

Lecture 01: Introduction to Spectral Graph Theory

January 19, 2021

Lecturer: Matthew Hirn

1 Introduction

The main book for this course is the draft of the book, *Spectral and Algebraic Graph Theory* by Daniel Spielman [1]; you can download it [here](#). As of this writing, I am using the version dated December 4, 2019. I will make every effort to keep my notation consistent with the book's notation.

As the course progresses we will incorporate other readings, particularly for topics on graph signal processing and graph convolutional (neural) networks. Stay tuned for those.

2 Graphs

An (unweighted, undirected) *graph* $G = (V, E)$ is a collection of vertices V and edges E . The edge set consists of unordered pairs of distinct vertices:

$$E \subset \{(a, b) : a, b \in V \text{ and } a \neq b \text{ and } (a, b) = (b, a)\}.$$

A better notation for an edge is probably $\{a, b\}$ (to emphasize that the order does not matter), but it is tradition to write (a, b) .

Many times we will want to add weights to the edges; these are called *weighted graphs* and are written $G = (V, E, w)$. Here $w : E \rightarrow \mathbb{R}$ gives the weight $w(a, b) = w(b, a)$ of each edge. Almost always the weights will be positive, i.e., $w(a, b) > 0$, but there might rare occasions when we want to allow for negative weights. We can (and will) view an unweighted graph as a weighted graph in which all the edge weights are equal to one.

Here is a simple example of a graph:

$$\begin{aligned} V &= \{1, 2, 3, 4, 5, 6\}, \\ E &= \{(1, 2), (1, 3), (2, 3), (3, 4), (4, 5), (4, 6)\}. \end{aligned} \tag{1}$$

Graphs have natural visual representations. We can draw dots for the vertices and lines for the edges. Figure 1 provides a drawing of the graph in (1). Here are some other abstract graphs that we will encounter in this course (see Figure 2 for drawings of them):

- The path graph:

$$\begin{aligned} V &= \{1, 2, \dots, n\}, \\ E &= \{(1, 2), (2, 3), \dots, (n-1, n)\}. \end{aligned}$$




Figure 1: A drawing of the graph from (1).

- The cycle graph:

$$V = \{1, 2, \dots, n\},$$

$$E = \{(1, 2), (2, 3), \dots, (n-1, n), (1, n)\}.$$

- The star graph:

$$V = \{1, 2, \dots, n\},$$

$$E = \{(1, 2), (1, 3), \dots, (1, n)\}.$$

We will use these abstract graphs to gain theoretical insights, but in terms of applications, the most interesting graphs come from real world data. In this case the vertices represents objects or things, and the edges indicate some form of relationship or similarity between two objects. Here are some examples:

- Friendship graphs (like Facebook): People are vertices, edges exist between pairs of people who are friends (see Figure 3).
- Airplane route graphs: Cities are vertices, and edges exist between pairs of cities for which there is a direct flight (see Figure 4).
- Molecular graphs: Atoms are vertices and edges exists between pairs of atoms that are bonded. The graph represents a single molecule (see Figure 5).
- Gene-gene interaction graphs: Genes are vertices and two genes are connected by an edge if there is physical relationship (e.g., one gene can bind to another) or a functional relationship (see Figure 6).

In fact, almost any data set can be represented as a weighted graph. Indeed, suppose one has data set

$$\text{data set} = \{a_1, \dots, a_n\},$$




Figure 2: Drawings of the path graph, the cycle graph, and the star graph.

and you are able to provide some notion of similarity between data points through a symmetric kernel function $k(a_i, a_j) = k(a_j, a_i)$ in which $k(a_i, a_j) = 0$ means either a_i and a_j are completely dissimilar or it is impossible to measure to their similarity and $k(a_i, a_j) > 0$ means a_i and a_j have some similarity, with larger values implying a greater degree of similarity. Then we can create a weighted graph $G = (V, E, w)$ in which

$$\begin{aligned} V &= \{a_1, \dots, a_n\}, \\ E &= \{(a_i, a_j) : k(a_i, a_j) > 0 \text{ and } i \neq j\}, \\ w(a_i, a_j) &= k(a_i, a_j). \end{aligned}$$

This is often a very useful way of thinking about a data set, and the techniques we develop in this course will help you analyze such data.

Remark 1. There are other types of graphs one can consider. These include:

- Pseudograph: Graphs with loops, i.e., we allow $(a, a) \in E$ for $a \in V$.
- Directed graphs, i.e., $(a, b) \neq (b, a)$. In other words, the edges have a directionality now.
- Mixed graphs: These are graphs with some undirected edges and some directed edges.




Figure 3: Facebook friendship graph of a particular person from <https://griffgraphs.wordpress.com/tag/social-network/>. Per the description on the website, red are high school friends, blue are college friends, yellow are his girlfriend's friends, purple are academic colleagues, and pink are friends met from traveling.

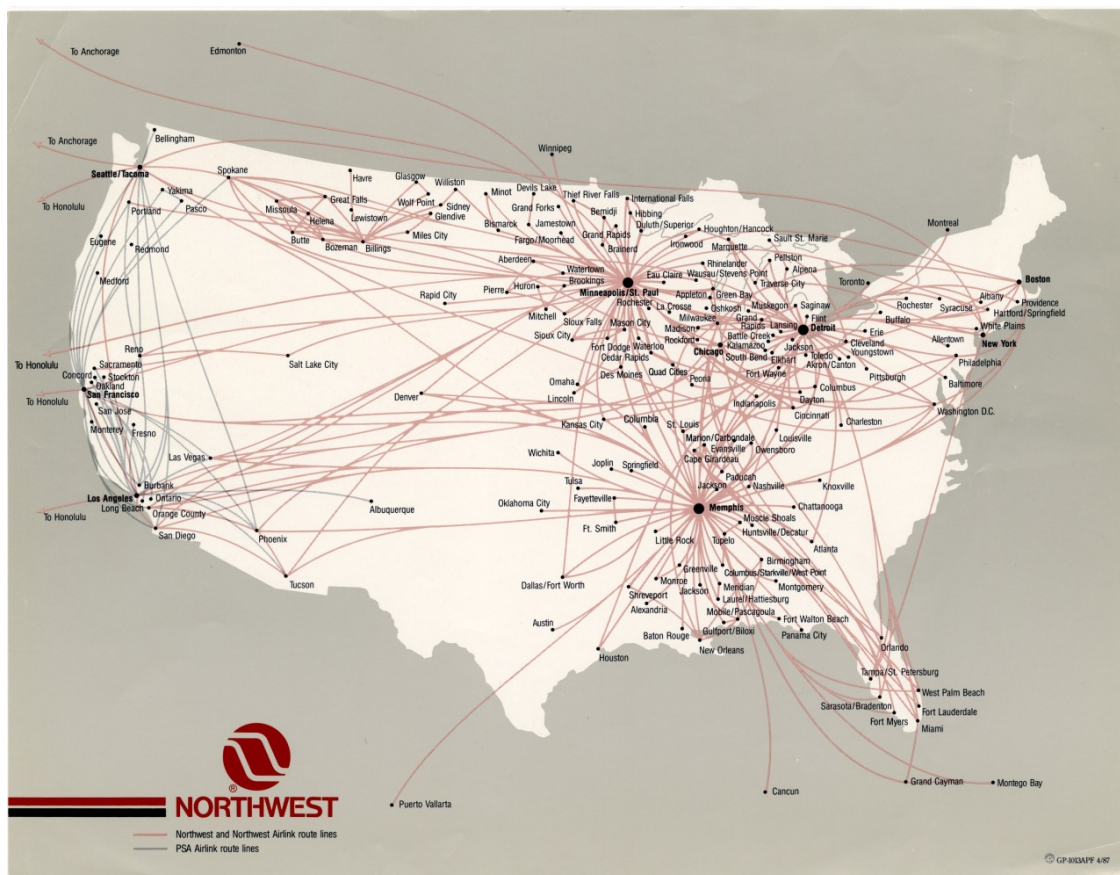


Figure 4: Northwest Airlines route map from 1987.




Figure 5: The Benzene molecule represented as a graph. The vertices marked with C represent carbon atoms, and the vertices marked with H represent hydrogen atoms; edges represent bonds.




Figure 6: Gene-gene interaction graph; provided by Prof. Arjun Krishnan of CMSE!

- Multigraphs: We allow for multiple distinct edges between the same pair of vertices.
- Hypergraphs: An edge can join more than two vertices.
- Simplicial complexes: These are structures that have vertices, edges, triangles, and their higher-dimensional counterparts.

For the most part in this course we will focus on undirected, (weighted) graphs with no loops. However, we might have an opportunity to look at some of the other types of graphs listed above, depending on time and interests.

Remark 2. From here on out, if we do not say so, we will assume $G = (V, E)$ is a graph with n vertices. If we need to label the vertices, often we will use $V = \{1, \dots, n\}$ (as above).

3 Matrices for graphs

Remark 3. We will denote functions or signals on the vertices of a graph G as $\mathbf{x} : V \rightarrow \mathbb{R}$. It will be very useful to think of these as functions, and so we will use $\mathbf{x}(a)$ to denote the value of \mathbf{x} at the vertex $a \in V$. However, it will also often be useful to think of \mathbf{x} as an $n \times 1$ column vector, in particular, so we can do linear algebra things (like matrix multiplication).

Remark 4. We will denote $n \times n$ matrices associated to G by bold uppercase letters, such as $\mathbf{M} = \mathbf{M}_G$. Entries of these matrices will be denoted by $\mathbf{M}(a, b)$, to emphasize that they depend on the vertices themselves, not the order in which we may write down the vertices.

We will associate different matrices \mathbf{M} to a graph G . Throughout much of the course, there will be three possible ways to think about such matrices:

1. As a spreadsheet that encodes the graph
2. As an operator that maps a function/vector \mathbf{x} on the vertices to a new function/vector $\mathbf{M}\mathbf{x}$.
3. As a quadratic form that maps a function/vector \mathbf{x} on the vertices to the number $\mathbf{x}^T \mathbf{M} \mathbf{x}$.

3.1 Spreadsheet: The adjacency matrix

The most obvious matrix to associate to a graph $G = (V, E)$ is its *adjacency matrix*, which is defined as

$$\mathbf{M}_G(a, b) := \begin{cases} 1 & (a, b) \in E, \\ 0 & (a, b) \notin E. \end{cases}$$

If the graph is weighted, that is $G = (V, E, w)$, then instead we use the weighted adjacency matrix:

$$\mathbf{M}_G(a, b) := \begin{cases} w(a, b) & (a, b) \in E, \\ 0 & (a, b) \notin E. \end{cases}$$

While the adjacency provides a nice way to encode a graph G , it is not terribly useful since it does not provide a natural operator or a quadratic form.

3.2 Operator: Random walks

One of the most natural operators associated to a graph is the *random walk operator*. The idea is the following. Suppose you start at some vertex $a \in V$. You are allowed to step to any other vertex $b \in V$ so long as $(a, b) \in E$; incidentally this set of vertices is called the neighborhood of a :

$$N(a) = N_G(a) := \{b \in V : (a, b) \in E\}.$$

However, you don't get to choose which vertex $b \in N(v)$ you step to, but rather you pick one at random. The probabilities are uniform in an unweighted graph and they are proportional to the weights in a weighted graph. The matrix that encodes the random walk is the random walk matrix.

To define the random walk matrix we need the notion of *degree*. The degree of a vertex in an unweighted graph is its number of neighbors:

$$\deg(a) := |N(a)|.$$

In a weighted graph we weight each of the neighbors according to their degree:

$$\deg(a) := \sum_{b \in N(v)} w(a, b).$$

Let \mathbf{d} denote the degree vector, i.e., $\mathbf{d}(a) := \deg(a)$, which we notice can be written as

$$\mathbf{d} = \mathbf{M}\mathbf{1},$$

where $\mathbf{1}$ is the vector of all ones (equivalently, the function that assigned $\mathbf{1}(a) = 1$ to every vertex $a \in V$). The *degree matrix* associated to a graph is the $n \times n$ diagonal matrix with \mathbf{d} on its diagonal:

$$\mathbf{D}(a, b) = \mathbf{D}_G(a, b) := \begin{cases} \mathbf{d}(a) & a = b, \\ 0 & a \neq b. \end{cases}$$

The random walk matrix is defined as:

$$\mathbf{W} = \mathbf{W}_G := \mathbf{M}_G \mathbf{D}_G^{-1}.$$

Let $\delta_a : V \rightarrow \mathbb{R}$ denote the function that assigns the value of one a and the value of zero to every other vertex in V , i.e.,

$$\delta_a(b) := \begin{cases} 1 & b = a, \\ 0 & b \neq a. \end{cases}$$

One can think of δ_a as a probability distribution on the vertices of G that indicates where we are going to start our random walk. It says that we are starting at the vertex a and

there is no chance we are starting anywhere else. Now let us suppose we take one step in our random walk. We want to know the probability of landing at each vertex in the graph. It will be given by:

$$\mathbf{W}\delta_a.$$

You can verify for yourself that $\mathbf{W}\delta_a(b)$ will only take nonzero values (that is, have non-zero probabilities) when $b \in N(a)$ and its entries will add up to one. If we want to know the probabilities of landing at each vertex in the graph after t steps of the random walk, it will be given by

$$\mathbf{W}^t\delta_a.$$

Spectral theory will play an important role here as it is a very useful tool by which to analyze repeated applications of an operator, i.e., powers \mathbf{W}^t .

Lecture 02: Introduction to Spectral Graph Theory (Part 2)

January 21, 2021

Lecturer: Matthew Hirn

3.3 Quadratic form: The graph Laplacian

The *graph Laplacian* is defined as

$$\mathbf{L} = \mathbf{L}_G := \mathbf{D}_G - \mathbf{M}_G.$$

It is a natural extension of the negative of the Laplacian operator Δ (from calculus) to graphs. To (partially) see this, recall that the Laplacian operator applied to a twice-differentiable function $f : \mathbb{R} \rightarrow \mathbb{R}$ computes the second derivative of f , i.e.,

$$\Delta f = f'' \text{ when } f : \mathbb{R} \rightarrow \mathbb{R}.$$

Recall further that the finite difference approximation of $f''(u)$ at $u \in \mathbb{R}$ is

$$f''(u) \approx \frac{f(u+h) - 2f(u) + f(u-h)}{h^2}. \quad (2)$$

Now consider the cycle graph with $n = 5$ vertices. Its graph Laplacian is:

$$\mathbf{L} = \begin{pmatrix} 2 & -1 & 0 & 0 & -1 \\ -1 & 2 & -1 & 0 & 0 \\ 0 & -1 & 2 & -1 & 0 \\ 0 & 0 & -1 & 2 & -1 \\ -1 & 0 & 0 & -1 & 2 \end{pmatrix}.$$

Now apply \mathbf{L} to \mathbf{x} and let us evaluate at an arbitrary vertex $a \in \{1, 2, 3, 4, 5\}$ of our cycle graph. We get:

$$-\mathbf{L}\mathbf{x}(a) = \mathbf{x}(a+1 \bmod 5) - 2\mathbf{x}(a) + \mathbf{x}(a-1 \bmod 5). \quad (3)$$

Notice that if you replace \mathbf{x} with f and set $h = 1$ and ignore the modular arithmetic, equation (3) looks very similar to (2).

The above discussion views \mathbf{L} as an operator (a sort of discrete differential operator in fact), but it will also be extremely useful to associate the following quadratic form to the graph Laplacian:

$$\mathbf{x}^T \mathbf{L} \mathbf{x} = \sum_{(a,b) \in E} w(a,b)(\mathbf{x}(a) - \mathbf{x}(b))^2. \quad (4)$$

Equation (4) measures the “smoothness” of \mathbf{x} with respect to the graph G . If \mathbf{x} changes drastically over several edges in $(a, b) \in E$ with a large weight $w(a, b)$, then (4) will be large. On the other hand, if \mathbf{x} only changes by small amounts over edges in the graph, then (4) will be small. Notice that changes in \mathbf{x} over pairs $a, b \in V$ such that $(a, b) \notin E$ do not contribute to (4).

4 Spectral Theory

Remark 5. We will denote the inner product between two vectors in two different ways, namely:

$$\langle \mathbf{x}, \mathbf{y} \rangle := \mathbf{x}^T \mathbf{y}.$$

Spectral theory refers to the study of *eigenvalues* and *eigenvectors* of matrices (or operators). Let us recall their definition. Let \mathbf{M} be an $n \times n$ matrix. An $n \times 1$ vector $\boldsymbol{\psi}$ is an eigenvector of \mathbf{M} with eigenvalue $\mu \in \mathbb{R}$ if

$$\mathbf{M}\boldsymbol{\psi} = \mu\boldsymbol{\psi},$$

and if $\boldsymbol{\psi}$ is not the all zeros vector. Here are some equivalent formulations of eigenvalues:

- μ is an eigenvalue if and only if $\mu\mathbf{I} - \mathbf{M}$ is a singular matrix (that is, not invertible).
- μ is an eigenvalue if and only if it is a root of the characteristic polynomial $p(z) = \det(z\mathbf{I} - \mathbf{M})$.

For general square matrices, it is hard to know when they admit eigenvectors and eigenvalues. However, for real valued symmetric matrices, we have the following.

Theorem 1 (Spectral Theorem). *Let \mathbf{M} be an $n \times n$ real valued, symmetric matrix. Then there exists n real eigenvalues $\mu_1, \dots, \mu_n \in \mathbb{R}$ (not necessarily distinct) and n real valued, orthonormal eigenvectors $\boldsymbol{\psi}_1, \dots, \boldsymbol{\psi}_n$ such that $\boldsymbol{\psi}_i$ is an eigenvector of \mathbf{M} with eigenvalue λ_i . That is:*

$$\begin{aligned} \mathbf{M}\boldsymbol{\psi}_i &= \mu_i\boldsymbol{\psi}_i, \\ \langle \boldsymbol{\psi}_i, \boldsymbol{\psi}_j \rangle &= \begin{cases} 1 & i = j \\ 0 & i \neq j \end{cases}. \end{aligned}$$

The eigenvalues of \mathbf{M} are unique but the eigenvectors are not. Indeed, if $\boldsymbol{\psi}$ is an eigenvector, then $-\boldsymbol{\psi}$ is also an eigenvector. Things can get even more complicated if some eigenvalues are repeated.

Many of the matrices we encounter in this course will be real valued and symmetric, and so the spectral theorem will apply. Others, such as the random walk matrix, will not be symmetric but they will be similar to a symmetric matrix, and so we will see that we can still use the Spectral Theorem to understand them.

Often the matrices we study will also be *positive semidefinite*. A matrix is positive semidefinite if it is symmetric and all of its eigenvalues are non-negative. A matrix is *positive definite* if it is symmetric and if all of its eigenvalues are positive. The next theorem proves the graph Laplacian is positive semidefinite.

Theorem 2. *The graph Laplacian is positive semidefinite.*

Proof. Let ψ be an eigenvector of \mathbf{L} with $\|\psi\| = 1$ and with eigenvalue λ . Then:

$$\psi^T \mathbf{L} \psi = \psi^T \lambda \psi = \lambda,$$

but on the other hand

$$\psi^T \mathbf{L} \psi = \sum_{(a,b) \in E} w(a,b) (\psi(a) - \psi(b))^2 \geq 0.$$

Furthermore, one can verify

$$\mathbf{L} \mathbf{1} = \mathbf{0},$$

where $\mathbf{0}$ is the vector of all zeros, meaning that $\lambda = 0$ is always an eigenvalue of \mathbf{L} with constant eigenvector. Thus \mathbf{L} is positive semidefinite and not positive definite. \square

Remark 6. We will always order the eigenvalues of \mathbf{L} in increasing order, i.e.,

$$0 = \lambda_1 \leq \lambda_2 \leq \lambda_3 \leq \cdots \leq \lambda_n.$$

I will also do my best to reserve the letter λ for eigenvalues of the graph Laplacian, and to use other (Greek) letters for eigenvalues of other matrices.

5 Previews

In this section we give previews of some of the things that will motivate us and some of the things we will study this semester.

5.1 Eigenvalue/eigenvector frequency

Eigenvalues and eigenvectors of the graph Laplacian have a natural notion of frequency, with smaller eigenvalues corresponding lower frequencies and larger eigenvalues corresponding to higher frequencies. We can use the path graph to illustrate this idea now; we will delve into it in more depth later. Let us take a look at the path graph on ten vertices; Figure 7 has a picture.

Here are the eigenvalues of the graph Laplacian of the path graph with $n = 10$ vertices: -5.16451812e-16, 9.78869674e-02, 3.81966011e-01, 8.24429495e-01, 1.38196601e+00, 2.00000000e+00, 2.61803399e+00, 3.17557050e+00, 3.61803399e+00, 3.90211303e+00. There is a pattern here that we will prove later, but for now will remain a mystery.



Figure 7: The path graph with $n = 10$ vertices.

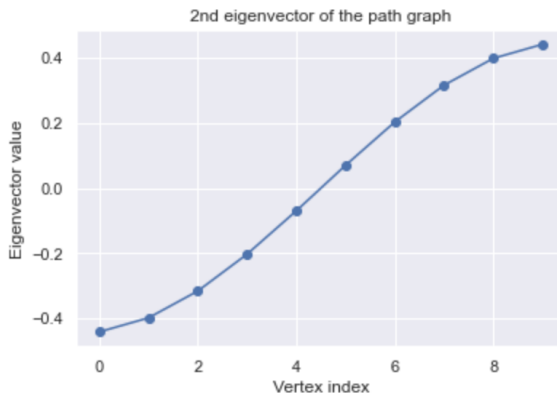


Figure 8: The second eigenvector of the graph Laplacian of the path graph with $n = 10$ vertices.

We know from the proof of Theorem 2 that the first eigenvalue is zero and its corresponding eigenvector is constant. Since that is not very interesting, let us move onto the second eigenvector, which is plotted in Figure 8. We can see the second eigenvector increases over the path graph.

Now let us plot the second, third, and fourth eigenvectors of the path graph to get a better feel for what is going on. Take a look at Figure 9. We see that as the eigenvalue increases, the eigenvectors oscillate at higher and higher frequencies. Put another way, the second eigenvector just increases, the third eigenvector goes down and then up, while the fourth eigenvector goes up, then down, then back up. These eigenvectors are like discrete versions of the modes of a vibrating string.




Figure 9: Left: The second eigenvector of the path graph. Middle: The third eigenvector of the path graph. Right: The fourth eigenvector of the path graph.

