious projects. Projects can involve implementing and experimenting with some of the algorithmic techniques in this area, giving an original presentation of material from the research literature that we didn't cover in class, and/or trying to obtain some new theoretical results. The project should provide good opportunities to connect the course material to your other interests and get some exposure to the frontier of research in spectral graph theory. The project will involve submitting topic ideas for feedback (due with problem set 2, 10/9, and revised as part of problem set 3, 10/23), a detailed project proposal (due 11/6, along with meeting with the course staff), a written paper (draft due in reading period, final version in exam period), and a project presentation (in exam period). We will post more details about the final project, including some directions to look for topics, early in the course.
- **Asynchronous Q&A and discussion:** We will use [Piazza](#) for Q&A, announcements, and discussions that are not attached to specific readings.
- **Equipment:** you should have a tablet that you can use for online brainstorming of mathematics with your classmates during breakout rooms in class and section and when working on problem sets and final projects. You can borrow an iPad and stylus from the Office of Undergraduate Education by calling (617) 495-7777.

Grading and Flexibility

The final grades for the course will be determined by roughly the following breakdown:

- 50% Problem Sets and Scribe Notes
- 25% Participation (in class meetings, section, office hours, Perusall, and/or Piazza)
- 25% Final Project

We recognize that the pandemic and remote learning may at times interfere with your ability to engage fully in the course. Thus, each of you will have 12 late days for the semester, of which you can use up to 3 days on any problem set, scribe note, or final project deadline. Do not use them carelessly at the start of the term; save them for unexpected circumstances when you might really need them. Any extensions beyond the 12 late days requires a note from your

resident dean (for undergraduates) or faculty advisor (for graduate students).

We will make recordings of all class meetings and sections available in case you are unable to attend or wish to rewatch them. If you must miss several meetings, please let us know so that you can fulfil the participation expectations in other ways.

Diversity and Inclusion

We would like to create a learning environment in our class that supports a diversity of thoughts, perspectives and experiences, and honors your identities (including race, gender, class, sexuality, socioeconomic status, religion, ability, etc.). We (like many people) are still in the process of learning about diverse perspectives and identities. If something was said in class (by anyone) that made you feel uncomfortable, please talk to us about it. If you feel like your performance in the class is being impacted by your experiences outside of class, please don't hesitate to come and talk with us. As a participant in course discussions, you should also strive to be open-minded and respectful of your classmates. (Statement modified from one by Dr. Monica Linden at Brown University.)

Health Accommodations

If you have a physical or mental health condition that affects your learning or classroom experience, please let us know as soon as possible so that we can do our best to support your learning (at minimum, providing all of the accommodations listed in your AEO letter if you have one). (Statement adapted from one by Prof. Krzysztof Gajos.)

Support Structures

Everyone can benefit from support during challenging times. Not only are we happy to listen and make accommodations as needed, we can also refer you to additional support structures on campus, including, but not limited to, the below.

- [Bureau of Study Counsel](#)
- [InTouch](#)
- [Counseling and Mental Health Services](#), 617-495-2042
- [Let's Talk](#)
- [Room 13](#), 617-495-4969

Collaboration Policy

Students are encouraged to discuss the course material and the homework problems with each other in small groups (2-3 people). Discussion of homework problems may include brainstorming and talking through possible solutions, but should not include one person telling the others how to solve the problem. In addition, each person must write up their solutions independently, and these write-ups should not be checked against each other or passed

around. While working on your problem sets, you should not refer to existing solutions, whether from other students, past offerings of this course, materials available on the internet, or elsewhere. All sources of ideas, including the names of any collaborators, must be listed on your submitted homework.

In general, we expect all students to abide by the Harvard College Honor Code. We view us all (teaching staff and students) as engaged in a *shared mission* of learning and discovery, not an adversarial process. The assignments we give and the rules we set for them (such as the collaboration policy) are designed with the aim of maximizing what you take away from the course. We trust that you will follow these rules, as doing so will maximize your own learning (and thus performance on exams) and will maintain a positive educational environment for everyone in the class. We welcome and will solicit feedback from you about what more we can do to support your learning.

Readings

Many of the readings in the course will be drawn from the following:

- Daniel A. Spielman. [Spectral and Algebraic Graph Theory](#).
- Luca Trevisan. [Lecture Notes on Graph Partitioning, Expanders, and Spectral Methods](#).
- Nisheeth Vishnoi. [Lx=b](#).

The introduction to the Spielman text also includes pointers to background material, for example on linear algebra. Another useful reference are the tutorial videos from the [Algorithmic Spectral Graph Theory Boot Camp](#) at the Simons Institute for the Theory of Computing.

Topics

(Bracketed items are contingent on time available.)

- Basics
 - Classes of graphs (from simple undirected graphs to weighted digraphs)
 - Associated matrices (adjacency, random-walk matrix, Laplacians)
 - Linear algebra (Spectral Theorem, Courant-Fischer, Perron-Frobenius)
- Markov chains
 - Stationary distributions
 - Mixing time and 2nd eigenvalue
 - Power method
 - Pagerank
 - Markov Chain Monte Carlo (MCMC) for sampling and counting
- Spectral partitioning and clustering
 - Cheeger's Inequality and spectral partitioning
 - Higher-order Cheeger and spectral clustering
 - [Spectral bounds on graph coloring]
 - [Partitioning of planar graphs]

- Expander graphs
 - Spectral vs. combinatorial definitions
 - Existence
 - Application to derandomization: Undirected S-T Connectivity in Logspace
 - [Expander decompositions and algorithmic applications]
- Resistor networks and spanning trees
 - Effective resistances and Schur complements
 - Matrix-Tree Theorem
 - Random spanning trees
- Spectral sparsification
 - Spectral approximation of matrices
 - Sparsification via random sampling
 - Sparsification via rank-1 updates
- Solving Laplacian systems
 - Preconditioned Richardson iterations
 - A nearly linear-time randomized algorithm
 - A nearly logarithmic-space deterministic algorithm
 - Applications to the power method and fast max-flow algorithms
- [High-dimensional expanders]
 - Basic theory
 - Application to sampling random spanning forests and matroid bases
- [Solving directed Laplacians]
 - Spectral approximation of digraphs
 - Algorithm(s) for Eulerian Laplacians
 - Reduction from general digraphs to Eulerian digraphs
- [Interlacing families]
 - Bipartite Ramanujan graphs of all sizes and degrees