Math 126 Problem 15: Opinion Dynamics

Spring 2027

Submission Instructions

Please upload your work to Gradescope by **11:59pm on Monday, April 5, 2027 (Pacific time).**

- You are encouraged to discuss ideas and work with your classmates. However, you *must acknowledge* your collaborators at the top of each solution on which you collaborated with others and you *must write* your solutions independently.
- Your solutions to theory questions must be written legibly, or typeset in LaTeX or markdown.
- I recommend that you write your solutions in a Python notebook using Google Colab.
- You should submit your solutions as a *single PDF* via the assignment on Gradescope.

Grading: The point of the problem set is for *you* to learn. To this end, I hope to disincentivize the use of LLMs by *not* grading your work for correctness. Instead, you will grade your own work by comparing it to my solutions. This self-grade is due the Friday *after* the problem set is due, also on Gradescope.
Original notebook: Google Colab.

Opinion Dynamics

Recently, many people have observed that politics feels more polarized. Some of this phenomenon is likely due to echo chambers and social media, but we can also gain insight through a more mathematical lens: A simplified model of how people form opinions.

Consider a graph, where the nodes represent people and the edges represent the friendships between them. Let $\mathbf{z}^{(0)} \in \mathbb{R}^n$ be a vector representing people's initial opinions on a topic, e.g., positive is for and negative is against. The DeGroot update procedure is arguably the most simple way to model opinion dynamics: At every iteration t, we build the new opinions $\mathbf{z}^{(t)}$ so that each person's opinion becomes the average of their friends' prior opinions $\mathbf{z}^{(t-1)}$.

Part A: Linear Algebra Representation

Define a matrix $\mathbf{A}$ so that

$$\mathbf{z}^{(t)} = \mathbf{A}\mathbf{z}^{(t-1)}.$$

Part B: Power Method

Using the `networkx` package in Python, load a small graph of your choice (e.g., the karate club graph). Build the matrix $\mathbf{A}$ for this graph. Then update a randomly initialized opinion vector 100 times.

Part C: Mean Centering and Normalization

As you may expect, averaging the opinions will make them converge to a scaling of the all ones vector (under some mild assumptions about the structure of the graph).

Instead, after each update step, mean center the opinions, and divide by their standard deviation. Mathematically, this effectively projects off the top eigenvector. Socially, this is like zooming in to focus only on the *differences* in opinion.

Now, plot the opinions after 100 updates. (It will help if you fix the position of the nodes in the visualization from one random run to the next.) What do you notice about the final opinions when you visualize them for different randomly initialized opinions?

Part D: Eigenvector View

Use `numpy` to compute the eigenvectors of $\mathbf{A}$. What is the first eigenvector? Overlay the second eigenvector on your graph visualization, how does it compare to the final opinion vector when you mean center and normalize?

Part E: Philosophy

Let's assume DeGroot dynamics are realistic, what could you conclude about the way people form opinions?