To drive the point home further, let us look at the eighth, ninth, and tenth eigenvectors, which correspond to the three eigenvectors with the three largest eigenvalues. They are plotted in Figure 10. We see the oscillate at very high frequencies! In particular, the last one goes up and down across each edge of the path graph.




Figure 10: Left: The eighth eigenvector of the path graph. Middle: The ninth eigenvector of the path graph. Right: The tenth eigenvector of the path graph.

Later in the course we will make this observation precise, first for the path graph, and then for general graphs. These results in turn will lead to lots of interesting things, such as clustering algorithms and graph signal processing, among others.

5.2 Graph signal processing

In the previous section we observed (anecdotally) that eigenvalue can be a proxy for the frequency of the eigenvector of the graph Laplacian. We can take this one step further by interpreting the eigenvectors of the graph Laplacian as the Fourier modes of the graph G . In doing so, we can create a mathematical language for signal processing on graphs, in which the aim is to extract information from a signal \mathbf{x} defined on a graph G .

One of the key things we will be able to do is to develop notions of graph filters that can be used to carry out this signal information extraction. Figure 11 shows an example in which numerically the manifold is approximated by a graph and the filters are derived from the resulting eigenvalues and eigenvectors of the graph Laplacian.

5.3 Graph convolutional networks

Graph convolutional networks merge signal processing on graphs with machine learning to carry out complex inference tasks. Spectral based graph convolutional networks are based upon graph signal processing and spectral graph theory.




Figure 11: Filters of two different sizes and locations on the person manifold, which numerically is approximated by a graph. Figure taken from [2].

5.4 Random walks on graphs

We already briefly saw random walks on graphs. We will study them in more detail. They also play a large role in machine learning on graphs and we will use them to develop models for filters on graphs that are pertinent to graph signal processing and graph convolutional networks.

5.5 Spectral clustering

Spectral graph theory also leads to a class of clustering algorithms called spectral clustering methods. Here is the idea, which we will study in more detail later in the course.

We already saw that the first eigenvalue of the graph Laplacian is zero, that is $\lambda_1 = 0$. It turns out if the graph is disconnected, then $\lambda_2 = 0$ as well. In fact, the multiplicity of the zero eigenvalue will tell us the number of connected components of the graph.

That is already pretty cool. However, we can say even more. Suppose that G is connected. Then we will show that necessarily $\lambda_2 > 0$. In fact, though, the magnitude of λ_2 tells us how well connected G is. We will make this statement quantitatively precise.

5.6 Graph embeddings/drawings


We can use the low frequency eigenvectors embed graphs in low dimensions, such as \mathbb{R}^2 for visualization. This can be a good way to get an initial impression of a graph you do not know much about. Here is one way we can embed a graph in two dimensions. For each

vertex $a \in V$, we map it to:

$$a \mapsto (\psi_2(a), \psi_3(a)) \in \mathbb{R}^2. \quad (5)$$

We do not use ψ_1 because it is the constant vector and will give us no information about the graph. We then plot the vertices using the coordinates $(\psi_2(a), \psi_3(a))$.

In Figure 12(a) is a visualization of the Minnesota graph, which represents the road network in the state of Minnesota. We compute its graph Laplacian and corresponding eigenvectors, and in Figure 12(b) use the 2nd and 3rd eigenvectors to embed the graph into \mathbb{R}^2 using the map (5). The embedding is not bad, particularly if you knew nothing about the graph.



(a) The Minnesota graph visualized using the physical geometry of the state (b) The embedding of the Minnesota graph into \mathbb{R}^2 using (5)

Figure 12: The Minnesota graph and its embedding

5.7 Graph approximation

We may discuss what it means for one graph to approximate another graph. Using this language, we can develop algorithms to approximate dense/arbitrary graphs with sparse graphs (that is, graphs with only a few edges), which reduces the memory required to store the graph on a computer and can make some graph processing algorithms more efficient. Depending on interest and time, we may also even discuss how to reduce the number of edges and vertices of a graph, while maintaining its spectral properties.

Lecture 03: Eigenvalues, Optimization, and Connectivity

January 26, 2021

Lecturer: Matthew Hirn

6 Eigenvalues and optimization: The Courant-Fischer Theorem

The *Rayleigh quotient* of a vector \mathbf{x} with respect to a matrix \mathbf{M} is

$$\text{Rayleigh quotient} := \frac{\mathbf{x}^T \mathbf{M} \mathbf{x}}{\mathbf{x}^T \mathbf{x}}.$$

An important fact is that the Rayleigh quotient is that the Rayleigh quotient of an eigenvector is its eigenvalue. That is, if $\mathbf{M}\boldsymbol{\psi} = \mu\boldsymbol{\psi}$, then

$$\frac{\boldsymbol{\psi}^T \mathbf{M} \boldsymbol{\psi}}{\boldsymbol{\psi}^T \boldsymbol{\psi}} = \mu.$$

Imagine now that you want to maximize the Rayleigh quotient of some symmetric matrix \mathbf{M} . The Courant-Fischer Theorem tells us that the maximum will be the largest eigenvalue of \mathbf{M} , and the vector that achieves this maximum will be the corresponding eigenvector. In fact, it will characterize every eigenvalue of \mathbf{M} .

Theorem 3 (Courant-Fischer Theorem). *Let \mathbf{M} be an $n \times n$ symmetric, real valued matrix with eigenvalues $\mu_1 \geq \mu_2 \geq \dots \geq \mu_n$. Then,*

$$\mu_k = \max_{\substack{S \subseteq \mathbb{R}^n \\ \dim(S)=k}} \min_{\substack{\mathbf{x} \in S \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T \mathbf{M} \mathbf{x}}{\mathbf{x}^T \mathbf{x}} = \min_{\substack{T \subseteq \mathbb{R}^n \\ \dim(T)=n-k+1}} \max_{\substack{\mathbf{x} \in T \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T \mathbf{M} \mathbf{x}}{\mathbf{x}^T \mathbf{x}},$$

where the outer max and min are over subspaces S and T of \mathbb{R}^n .

We will need the following lemma to prove Theorem 3.

Lemma 4. *Let \mathbf{M} be an $n \times n$ symmetric, real valued matrix with eigenvalues μ_1, \dots, μ_n and corresponding eigenvectors $\boldsymbol{\psi}_1, \dots, \boldsymbol{\psi}_n$. Then for any $\mathbf{x} \in \mathbb{R}^n$,*

$$\mathbf{x}^T \mathbf{M} \mathbf{x} = \sum_{i=1}^n \mu_i |\langle \mathbf{x}, \boldsymbol{\psi}_i \rangle|^2.$$

Proof. Since \mathbf{M} is a real valued, symmetric matrix we can apply the Spectral Theorem (Theorem 1). Let $\boldsymbol{\psi}_1, \dots, \boldsymbol{\psi}_n$ be the orthonormal eigenvectors of \mathbf{M} with corresponding eigenvalues μ_1, \dots, μ_n . Since we have n such eigenvectors, they form an orthonormal basis (ONB) for \mathbb{R}^n , meaning that we may write

$$\mathbf{x} = \sum_{i=1}^n \langle \mathbf{x}, \boldsymbol{\psi}_i \rangle \boldsymbol{\psi}_i \quad \text{and} \quad \|\mathbf{x}\|^2 = \sum_{i=1}^n |\langle \mathbf{x}, \boldsymbol{\psi}_i \rangle|^2. \quad (6)$$

Equation (6) is a standard fact from Linear Algebra that we will use a lot in this course; if you do not remember it and/or do not remember why it is true, please go find your favorite Linear Algebra book and look it up ☺.

The lemma will follow from (6) and the following calculation:

$$\begin{aligned} \mathbf{x}^T \mathbf{M} \mathbf{x} &= \langle \mathbf{x}, \mathbf{M} \mathbf{x} \rangle \\ &= \left\langle \sum_{i=1}^n \langle \mathbf{x}, \boldsymbol{\psi}_i \rangle \boldsymbol{\psi}_i, \mathbf{M} \sum_{j=1}^n \langle \mathbf{x}, \boldsymbol{\psi}_j \rangle \boldsymbol{\psi}_j \right\rangle \\ &= \left\langle \sum_{i=1}^n \langle \mathbf{x}, \boldsymbol{\psi}_i \rangle \boldsymbol{\psi}_i, \sum_{j=1}^n \langle \mathbf{x}, \boldsymbol{\psi}_j \rangle \mathbf{M} \boldsymbol{\psi}_j \right\rangle \\ &= \left\langle \sum_{i=1}^n \langle \mathbf{x}, \boldsymbol{\psi}_i \rangle \boldsymbol{\psi}_i, \sum_{j=1}^n \langle \mathbf{x}, \boldsymbol{\psi}_j \rangle \mu_j \boldsymbol{\psi}_j \right\rangle \\ &= \sum_{i,j=1}^n \mu_j \langle \mathbf{x}, \boldsymbol{\psi}_i \rangle \langle \mathbf{x}, \boldsymbol{\psi}_j \rangle \underbrace{\langle \boldsymbol{\psi}_i, \boldsymbol{\psi}_j \rangle}_{\delta(i-j)} \\ &= \sum_{i=1}^n \mu_i |\langle \mathbf{x}, \boldsymbol{\psi}_i \rangle|^2. \end{aligned}$$

□

Proof of Theorem 3. We will prove the second formulation, as the proof of the first formulation can be found in [1, Chapter 2]. To start, let

$$T = \text{span}\{\boldsymbol{\psi}_k, \dots, \boldsymbol{\psi}_n\}.$$

We can expand any $\mathbf{x} \in T$ as

$$\mathbf{x} = \sum_{i=k}^n \langle \mathbf{x}, \boldsymbol{\psi}_i \rangle \boldsymbol{\psi}_i.$$

Furthermore, using Lemma 4 and (6),

$$\frac{\mathbf{x}^T \mathbf{M} \mathbf{x}}{\mathbf{x}^T \mathbf{x}} = \frac{\sum_{i=k}^n \mu_i |\langle \mathbf{x}, \boldsymbol{\psi}_i \rangle|^2}{\sum_{j=k}^n |\langle \mathbf{x}, \boldsymbol{\psi}_j \rangle|^2} \leq \frac{\mu_k \sum_{i=k}^n |\langle \mathbf{x}, \boldsymbol{\psi}_i \rangle|^2}{\sum_{j=k}^n |\langle \mathbf{x}, \boldsymbol{\psi}_j \rangle|^2} = \mu_k.$$

Therefore,

$$\max_{\substack{\mathbf{x} \in T \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T \mathbf{M} \mathbf{x}}{\mathbf{x}^T \mathbf{x}} \leq \mu_k.$$

Thus we have an upper bound.

To prove equality, we will show that for all subspaces T of dimension $n - k + 1$ we have

$$\max_{\substack{\mathbf{x} \in T \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T \mathbf{M} \mathbf{x}}{\mathbf{x}^T \mathbf{x}} \geq \mu_k.$$

To do so, let

$$S = \text{span}\{\boldsymbol{\psi}_1, \dots, \boldsymbol{\psi}_k\}.$$

Since S has dimension k , any T of dimension $n - k + 1$ has an intersection with S of dimension at least one, i.e., $\dim(S \cap T) \geq 1$. Therefore

$$\max_{\substack{\mathbf{x} \in T \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T \mathbf{M} \mathbf{x}}{\mathbf{x}^T \mathbf{x}} \geq \max_{\substack{\mathbf{x} \in S \cap T \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T \mathbf{M} \mathbf{x}}{\mathbf{x}^T \mathbf{x}} \geq \min_{\substack{\mathbf{x} \in S \cap T \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T \mathbf{M} \mathbf{x}}{\mathbf{x}^T \mathbf{x}} \geq \min_{\substack{\mathbf{x} \in S \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T \mathbf{M} \mathbf{x}}{\mathbf{x}^T \mathbf{x}}.$$

Since any $\mathbf{x} \in S$ can be written as

$$\mathbf{x} = \sum_{i=1}^k \langle \mathbf{x}, \boldsymbol{\psi}_i \rangle \boldsymbol{\psi}_i,$$

we have, using Lemma 4 and (6) again, for any $\mathbf{x} \in S$,

$$\frac{\mathbf{x}^T \mathbf{M} \mathbf{x}}{\mathbf{x}^T \mathbf{x}} = \frac{\sum_{i=1}^k \mu_i |\langle \mathbf{x}, \boldsymbol{\psi}_i \rangle|^2}{\sum_{j=1}^k |\langle \mathbf{x}, \boldsymbol{\psi}_j \rangle|^2} \geq \frac{\mu_k \sum_{i=1}^k |\langle \mathbf{x}, \boldsymbol{\psi}_i \rangle|^2}{\sum_{j=1}^k |\langle \mathbf{x}, \boldsymbol{\psi}_j \rangle|^2} = \mu_k.$$

□

7 The Laplacian and connectivity

Recall the graph Laplacian is defined as $\mathbf{L} = \mathbf{D} - \mathbf{M}$. Let us think of \mathbf{L} as an operator and compute $\mathbf{L}\mathbf{x}(a)$ at a vertex $a \in V$:

$$\begin{aligned} \mathbf{L}\mathbf{x}(a) &= \mathbf{D}\mathbf{x}(a) - \mathbf{M}\mathbf{x}(a) \\ &= \mathbf{d}(a)\mathbf{x}(a) - \sum_{b \in N(a)} w(a, b)\mathbf{x}(b) \\ &= \left(\sum_{b \in N(a)} w(a, b) \right) \mathbf{x}(a) - \sum_{b \in N(a)} w(a, b)\mathbf{x}(b) \\ &= \sum_{b \in N(a)} w(a, b)(\mathbf{x}(a) - \mathbf{x}(b)). \end{aligned} \tag{7}$$

Equation (7) gives us an alternate way of writing $\mathbf{L}\mathbf{x}(a)$. We immediately see from (7) that

$$\mathbf{L}\mathbf{1} = \mathbf{0},$$

which verifies this fact that we had used in the proof of Theorem 2. Thus $\lambda_1 = 0$ is always the smallest eigenvalue of \mathbf{L} and it has corresponding eigenvector $\mathbf{1}$. The next theorem shows that whether λ_2 is zero or not corresponds to whether G is connected or not; it is our first theorem relating spectral properties of \mathbf{L} to the structure of G .

Theorem 5. *Let $G = (V, E, w)$ be a weighted graph, and let $0 = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ be the eigenvalues of its graph Laplacian \mathbf{L} . Then, $\lambda_2 > 0$ if and only if G is connected.*

Proof. “If and only if” means we need to prove two things:

1. $\lambda_2 > 0$ implies G is connected;
2. G is connected implies $\lambda_2 > 0$.

Let us start with the first statement. First note that it is equivalent to “ G is disconnected implies $\lambda_2 = 0$.” Let us prove this statement. Since G is disconnected, it can be written as the union of two subgraphs, $G = G_1 \cup G_2$, where there are no edges going between G_1 and G_2 ; see also Figure 13. Order the vertices of G so that ones are from G_1 and the last ones




Figure 13: Since G is disconnected, it can be partitioned into two subgraphs G_1 and G_2 that have no edges between them.

are from G_2 . Since there are no edges going between G_1 and G_2 , we can write the graph Laplacian of G as:

$$\mathbf{L}_G = \begin{pmatrix} \mathbf{L}_{G_1} & \mathbf{0} \\ \mathbf{0} & \mathbf{L}_{G_2} \end{pmatrix},$$

where in the above equation $\mathbf{0}$ represents a sub-matrix of zeros ($\mathbf{0}$ will alternate between meaning a vectors of zeros and a matrix of zeros; the context should make clear which interpretation to use). Therefore \mathbf{L}_G has two independent eigenvectors both with eigenvalue zero:

$$\boldsymbol{\psi}_1 = \begin{pmatrix} \mathbf{1} \\ \mathbf{0} \end{pmatrix} \quad \text{and} \quad \boldsymbol{\psi}_2 = \begin{pmatrix} \mathbf{0} \\ \mathbf{1} \end{pmatrix},$$

where $\boldsymbol{\psi}_1$ is constant on G_1 and zero on G_2 and $\boldsymbol{\psi}_2$ is zero on G_1 and constant on G_2 .

Now for the second statement. Let $\boldsymbol{\psi}$ be an eigenvector of \mathbf{L} with eigenvalue zero,

$$\mathbf{L}\boldsymbol{\psi} = \mathbf{0}.$$

Therefore:

$$\boldsymbol{\psi}^T \mathbf{L}\boldsymbol{\psi} = \sum_{(a,b) \in E} w(a,b)(\boldsymbol{\psi}(a) - \boldsymbol{\psi}(b))^2 = 0.$$

Thus, for every $a, b \in V$ such that $(a, b) \in E$ (i.e., every pair of vertices connected by an edge), we have $\boldsymbol{\psi}(a) = \boldsymbol{\psi}(b)$. Now fix an $a \in V$ and let $b \in V$ be any vertex in G . Since G is connected there exists a path in G from a to b ; that is, there is a sequence of vertices:

$$a = v_1, v_2, \dots, v_m = b \quad \text{such that} \quad (v_i, v_{i+1}) \in E, \quad \forall 1 \leq i \leq m-1.$$

Therefore:

$$\boldsymbol{\psi}(a) = \boldsymbol{\psi}(v_1) = \boldsymbol{\psi}(v_2) = \dots = \boldsymbol{\psi}(v_m) = \boldsymbol{\psi}(b).$$

Thus $\boldsymbol{\psi}$ is constant on G and we conclude that the only eigenvectors of \mathbf{L} with eigenvalue zero are constant vectors. Thus there is only one independent eigenvector with eigenvalue zero, and so $\lambda_2 > 0$. \square

Remark 7. In fact we can strengthen Theorem 5 and show that the multiplicity of the zero eigenvalue is equal to the number of connected components of G .

Using similar techniques we can also prove that an eigenvector $\boldsymbol{\psi}$ of zero eigenvalue must be constant on each of the connected components of G .

Theorem 6. Let $G = (V, E, w)$ be a weighted graph and let $\boldsymbol{\psi}$ be an eigenvector of its graph Laplacian with eigenvalue 0, i.e., $\mathbf{L}\boldsymbol{\psi} = \mathbf{0}$. Then $\boldsymbol{\psi}$ must be constant on each of the connected components of G .

Proof. As in the proof of Theorem 5, since $\mathbf{L}\boldsymbol{\psi} = \mathbf{0}$ we have

$$0 = \boldsymbol{\psi}^T \mathbf{L}\boldsymbol{\psi} = \sum_{(a,b) \in E} w(a,b)(\boldsymbol{\psi}(a) - \boldsymbol{\psi}(b))^2,$$

and it follows that $\boldsymbol{\psi}(a) = \boldsymbol{\psi}(b)$ for all $(a, b) \in E$. Now let $a, b \in V$ and suppose that a and b are in the same connected component of G . Since they are in the same connected component, there is a path from a to b , i.e.,

$$a = v_1, v_2, \dots, v_m = b \quad \text{such that} \quad (v_i, v_{i+1}) \in E, \quad \forall 1 \leq i \leq m-1.$$

Therefore:

$$\boldsymbol{\psi}(a) = \boldsymbol{\psi}(v_1) = \boldsymbol{\psi}(v_2) = \dots = \boldsymbol{\psi}(v_m) = \boldsymbol{\psi}(b),$$

and it follows that $\boldsymbol{\psi}$ is constant on the connected component. \square

Remark 8. If we want to cluster our data in terms of connected components, Remark 7 and Theorem 6 tells us how to do so. We can compute the graph Laplacian \mathbf{L} and then compute its eigenvalues and eigenvectors. We then check the multiplicity of the zero eigenvalue, which gives us the number of connected components. If the multiplicity of the zero eigenvalue is k , we can use ψ_1, \dots, ψ_k to find the connected components of G , although it might not be so easy that each eigenvector corresponds to the indicator function on a connected component. Nevertheless, since each ψ_1, \dots, ψ_k has to be constant on each connected component, and since the eigenvectors can be taken to be orthogonal, we can combine the information from each eigenvector to figure out the clusters.

Here is an example with four clusters. Suppose our graph G has 20 vertices and 4 connected components G_1, G_2, G_3 , and G_4 , all of which have 5 vertices. Then the following four eigenvectors are orthogonal and all have eigenvalue zero:

$$\psi_1 = \mathbf{1}, \quad \psi_2 = \mathbf{1}_{G_1 \cup G_2} - \mathbf{1}_{G_3 \cup G_4}, \quad \psi_3 = \mathbf{1}_{G_1} - \mathbf{1}_{G_2}, \quad \psi_4 = \mathbf{1}_{G_3} - \mathbf{1}_{G_4}.$$

Clearly ψ_1 tells us nothing. The second eigenvector, ψ_2 , is useful and separates $G_1 \cup G_2$ from $G_3 \cup G_4$, but does not separate G_1 from G_2 nor G_3 from G_4 . The third eigenvector, though, separates G_1 from G_2 , but still tells us nothing for separating G_3 from G_4 . However, the last eigenvector, ψ_4 , allows us to separate G_3 from G_4 .

A way to automate the example is the following. Suppose the multiplicity of the zero eigenvalue is k . Then we will use ψ_1, \dots, ψ_k to embed each vertex into \mathbb{R}^k via:

$$a \mapsto (\psi_1(a), \dots, \psi_k(a)) \in \mathbb{R}^k.$$

Those vertices with the same k -tuples are in the same cluster, and those with different k -dimensional embeddings are in different clusters. In the above example we have:

$$\begin{aligned} a \in G_1 &\mapsto (1, 1, 1, 0) \\ a \in G_2 &\mapsto (1, 1, -1, 0) \\ a \in G_3 &\mapsto (1, -1, 0, 1) \\ a \in G_4 &\mapsto (1, -1, 0, -1) \end{aligned}$$

Remark 9. Later on in the course we will show, quantitatively, that even if G is connected, the size of $\lambda_2 > 0$ will indicate how well connected G is! We can use this idea to cluster G even when it has one connected component.

Lecture 04: The Complete Graph and Drawing Graphs

January 28, 2021

Lecturer: Matthew Hirn

8 The complete graph

Throughout the course, but especially in the first part of it, we will compute the eigenvalues (and eigenvectors) of particular graphs. We begin doing so in this section, and we start with the complete graph.

The *complete graph* on n vertices, often denoted by $K_n = (V, E)$, is defined as:

$$\begin{aligned} V &= \{1, \dots, n\}, \\ E &= \{(a, b) : a, b \in V \text{ and } a \neq b\}. \end{aligned}$$

In other words, the complete graph contains every possible edge; see Figure 14 for a picture of K_5 .




Figure 14: The complete graph on $n = 5$ vertices, K_5 .

Theorem 7. *The graph Laplacian of K_n has eigenvalue 0 with multiplicity 1 and eigenvalue n with multiplicity $n - 1$.*

Proof. Since K_n is connected, we know from Theorem 5 that $\lambda_1 = 0$ and $\lambda_2 > 0$; therefore, the eigenvalue 0 has multiplicity 1, and we know its eigenvector is $\mathbf{1}$.

To compute the non-zero eigenvalues, let $\psi \neq \mathbf{0}$ be any vector orthogonal to $\mathbf{1}$. Thus:

$$\langle \psi, \mathbf{1} \rangle = \sum_{a \in V} \psi(a) = 0 \implies \psi(a) = - \sum_{\substack{b \in V \\ b \neq a}} \psi(b).$$

Let us now compute $L\psi(a)$ for an arbitrary vertex $a \in V$. Using (7) and the above equation, we get:

$$L\psi(a) = \sum_{\substack{b \in V \\ b \neq a}} (\psi(a) - \psi(b)) = (n-1)\psi(a) - \sum_{\substack{b \in V \\ b \neq a}} \psi(b) = n\psi(a).$$

Since a was arbitrary we have $\mathbf{L}\boldsymbol{\psi} = n\boldsymbol{\psi}$ for any vector orthogonal to $\mathbf{1}$. In other words, every vector orthogonal to $\mathbf{1}$ is an eigenvector of \mathbf{L} with eigenvalue n , and so the eigenvalue n has multiplicity $n - 1$. \square

9 Drawing with Laplacian eigenvectors

In this section we give mathematical justification to embedding a connected graph into \mathbb{R}^2 using the first two non-trivial eigenvectors of its graph Laplacian. Recall the map was given by $a \mapsto (\boldsymbol{\psi}_2(a), \boldsymbol{\psi}_3(a)) \in \mathbb{R}^2$, and an example was given in Figure 12.

Let us begin with the simpler problem of embedding a graph into \mathbb{R} . Let $\mathbf{x} : V \rightarrow \mathbb{R}$ be a candidate function on the vertices that we are going to use to map G into \mathbb{R} via $a \mapsto \mathbf{x}(a)$. Our modeling assumption is that we would like vertices that are neighbors to be close to one another in the embedding. So a natural candidate is:

$$\mathbf{x} = \arg \inf_{\mathbf{z} \in \mathbb{R}^n} \mathbf{z}^T \mathbf{L} \mathbf{z} = \sum_{(a,b) \in E} (\mathbf{z}(a) - \mathbf{z}(b))^2.$$

However, there are some issues. The first is that there is nothing to prevent us from picking $\mathbf{x} = \mathbf{0}$, the all zeros vector, which maps every vertex to 0. We can remedy this by enforcing that \mathbf{x} have unit norm, i.e.,

$$\mathbf{x} = \arg \inf_{\substack{\mathbf{z} \in \mathbb{R}^n \\ \|\mathbf{z}\|=1}} \mathbf{z}^T \mathbf{L} \mathbf{z}.$$

This new version still does not work, though, as we can pick $\mathbf{x} = (1/\sqrt{n})\mathbf{1}$; that is the constant vector. This embedding still maps every vertex to the same point in \mathbb{R} , which is not a very interesting or useful embedding. We can fix this problem by adding the constraint that \mathbf{x} be orthogonal to $\mathbf{1}$:

$$\mathbf{x} = \arg \inf_{\substack{\mathbf{z} \in \mathbb{R}^n \\ \|\mathbf{z}\|=1 \\ \langle \mathbf{z}, \mathbf{1} \rangle = 0}} \mathbf{z}^T \mathbf{L} \mathbf{z}. \quad (8)$$

This final formulation will give us something interesting. In fact, we already know what it will give us! Using the Courant-Fischer Theorem (Theorem 3), we know that the value of the argument of (8) is λ_2 and that $\mathbf{x} = \boldsymbol{\psi}_2$ will be the solution.

But our original goal was to embed G into \mathbb{R}^2 . To do so, we need two coordinate functions, $\mathbf{x} : V \rightarrow \mathbb{R}$ and $\mathbf{y} : V \rightarrow \mathbb{R}$, that will give us an embedding $a \mapsto (\mathbf{x}(a), \mathbf{y}(a)) \in \mathbb{R}^2$. Using again the modeling assumption that we want neighboring vertices to be close to one another, we seek to minimize the sum of the squares of the lengths between neighboring vertices in the embedding:

$$\mathbf{x}, \mathbf{y} = \arg \inf_{\mathbf{z}_1, \mathbf{z}_2 \in \mathbb{R}^n} \sum_{(a,b) \in E} \left\| \begin{pmatrix} \mathbf{z}_1(a) \\ \mathbf{z}_2(a) \end{pmatrix} - \begin{pmatrix} \mathbf{z}_1(b) \\ \mathbf{z}_2(b) \end{pmatrix} \right\|^2.$$

We first notice that this minimization problem can also be written in terms of the graph Laplacian, as:

$$\begin{aligned} \sum_{(a,b) \in E} \left\| \begin{pmatrix} \mathbf{z}_1(a) \\ \mathbf{z}_2(a) \end{pmatrix} - \begin{pmatrix} \mathbf{z}_1(b) \\ \mathbf{z}_2(b) \end{pmatrix} \right\|^2 &= \sum_{(a,b) \in E} (\mathbf{z}_1(a) - \mathbf{z}_1(b))^2 + (\mathbf{z}_2(a) - \mathbf{z}_2(b))^2, \\ &= \mathbf{z}_1^T \mathbf{L} \mathbf{z}_1 + \mathbf{z}_2^T \mathbf{L} \mathbf{z}_2. \end{aligned}$$

As before we need to impose constraints on \mathbf{x} and \mathbf{y} to avoid degenerate solutions. We impose the unit norm constraint,

$$\|\mathbf{x}\| = \|\mathbf{y}\| = 1,$$

and the orthogonality constraint with respect to the constant vector,

$$\langle \mathbf{x}, \mathbf{1} \rangle = \langle \mathbf{y}, \mathbf{1} \rangle = 0.$$

However, since we have two coordinates now, we have another type of degenerate solution in which $\mathbf{x} = \mathbf{y} = \psi_2$. In order to avoid this solution, we impose that \mathbf{x} be orthogonal to \mathbf{y} , i.e., $\langle \mathbf{x}, \mathbf{y} \rangle = 0$, which leads to the optimization problem:

$$\mathbf{x}, \mathbf{y} = \arg \inf_{\substack{\mathbf{z}_1, \mathbf{z}_2 \in \mathbb{R}^n \\ \|\mathbf{z}_1\| = \|\mathbf{z}_2\| = 1 \\ \langle \mathbf{z}_1, \mathbf{1} \rangle = \langle \mathbf{z}_2, \mathbf{1} \rangle = 0 \\ \langle \mathbf{z}_1, \mathbf{z}_2 \rangle = 0}} \sum_{(a,b) \in E} \left\| \begin{pmatrix} \mathbf{z}_1(a) \\ \mathbf{z}_2(a) \end{pmatrix} - \begin{pmatrix} \mathbf{z}_1(b) \\ \mathbf{z}_2(b) \end{pmatrix} \right\|^2. \quad (9)$$

A natural candidate for the solution of (9) is $\mathbf{x} = \psi_2$ and $\mathbf{y} = \psi_3$.

More generally, suppose we want to embed G into \mathbb{R}^k with k coordinate functions $\mathbf{x}_1, \dots, \mathbf{x}_k$ that are orthonormal and are orthogonal to $\mathbf{1}$ and that minimize $\sum_{i=1}^k \mathbf{z}_i^T \mathbf{L} \mathbf{z}_i$, i.e.,

$$\mathbf{x}_1, \dots, \mathbf{x}_k = \arg \inf_{\mathbf{z}_1, \dots, \mathbf{z}_k} \sum_{i=1}^k \mathbf{z}_i^T \mathbf{L} \mathbf{z}_i \quad \text{subject to} \quad \langle \mathbf{z}_i, \mathbf{z}_j \rangle = \delta(i-j) \text{ and } \langle \mathbf{z}_i, \mathbf{1} \rangle = 0. \quad (10)$$

As in the two-dimensional case, a natural candidate is to take $\mathbf{x}_i = \psi_{i+1}$. Since $\psi_i^T \mathbf{L} \psi_i = \lambda_i$, we see that the value of the argument of (10) will be $\sum_{i=2}^{k+1} \lambda_i$. The following theorem says this is the best we can do.

Theorem 8. *Let $G = (V, E)$ be a graph and let $0 = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ be the eigenvalues of its graph Laplacian \mathbf{L} with associated orthonormal eigenvectors ψ_1, \dots, ψ_n . Let $\mathbf{x}_1, \dots, \mathbf{x}_k \in \mathbb{R}^n$ be orthonormal vectors that are all orthogonal to $\mathbf{1}$. Then*

$$\sum_{i=1}^k \mathbf{x}_i^T \mathbf{L} \mathbf{x}_i \geq \sum_{i=2}^{k+1} \lambda_i,$$

and furthermore, one achieves equality if and only if

$$\langle \mathbf{x}_i, \psi_j \rangle = 0 \text{ for all } j \text{ such that } \lambda_j > \lambda_{k+1}.$$

Proof. It is a standard fact from Linear Algebra that a set of orthonormal vectors $\mathbf{x}_1, \dots, \mathbf{x}_k \in \mathbb{R}^n$ can be completed to form an orthonormal basis $\mathbf{x}_1, \dots, \mathbf{x}_n$. We now have two ONBs, $\mathbf{x}_1, \dots, \mathbf{x}_n$ and $\boldsymbol{\psi}_1, \dots, \boldsymbol{\psi}_n$. Using (6),

$$\sum_{i=1}^n |\langle \mathbf{x}_i, \boldsymbol{\psi}_j \rangle|^2 = 1 \quad \text{and} \quad \sum_{j=1}^n |\langle \mathbf{x}_i, \boldsymbol{\psi}_j \rangle|^2 = 1. \quad (11)$$

Also, let $\boldsymbol{\psi}_1$ be the constant vector.

Since $\langle \mathbf{x}_i, \boldsymbol{\psi}_1 \rangle = 0$ for $1 \leq i \leq k$, we have using Lemma 4 and (11):

$$\begin{aligned} \forall 1 \leq i \leq k, \quad \mathbf{x}_i^T \mathbf{L} \mathbf{x}_i &= \sum_{j=2}^n \lambda_j |\langle \mathbf{x}_i, \boldsymbol{\psi}_j \rangle|^2 \\ &= \lambda_{k+1} + \sum_{j=2}^n (\lambda_j - \lambda_{k+1}) |\langle \mathbf{x}_i, \boldsymbol{\psi}_j \rangle|^2 \\ &= \lambda_{k+1} + \sum_{j=2}^{k+1} \underbrace{(\lambda_j - \lambda_{k+1})}_{\leq 0} |\langle \mathbf{x}_i, \boldsymbol{\psi}_j \rangle|^2 + \sum_{j=k+2}^n \underbrace{(\lambda_j - \lambda_{k+1})}_{\geq 0} |\langle \mathbf{x}_i, \boldsymbol{\psi}_j \rangle|^2 \\ &\geq \lambda_{k+1} + \sum_{j=2}^{k+1} (\lambda_j - \lambda_{k+1}) |\langle \mathbf{x}_i, \boldsymbol{\psi}_j \rangle|^2. \end{aligned}$$

Notice that the inequality will be an equality if and only if $\langle \mathbf{x}_i, \boldsymbol{\psi}_j \rangle = 0$ for all j such that $\lambda_j > \lambda_{k+1}$.

Now let us sum over $1 \leq i \leq k$ and use the previous calculation along with (11):

$$\begin{aligned} \sum_{i=1}^k \mathbf{x}_i^T \mathbf{L} \mathbf{x}_i &\geq k\lambda_{k+1} + \sum_{j=2}^{k+1} (\lambda_j - \lambda_{k+1}) \sum_{i=1}^k |\langle \mathbf{x}_i, \boldsymbol{\psi}_j \rangle|^2 \\ &= k\lambda_{k+1} + \sum_{j=2}^{k+1} \underbrace{(\lambda_j - \lambda_{k+1})}_{\leq 0} \underbrace{\left(1 - \sum_{i=k+1}^n |\langle \mathbf{x}_i, \boldsymbol{\psi}_j \rangle|^2\right)}_{0 \leq \dots \leq 1} \\ &\geq k\lambda_{k+1} + \sum_{j=2}^{k+1} (\lambda_j - \lambda_{k+1}) \\ &= \sum_{j=2}^{k+1} \lambda_j. \end{aligned} \quad (12)$$

The last step is to show that the final inequality, (12), is an equality if and only if $\langle \mathbf{x}_i, \boldsymbol{\psi}_j \rangle = 0$ for all j such that $\lambda_j > \lambda_{k+1}$. I leave that for you to verify on your own. \square

Theorem 8 shows that the eigenvector embedding of a graph G is optimal in the sense that we defined in this section. It thus gives a mathematical justification for this embedding

method. However, one must be careful to remember that the embedding will only be good if the quantity we are minimizing makes sense for the particular graph you are studying. There are lots of other graph embedding/data visualization algorithms that minimize (or maximize) other quantities.

Lecture 05: Product Graphs and Star Graphs

February 2, 2021


Lecturer: Matthew Hirn

10 Product graphs

Let $G = (V, E, v)$ and $H = (W, F, w)$ be two weighted graphs. Define the *product graph* $G \times H$ as the graph with vertex set $V \times W$ and with edge set $E_{G \times H}$:

- $((a_1, b), (a_2, b))$ with weight $w_{G \times H}((a_1, b), (a_2, b)) = v(a_1, a_2)$ where $(a_1, a_2) \in E$; and
- $((a, b_1), (a, b_2))$ with weight $w_{G \times H}((a, b_1), (a, b_2)) = w(b_1, b_2)$ where $(b_1, b_2) \in F$.

Let P_n be the path graph on n vertices. The graph $P_m \times P_n$ is the $m \times n$ *grid graph*; see Figure 15 for a picture.

Figure 15: The 8×12 grid graph.

Let's draw the grid graph in \mathbb{R}^2 using the eigenvector embedding of Section 9; the embedding of the 8×12 grid graph is given in Figure 16. Comparing to Figure 15, the eigenvector embedding is a remarkably good drawing of the grid graph given that it used nothing specific to the grid graph. The reason for this is that the eigenvectors of a product graph $G \times H$ are the product of the eigenvectors of the two graphs G and H . The next theorem explains.

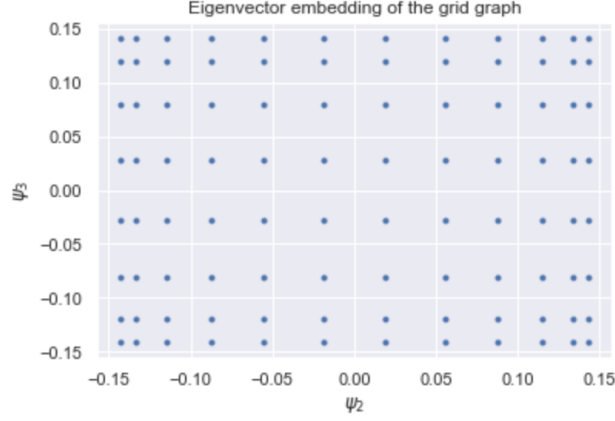


Figure 16: The eigenvector embedding of the 8×12 grid graph.

Theorem 9. Let $G = (V, E, v)$ and $H = (W, F, w)$ be weighted graphs with graph Laplacian eigenvalues $\lambda_1, \dots, \lambda_n$ and μ_1, \dots, μ_m , and graph Laplacian eigenvectors $\alpha_1, \dots, \alpha_n$ and β_1, \dots, β_m , respectively. Then $\mathbf{L}_{G \times H}$ has eigenvalues

$$\{\lambda_i + \mu_j : 1 \leq i \leq n, 1 \leq j \leq m\},$$

with eigenvectors

$$\psi_{i,j}(a, b) = \alpha_i(a)\beta_j(b), \quad \forall 1 \leq i \leq n, 1 \leq j \leq m.$$

Proof. We drop the (i, j) sub-indices but otherwise everything is defined as in the statement of theorem. Recall from (7) that

$$\mathbf{L}\mathbf{x}(c) = \sum_{d \in N(c)} \tilde{w}(c, d)(\mathbf{x}(c) - \mathbf{x}(d)),$$

for any graph $\tilde{G} = (\tilde{V}, \tilde{E}, \tilde{w})$. Let us apply this fact and make the following calculation:

$$\begin{aligned} \mathbf{L}_{G \times H}\psi(a, b) &= \sum_{(c, d) \in N(a, b)} w_{G \times H}((a, b), (c, d))(\psi(a, b) - \psi(c, d)) \\ &= \sum_{(a, a_1) \in E} v(a, a_1)(\psi(a, b) - \psi(a_1, b)) + \sum_{(b, b_1) \in F} w(b, b_1)(\psi(a, b) - \psi(a, b_1)) \\ &= \sum_{(a, a_1) \in E} v(a, a_1)(\alpha(a)\beta(b) - \alpha(a_1)\beta(b)) + \sum_{(b, b_1) \in F} w(b, b_1)(\alpha(a)\beta(b) - \alpha(a)\beta(b_1)) \\ &= \beta(b) \sum_{a_1 \in N(a)} v(a, a_1)(\alpha(a) - \alpha(a_1)) + \alpha(a) \sum_{b_1 \in N(b)} w(b, b_1)(\beta(b) - \beta(b_1)) \\ &= \beta(b)\mathbf{L}_G\alpha(a) + \alpha(a)\mathbf{L}_H\beta(b) \\ &= \beta(b)\lambda\alpha(a) + \alpha(a)\mu\beta(b) \\ &= (\lambda + \mu)\alpha(a)\beta(b) = (\lambda + \mu)\psi(a, b). \end{aligned}$$

□

11 The star graph

In your first homework you were asked to make a conjecture regarding the eigenvalues of the star graph on n vertices. Let us now prove that they are 0, 1 (with multiplicity $n - 2$), and n . Recall the star graph on n vertices, which we will denote by $S_n = (V, E)$, is defined as

$$\begin{aligned} V &= \{1, \dots, n\}, \\ E &= \{(1, a) : 2 \leq a \leq n\}. \end{aligned}$$

We will first need the following lemma, which is of interest even for general graphs.

Lemma 10. *Let $G = (V, E)$ be a graph and let $a, b \in V$ be vertices of degree one that are both connected to another vertex $c \in V$. Then, the vector $\psi = \delta_a - \delta_b$ is an eigenvector of L with eigenvalue 1.*

Proof. We will use (7) and calculate $L\psi(v)$ for each vertex $v \in V$. Applying (7) we have:

$$L\psi(v) = \sum_{u \in N(v)} (\psi(v) - \psi(u)) = \sum_{u \in N(v)} (\delta_a(v) - \delta_b(v) - \delta_a(u) + \delta_b(u)). \quad (13)$$

Now we have four cases.

1. $v = a$. Then c is only the neighbor of $v = a$ and (13) is equal to $(\delta_a(a) - \delta_b(a) - \delta_a(c) + \delta_b(c)) = 1 = \psi(a)$.
2. $v = b$. Again c is the only neighbor of $v = b$ and (13) is equal to $(\delta_a(b) - \delta_b(b) - \delta_a(c) + \delta_b(c)) = -1 = \psi(b)$.
3. $v = c$. In this case a and b are neighbors of $v = c$, and c may have other neighbors too. We write (13) as:

$$\begin{aligned} (13) &= (\delta_a(c) - \delta_b(c) - \delta_a(a) + \delta_b(a)) \\ &\quad + (\delta_a(c) - \delta_b(c) - \delta_a(b) + \delta_b(b)) \\ &\quad + \sum_{\substack{u \in N(c) \\ u \neq a, b}} (\delta_a(c) - \delta_b(c) - \delta_a(u) + \delta_b(u)) \\ &= -1 + 1 + 0 = 0 = \psi(c). \end{aligned}$$

4. $v \neq a, b, c$. In this case we may write (13) as

$$(13) = \sum_{\substack{u \in N(v) \\ u \neq a, b}} (\delta_a(v) - \delta_b(v) - \delta_a(u) + \delta_b(u)) = 0 = \psi(v).$$

□

As a corollary we have the following lemma which will also be useful.

Lemma 11. *Let $G = (V, E)$ be a graph, let $a, b \in V$ be vertices of degree one that are both connected to another vertex $c \in V$, and let ϕ be an eigenvector of \mathbf{L} with eigenvalue $\lambda \neq 1$. Then $\phi(a) = \phi(b)$.*

Proof. By Lemma 10, $\psi = \delta_a - \delta_b$ is an eigenvector of \mathbf{L} with eigenvalue 1. Furthermore, since ϕ is also an eigenvector of \mathbf{L} but with eigenvalue $\lambda \neq 1$, we know by Exercise 1 of Homework 01 that $\langle \phi, \psi \rangle = 0$. Therefore:

$$0 = \langle \phi, \psi \rangle = \sum_{v \in V} \phi(v) \psi(v) = \sum_{v \in V} \phi(v) (\delta_a(v) - \delta_b(v)) = \phi(a) - \phi(b).$$

□

Now we can prove the following theorem about the star graph.

Theorem 12. *The star graph S_n has eigenvalue 0 with multiplicity 1, eigenvalue 1 with multiplicity $n - 2$, and eigenvalue n with multiplicity 1.*

Proof. Since the star graph is connected we know it has eigenvalue 0 with multiplicity 1 and the eigenvector is $\mathbf{1}$. Notice that a and $a + 1$, for $2 \leq a \leq n - 1$, are vertices of S_n of degree 1 both connected to vertex $c = 1$. Therefore

$$\psi_a = \delta_a - \delta_{a+1}, \quad \forall 2 \leq a \leq n - 1,$$

are eigenvectors of \mathbf{L} , each with eigenvalue 1. Even though $\{\psi_a : 2 \leq a \leq n - 1\}$ are not orthogonal, they are independent, and so the eigenvalue 1 must have multiplicity at least $n - 2$. Thus we just need to determine λ_n .

In your current homework (Exercise 2 of Homework 02), you are asked to show that the trace of a symmetric, real-valued $n \times n$ matrix is equal to the sum of its eigenvalue. Let us apply this to \mathbf{L} :

$$n - 2 + \lambda_n = \sum_{i=1}^n \lambda_i = \text{Tr}(\mathbf{L}) = \text{Tr}(\mathbf{D} - \mathbf{M}) = \text{Tr}(\mathbf{D}) = \sum_{a \in V} \deg(a). \quad (14)$$

In S_n we have one vertex of degree $n - 1$ and $n - 1$ vertices of degree 1. Therefore (14) reads:

$$n - 2 + \lambda_n = 2n - 2 \implies \lambda_n = n.$$

That completes the theorem, but as a bonus we can compute the eigenvector associated to λ_n as well. By Lemma 11 we know that $\psi_n(a) = \psi_n(a + 1)$ for all $2 \leq a \leq n - 1$ which means that $\psi_n(a)$ is constant over $2 \leq a \leq n$ (the points of the star). Let us set $\psi_n(a) = 1$ for $2 \leq a \leq n$. To determine $\psi_n(1)$, we note that on the other hand, ψ_n must also be orthogonal $\mathbf{1}$. Therefore:

$$0 = \langle \psi_n, \mathbf{1} \rangle = \sum_{a \in V} \psi_n(a) = n - 1 + \psi_n(1),$$

and we see that $\psi_n(1) = -(n - 1)$. □

Lecture 06: Test Vectors and Comparing Graphs

February 4, 2021

Lecturer: Matthew Hirn

12 Upper bounding the second eigenvalue via test vectors

In this section we describe a general technique for getting upper bounds on λ_2 . We first prove the following theorem that gives a useful characterization of λ_2 (note that this theorem more or less follows from Section 9, but we state and prove it separately here since it is so useful).

Theorem 13. *Let $G = (V, E, w)$ be a weighted graph with graph Laplacian eigenvalues $0 = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$. Then*

$$\lambda_2 = \min_{\substack{\mathbf{x} \in \mathbb{R}^n \\ \langle \mathbf{x}, \mathbf{1} \rangle = 0}} \frac{\mathbf{x}^T \mathbf{L} \mathbf{x}}{\mathbf{x}^T \mathbf{x}}.$$

Proof. Let ψ_1, \dots, ψ_n be orthonormal eigenvectors of \mathbf{L} with $\psi_1 = (1/\sqrt{n})\mathbf{1}$. Thus $\langle \psi_2, \mathbf{1} \rangle = 0$ and

$$\frac{\psi_2^T \mathbf{L} \psi_2}{\psi_2^T \psi_2} = \lambda_2.$$

Therefore

$$\min_{\substack{\mathbf{x} \in \mathbb{R}^n \\ \langle \mathbf{x}, \mathbf{1} \rangle = 0}} \frac{\mathbf{x}^T \mathbf{L} \mathbf{x}}{\mathbf{x}^T \mathbf{x}} \leq \lambda_2.$$

Now let $\mathbf{x} \in \mathbb{R}^n$ with $\langle \mathbf{x}, \mathbf{1} \rangle = 0$. Using Lemma 4 we have

$$\mathbf{x}^T \mathbf{L} \mathbf{x} = \sum_{i=2}^n \lambda_i |\langle \mathbf{x}, \psi_i \rangle|^2 \geq \lambda_2 \sum_{i=2}^n |\langle \mathbf{x}, \psi_i \rangle|^2 = \lambda_2 \|\mathbf{x}\|^2 = \lambda_2 \mathbf{x}^T \mathbf{x}.$$

Thus

$$\min_{\substack{\mathbf{x} \in \mathbb{R}^n \\ \langle \mathbf{x}, \mathbf{1} \rangle = 0}} \frac{\mathbf{x}^T \mathbf{L} \mathbf{x}}{\mathbf{x}^T \mathbf{x}} \geq \min_{\substack{\mathbf{x} \in \mathbb{R}^n \\ \langle \mathbf{x}, \mathbf{1} \rangle = 0}} \frac{\lambda_2 \mathbf{x}^T \mathbf{x}}{\mathbf{x}^T \mathbf{x}} = \lambda_2.$$

□

Theorem 13 is useful for upper bounding λ_2 in the following way. It shows that any $\mathbf{x} \in \mathbb{R}^n$ with $\langle \mathbf{x}, \mathbf{1} \rangle = 0$ provides an upper bound for λ_2 since

$$\lambda_2 \leq \frac{\mathbf{x}^T \mathbf{L} \mathbf{x}}{\mathbf{x}^T \mathbf{x}}.$$

If we can guess an \mathbf{x} that is not too far from ψ_2 , then this upper bound might be a good approximation of λ_2 . Such an \mathbf{x} is called a *test vector*.

Let us apply this technique to the path graph on n vertices, which we will denote by $P_n = (V, E)$, where

$$\begin{aligned} V &= \{1, \dots, n\}, \\ E &= \{(a, a+1) : 1 \leq a \leq n-1\}. \end{aligned}$$

Recall in Figure 8 we plotted the second eigenvector of P_{10} and saw it was an increasing function of the vertex index a . Recall as well that if ψ_2 is an eigenvector, then $-\psi_2$ is also an eigenvector; the latter, in this case, would be a decreasing function of the vertex index. Even though ψ_2 for P_{10} is clearly not an affine function, we choose as our test vector

$$\mathbf{x}(a) = (n+1) - 2a, \quad 1 \leq a \leq n.$$

Figure 17 plots \mathbf{x} against $-\psi_2$ and shows, visually, that it is a good guess!

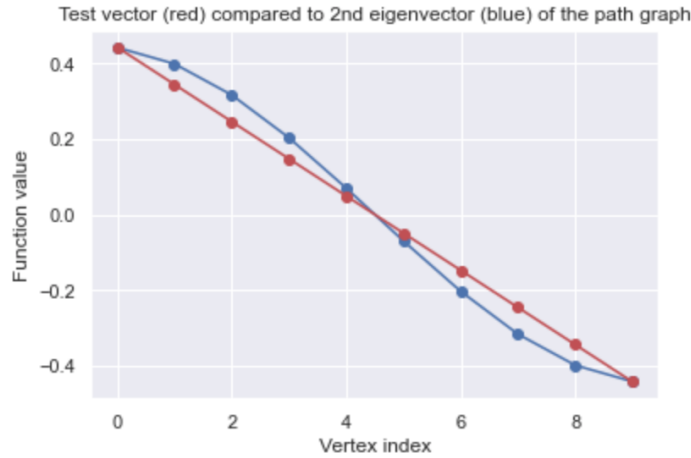


Figure 17: The second eigenvector, $-\psi_2$ (blue), of P_{10} , along with the test vector \mathbf{x} (red).

We also see that

$$\langle \mathbf{x}, \mathbf{1} \rangle = \sum_{a \in V} \mathbf{x}(a) = n(n+1) - 2 \sum_{a=1}^n a = n(n+1) - 2 \frac{n(n+1)}{2} = 0,$$

and so indeed we can use \mathbf{x} to upper bound $\lambda_2(P_n)$. Thus,

$$\begin{aligned}\lambda_2(P_n) &\leq \frac{\sum_{a=1}^{n-1} (\mathbf{x}(a) - \mathbf{x}(a+1))^2}{\sum_{a=1}^n \mathbf{x}(a)^2} \\ &= \frac{\sum_{a=1}^{n-1} 2^2}{\sum_{a=1}^n (n+1-2a)^2} \\ &= \frac{4(n-1)}{n(n+1)(n-1)/3} \\ &= \frac{12}{n(n+1)}.\end{aligned}$$

Remark 10. Theorem 13 can be thought of as a specific case of the Courant-Fischer Theorem (Theorem 3). We could try to use the same theorem to get a lower bound on λ_2 , but we would need to use the following formulation of λ_2 :

$$\lambda_2 = \max_{\substack{S \subseteq \mathbb{R}^n \\ \dim(S)=n-1}} \min_{\substack{\mathbf{x} \in S \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T \mathbf{L} \mathbf{x}}{\mathbf{x}^T \mathbf{x}}.$$

This is similar to Theorem 13; indeed $S_0 = \{\mathbf{x} \in \mathbb{R}^n : \langle \mathbf{x}, \mathbf{1} \rangle = 0\}$ is a subspace of \mathbb{R}^n of dimension $n-1$. However, to get a lower bound on λ_2 we would need a lower bound (not an upper bound) on

$$\min_{\substack{\mathbf{x} \in S \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T \mathbf{L} \mathbf{x}}{\mathbf{x}^T \mathbf{x}},$$

over a set S of dimension $n-1$, which is not easy to do. The next section will give an alternative method.

13 Comparing graphs and the Loewner partial order

In this section we will discuss how to compare graphs through their Laplacian quadratic forms. This will be immediately useful for bounding eigenvalues of graphs from below. Later on, it will also be useful for defining what it means for one graph to approximate another graph.

For a symmetric matrix \mathbf{A} we will write

$$\mathbf{A} \succcurlyeq \mathbf{0},$$

if \mathbf{A} is positive semidefinite (recall, this means all the eigenvalues of \mathbf{A} are non-negative). Thus,

$$\mathbf{A} \succcurlyeq \mathbf{0} \iff \mathbf{x}^T \mathbf{A} \mathbf{x} \geq 0, \quad \forall \mathbf{x} \in \mathbb{R}^n.$$

Similarly, we will write

$$\mathbf{A} \succcurlyeq \mathbf{B} \iff \mathbf{A} - \mathbf{B} \succcurlyeq \mathbf{0} \iff \mathbf{x}^T \mathbf{A} \mathbf{x} \geq \mathbf{x}^T \mathbf{B} \mathbf{x}, \quad \forall \mathbf{x} \in \mathbb{R}^n.$$

The relation \succsim is the *Loewner partial order*. We say partial order because there will be pairs of symmetric matrices (\mathbf{A}, \mathbf{B}) such that $\mathbf{A} \not\sucsim \mathbf{B}$ and $\mathbf{A} \not\precsim \mathbf{B}$. On the other hand, for pairs of symmetric matrices for which it does apply, it acts like an order. For example,

$$\mathbf{A} \succsim \mathbf{B} \text{ and } \mathbf{B} \succsim \mathbf{C} \implies \mathbf{A} \succsim \mathbf{C},$$

and

$$\mathbf{A} \succsim \mathbf{B} \implies \mathbf{A} + \mathbf{C} \succsim \mathbf{B} + \mathbf{C}.$$

We will overload this notation so we can compare graphs as well. In particular, for two graph G and H with the same number of vertices we define:

$$G \succsim H \iff \mathbf{L}_G \succsim \mathbf{L}_H.$$

Recall that $H = (V, F, w)$ is a subgraph of $G = (V, E, w)$ if $F \subseteq E$. Here is our first little result using this notation.

Theorem 14. *Let H be a subgraph of G . Then $G \succsim H$.*

Proof. We compute the graph Laplacian quadratic form:

$$\mathbf{x}^T \mathbf{L}_G \mathbf{x} = \sum_{(a,b) \in E} w(a,b)(\mathbf{x}(a) - \mathbf{x}(b))^2 \geq \sum_{(a,b) \in F} w(a,b)(\mathbf{x}(a) - \mathbf{x}(b))^2 = \mathbf{x}^T \mathbf{L}_H \mathbf{x},$$

where in the inequality we used the fact that $F \subseteq E$. □

Define $c \cdot G$ as the same graph as $G = (V, E, w)$ but in which every edge weight is multiplied by c , i.e.,

$$c \cdot G := (V, E, c \cdot w).$$

The following theorem shows the usefulness of the Loewner partial order as it relates to the eigenvalues of graph Laplacians.

Theorem 15. *If G and H are graphs such that*

$$G \succsim c \cdot H,$$

then

$$\lambda_k(G) \geq c \cdot \lambda_k(H), \quad \forall 1 \leq k \leq n.$$

Proof. Using the definition of $G \succsim c \cdot H$ we have:

$$\begin{aligned} \mathbf{x}^T \mathbf{L}_G \mathbf{x} &\geq c \cdot \mathbf{x}^T \mathbf{L}_H \mathbf{x}, \quad \forall \mathbf{x} \in \mathbb{R}^n \\ \implies \frac{\mathbf{x}^T \mathbf{L}_G \mathbf{x}}{\mathbf{x}^T \mathbf{x}} &\geq c \cdot \frac{\mathbf{x}^T \mathbf{L}_H \mathbf{x}}{\mathbf{x}^T \mathbf{x}}, \quad \forall \mathbf{x} \in \mathbb{R}^n, \mathbf{x} \neq \mathbf{0} \\ \implies \max_{\substack{\mathbf{x} \in T \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T \mathbf{L}_G \mathbf{x}}{\mathbf{x}^T \mathbf{x}} &\geq c \cdot \max_{\substack{\mathbf{x} \in T \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T \mathbf{L}_H \mathbf{x}}{\mathbf{x}^T \mathbf{x}}, \quad \forall T \subseteq \mathbb{R}^n, \dim(T) = k. \end{aligned} \tag{15}$$

Now define $T_0 \subseteq \mathbb{R}^n$ as

$$T_0 := \arg \min_{\substack{T \subseteq \mathbb{R}^n \\ \dim(T)=k}} \max_{\substack{\mathbf{x} \in T \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T \mathbf{L}_G \mathbf{x}}{\mathbf{x}^T \mathbf{x}}.$$

By the Courant-Fischer Theorem (Theorem 3) and (15) we have:

$$\lambda_k(G) = \max_{\substack{\mathbf{x} \in T_0 \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T \mathbf{L}_G \mathbf{x}}{\mathbf{x}^T \mathbf{x}} \geq c \cdot \max_{\substack{\mathbf{x} \in T_0 \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T \mathbf{L}_H \mathbf{x}}{\mathbf{x}^T \mathbf{x}} \geq c \cdot \min_{\substack{T \subseteq \mathbb{R}^n \\ \dim(T)=k}} \max_{\substack{\mathbf{x} \in T \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T \mathbf{L}_H \mathbf{x}}{\mathbf{x}^T \mathbf{x}} = c \cdot \lambda_k(H).$$

□

Remark 11. Later on will want to approximate one graph G by another graph H . We will say that H is a c -approximation of G if

$$c \cdot H \succcurlyeq G \succcurlyeq (1/c) \cdot H.$$

14 The path inequality

In order to use Theorem 15 we need to be able to show $G \succcurlyeq c \cdot H$ for pairs of graphs G and H . We will start by doing so for some simple graphs, and then extend to more general graphs.

To that end, recall P_n is the unweighted path graph on n vertices, and define $G_{1,n}$ as the unweighted graph on n vertices with the single edge $(1, n)$.

Theorem 16. For $n \geq 2$,

$$(n-1) \cdot P_n \succcurlyeq G_{1,n}. \quad (16)$$

Proof. Equation (16) means that we need to show:

$$\forall \mathbf{x} \in \mathbb{R}^n, \quad (n-1) \sum_{a=1}^{n-1} (\mathbf{x}(a+1) - \mathbf{x}(a))^2 \geq (\mathbf{x}(n) - \mathbf{x}(1))^2. \quad (17)$$

Define $\Delta \in \mathbb{R}^{n-1}$ as

$$\Delta(a) := \mathbf{x}(a+1) - \mathbf{x}(a), \quad \forall 1 \leq a \leq n-1.$$

Notice now that (17) can be rewritten as

$$(n-1) \sum_{a=1}^{n-1} \Delta(a)^2 \geq (\mathbf{x}(n) - \mathbf{x}(1))^2 = \left(\sum_{a=1}^{n-1} \Delta(a) \right)^2. \quad (18)$$

Recall the Cauchy-Schwartz inequality $|\langle \mathbf{x}, \mathbf{y} \rangle| \leq \|\mathbf{x}\| \|\mathbf{y}\|$. Equation (18) will follow from it with $\mathbf{x} = \Delta$ and $\mathbf{y} = \mathbf{1}$. Indeed,

$$\left(\sum_{a=1}^{n-1} \Delta(a) \right)^2 = |\langle \mathbf{1}, \Delta \rangle|^2 \leq \|\mathbf{1}\|^2 \|\Delta\|^2 = (n-1) \sum_{a=1}^{n-1} \Delta(a)^2.$$

□

Now we will prove a lower bound for $\lambda_2(P_n)$ using Theorem 16. We will do so by proving that

$$P_n \succcurlyeq c_n \cdot K_n,$$

for a constant c_n that we will compute (recall K_n is the complete graph on n vertices). First, we will need a new way of writing L_{K_n} .

Let $G_{a,b}$ be the graph on n vertices that contains only the edge (a, b) . Its graph Laplacian, $L_{G_{a,b}}$, is a matrix full of zeros except for the 2×2 sub-matrix at the intersection of the rows and columns indexed by a and b , where it is

$$\begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}.$$

One can write $L_{G_{a,b}}$ as

$$L_{G_{a,b}} = (\delta_a - \delta_b)(\delta_a - \delta_b)^T.$$

Now let $G = (V, E, w)$ be a general weighted graph. One can verify that

$$L_G = \sum_{(a,b) \in E} w(a,b) L_{G_{a,b}}. \quad (19)$$

Using (19), we have

$$L_{K_n} = \sum_{1 \leq a < b \leq n} L_{G_{a,b}}.$$

Let $P_{a,b}$, for $a < b$, be the subgraph of P_n that keeps the edges that are part of the path from a to b , and removes the other edges. Note that $P_{a,b}$ has $(b-a)$ edges. Thus, using Theorem 16 we have

$$G_{a,b} \preccurlyeq (b-a)P_{a,b}.$$

Furthermore, using Theorem 14,

$$P_{a,b} \preccurlyeq P_n \implies G_{a,b} \preccurlyeq (b-a)P_{a,b} \preccurlyeq (b-a)P_n.$$

Using this inequality we obtain:

$$L_{K_n} = \sum_{1 \leq a < b \leq n} L_{G_{a,b}} \preccurlyeq \sum_{1 \leq a < b \leq n} (b-a)L_{P_n} = \left(\sum_{1 \leq a < b \leq n} (b-a) \right) L_{P_n}.$$

One can calculate

$$\sum_{1 \leq a < b \leq n} (b-a) = \frac{1}{6}n(n^2-1),$$

and so,

$$L_{K_n} \preccurlyeq \frac{n(n^2-1)}{6} L_{P_n}.$$

To obtain an estimate for $\lambda_2(P_n)$, we apply Theorem 7 (which proves $\lambda_2(K_n) = n$) and Theorem 15 to obtain:

$$n = \lambda_2(K_n) \leq \frac{n(n^2 - 1)}{6} \lambda_2(P_n) \implies \lambda_2(P_n) \geq \frac{6}{n^2 - 1}.$$

If we combine this lower bound with the upper bound we obtained in Section 12, we get

$$\frac{6}{(n-1)(n+1)} \leq \lambda_2(P_n) \leq \frac{12}{n(n+1)}.$$

We will see soon this is a really good estimate of $\lambda_2(P_n)$!

Lecture 07: The Cycle Graph and the Path Graph

February 9, 2021

Lecturer: Matthew Hirn

15 The cycle graph

In order to compute the eigenvectors and eigenvalues of the path graph, P_n , it will be useful, and easier, to compute the eigenvectors and eigenvalues of the cycle graph. Recall the cycle graph on n vertices, which we will denote by $C_n = (V, E)$, is defined as:

$$\begin{aligned} V &= \{1, \dots, n\}, \\ E &= \{(1, 2), (2, 3), \dots, (n-1, n), (n, 1)\}. \end{aligned}$$

The following theorem shows that the eigenvectors of the cycle graph are the standard, real valued Fourier modes, namely cosine and sine functions at different discrete frequencies.

Theorem 17. *The graph Laplacian of the cycle graph, C_n , has eigenvectors*

$$\begin{aligned} \mathbf{x}_k(a) &= \cos(2\pi ka/n), \quad 0 \leq k \leq n/2, \\ \mathbf{y}_k(a) &= \sin(2\pi ka/n), \quad 1 \leq k \leq \lceil n/2 - 1 \rceil. \end{aligned}$$

The eigenvectors \mathbf{x}_k and \mathbf{y}_k have eigenvalue

$$\mu_k = 2 - 2 \cos(2\pi k/n).$$

Proof. We prove, via direct calculation, the result for \mathbf{y}_k . A similar calculation works for \mathbf{x}_k , which can be found in [1, Chapter 6.5]. We are going to use the following trig identity:

$$\sin(\alpha \pm \beta) = \sin \alpha \cos \beta \pm \cos \alpha \sin \beta,$$

which implies

$$\sin(\alpha - \beta) + \sin(\alpha + \beta) = 2 \sin \alpha \cos \beta.$$

Now we make the following calculation for each $a \in V$:

$$\begin{aligned}
\mathbf{L}_{C_n} \mathbf{y}_k(a) &= 2\mathbf{y}_k(a) - \mathbf{y}_k(a-1 \bmod n) - \mathbf{y}_k(a+1 \bmod n) \\
&= 2 \sin\left(\frac{2\pi ka}{n}\right) - \sin\left(\frac{2\pi k(a-1)}{n}\right) - \sin\left(\frac{2\pi k(a+1)}{n}\right) \\
&= 2 \sin\left(\frac{2\pi ka}{n}\right) - \sin\left(\frac{2\pi ka}{n} - \frac{2\pi k}{n}\right) - \sin\left(\frac{2\pi ka}{n} + \frac{2\pi k}{n}\right) \\
&= 2 \sin\left(\frac{2\pi ka}{n}\right) - 2 \sin\left(\frac{2\pi ka}{n}\right) \cos\left(\frac{2\pi k}{n}\right) \\
&= \left(2 - 2 \cos\left(\frac{2\pi k}{n}\right)\right) \sin\left(\frac{2\pi ka}{n}\right) \\
&= (2 - 2 \cos(2\pi k/n)) \mathbf{y}_k(a).
\end{aligned}$$

□

Remark 12. If we order the eigenvalues of \mathbf{L}_{C_n} in increasing order, $0 = \lambda_1 < \lambda_2 \leq \lambda_3 \leq \dots \leq \lambda_n$ we have:

$$\begin{aligned}
\lambda_1 &= 0 \quad \text{with} \quad \boldsymbol{\psi}_1(a) = \mathbf{x}_0(a) = \mathbf{1}(a) \\
\lambda_2 &= 2 - 2 \cos(2\pi/n) \quad \text{with} \quad \boldsymbol{\psi}_2(a) = \mathbf{x}_1(a) = \cos(2\pi a/n) \\
\lambda_3 &= 2 - 2 \cos(2\pi/n) \quad \text{with} \quad \boldsymbol{\psi}_3(a) = \mathbf{y}_1(a) = \sin(2\pi a/n),
\end{aligned}$$

and $\lambda_4 > \lambda_3$. Thus the eigenvector embedding of the cycle graph is

$$a \mapsto (\boldsymbol{\psi}_2(a), \boldsymbol{\psi}_3(a)) = (\cos(2\pi a/n), \sin(2\pi a/n)).$$

This is a really good embedding, as it gives evenly spaced samples on the unit circle! See also Exercise 5 from Homework 02.

16 The path graph

Now that we have computed the eigenvalues and eigenvectors of the cycle graph, C_n , we are in a good position to compute the eigenvalues and eigenvectors of the path graph, P_n .

Theorem 18. *The graph Laplacian of the path graph, P_n , has eigenvalues*

$$\lambda_{k+1} = 2 - 2 \cos(\pi k/n), \quad 0 \leq k < n,$$

with eigenvectors

$$\boldsymbol{\psi}_{k+1}(a) = \cos\left(\frac{\pi ka}{n} - \frac{\pi k}{2n}\right), \quad 0 \leq k < n.$$

Proof. Notice the values λ_{k+1} are the eigenvalues of C_{2n} , the cycle graph on $2n$ vertices. This is not an accident and we will use C_{2n} to prove path graph result. The idea is to view P_n as the quotient of C_{2n} . In particular:

$$P_n = C_{2n} / \sim,$$

where \sim identifies $a \in V(C_{2n})$ with $2n + 1 - a \in V(C_{2n})$, yielding the vertex $a \in V(P_n)$. The edges $(1, 2n), (n, n + 1) \in E(C_{2n})$ are eliminated under this equivalence relation, and the edges $(a, a + 1), (2n - a, 2n + 1 - a) \in E(C_{2n})$ are identified as the same edge for $1 \leq a < n$, yielding the edge $(a, a + 1) \in E(P_n)$. See Figure 18 for a picture illustrating this equivalence relation.




Figure 18: Illustration of the equivalence relation $P_n = C_{2n} / \sim$ for $n = 4$. Pairs of lines of the same color identify two vertices of C_{2n} with one vertex in P_n .

Now consider the function $\phi_{k+1} : V(C_{2n}) \rightarrow \mathbb{R}$ defined as

$$\phi_{k+1}(a) := \cos \left(\frac{\pi k a}{n} - \frac{\pi k}{2n} \right), \quad 1 \leq a \leq 2n.$$

Notice that $\phi_{k+1}(a) = \psi_{k+1}(a)$ for $1 \leq a \leq n$. In fact, though, ϕ_{k+1} is well defined on P_n under the equivalence relation \sim (not just by restricting to $1 \leq a \leq n$). In other words,

$$\phi_{k+1}(a) = \phi_{k+1}(2n + 1 - a), \quad \forall 1 \leq a \leq n. \quad (20)$$

Indeed,

$$\begin{aligned}
\phi_{k+1}(2n+1-a) &= \cos\left(\frac{\pi k(2n+1-a)}{n} - \frac{\pi k}{2n}\right) \\
&= \cos\left(2\pi k + \frac{\pi k}{n} - \frac{\pi ka}{n} - \frac{\pi k}{2n}\right) \\
&= \cos\left(-\frac{\pi ka}{n} + \frac{\pi k}{2n}\right) \\
&= \cos\left(\frac{\pi ka}{n} - \frac{\pi k}{2n}\right) \\
&= \phi_{k+1}(a).
\end{aligned}$$

Now we will use the trig identity

$$\cos(\alpha \pm \beta) = \cos(\alpha)\cos(\beta) \mp \sin\alpha\sin\beta,$$

to write ϕ_{k+1} in terms of the eigenvectors \mathbf{x}_k , \mathbf{y}_k of $\mathbf{L}_{C_{2n}}$, which we computed in Theorem 17. We have

$$\begin{aligned}
\phi_{k+1}(a) &= \cos\left(\frac{2\pi ka}{2n} - \frac{\pi k}{2n}\right) \\
&= \cos\left(\frac{2\pi ka}{2n}\right)\cos\left(\frac{\pi k}{2n}\right) + \sin\left(\frac{2\pi ka}{2n}\right)\sin\left(\frac{\pi k}{2n}\right) \\
&= \cos\left(\frac{\pi k}{2n}\right)\mathbf{x}_k(a) + \sin\left(\frac{\pi k}{2n}\right)\mathbf{y}_k(a).
\end{aligned}$$

Thus ϕ_{k+1} is a linear combination of \mathbf{x}_k and \mathbf{y}_k . Since \mathbf{x}_k and \mathbf{y}_k have the same eigenvalue λ_{k+1} , this means that ϕ_{k+1} is an eigenvector of $\mathbf{L}_{C_{2n}}$ with eigenvalue λ_{k+1} as well, i.e.,

$$\mathbf{L}_{C_{2n}}\phi_{k+1} = \lambda_{k+1}\phi_{k+1}. \quad (21)$$

We can now show that ψ_{k+1} is an eigenvector of \mathbf{L}_{P_n} with eigenvalue λ_{k+1} . Using (20) and (21), we make the following calculation:

$$\begin{aligned}
\forall 1 < a < n, \quad \mathbf{L}_{P_n}\psi_{k+1}(a) &= 2\psi_{k+1}(a) - \psi_{k+1}(a-1) - \psi_{k+1}(a+1) \\
&= \frac{1}{2}\left(2\phi_{k+1}(a) - \phi_{k+1}(a-1) - \phi_{k+1}(a+1) + \dots \right. \\
&\quad \left. \dots + 2\phi_{k+1}(2n+1-a) - \phi_{k+1}(2n+1-(a-1)) - \phi_{k+1}(2n+1-(a+1))\right) \\
&= \frac{1}{2}\left(\mathbf{L}_{C_{2n}}\phi_{k+1}(a) + \mathbf{L}_{C_{2n}}\phi_{k+1}(2n+1-a)\right) \\
&= \frac{1}{2}\left(\lambda_{k+1}\phi_{k+1}(a) + \lambda_{k+1}\phi_{k+1}(2n+1-a)\right) \\
&= \lambda_{k+1}\psi_{k+1}(a).
\end{aligned}$$

For $a = 1$ we have:

$$\begin{aligned}\mathbf{L}_{P_n} \psi_{k+1}(1) &= \psi_{k+1}(1) - \psi_{k+1}(2) = 2\psi_{k+1}(1) - \psi_{k+1}(2) - \psi_{k+1}(1) \\ &= 2\phi_{k+1}(1) - \phi_{k+1}(2) - \phi_{k+1}(2n) \\ &= \mathbf{L}_{C_{2n}} \phi_{k+1}(1) = \lambda_{k+1} \phi_{k+1}(1) = \lambda_{k+1} \psi_{k+1}(1).\end{aligned}$$

A similar calculation works for $a = n$ as well. □

Remark 13. In Sections 12 and 14 we estimated the second eigenvalue of \mathbf{L}_{P_n} as

$$\frac{6}{(n-1)(n+1)} \leq \lambda_2(P_n) \leq \frac{12}{n(n+1)}.$$

Theorem 18 proves this is a very good estimate. Indeed, recall the Taylor series of $\cos(u)$ is

$$\cos(u) = 1 - \frac{u^2}{2} + O(u^4).$$

Thus,

$$\lambda_2(P_n) = 2(1 - \cos(\pi/n)) = 2 \left(1 - 1 + \frac{1}{2} \frac{\pi^2}{n^2} + O(n^{-4}) \right) = \frac{\pi^2}{n^2} + O(n^{-4}).$$

Remark 14. Recall in Figure 9 and Figure 10 we made the empirical observation that the eigenvectors of the path graph on 10 vertices increase in frequency with increasing eigenvalue. This theorem proves that observation was not an accident, and that it holds for any path graph P_n . Indeed, since

$$\psi_k(a) = \cos \left(\frac{\pi(k-1)a}{n} - \frac{\pi(k-1)}{2n} \right), \quad 1 \leq k \leq n,$$

we see the frequency of ψ_k is $\pi(k-1)/n$, which increases as k increases. Using this formula for ψ_k one can show that if there is no vertex $a \in V_{P_n}$ for which $\psi_k(a) = 0$, then there are exactly $k-1$ edges $(a, b) \in E_{P_n}$ for which $\psi_k(a)\psi_k(b) < 0$, which corresponds to when ψ_k changes sign over an edge. In fact, we will be able to generalize this result to weighted path graphs, and we will even be able to formulate an analogous result for general graphs. These results will clarify what we mean by the frequency of an eigenvector of the graph Laplacian for a general graph, and they will show that the frequency of these eigenvectors increases with increasing eigenvalue. We will work towards this result over the next few lectures, along the way covering items of independent interest as well.

Lecture 08: The Adjacency Matrix and Eigenvalue Interlacing

February 11, 2021

Lecturer: Matthew Hirn

17 The adjacency matrix

Thus far in the course we have taken the point of view that the adjacency matrix is a “spreadsheet” and that the matrix of real interest is the graph Laplacian, which can be viewed as an operator and/or quadratic form. In this section we revisit the adjacency matrix and show that it has some interesting properties as well.

Let $G = (V, E, w)$ be a weighted graph and let \mathbf{M} be its adjacency matrix. As an operator, \mathbf{M} acts as follows:

$$\mathbf{M}\mathbf{x}(a) = \sum_{b \in N(a)} w(a, b)\mathbf{x}(b),$$

that is $\mathbf{M}\mathbf{x}(a)$ replaces $\mathbf{x}(a)$ with a weighted sum of the values of $\mathbf{x}(b)$ for b in the neighborhood of a . This is actually a pretty natural thing to do in machine learning problems on graphs, for example in node classification. Suppose you do not know $\mathbf{x}(a)$ but you do know $\mathbf{x}(b)$ for all $b \in N(a)$; a pretty good guess for $\mathbf{x}(a)$ then might be to take a weighted average of the values of $\mathbf{x}(b)$ for $b \in N(a)$, which is exactly what $\mathbf{M}\mathbf{x}(a)$ does.

We also remark that the quadratic form of \mathbf{M} is:

$$\mathbf{x}^T \mathbf{M} \mathbf{x} = \sum_{a \in V} \sum_{b \in N(a)} w(a, b)\mathbf{x}(a)\mathbf{x}(b) = 2 \sum_{(a, b) \in E} w(a, b)\mathbf{x}(a)\mathbf{x}(b).$$

Since \mathbf{M} is real valued and symmetric, as an operator, \mathbf{M} has n eigenvalues and n orthonormal eigenvectors,

$$\mathbf{M}\phi_i = \mu_i \phi_i, \quad 1 \leq i \leq n.$$

We will order the eigenvalues in decreasing order,

$$\mu_1 \geq \mu_2 \geq \cdots \geq \mu_n.$$

In doing so, μ_i will correspond to λ_i , the i^{th} eigenvalue of \mathbf{L} , since we have ordered them as $\lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_n$. Indeed, consider the case of a d -regular graph. Then $\mathbf{D} = d\mathbf{I}$ and so

$$\mathbf{L} = \mathbf{D} - \mathbf{M} = d\mathbf{I} - \mathbf{M}.$$

In this case we have

$$\lambda_i = d - \mu_i,$$

and we see, for example, that $\mu_1 = d$ since $\lambda_1 = 0$, and furthermore that $\phi_1 = \psi_1 = \mathbf{1}$.

But this required G being d -regular, which most graphs are not. For a general graph $G = (V, E)$, let

$$d_{\max} = d_{\max}(G) = \max_{a \in V} \deg(a),$$

and define

$$d_{\text{avg}} = d_{\text{avg}}(G) = \frac{1}{|V|} \sum_{a \in V} \deg(a).$$

For an arbitrary graph we can bound μ_1 using these two quantities.

Theorem 19. *Let $G = (V, E)$ be a graph. Then*

$$d_{\text{avg}} \leq \mu_1 \leq d_{\max}.$$

Proof. From the Courant-Fischer Theorem (Theorem 3) we know that

$$\mu_1 = \max_{\substack{\mathbf{x} \in \mathbb{R}^n \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T \mathbf{M} \mathbf{x}}{\mathbf{x}^T \mathbf{x}}.$$

Let us use the test vector $\mathbf{x} = \mathbf{1}$. We have:

$$\mu_1 \geq \frac{\mathbf{1}^T \mathbf{M} \mathbf{1}}{\mathbf{1}^T \mathbf{1}} = \frac{\sum_{a \in V} \sum_{b \in N(a)} \mathbf{1}(a) \mathbf{1}(b)}{n} = \frac{1}{n} \sum_{a \in V} \deg(a) = d_{\text{avg}}.$$

To prove the upper bound, let ϕ_1 be the eigenvector of \mathbf{M} with eigenvalue μ_1 . Let $a \in V$ be the vertex at which ϕ_1 takes its maximum value, so that

$$\phi_1(a) \geq \phi_1(b), \quad b \in V.$$

We may assume that $\phi_1(a) > 0$; if not, replace ϕ_1 with $-\phi_1$ if ϕ_1 is strictly negative. We have

$$\mu_1 = \frac{\mathbf{M} \phi_1(a)}{\phi_1(a)} = \frac{\sum_{b \in N(a)} \phi_1(b)}{\phi_1(a)} = \sum_{b \in N(a)} \frac{\phi_1(b)}{\phi_1(a)} \leq \sum_{b \in N(a)} 1 = \deg(a) \leq d_{\max}. \quad (22)$$

□

It turns out that if $\mu_1 = d_{\max}$, then G must be d_{\max} -regular. Here is the result.

Corollary 20. *If G is connected and $\mu_1 = d_{\max}$, then G is d_{\max} -regular.*

Proof. If $\mu_1 = d_{\max}$ then we must have equality in the inequalities in (22). Thus we must have

$$\deg(a) = d_{\max},$$

and

$$\phi_1(b) = \phi_1(a), \quad \forall b \in N(a).$$

Thus ϕ_1 takes its maximum value at a and all $b \in N(a)$. Now we can apply equation (22) to each $b \in N(a)$. Since $\mu_1 = d_{\max}$, we again must have equality where there are inequalities, which means $\deg(b) = d_{\max}$ for all $b \in N(a)$ and $\phi_1(c) = \phi_1(b)$ for all $c \in N(b)$ and all $b \in N(a)$. If we keep repeating the argument, since G is connected, we will conclude that $\deg(v) = d_{\max}$ for all $v \in V$. \square

18 Eigenvalue interlacing

We now state and prove some useful results regarding the eigenvalues of a real-valued, symmetric matrix and the eigenvalues of some of its sub-matrices. This will be immediately useful for strengthening Theorem 19.

For a weighted graph $G = (V, E, w)$ and subset $S \subset V$ we define the *subgraph induced by S* , written as $G(S)$, to be:

$$\begin{aligned} G(S) &:= (S, E(S), w|_{E(S)}), \\ E(S) &:= \{(a, b) \in E : a, b \in S\}. \end{aligned} \tag{23}$$

Relatedly, for a symmetric matrix \mathbf{A} whose rows and columns are indexed by a vertex set V , for $S \subset V$ we write $\mathbf{A}(S)$ to be the symmetric sub-matrix obtained by keeping only the rows and columns indexed by S . We remark that

$$\mathbf{M}_G(S) = \mathbf{M}_{G(S)},$$

but on the other hand, in general $\mathbf{D}_G(S) \neq \mathbf{D}_{G(S)}$ and thus $\mathbf{L}_G(S) \neq \mathbf{L}_{G(S)}$.

Theorem 21 (Cauchy's Interlacing Theorem). *Let \mathbf{A} be an $n \times n$ real-valued, symmetric matrix and let \mathbf{B} be an $(n-1) \times (n-1)$ sub-matrix of \mathbf{A} obtained by deleting the same row and column of \mathbf{A} . Let $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_n$ be the eigenvalues of \mathbf{A} and let $\beta_1 \geq \beta_2 \geq \dots \geq \beta_{n-1}$ be the eigenvalues of \mathbf{B} . Then:*

$$\alpha_1 \geq \beta_1 \geq \alpha_2 \geq \beta_2 \geq \dots \geq \alpha_{n-1} \geq \beta_{n-1} \geq \alpha_n.$$

Proof. Without loss of generality we may assume that \mathbf{B} is obtained from \mathbf{A} by removing the first row and column of \mathbf{A} . We will use the Courant-Fischer Theorem (Theorem 3) to prove this theorem. Applying it to \mathbf{B} we have:

$$\beta_k = \max_{\substack{S \subseteq \mathbb{R}^{n-1} \\ \dim(S)=k}} \min_{\substack{\mathbf{y} \in S \\ \mathbf{y} \neq \mathbf{0}}} \frac{\mathbf{y}^T \mathbf{B} \mathbf{y}}{\mathbf{y}^T \mathbf{y}} = \max_{\substack{S \subseteq \mathbb{R}^{n-1} \\ \dim(S)=k}} \min_{\substack{\mathbf{y} \in S \\ \mathbf{y} \neq \mathbf{0}}} \frac{\begin{pmatrix} 0 \\ \mathbf{y} \end{pmatrix}^T \mathbf{A} \begin{pmatrix} 0 \\ \mathbf{y} \end{pmatrix}}{\mathbf{y}^T \mathbf{y}}. \tag{24}$$

Define \mathcal{S} as the set of sets of the following form:

$$\mathcal{S} = \{S \subseteq \mathbb{R}^n : \dim(S) = k \text{ and every } \mathbf{x} \in S \text{ is of the form } \mathbf{x} = (0 \ \mathbf{y})^T \text{ for some } \mathbf{y} \in \mathbb{R}^{n-1}\}.$$

The right hand side of (24) can then be written as

$$(24) = \max_{S \in \mathcal{S}} \min_{\substack{\mathbf{x} \in S \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{\mathbf{x}^T \mathbf{x}} \leq \max_{\substack{S \subseteq \mathbb{R}^n \\ \dim(S)=k}} \min_{\substack{\mathbf{x} \in S \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{\mathbf{x}^T \mathbf{x}} = \alpha_k.$$

A similar proof works to prove that $\beta_k \geq \alpha_{k+1}$. Indeed, applying the other version of the Courant-Fischer Theorem we have

$$\begin{aligned} \beta_k &= \min_{\substack{T \subseteq \mathbb{R}^{n-1} \\ \dim(T)=n-1-k+1}} \max_{\substack{\mathbf{y} \in T \\ \mathbf{y} \neq \mathbf{0}}} \frac{\mathbf{y}^T \mathbf{B} \mathbf{y}}{\mathbf{y}^T \mathbf{y}} \\ &= \min_{\substack{T \subseteq \mathbb{R}^{n-1} \\ \dim(T)=n-k}} \max_{\substack{\mathbf{y} \in T \\ \mathbf{y} \neq \mathbf{0}}} \frac{\begin{pmatrix} 0 \\ \mathbf{y} \end{pmatrix}^T \mathbf{A} \begin{pmatrix} 0 \\ \mathbf{y} \end{pmatrix}}{\mathbf{y}^T \mathbf{y}} \geq \min_{\substack{T \subseteq \mathbb{R}^n \\ \dim(T)=n-k}} \max_{\substack{\mathbf{x} \in T \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{\mathbf{x}^T \mathbf{x}} = \alpha_{k+1}. \end{aligned}$$

□

Using Theorem 21 we have the following corollary.

Corollary 22. *Let $G = (V, E, w)$ and let $S \subseteq V$ with $|S| = n - k$. Then*

$$\mu_i(G) \geq \mu_i(G(S)), \quad 1 \leq i \leq n - k.$$

Proof. We will do a proof by induction. For the base cases we will take $k = 0, 1$. When $k = 0$ we have $S = V$ and the result is obvious. When $k = 1$ then $|S| = n - 1$ so $\mathbf{M}_{G(S)} = \mathbf{M}_G(S)$ removes the same row and column from \mathbf{M}_G . Therefore we can apply the Cauchy Interlacing Theorem (Theorem 21) and the result follows.

For the inductive step, assume the result is true for $0 \leq k \leq \ell$ and let us prove it for $k = \ell + 1$. Let $S \subset V$ be any subset with $|S| = n - (\ell + 1)$. For any such subset, there exists a subset \tilde{S} such that $S \subset \tilde{S}$ and $|\tilde{S}| = n - \ell$. Thus by the inductive hypothesis

$$\mu_i(G) \geq \mu_i(G(\tilde{S})).$$

On the other hand, we also have

$$\mathbf{M}_{G(S)} = \mathbf{M}_{G(\tilde{S})}(S),$$

and S only removes one row/column from \tilde{S} , so we can apply the Cauchy Interlacing Theorem to conclude that

$$\mu_i(G) \geq \mu_i(G(\tilde{S})) \geq \mu_i(G(S)).$$

□

Now we can use these results to improve our estimate of μ_1 .

Theorem 23. *Let $G = (V, E)$ be a graph. Then,*

$$\mu_1(G) \geq d_{\text{avg}}(G(S)), \quad \forall S \subseteq V.$$

Proof. By Corollary 22 we know that

$$\mu_1(G) \geq \mu_1(G(S)).$$

By Theorem 19 we have:

$$\mu_1(G(S)) \geq d_{\text{avg}}(G(S)).$$

□

19 Perron-Frobenius Theory for symmetric matrices

We will next prove the Perron-Frobenius Theorem for the adjacency matrix of a weighted, connected graph G . It will state, among other things, that even though ϕ_1 does not have to be the constant vector, it does have to be strictly positive. Let us first make the following remark on μ_n , the smallest eigenvalue of \mathbf{M} .

Remark 15. The eigenvalue μ_n must be negative so long as G has at least one edge. Indeed, by Theorem 19, we know that $\mu_1 \geq d_{\text{avg}}$. We also know, from Exercise 2 of Homework 2, that

$$0 = \text{Tr}(\mathbf{M}) = \sum_{i=1}^n \mu_i.$$

Theorem 24 (Perron-Frobenius Theorem). *Let $G = (V, E, w)$ be a connected graph, let \mathbf{M} be the adjacency matrix of G , and let $\mu_1 \geq \mu_2 \geq \dots \geq \mu_n$ be the eigenvalues of \mathbf{M} . Then*

1. *One can take $\phi_1(a) > 0$ for all $a \in V$*
2. $\mu_1 \geq -\mu_n$
3. $\mu_1 > \mu_2$

To prove Theorem 24 we will need two other results. The first one deals with maximizing and minimizing the Rayleigh quotient. We know from the Courant-Fischer Theorem (Theorem 3) that the maximum is the largest eigenvalue and the minimum is the smallest eigenvalue, and that one way to achieve these values is by evaluating the Rayleigh quotient at the corresponding eigenvector. It turns out this is the only way to achieve these values.

Theorem 25. Let \mathbf{A} be a real-valued, symmetric matrix and let μ_1 be the largest eigenvalue of \mathbf{A} . If

$$\mu_1 = \frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{\mathbf{x}^T \mathbf{x}},$$

for some $\mathbf{x} \in \mathbb{R}^n$, then $\mathbf{A} \mathbf{x} = \mu_1 \mathbf{x}$. The same result holds for μ_n , the minimum eigenvalue of \mathbf{A} .

Proof. By the Courant-Fischer Theorem (Theorem 3) we know that

$$\mu_1 = \max_{\substack{\mathbf{y} \in \mathbb{R}^n \\ \mathbf{y} \neq \mathbf{0}}} R(\mathbf{y}) = \max_{\substack{\mathbf{y} \in \mathbb{R}^n \\ \mathbf{y} \neq \mathbf{0}}} \frac{\mathbf{y}^T \mathbf{A} \mathbf{y}}{\mathbf{y}^T \mathbf{y}},$$

where $R(\mathbf{y})$ is the Rayleigh quotient of \mathbf{y} . Since $R(\mathbf{x})$ is the maximum of $R : \mathbb{R}^n \rightarrow \mathbb{R}$, we know that

$$\nabla R(\mathbf{x}) = \mathbf{0}.$$

Thus let us compute the gradient of R . We first note that

$$\mathbf{y}^T \mathbf{y} = \sum_{b \in V} \mathbf{y}(b)^2,$$

and so

$$\frac{\partial}{\partial \mathbf{y}(a)} \mathbf{y}^T \mathbf{y} = \frac{\partial}{\partial \mathbf{y}(a)} \sum_{b \in V} \mathbf{y}(b)^2 = 2\mathbf{y}(a).$$

Thus

$$\nabla(\mathbf{y}^T \mathbf{y}) = 2\mathbf{y}.$$

Similarly,

$$\nabla(\mathbf{y}^T \mathbf{A} \mathbf{y}) = 2\mathbf{A} \mathbf{y}.$$

Therefore,

$$\nabla R(\mathbf{y}) = \frac{(\mathbf{y}^T \mathbf{y})(2\mathbf{A} \mathbf{y}) - (\mathbf{y}^T \mathbf{A} \mathbf{y})(2\mathbf{y})}{(\mathbf{y}^T \mathbf{y})^2}.$$

Thus $\nabla R(\mathbf{x}) = \mathbf{0}$ implies that

$$(\mathbf{x}^T \mathbf{x}) \mathbf{A} \mathbf{x} = (\mathbf{x}^T \mathbf{A} \mathbf{x}) \mathbf{x} \implies \mathbf{A} \mathbf{x} = \frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{\mathbf{x}^T \mathbf{x}} \mathbf{x} = \mu_1 \mathbf{x}.$$

The proof for μ_n is similar. □

Lecture 09: Homework 03 Solutions

February 16, 2021

Lecturer: Matthew Hirn

Exercise 2. Recall the barbell graph $B_{m,n}$ depicted in Figure 19.




Figure 19: The barbell graph, $B_{m,n}$.

For Exercise 2 of Homework 3 (finding an upper bound for $\lambda_2(B_{m,n})$), we can first do some numerical investigation. Let us implement a function to return the adjacency matrix of the barbell graph $B_{m,n}$ so that we can compute its second eigenvector. Here in Figure 20 is a plot of ψ_2 for $B_{4,6}$.




Figure 20: Plot of the ψ_2 of $B_{4,6}$, with the path vertices in the middle, one complete graph on the left, and one complete graph on the right.

We see that ψ_2 is nearly constant on each K_n , with the exception of the vertices shared by P_m . We also see that ψ_2 is nearly linear on P_m . Therefore we propose to use a test vector

\mathbf{x} that is constant on each K_n and on P_m is the same test vector we used when analyzing the path graph. If we let the first m vertices of $B_{m,n}$ be those vertices on the path graph, then we have

$$\mathbf{x}(a) := \begin{cases} m-1 & a \in K_n \text{ (for the } K_n \text{ that includes } a=1) \\ (m+1)-2a & a \in P_m \\ -(m-1) & a \in K_n \text{ for the } K_n \text{ that includes } a=m \end{cases}.$$

See Figure 21 for a comparison between ψ_2 and \mathbf{x} ; one can see they are in good agreement.




Figure 21: The test vector \mathbf{x} , in orange, for the barbell graph $B_{4,6}$, with ψ_2 (rescaled) in blue for comparison.

Now we can get an upper bound on $\lambda_2(B_{m,n})$ using \mathbf{x} and computing:

$$\lambda_2(B_{m,n}) \leq \frac{\mathbf{x}^T \mathbf{L}_{B_{m,n}} \mathbf{x}}{\mathbf{x}^T \mathbf{x}}.$$

Computing the right hand side of the above gives the desired upper bound.

Exercise 3. The lower bound follows from this more general result. Let $G = (V, E)$ be a connected graph. For any $a, b \in V$, let $P_{a,b}(G)$ be a shortest path from a to b in G . Define $\text{len}(P_{a,b}(G))$, the length of $P_{a,b}(G)$, to be the number of edges on $P_{a,b}(G)$. Define the *diameter* of G as:

$$\text{diam}(G) := \max\{\text{len}(P_{a,b}(G)) : a, b \in V\}.$$

One can then show

$$\lambda_2(G) \geq \frac{2}{(n-1)\text{diam}(G)}.$$

Here is the proof.

Proof. Let $a, b \in V$. Using the path inequality (Theorem 16), as well as Theorem 14, we have:

$$G_{a,b} \preceq \text{len}(P_{a,b}(G))P_{a,b}(G) \preceq \text{diam}(G)P_{a,b}(G) \preceq \text{diam}(G)G.$$

Now summing over all $1 \leq a < b \leq n$, we obtain:

$$K_n = \sum_{1 \leq a < b \leq n} G_{a,b} \frac{n(n-1)}{2} \text{diam}(G) G.$$

Applying Theorem 15, we have:

$$n = \lambda_2(K_n) \leq \frac{n(n-1)}{2} \text{diam}(G) \lambda_2(G) \implies \lambda_2(G) \geq \frac{2}{(n-1) \text{diam}(G)}.$$

□

Lecture 09: Perron-Frobenius Theory

February 16, 2021

Lecturer: Matthew Hirn

Our goal in this lecture is to prove Perron-Frobenius Theorem, which we stated in the last lecture for the adjacency matrix \mathbf{M} of a graph G , but which in fact has a more general statement for symmetric, non-negative matrices. We will provide the more general statement shortly.

First, let \mathbf{A} be an $n \times n$ symmetric matrix with non-negative entries. Define $\text{diag}(\mathbf{A})$ as the $n \times n$ matrix with the diagonal of \mathbf{A} on its diagonal, and zeroes elsewhere. Note that we can decompose \mathbf{A} as:

$$\mathbf{A} = \text{diag}(\mathbf{A}) + \mathbf{M},$$

where \mathbf{M} is an $n \times n$, symmetric matrix with zeros down its diagonal. We use the notation \mathbf{M} because, indeed, \mathbf{M} defines an adjacency matrix of a weighted graph $G(\mathbf{A}) = (V, E, w)$ with vertices $V = \{1, \dots, n\}$, edges

$$E = \{(a, b) \in V \times V : \mathbf{M}(a, b) > 0\},$$

and weights $w(a, b) = \mathbf{M}(a, b)$. We refer to $G(\mathbf{A})$ as the *graph induced by \mathbf{A}* . We are now ready to state the more general version of the Perron-Frobenius Theorem.

Theorem 26 (Perron-Frobenius Theorem, more general version). *Let \mathbf{A} be an $n \times n$ symmetric matrix with non-negative entries, with eigenvalues $\mu_1 \geq \mu_2 \geq \dots \geq \mu_n$, and with corresponding orthonormal eigenvectors $\phi_1, \phi_2, \dots, \phi_n$. If the graph $G(\mathbf{A})$ induced by \mathbf{A} is connected, then*

1. *One can take $\phi_1(a) > 0$ for all $a \in V$*
2. *$\mu_1 \geq |\mu_n|$*
3. *$\mu_1 > \mu_2$*

As mentioned last time, we will need two results to prove Theorem 24. The first of these was the following theorem.

Theorem 25. *Let \mathbf{A} be a real-valued, symmetric matrix and let μ_1 be the largest eigenvalue of \mathbf{A} . If*

$$\mu_1 = \frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{\mathbf{x}^T \mathbf{x}},$$

for some $\mathbf{x} \in \mathbb{R}^n$, then $\mathbf{A} \mathbf{x} = \mu_1 \mathbf{x}$. The same result holds for μ_n , the minimum eigenvalue of \mathbf{A} .

We proved Theorem 25 in the last lecture. The second result we will need is the following lemma.

Lemma 27. *Let \mathbf{A} be an $n \times n$ symmetric matrix with non-negative entries such that the graph $G(\mathbf{A})$ induced by \mathbf{A} is connected. Let ϕ be an eigenvector of \mathbf{A} such that $\phi(a) \geq 0$ for all $a \in V$. Then $\phi(a) > 0$ for all $a \in V$.*

Proof. Suppose ϕ is not strictly positive. Then there exists some vertex $a \in V$ for which $\phi(a) = 0$. Since $G = G(\mathbf{A}) = (V, E, w)$ is connected, there must also be a vertex $b \in V$ (possibly with $b = a$) such that $\phi(b) = 0$ and for some $c \in N(b)$, $\phi(c) > 0$. Let μ be the eigenvalue of ϕ . Then:

$$0 = \mu\phi(b) = \mathbf{A}\phi(b) = \mathbf{A}(b, b)\phi(b) + \sum_{v \in N(b)} \mathbf{A}(b, v)\phi(v) \geq \mathbf{A}(b, c)\phi(c) > 0,$$

since $(b, c) \in E$ implies that $\mathbf{A}(b, c) > 0$. However, this is a contradiction. \square

Now we can prove the Perron-Frobenius Theorem.

Proof of Theorem 24. Let $G = G(\mathbf{A}) = (V, E, w)$ be the graph induced by \mathbf{A} .

First part: Let ϕ_1 be an eigenvector of \mathbf{A} with eigenvalue μ_1 and with unit norm, $\|\phi_1\| = 1$. Set

$$\mathbf{y}(a) := |\phi_1(a)|, \quad a \in V.$$

Notice that $\|\mathbf{y}\| = \|\phi_1\| = 1$ as well, so for both vectors their Rayleigh quotient has denominator equal to one. We have:

$$\begin{aligned} \mu_1 &= \phi_1^T \mathbf{A} \phi_1 = \sum_{a \in V} \mathbf{A}(a, a) \phi_1(a)^2 + 2 \sum_{(a, b) \in E} \mathbf{A}(a, b) \phi_1(a) \phi_1(b) \\ &\leq \sum_{a \in V} \mathbf{A}(a, a) |\phi_1(a)|^2 + 2 \sum_{(a, b) \in E} \mathbf{A}(a, b) |\phi_1(a)| |\phi_1(b)| \\ &= \sum_{a \in V} \mathbf{A}(a, a) \mathbf{y}(a)^2 + 2 \sum_{(a, b) \in E} \mathbf{A}(a, b) \mathbf{y}(a) \mathbf{y}(b) \\ &= \mathbf{y}^T \mathbf{A} \mathbf{y}. \end{aligned}$$

But by the Courant-Fischer Theorem (Theorem 3),

$$\mu_1 = \max_{\substack{\mathbf{x} \in \mathbb{R}^n \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{\mathbf{x}^T \mathbf{x}},$$

which means that we must have $\mu_1 = \mathbf{y}^T \mathbf{A} \mathbf{y}$. But then, by Theorem 25, \mathbf{y} must be an eigenvector of \mathbf{A} . Finally, applying Lemma 27, we see that \mathbf{y} must be strictly positive.

Second part: Let ϕ_n be an eigenvector of \mathbf{A} with eigenvalue μ_n and with unit norm, $\|\phi_n\| = 1$. Similarly to the first part, set

$$\mathbf{y}(a) := |\phi_n(a)|, \quad a \in V.$$

We have:

$$\begin{aligned}
|\mu_n| &= |\phi_n^T \mathbf{A} \phi_n| = \left| \sum_{a \in V} \mathbf{A}(a, a) \phi_n(a)^2 + 2 \sum_{(a, b) \in E} \mathbf{A}(a, b) \phi_n(a) \phi_n(b) \right| \\
&\leq \sum_{a \in V} \mathbf{A}(a, a) |\phi_n(a)|^2 + 2 \sum_{(a, b) \in E} \mathbf{A}(a, b) |\phi_n(a)| |\phi_n(b)| \\
&= \sum_{a \in V} \mathbf{A}(a, a) \mathbf{y}(a)^2 + 2 \sum_{(a, b) \in E} \mathbf{A}(a, b) \mathbf{y}(a) \mathbf{y}(b) \\
&= \mathbf{y}^T \mathbf{A} \mathbf{y} \\
&\leq \mu_1.
\end{aligned}$$

Third part: Take ϕ_1 to be the normalized eigenvector of \mathbf{A} with eigenvalue μ_1 that is strictly positive (using part 1). Let ϕ_2 be an eigenvector of \mathbf{A} with eigenvalue μ_2 , unit norm, $\|\phi_2\| = 1$, and that is orthogonal to ϕ_1 . Following the same method as the previous two parts, set

$$\mathbf{y}(a) := |\phi_2(a)|, \quad a \in V.$$

We have that

$$\begin{aligned}
\mu_2 &= \phi_2^T \mathbf{A} \phi_2 = \sum_{a \in V} \mathbf{A}(a, a) \phi_2(a)^2 + 2 \sum_{(a, b) \in E} \mathbf{A}(a, b) \phi_2(a) \phi_2(b) \\
&\leq \sum_{a \in V} \mathbf{A}(a, a) |\phi_2(a)|^2 + 2 \sum_{(a, b) \in E} \mathbf{A}(a, b) |\phi_2(a)| |\phi_2(b)| \\
&= \sum_{a \in V} \mathbf{A}(a, a) \mathbf{y}(a)^2 + 2 \sum_{(a, b) \in E} \mathbf{A}(a, b) \mathbf{y}(a) \mathbf{y}(b) \\
&= \mathbf{y}^T \mathbf{A} \mathbf{y} \\
&\leq \mu_1.
\end{aligned}$$

Of course we already knew $\mu_1 \geq \mu_2$, but the important parts of the above calculation are the steps in between. Now suppose that $\mu_1 = \mu_2$. Then, by the above calculation, $\mathbf{y}^T \mathbf{A} \mathbf{y} = \mu_1$. Thus, by Theorem 25, \mathbf{y} must be an eigenvector of \mathbf{A} with eigenvalue μ_1 . Since \mathbf{y} is a non-negative eigenvector of \mathbf{A} , we can apply Lemma 27 to conclude that \mathbf{y} is strictly positive, which in turn implies that $\phi_2(a) \neq 0$ for all $a \in V$.

Now, since $\langle \phi_1, \phi_2 \rangle = 0$ and $\phi_1(a) > 0$ for all $a \in V$, it must be that ϕ_2 has some positive entries and some negative entries. Since G is connected and since $\phi_2(a) \neq 0$ for all $a \in V$, it must be that there exists an $(u, v) \in E$ such that

$$\phi_2(u) < 0 < \phi_2(v).$$

But if that is the case, then

$$\begin{aligned}\mu_2 &= \sum_{a \in V} \mathbf{A}(a, a) \phi_2(a)^2 + 2 \sum_{\substack{(a,b) \in E \\ (a,b) \neq (u,v)}} \mathbf{A}(a, b) \phi_2(a) \phi_2(b) + \underbrace{2\mathbf{A}(u, v) \phi_2(u) \phi_2(v)}_{<0} \\ &< \sum_{a \in V} \mathbf{A}(a, a) |\phi_2(a)|^2 + 2 \sum_{(a,b) \in E} \mathbf{A}(a, b) |\phi_2(a)| |\phi_2(b)| = \mu_1.\end{aligned}$$

Thus $\mu_2 < \mu_1$, which contradicts the assumption that $\mu_2 = \mu_1$; we thus conclude that indeed $\mu_2 < \mu_1$. \square

Remark 16. Let $\mu_1 \geq \mu_2 \geq \dots \geq \mu_n$ be the eigenvalues of the adjacency matrix, \mathbf{M} , of a graph $G = (V, E)$. One can show if G is connected, then $\mu_1 = -\mu_n$ if and only if G is bipartite. The details can be found in [1, Chapter 4]

We can extend the Perron-Frobenius Theorem to the graph Laplacian (this was already clear) and graph Laplacian like matrices. Here is the result.

Corollary 28. Let $\tilde{\mathbf{L}}$ be a symmetric, $n \times n$ matrix with non-positive off-diagonal entries, so that

$$\tilde{\mathbf{L}} = \text{diag}(\tilde{\mathbf{L}}) - \mathbf{M},$$

where \mathbf{M} is a symmetric, real-valued matrix with non-negative entries and zeroes down its diagonal. Suppose the graph G induced by \mathbf{M} is connected. Let $\tilde{\lambda}_1$ be the smallest eigenvalue of $\tilde{\mathbf{L}}$ with eigenvector $\tilde{\psi}_1$. Then $\tilde{\psi}_1$ may be taken to be strictly positive, and $\tilde{\lambda}_1$ has multiplicity one.

Proof. Set

$$\mathbf{A} = \sigma \mathbf{I} - \tilde{\mathbf{L}} = (\sigma \mathbf{I} - \text{diag}(\tilde{\mathbf{L}})) + \mathbf{M}.$$

If we set σ such that

$$\sigma \geq \max_{a \in V} \tilde{\mathbf{L}}(a, a),$$

then the matrix \mathbf{A} will be a symmetric matrix with non-negative entries such that the graph $G(\mathbf{A})$ is connected. We may therefore apply the Perron-Frobenius Theorem (the general version, Theorem 26) to the matrix \mathbf{A} , and conclude that its largest eigenvalue, μ_1 , has multiplicity one, and the eigenvector associated to μ_1 , ϕ_1 , may be taken to be strictly positive. But any eigenvector ϕ of \mathbf{A} is an eigenvector of $\tilde{\mathbf{L}}$, and if $\mathbf{A}\phi = \mu\phi$ then $\tilde{\mathbf{L}}\phi = (\sigma - \mu)\phi$, i.e., $\tilde{\lambda} = \sigma - \mu$ is the corresponding eigenvalue of $\tilde{\mathbf{L}}$. Thus $\tilde{\lambda}_1 = \sigma - \mu_1$ has multiplicity one. \square

Lecture 10: The Weighted Path Graph

February 18, 2021

Lecturer: Matthew Hirn

20 The weighted path graph

In Theorem 18 we proved that the eigenvectors of the path graph, P_n , have the formula

$$\psi_k(a) = \cos\left(\frac{\pi(k-1)a}{n} - \frac{\pi(k-1)}{2n}\right).$$

As such, one can show that ψ_k changes sign exactly $k-1$ times. We also observed this phenomenon, empirically, in Figures 9 and 10.

In fact, this result holds for weighted path graphs too! In Figure 22 we plot the 2nd, 3rd, and 4th eigenvectors of a weighted path graph with random weights and with $n = 10$ vertices; Figure 23 plots the 8th, 9th, and 10th eigenvectors of the same weighted path graph. We can see they are “deformed” versions of the unweighted path graph, but nevertheless, they change sign the same number of times.




Figure 22: The 2nd, 3rd, and 4th eigenvectors of a weighted path graph on $n = 10$ vertices.

Let us now work towards proving this result. We will begin with a simple result, which will help motivate the one that comes next.

Theorem 29. *If \mathbf{A} is an $n \times n$ positive semi-definite matrix, then so is \mathbf{BAB}^T for every $n \times n$ matrix \mathbf{B} .*

Proof. For any $\mathbf{x} \in \mathbb{R}^n$, we have

$$\mathbf{x}^T \mathbf{BAB}^T \mathbf{x} = (\mathbf{B}^T \mathbf{x})^T \mathbf{A} (\mathbf{B}^T \mathbf{x}) \geq 0,$$

since \mathbf{A} is positive semi-definite. □




Figure 23: The 8th, 9th, and 10th eigenvectors of a weighted path graph on $n = 10$ vertices.

Now we will generalize this result.

Theorem 30 (Sylvester's Law of Inertia). *Let \mathbf{A} be a real-valued, symmetric matrix and \mathbf{B} be an invertible matrix. Then the matrix \mathbf{BAB}^T has the same number of positive, negative, and zero eigenvalues as \mathbf{A} .*

Proof. We first note that \mathbf{A} and \mathbf{BAB}^T have the same rank (remember, \mathbf{B} is invertible), and so they have the same number of zero eigenvalues.

Let us prove that \mathbf{A} has at least as many positive eigenvalues as \mathbf{BAB}^T . To that end, let $\gamma_1 \geq \gamma_2 \geq \dots \geq \gamma_k > 0$ be the positive eigenvalues of \mathbf{BAB}^T with eigenvectors $\mathbf{y}_1, \dots, \mathbf{y}_k$. Let Y_k be the span of these eigenvectors,

$$Y_k := \text{span}\{\mathbf{y}_1, \dots, \mathbf{y}_k\}.$$

Now define S_k as

$$S_k := \text{span}\{\mathbf{B}^T \mathbf{y}_1, \dots, \mathbf{B}^T \mathbf{y}_k\}.$$

Since \mathbf{B} is invertible, $\dim(S_k) = k$. Let $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_n$ be the eigenvalues of \mathbf{A} . Using the Courant-Fischer Theorem (Theorem 3) and Lemma 4, we have

$$\alpha_k = \max_{\substack{S \subseteq \mathbb{R}^n \\ \dim(S)=k}} \min_{\substack{\mathbf{x} \in S \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{\mathbf{x}^T \mathbf{x}} \geq \min_{\substack{\mathbf{x} \in S_k \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{\mathbf{x}^T \mathbf{x}} = \min_{\substack{\mathbf{y} \in Y_k \\ \mathbf{y} \neq \mathbf{0}}} \frac{\mathbf{y}^T \mathbf{B} \mathbf{A} \mathbf{B}^T \mathbf{y}}{\mathbf{y}^T \mathbf{B} \mathbf{B}^T \mathbf{y}} = \min_{\substack{\mathbf{y} \in Y_k \\ \mathbf{y} \neq \mathbf{0}}} \frac{\sum_{i=1}^k \gamma_i |\langle \mathbf{y}, \mathbf{y}_i \rangle|^2}{\|\mathbf{B}^T \mathbf{y}\|^2} > 0.$$

Thus \mathbf{A} has at least k positive eigenvalues. Similarly, one can prove that \mathbf{A} has at least as many negative eigenvalues as \mathbf{BAB}^T . The proof is therefore complete. \square

Remark 17. In your homework you were asked to prove that \mathbf{BAB}^{-1} has the same eigenvalues as \mathbf{A} . In the above, the matrix \mathbf{BAB}^T is not similar to \mathbf{A} , rather it is congruent; that is, we have applied a change of basis to \mathbf{A} .

Before we can prove that the k^{th} eigenvector of a weighted path graph will change sign $k - 1$ times, we need one more lemma. In the following lemma we will consider weighted path graphs with *negative* edge weights, and their corresponding graph Laplacian matrices. While this may seem strange, we will see its usefulness shortly.

Lemma 31. *Let $G = (V, E, w)$ be a weighted path graph that can have negative (but not zero) weights $w(a-1, a)$ ($2 \leq a \leq n$), and let \mathbf{L} be the graph Laplacian of G :*

$$\mathbf{L} = \sum_{a=2}^n w(a-1, a) \mathbf{L}_{G_{a-1, a}}. \quad (25)$$

Then \mathbf{L} has one zero eigenvalue, the number of negative eigenvalues of \mathbf{L} equals the number of negative edge weights, and the number of positive eigenvalues of \mathbf{L} equals the number of positive edge weights.

Proof. Recall,

$$\mathbf{x}^T \mathbf{L} \mathbf{x} = \sum_{a=2}^n w(a-1, a) (\mathbf{x}(a) - \mathbf{x}(a-1))^2.$$

We will make a change of variables to diagonalize \mathbf{L} . To that end, for any $\mathbf{x} \in \mathbb{R}^n$ define $\mathbf{y} \in \mathbb{R}^n$ as

$$\mathbf{y}(a) = \begin{cases} \mathbf{x}(1) & a = 1 \\ \mathbf{x}(a) - \mathbf{x}(a-1) & 2 \leq a \leq n \end{cases}.$$

Clearly we have

$$\mathbf{x}(a) = \sum_{b=1}^a \mathbf{y}(b).$$

Therefore, if we set \mathbf{B} to be the lower triangular matrix with all ones on the diagonal and all ones in the lower triangle, i.e.,

$$\mathbf{B} = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ 1 & 1 & 0 & \cdots & 0 \\ 1 & 1 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & 1 & \cdots & 1 \end{pmatrix}$$

then

$$\mathbf{x} = \mathbf{B} \mathbf{y}.$$

Clearly \mathbf{B} has full rank, and so, by Sylvester's Law of Inertia (Theorem 30), we know that $\mathbf{B}^T \mathbf{L} \mathbf{B}$ has the same number of positive, negative, and zero eigenvalues as \mathbf{L} . Furthermore, since $\mathbf{y}(a) = \mathbf{x}(a) - \mathbf{x}(a-1)$ for $2 \leq a \leq n$ and $\mathbf{x} = \mathbf{B} \mathbf{y}$,

$$\mathbf{y}^T \mathbf{B}^T \mathbf{L} \mathbf{B} \mathbf{y} = \sum_{a=2}^n w(a-1, a) \mathbf{y}(a)^2.$$

We see that the quadratic form of $\mathbf{B}^T \mathbf{L} \mathbf{B}$ has no cross terms. Therefore, $\mathbf{B}^T \mathbf{L} \mathbf{B}$ is a diagonal matrix and

$$\mathbf{B}^T \mathbf{L} \mathbf{B}(a, a) = \begin{cases} 0 & a = 1 \\ w(a-1, a) & 2 \leq a \leq n \end{cases}.$$

Thus $\mathbf{B}^T \mathbf{L} \mathbf{B}$ has one zero eigenvalue and the weights of the path graph are its other eigenvalues. Since $\mathbf{B}^T \mathbf{L} \mathbf{B}$ has the same number of zero, negative, and positive eigenvalues as \mathbf{L} , the result follows immediately. \square

Now we are at last ready to state our result on the weighted path graph.

Theorem 32. *Let $G = (V, E, w)$ be a weighted path graph on n vertices. Let \mathbf{L} be its graph Laplacian with eigenvalues $0 = \lambda_1 < \lambda_2 < \dots < \lambda_n$, and let ψ_k be the k^{th} eigenvector of \mathbf{L} with eigenvalue λ_k . If $\psi_k(a) \neq 0$ for all $a \in V$, then there are exactly $k - 1$ edges $(a - 1, a) \in E$ for which $\psi_k(a - 1)\psi_k(a) < 0$.*

Proof. Let Ψ_k denote the diagonal $n \times n$ matrix with ψ_k on its diagonal. Consider the matrix:

$$\mathbf{A} = \Psi_k(\mathbf{L} - \lambda_k \mathbf{I})\Psi_k.$$

The matrix $\mathbf{L} - \lambda_k \mathbf{I}$ has one zero eigenvalue and $k - 1$ negative eigenvalues. Since $\psi_k(a) \neq 0$ for all $a \in V$, the matrix Ψ_k is invertible; note also, $\Psi_k^T = \Psi_k$. Thus, by Sylvester's Law of Inertia (Theorem 30), we know that \mathbf{A} has one zero eigenvalue and $k - 1$ negative eigenvalues as well.

I now claim that

$$\mathbf{A} = \sum_{a=2}^n w(a - 1, a) \psi_k(a - 1) \psi_k(a) \mathbf{L}_{G_{a-1, a}}. \quad (26)$$

If (26) is true, we are nearly finished. Indeed, compare (26) with (25). We see that \mathbf{A} is the graph Laplacian of a second weighted path graph with weights $w(a - 1, a) \psi_k(a - 1) \psi_k(a)$ over the edges in the path. Since \mathbf{A} has $k - 1$ negative eigenvalues, Lemma 31 says this second weighted path graph must have $k - 1$ negative edge weights. Since $w(a - 1, a) > 0$, it must be that $\psi_k(a - 1) \psi_k(a) < 0$ over $k - 1$ edges $(a - 1, a) \in E$. That completes the proof, aside from verifying (26), which we will do in the next lecture. \square

Lecture 11: Fiedler's Nodal Domain Theorem

February 23, 2021

Lecturer: Matthew Hirn

In the last lecture we partially proved the following result.

Theorem 32. *Let $G = (V, E, w)$ be a weighted path graph on n vertices. Let \mathbf{L} be its graph Laplacian with eigenvalues $0 = \lambda_1 < \lambda_2 < \dots < \lambda_n$, and let ψ_k be the k^{th} eigenvector of \mathbf{L} with eigenvalue λ_k . If $\psi_k(a) \neq 0$ for all $a \in V$, then there are exactly $k - 1$ edges $(a - 1, a) \in E$ for which $\psi_k(a - 1)\psi_k(a) < 0$.*

We now want to complete the proof. Recall we defined the matrix \mathbf{A} as

$$\mathbf{A} := \mathbf{\Psi}_k(\mathbf{L} - \lambda_k \mathbf{I})\mathbf{\Psi}_k,$$

where $\mathbf{\Psi}_k$ is the $n \times n$ matrix with ψ_k down its diagonal and zeros elsewhere. To complete the proof of Theorem 32, we need to show that

$$\mathbf{A} = \sum_{a=2}^n w(a-1, a) \psi_k(a-1) \psi_k(a) \mathbf{L}_{G_{a-1, a}},$$

which was labeled as equation (26) in the previous lecture. Let us do that now.

Proof of equation (26). We will prove the alternate formula for \mathbf{A} in two steps. First, we will show (26) holds for the off diagonal entries; then, we will show it holds for the diagonal entries. For the off diagonal entries we first note that

$$\mathbf{A} = \mathbf{\Psi}_k(\mathbf{L} - \lambda_k \mathbf{I})\mathbf{\Psi}_k = \mathbf{\Psi}_k \mathbf{L} \mathbf{\Psi}_k - \lambda_k \mathbf{\Psi}_k^2.$$

Since $\mathbf{\Psi}_k$ is a diagonal matrix, the off diagonal entries of \mathbf{A} are contained solely in $\mathbf{\Psi}_k \mathbf{L} \mathbf{\Psi}_k$. So, we will compute the entries of this matrix. First we compute:

$$\mathbf{\Psi}_k \mathbf{L}(u, v) = \sum_{c \in V} \mathbf{\Psi}_k(u, c) \mathbf{L}(c, v) = \mathbf{\Psi}_k(u, u) \mathbf{L}(u, v) = \psi_k(u) \mathbf{L}(u, v).$$

Then we have:

$$\mathbf{\Psi}_k \mathbf{L} \mathbf{\Psi}_k(u, v) = \sum_{c \in V} \mathbf{\Psi}_k \mathbf{L}(u, c) \mathbf{\Psi}_k(c, v) = \mathbf{\Psi}_k \mathbf{L}(u, v) \mathbf{\Psi}_k(v, v) = \psi_k(u) \psi_k(v) \mathbf{L}(u, v).$$

Therefore, since $\mathbf{L}(u, v) = -w(u, v)$ for $(u, v) \in E$ and $\mathbf{L}(u, v) = 0$ for $(u, v) \notin E$, we have,

$$\mathbf{A}(u, v) = \mathbf{\Psi}_k \mathbf{L} \mathbf{\Psi}_k(u, v) = \begin{cases} -w(u, v) \psi_k(u) \psi_k(v) & (u, v) \in E \\ 0 & (u, v) \notin E \end{cases}.$$

On the other hand, for $u \neq v$,

$$\mathbf{L}_{G_{a,b}}(u, v) = \begin{cases} -1 & (a, b) = (u, v) \\ 0 & (a, b) \neq (u, v) \end{cases},$$

and thus for $u \neq v$,

$$\begin{aligned} \sum_{(a,b) \in E} w(a, b) \psi_k(a) \psi_k(b) \mathbf{L}_{G_{a,b}}(u, v) &= \begin{cases} -w(u, v) \psi_k(u) \psi_k(v) & (u, v) \in E \\ 0 & (u, v) \notin E \end{cases} \\ &= \mathbf{A}(u, v). \end{aligned}$$

For the diagonal entries, we will show the row sums of \mathbf{A} are the same as the row sums of the right hand side of (26). Since we know the off diagonal entries are the same, this will show the diagonal entries are the same too. We compute the row sums of \mathbf{A} via:

$$\mathbf{A}\mathbf{1} = \Psi_k(\mathbf{L} - \lambda_k \mathbf{I})\Psi_k \mathbf{1} = \Psi_k(\mathbf{L} - \lambda_k \mathbf{I})\psi_k = \Psi_k(\mathbf{L}\psi_k - \lambda_k \psi_k) = \Psi_k(\lambda_k \psi_k - \lambda_k \psi_k) = \mathbf{0}.$$

On the other hand, $\mathbf{L}_{G_{a,b}} \mathbf{1} = \mathbf{0}$ for all $(a, b) \in E$, so the row sums of the right hand side of (26) are zeros too. The proof is complete. \square

Remark 18. In Theorem 32, we assumed λ_k has multiplicity one for each $1 \leq k \leq n$, but in fact one can prove this must be true, and thus this assumption can be removed.

Remark 19. We also assume that $\psi_k(a) \neq 0$ for all $a \in V$ in Theorem 32. We can also relax this assumption and show that the eigenvector changes sign $k - 1$ times over the path.

21 Nodal domains

We are now going to prove Fiedler's Nodal Domain Theorem, which will generalize the result on the weighted path graph to any connected, weighted graph! In doing so, our analysis and interpretation of what the “frequency” of an eigenvector means will be complete.

Let $G = (V, E, w)$ be a weighted, connected graph, and let us first define a nodal domain. For $S \subseteq V$, recall that $G(S)$ is the subgraph induced by S ; see equation (23) for the definition. For each eigenvector ψ_k of the graph Laplacian of G , $1 \leq k \leq n$, define the set $W_k \subseteq S$ as

$$W_k := \{a \in V : \psi_k(a) \geq 0\}.$$

The subgraph induced by W_k , $G(W_k)$, is the k^{th} *nodal domain* for the graph G . Note, it will depend on the choice of ψ_k , but we will be able to say something about the number of connected components of $G(W_k)$ irrespective the choice of ψ_k .

To gain some intuition, let us go back to the weighted path graph. We know its eigenvector ψ_k will change sign exactly $k - 1$ times. Thus ψ_2 will change sign once, and $G(W_2)$ will have one connected component. ψ_3 will change sign twice, and depending on the sign of ψ_3 , $G(W_3)$ will have either one or two connected components. Similarly, $G(W_4)$ will have two

connected components, and $G(W_5)$ will have two or three connected components. In other words, the number of connected components of $G(W_k)$ for the path graph will increase with the eigenvalue index. One may also wish to look back at Figures 9, 10, 22, and 23.

We see from the path graph example that measuring the number of connected components is a good proxy for the frequency of the eigenvector. Indeed, the more connected components $G(W_k)$ has, the more ψ_k must change its sign over the edges of G . For general weighted graphs, we will take the number of connected components of $G(W_k)$ as our measure of the frequency of ψ_k . Fiedler's Nodal Domain Theorem will show that the number of these connected components, and hence the frequency of ψ_k , is upper bounded by the index of the eigenvector. Its proof will leverage two of our previous big results, namely the Cauchy Eigenvalue Interlacing Theorem (Theorem 21) and the Perron-Frobenius Theorem (Theorem 26 and Corollary 28).

Theorem 33 (Fiedler's Nodal Domain Theorem). *Let $G = (V, E, w)$ be a connected graph. Then for $k \geq 2$, $G(W_k)$ has at most $k - 1$ connected components.*

Proof. Since $\psi_1 = \mathbf{1}$ and $\langle \psi_1, \psi_k \rangle = 0$ for all $k \geq 2$, we know that ψ_k must have positive and negative entries and $W_k \neq \emptyset$.

Suppose that $G(W_k)$ has t connected components. Let us order the vertices of G so that the vertices in one connected component of $G(W_k)$ come first, the vertices in another connected component of $G(W_k)$ come second, and so on so forth, and the vertices $a \in V$ for which $\psi_k(a) < 0$ come last. Then the graph Laplacian of G can be written as:

$$L = \begin{pmatrix} B_1 & 0 & 0 & \cdots & 0 & C_1 \\ 0 & B_2 & 0 & \cdots & 0 & C_2 \\ 0 & 0 & B_3 & \cdots & 0 & C_3 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & B_t & C_t \\ C_1^T & C_2^T & C_3^T & \cdots & C_t^T & A \end{pmatrix},$$

where:

- B_i encodes the vertices and edges in the i^{th} connected component of $G(W_k)$. Note that each B_i is symmetric with non-positive entries on the off-diagonal, and hence the Perron-Frobenius Theorem (specifically Corollary 28) applies to each one.
- C_i encodes the edges between the vertices of the i^{th} connected component of $G(W_k)$ and those vertices $a \in V$ for which $\psi_k(a) < 0$. Note that every entry of C_i is non-positive and at least one entry of C_i must be non-zero (thus negative) since otherwise G would not be connected.
- A encodes the edges between the vertices $a \in V$ for which $\psi_k(a) < 0$.

We can also decompose ψ_k using the same blocks of vertices:

$$\psi_k = \begin{pmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \\ \mathbf{x}_3 \\ \vdots \\ \mathbf{x}_t \\ \mathbf{y} \end{pmatrix},$$

where

- \mathbf{x}_i is a vector of length the number of vertices in the i^{th} connected component of $G(W_k)$, and the entries of \mathbf{x}_i are non-negative.
- \mathbf{y} is a vector of length the number of vertices $a \in V$ for which $\psi_k(a) < 0$, and every entry of \mathbf{y} is negative.

We will now show that the smallest eigenvalue of each \mathbf{B}_i is strictly less than λ_k , the eigenvalue of ψ_k . First note that we have:

$$\begin{pmatrix} \mathbf{B}_1 & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{C}_1 \\ \mathbf{0} & \mathbf{B}_2 & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{C}_2 \\ \mathbf{0} & \mathbf{0} & \mathbf{B}_3 & \cdots & \mathbf{0} & \mathbf{C}_3 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{B}_t & \mathbf{C}_t \\ \mathbf{C}_1^T & \mathbf{C}_2^T & \mathbf{C}_3^T & \cdots & \mathbf{C}_t^T & \mathbf{A} \end{pmatrix} \begin{pmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \\ \mathbf{x}_3 \\ \vdots \\ \mathbf{x}_t \\ \mathbf{y} \end{pmatrix} = \begin{pmatrix} \mathbf{B}_1\mathbf{x}_1 + \mathbf{C}_1\mathbf{y} \\ \mathbf{B}_2\mathbf{x}_2 + \mathbf{C}_2\mathbf{y} \\ \mathbf{B}_3\mathbf{x}_3 + \mathbf{C}_3\mathbf{y} \\ \vdots \\ \mathbf{B}_t\mathbf{x}_t + \mathbf{C}_t\mathbf{y} \\ \sum_{i=1}^t \mathbf{C}_i^T \mathbf{x}_i + \mathbf{A}\mathbf{y} \end{pmatrix} = \lambda_k \begin{pmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \\ \mathbf{x}_3 \\ \vdots \\ \mathbf{x}_t \\ \mathbf{y} \end{pmatrix},$$

and in particular,

$$\mathbf{B}_i\mathbf{x}_i + \mathbf{C}_i\mathbf{y} = \lambda_k\mathbf{x}_i, \quad 1 \leq i \leq t. \quad (27)$$

Now, every entry of \mathbf{C}_i is non-positive, at least one entry of \mathbf{C}_i is negative, and every entry of \mathbf{y} is negative; therefore, every entry of $\mathbf{C}_i\mathbf{y}$ is non-negative and at least one entry of $\mathbf{C}_i\mathbf{y}$ is positive. It follows that $\mathbf{x}_i \neq \mathbf{0}$, since if this were the case we would have $\mathbf{C}_i\mathbf{y} = \mathbf{0}$, which contradicts the fact that at least one entry of $\mathbf{C}_i\mathbf{y}$ is positive.

Furthermore, using the fact that every entry of \mathbf{x}_i is non-negative, plus our conclusions regarding $\mathbf{C}_i\mathbf{y}$, as well as (27), we have

$$\lambda_k\mathbf{x}_i(a) - \mathbf{C}_i\mathbf{y}(a) \leq \lambda_k\mathbf{x}_i(a) \implies \mathbf{B}_i\mathbf{x}_i(a) \leq \lambda_k\mathbf{x}_i(a), \quad \forall a.$$

We deduce that

$$\mathbf{x}_i^T \mathbf{B}_i \mathbf{x}_i \leq \lambda_k \mathbf{x}_i^T \mathbf{x}_i. \quad (28)$$

Now, by the Courant-Fischer Theorem (Theorem 3) and (28), we know that

$$\lambda_1(\mathbf{B}_i) = \min_{\mathbf{x} \neq \mathbf{0}} \frac{\mathbf{x}^T \mathbf{B}_i \mathbf{x}}{\mathbf{x}^T \mathbf{x}} \leq \frac{\mathbf{x}_i^T \mathbf{B}_i \mathbf{x}_i}{\mathbf{x}_i^T \mathbf{x}_i} \leq \lambda_k.$$

We have two cases:

1. Suppose \mathbf{x}_i has at least one zero entry and $\mathbf{x}_i^T \mathbf{B}_i \mathbf{x}_i = \lambda_k \mathbf{x}_i^T \mathbf{x}_i$ (if it is a strict inequality we done). Applying the Perron-Frobenius Theorem (Corollary 28), though, we conclude that λ_k cannot be $\lambda_1(\mathbf{B}_i)$ since the eigenvector associated to $\lambda_1(\mathbf{B}_i)$ cannot have any zero entries. Thus $\lambda_1(\mathbf{B}_i) < \lambda_k$.
2. Suppose \mathbf{x}_i has all positive entries. Since $\mathbf{C}_i \mathbf{y}$ is non-negative with at least one positive entry, we have that $\mathbf{x}_i^T \mathbf{C}_i \mathbf{y} > 0$. Therefore, using (27),

$$\mathbf{x}_i^T \mathbf{B}_i \mathbf{x}_i = \lambda_k \mathbf{x}_i^T \mathbf{x}_i - \mathbf{x}_i^T \mathbf{C}_i \mathbf{y} < \lambda_k \mathbf{x}_i^T \mathbf{x}_i,$$

and thus $\lambda_1(\mathbf{B}_i) < \lambda_k$.

We have proved that

$$\lambda_1(\mathbf{B}_i) < \lambda_k, \quad 1 \leq i \leq t.$$

It follows that the matrix

$$\mathbf{B} = \begin{pmatrix} \mathbf{B}_1 & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{B}_2 & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{B}_t \end{pmatrix},$$

has at least t eigenvalues less than λ_k . We now want to apply the Cauchy Interlacing Theorem to \mathbf{L} and the sub-matrix \mathbf{B} . However, \mathbf{B} is not a principal sub-matrix of \mathbf{L} so we need to verify that we can.

First, let $\tilde{\mathbf{A}}$ be an $n \times n$ matrix and let $\tilde{\mathbf{B}}$ be an $(n-1) \times (n-1)$ principal sub-matrix of $\tilde{\mathbf{A}}$. This time, let us order the eigenvalues of $\tilde{\mathbf{A}}$ and $\tilde{\mathbf{B}}$ in increasing order:

$$\lambda_1(\tilde{\mathbf{A}}) \leq \lambda_2(\tilde{\mathbf{A}}) \leq \cdots \leq \lambda_n(\tilde{\mathbf{A}}) \quad \text{and} \quad \lambda_1(\tilde{\mathbf{B}}) \leq \lambda_2(\tilde{\mathbf{B}}) \leq \cdots \leq \lambda_{n-1}(\tilde{\mathbf{B}}).$$

Applying the Cauchy Interlacing Theorem (Theorem 21), we know that

$$\lambda_1(\tilde{\mathbf{A}}) \leq \lambda_1(\tilde{\mathbf{B}}) \leq \lambda_2(\tilde{\mathbf{A}}) \leq \lambda_2(\tilde{\mathbf{B}}) \leq \cdots \leq \lambda_{n-1}(\tilde{\mathbf{A}}) \leq \lambda_{n-1}(\tilde{\mathbf{B}}) \leq \lambda_n(\tilde{\mathbf{A}}).$$

In fact, even if $\tilde{\mathbf{B}}$ is an $(n-\ell) \times (n-\ell)$ sub-matrix of $\tilde{\mathbf{A}}$ obtained by removing the same rows and columns from $\tilde{\mathbf{A}}$, we can prove that

$$\lambda_i(\tilde{\mathbf{A}}) \leq \lambda_i(\tilde{\mathbf{B}}), \quad 1 \leq i \leq n-\ell. \quad (29)$$

The proof is nearly identical to the proof of Corollary 22, and in particular uses the Cauchy Interlacing Theorem.

Let us apply (29) with $\tilde{\mathbf{A}} = \mathbf{L}$ and $\tilde{\mathbf{B}} = \mathbf{B}$. We conclude that

$$\lambda_i(\mathbf{L}) \leq \lambda_i(\mathbf{B}) < \lambda_k(\mathbf{L}), \quad 1 \leq i \leq t.$$

Since it is only possible for $\lambda_i(\mathbf{L}) < \lambda_k(\mathbf{L})$ for $1 \leq i \leq k-1$, we conclude that $t \leq k-1$. \square

Remark 20. The theorem is sharp. Indeed, consider the star graph $S_5 = (V, E)$ on five vertices such that $V = \{1, \dots, 5\}$ with $a = 1$ being the hub vertex. One can verify that ψ defined as $\psi(1) = 0$, $\psi(2) = \psi(3) = \psi(4) = 1$, and $\psi(5) = -3$ is an eigenvector with eigenvalue $\lambda_2 = 1$ (use Lemma 10). On the other hand, if we consider the set $\{a \in V : \psi(a) > 0\}$, we get three connected components.

Lecture 12: Graph Signal Processing, Part I

February 25, 2021

Lecturer: Matthew Hirn

22 Graph signal processing

The main reference for this section is [3].

Now that we have completed our understanding of eigenvector frequency, we can take a small tangent and discuss *graph signal processing*. In graph signal processing, one is given a graph $G = (V, E, w)$ and a signal $\mathbf{x} : V \rightarrow \mathbb{R}$. The goal is to extract information from the signal \mathbf{x} with respect to the graph G on which it is defined. Unlike our work thus far, in which our sole goal was to extract information about the graph G , and the use of signal \mathbf{x} , such as test vectors, was a means to that end, here the signal \mathbf{x} is the object of interest.

Here are some examples of this setting:

- The vertices of the graph consist of words and one places edges between words that have similar meaning and which are used together. We can represent a text document as a signal \mathbf{x} on this graph by assigning $\mathbf{x}(a)$ (remember $a \in V$ is a word) to be the number of times the word appears in the document. By extracting information from the signal \mathbf{x} , we extract information about the document it represents.
- The graph represents a discrete sampling of the earth (so a discrete sampling of the sphere in which we connect nearby samples with edges) or some other graph representation of a geographic region, such as the Minnesota graph. The signal \mathbf{x} could represent a phenomenon over the region, such as weather patterns, or the spread of a disease (such as COVID-19). Different signals could represent different time snapshots of the weather or disease spread, respectively.
- Gene-gene interaction network in which nodes are genes and edges are placed between genes with a physical and/or functional relationship. A signal $\mathbf{x} : V \rightarrow \mathbb{R}$ could correspond to single cell sequencing of the genes for a particular individual. We could have multiple signals, some from healthy individuals, and some from individuals with a particular disease.

22.1 The graph Fourier transform

In classical signal processing one often analyzes a signal through its *Fourier transform*, which measures the frequencies present in the signal. Suppose we have a signal $f : \mathbb{R} \rightarrow \mathbb{R}$, such as

a time series. We measure the frequencies of f by computing its Fourier transform. Define $e_\omega : \mathbb{R} \rightarrow \mathbb{C}$ as

$$e_\omega(t) := e^{i\omega t} = \cos(\omega t) + i \sin(\omega t), \quad \omega \in \mathbb{R}, \quad t \in \mathbb{R},$$

and define the Fourier transform of f , denote by $\hat{f} : \mathbb{R} \rightarrow \mathbb{C}$, as

$$\hat{f}(\omega) := \langle f, e_\omega \rangle = \int_{-\infty}^{+\infty} f(t) e^{-i\omega t} dt. \quad (30)$$

If

$$\|f\|_1 := \int_{-\infty}^{+\infty} |f(t)| dt < \infty,$$

then (30) is finite for all $\omega \in \mathbb{R}$ and the Fourier transform is well defined. If $\|\hat{f}\|_1 < \infty$ as well, then one can show that \hat{f} determines f and we have the following Fourier inversion result:

$$f(t) = \int_{-\infty}^{+\infty} \hat{f}(\omega) e^{i\omega t} d\omega = \int_{-\infty}^{+\infty} \langle f, e_\omega \rangle e_\omega(t) d\omega. \quad (31)$$

The second way of writing (31) is meant to recall an orthonormal basis expansion of a vector. Indeed, if one replaces the integral with a finite sum, it is exactly that. The function e_ω is an oscillating function that vibrates at the frequency ω , and the Fourier transform measures the overlap of f with e_ω . In light of (31), we see that knowing this frequency representation of f is equivalent to knowing f itself, and each Fourier coefficient $\hat{f}(\omega)$ encodes the proportion of f that vibrates at the frequency ω .

We now ask the question as to how the complex sinusoidal functions e_ω arise. To that end, define the *Laplacian* Δ as the operator that maps a twice-differentiable function $f : \mathbb{R} \rightarrow \mathbb{R}$ to the negative of its second derivative:

$$\Delta f := -f''.$$

Notice, the functions e_ω are eigenfunctions of Δ . Indeed,

$$\Delta e_\omega(t) = -e''_\omega(t) = -\frac{d^2}{dt^2} e^{i\omega t} = -(i\omega)^2 e^{i\omega t} = \omega^2 e_\omega(t), \quad (32)$$

and so the eigenvalue associated to e_ω is ω^2 . These eigenfunctions are orthogonal (in the sense of distributions), and the Fourier inversion formula (31) implies that any other eigenfunction would not be linearly independent.

As discussed back in Section 3.3 (Lecture 02), the graph Laplacian \mathbf{L} is the analogue of the Laplacian for signals \mathbf{x} defined on a graph $G = (V, E, w)$. The above discussion shows, then, that the eigenvectors of \mathbf{L} are the analogues of the sinusoidal functions e_k . Indeed, these eigenvectors $\{\psi_k\}_{k=1}^n$ can be thought of as the harmonics of the graph G . Furthermore, by our recent work on the cycle graph (Theorem 17), the path graph (Theorem 18), the weighted path graph (Theorem 32), and Fiedler's Nodal Domain Theorem for general weighted graph (Theorem 33), we know that the index k of ψ_k is a proxy for the frequency of the eigenvector.

Therefore, we define the *graph Fourier transform* of a signal $\mathbf{x} : V \rightarrow \mathbb{R}$, denoted by $\hat{\mathbf{x}} \in \mathbb{R}^n$, as

$$\hat{\mathbf{x}}(k) := \langle \mathbf{x}, \boldsymbol{\psi}_k \rangle, \quad 1 \leq k \leq n.$$

Since $\{\boldsymbol{\psi}_k\}_{k=1}^n$ forms an ONB for \mathbb{R}^n , we have the Fourier inverse:

$$\mathbf{x} = \sum_{k=1}^n \hat{\mathbf{x}}(k) \boldsymbol{\psi}_k = \sum_{k=1}^n \langle \mathbf{x}, \boldsymbol{\psi}_k \rangle \boldsymbol{\psi}_k. \quad (33)$$

Thus $\hat{\mathbf{x}}$ provides an alternate way of encoding the signal \mathbf{x} through its frequency responses with respect to the harmonics of the graph G ; see also Figure 24. We also observe, in light of (32), that λ_k can be interpreted as the frequency squared of the eigenvector $\boldsymbol{\psi}_k$.




Figure 24: (a) The Minnesota graph and a signal \mathbf{x} defined on the vertices of it, with values $\mathbf{x}(a)$ indicated by the blue spikes at each vertex. (b) The graph Fourier transform, $\hat{\mathbf{x}}(k)$, of the signal \mathbf{x} , plotted as a function of the eigenvalues λ_k . In this case $\hat{\mathbf{x}}(k) = e^{-5\lambda_k}$ and \mathbf{x} is obtained using the graph Fourier inverse formula (33). Figure taken from [3, Figure 4].

22.2 Smoothness and decay

In signal processing one relates the smoothness of a signal f to the decay of its Fourier transform \hat{f} . Informally, the smoother the signal f (as measured by the number of times one can differentiate f), the faster its Fourier transform $\hat{f}(\omega)$ will decay as $|\omega| \rightarrow \infty$, and vice versa.

We can make a similar observation for graph signals \mathbf{x} . Recall we can use the graph Laplacian quadratic form to measure the smoothness of \mathbf{x} , and by Lemma 4 we can write it using the Fourier transform of \mathbf{x} :

$$\mathbf{x}^T \mathbf{L} \mathbf{x} = \sum_{(a,b) \in E} w(a,b) (\mathbf{x}(a) - \mathbf{x}(b))^2 = \sum_{k=1}^n \lambda_k |\langle \mathbf{x}, \boldsymbol{\psi}_k \rangle|^2 = \sum_{k=1}^n \lambda_k |\hat{\mathbf{x}}(k)|^2.$$

Recalling that $0 = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$, we see that if $\mathbf{x}^T \mathbf{L} \mathbf{x}$ is small (and hence \mathbf{x} is smooth), then $|\hat{\mathbf{x}}(k)|$ will be small for large k , and conversely, if $|\hat{\mathbf{x}}(k)|$ decays as k gets larger, then $\mathbf{x}^T \mathbf{L} \mathbf{x}$ will be small and hence \mathbf{x} will be smooth.

It is important to note how strongly this notion of smoothness depends on the underlying graph structure. Indeed, if two graphs $G_1 = (V, E_1)$ and $G_2 = (V, E_2)$ have the same vertices V but different edge sets E_1 and E_2 , a signal \mathbf{x} can be defined on either graph. However, the smoothness of \mathbf{x} with respect to G_1 may be very different than the smoothness of \mathbf{x} with respect to G_2 ; see Figure 25 for an example.

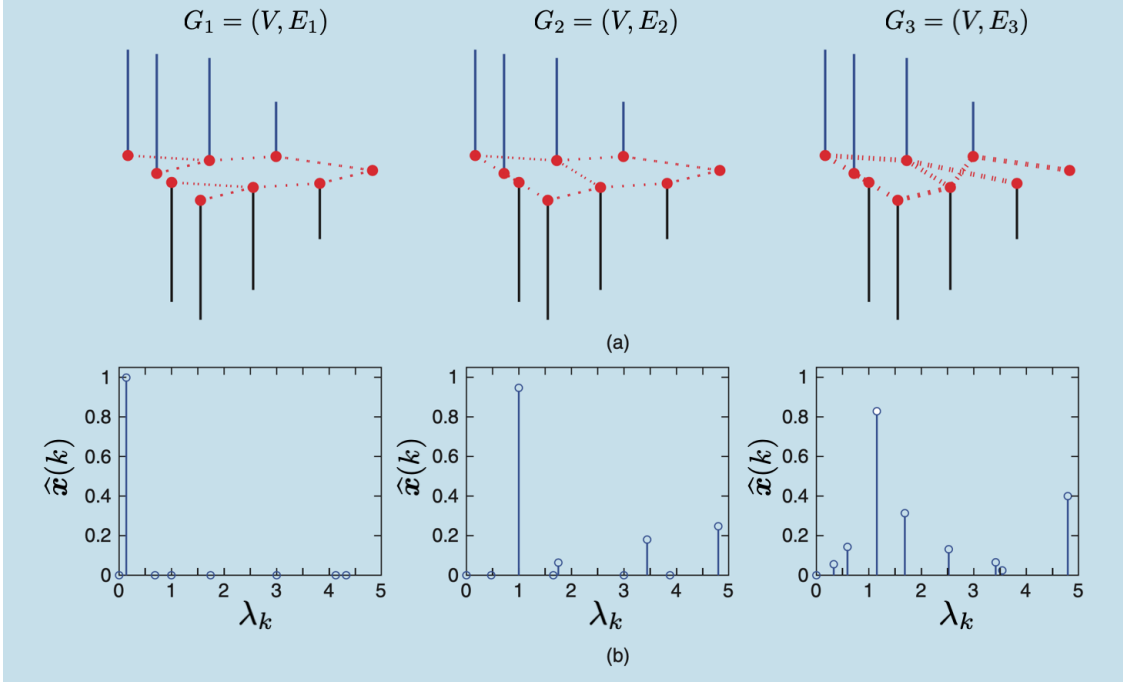


Figure 25: (a) Three graphs $G_i = (V, E_i)$, $1 \leq i \leq 3$, with a common vertex set but different edge sets. A common signal $\mathbf{x} : V \rightarrow \mathbb{R}$ is plotted on each graph. By considering the edges in each graph, we can visually see the smoothness of \mathbf{x} will change from graph to graph, with decreasing smoothness from left to right. We have $\mathbf{x}^T \mathbf{L}_{G_1} \mathbf{x} = 0.14$, $\mathbf{x}^T \mathbf{L}_{G_2} \mathbf{x} = 1.31$, and $\mathbf{x}^T \mathbf{L}_{G_3} \mathbf{x} = 1.81$, indicating quantitatively that \mathbf{x} is smoothest on G_1 and is the least smooth on G_3 . (b) The graph Fourier transform, $\hat{\mathbf{x}}(k)$, with respect to each graph G_i . In the case of G_1 , on which \mathbf{x} is the smoothest, we see that $\hat{\mathbf{x}}(k)$ is concentrated on $k = 2$. On the other hand, in the case of G_3 , on which \mathbf{x} is the least smooth, $|\hat{\mathbf{x}}(k)|$ is small for small k and is larger for some intermediate and larger values of k . The graph G_2 is an intermediate case, which is reflected by $\hat{\mathbf{x}}$ on this graph. Figure taken from [3, Figure S1].

Lecture 13: Graph Signal Processing, Part II

March 4, 2021

Lecturer: Matthew Hirn

22.3 Filtering and graph convolution

One of the primary ways signal processing algorithms extract information from a signal $f : \mathbb{R} \rightarrow \mathbb{R}$ is by *filtering* f with a filter $h : \mathbb{R} \rightarrow \mathbb{C}$. This filtering operation is accomplished with a *convolution*, which is defined as:

$$f * h(t) := \int_{-\infty}^{+\infty} f(u)h(t-u) du. \quad (34)$$

See, for example, Figure 26.




Figure 26: (a) A gray scale image of a butterfly. (b) The filtered butterfly image, using a small oscillating filter with zero average that detects edges. (c) The filter used, in which green is positive and pink is negative (note: the filter is not to scale and is enlarged for visualization purposes).

By the Fourier convolution theorem,

$$\widehat{f * h}(\omega) = \widehat{f}(\omega)\widehat{h}(\omega),$$

and thus

$$f * h(t) = \int_{-\infty}^{+\infty} \widehat{f * h}(\omega) e^{i\omega t} d\omega = \int_{-\infty}^{+\infty} \widehat{f}(\omega)\widehat{h}(\omega) e^{i\omega t} d\omega. \quad (35)$$

In other words, convolution in time/space corresponds to multiplication in frequency. In classical signal processing, filters h are often defined in frequency, meaning that one specifies

$\widehat{h}(\omega)$ and then obtains $h(t)$ through the Fourier inverse (31). Additionally, such filters are often designed to extract a localized range of frequencies from f , that is, $\widehat{h}(\omega)$ has a localized support so that $\widehat{f * h}(\omega)$ corresponds to $\widehat{f}(\omega)$ in the support of $\widehat{h}(\omega)$. Such filters h are often localized in time/space as well (if, for example, \widehat{h} is smooth), and as such in space we see that $f * h(t)$ replaces $f(t)$ with a local aggregation of information in the signal f around t .

On a graph $G = (V, E, w)$, we must be more creative to define convolution since the notion of a translation does not make sense except on certain types of graphs (such as the path graph and the cycle graph). However, now that we have the notion of a graph Fourier transform, one option is to define convolution in a way analogous to (35). To that end, let $\mathbf{x} : V \rightarrow \mathbb{R}$ be a graph signal and $\mathbf{h} : V \rightarrow \mathbb{R}$ another graph signal (which we will often think of as a filter). Recall we let $0 = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ be the eigenvalue of the graph Laplacian \mathbf{L} , and let $\boldsymbol{\psi}_1, \boldsymbol{\psi}_2, \dots, \boldsymbol{\psi}_n$ be corresponding orthonormal eigenvectors. Define the convolution of \mathbf{x} and \mathbf{h} as:

$$\mathbf{x} * \mathbf{h}(a) := \sum_{k=1}^n \widehat{\mathbf{x}}(k) \widehat{\mathbf{h}}(k) \boldsymbol{\psi}_k(a). \quad (36)$$

Equation (36) resembles (35), and indeed (35) is the motivation for defining graph convolution as (36). It is important to note, though, that (35) is a consequence of the definition of convolution (34) for Euclidean supported signals, whereas we are taking (36) as the definition of convolution for graph supported signals.

If we set

$$\mathbf{H} := \begin{pmatrix} \widehat{\mathbf{h}}(1) & 0 & \dots & 0 \\ 0 & \widehat{\mathbf{h}}(2) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \widehat{\mathbf{h}}(n) \end{pmatrix},$$

and we let $\boldsymbol{\Psi}$ be the $n \times n$ matrix such that the k^{th} column of $\boldsymbol{\Psi}$ is $\boldsymbol{\psi}_k$, then one can rewrite (36) as

$$\mathbf{x} * \mathbf{h} = \boldsymbol{\Psi} \mathbf{H} \boldsymbol{\Psi}^T \mathbf{x}, \quad (37)$$

which resembles the decomposition of \mathbf{L} as $\mathbf{L} = \boldsymbol{\Psi} \boldsymbol{\Lambda} \boldsymbol{\Psi}^T$ (see Exercise 2 of Homework 1).

Indeed, it is often the case, like in classical signal processing, that we specify a graph filter \mathbf{h} through its graph Fourier transform $\widehat{\mathbf{h}}$. Furthermore, it is often useful to do so through a function $g : [0, \infty) \rightarrow \mathbb{R}$ that acts on the eigenvalues of \mathbf{L} , so that

$$\widehat{\mathbf{h}}(k) = g(\lambda_k).$$

In this case, if we define

$$g(\boldsymbol{\Lambda}) := \begin{pmatrix} g(\lambda_1) & 0 & \dots & 0 \\ 0 & g(\lambda_2) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & g(\lambda_n) \end{pmatrix}, \quad (38)$$

then

$$\mathbf{x} * \mathbf{h} = \Psi g(\Lambda) \Psi^T \mathbf{x}. \quad (39)$$

There are several advantages to formulating the graph filter \mathbf{h} as $\widehat{\mathbf{h}}(k) = g(\lambda_k)$. Here are some of them:

- In some cases, the regularity of g can be used to imply localization of \mathbf{h} on the graph G (more on this below).
- By defining the filter \mathbf{h} through a function $g : [0, \infty) \rightarrow \mathbb{R}$ that is valid for any eigenvalue $\lambda \in [0, \infty)$, we can transfer the filter \mathbf{h} between multiple graphs. This will be useful for graph neural networks trained on multiple graphs (e.g., for graph classification).
- We will be able to define the dilation of the filter \mathbf{h} in frequency, which on the graph will make its support smaller and larger depending on the dilation factor. This in turn will allow us to define multiscale graph signal filtering operators.

On the other hand, there are some disadvantages as well. Chief among them is that if an eigenvalue $\lambda_k = \lambda_{k+1} = \dots = \lambda_{k+\ell}$ has multiplicity more than one, the Fourier transform of the filter, $\widehat{\mathbf{h}}$, must take the same value across the corresponding indices, $\widehat{\mathbf{h}}(k) = \widehat{\mathbf{h}}(k+1) = \dots = \widehat{\mathbf{h}}(k+\ell)$. Such filters lack the ability to encode local directionality in the graph G .

22.4 Spatial graph filtering

Let $G = (V, E)$ and let $a, b \in V$. Define the *graph distance* between a and b , denoted by $\text{dist}(a, b) = \text{dist}_G(a, b)$, to be the length of the shortest path from a to b . Previously we defined $N(a)$ to be all the neighbors of G ; we now generalize it to the *m -hop neighborhood* of a :

$$N_m(a) := \{b \in V : b \neq a \text{ and } \text{dist}(a, b) \leq m\}.$$

Note that $N_1(a) = N(a)$.

A spatial graph signal filtering is obtained through a matrix \mathbf{A}_m (not necessarily symmetric) such that the diagonal entries $\mathbf{A}_m(a, a)$ can be anything, and the off-diagonal entries $\mathbf{A}_m(a, b) = 0$ for all $b \notin N_m(a)$. The filtering of \mathbf{x} by \mathbf{A}_m computes:

$$\mathbf{A}_m \mathbf{x}(a) = \mathbf{A}_m(a, a) \mathbf{x}(a) + \sum_{b \in N_m(a)} \mathbf{A}_m(a, b) \mathbf{x}(b).$$

From the above formula, we see that \mathbf{A}_m replaces $\mathbf{x}(a)$ with a weighted combination of $\mathbf{x}(a)$ plus $\mathbf{x}(b)$ for all $b \in V$ within m steps (or hops) of a . The following theorem relates certain types of spectral filters based on polynomials to spatial m -hop filters.

Theorem 34. *Let $G = (V, E)$ and let $p : \mathbb{R} \rightarrow \mathbb{R}$ be a polynomial of order m , i.e.,*

$$p(t) = \sum_{j=0}^m c_j t^j. \quad (40)$$

Define a filter $\mathbf{h} : V \rightarrow \mathbb{R}$ through its graph Fourier transform as

$$\widehat{\mathbf{h}}(k) = p(\lambda_k).$$

Then \mathbf{h} defines an m -hop filter \mathbf{A}_m .

To prove Theorem 34 we need two lemmas. First, for a polynomial p as in (40), and for a matrix \mathbf{B} , define

$$p(\mathbf{B}) := \sum_{j=0}^m c_j \mathbf{B}^j.$$

Note this definition is consistent with the definition of $g(\mathbf{\Lambda})$ in (38) when $g = p$ and $\mathbf{B} = \mathbf{\Lambda}$.

Lemma 35. Let \mathbf{B} be a symmetric, $n \times n$ matrix with eigenvalues $\lambda_1, \dots, \lambda_n$ (here, even though we use λ_k to denote the eigenvalues, \mathbf{B} does not have to be a graph Laplacian) and corresponding orthonormal eigenvectors ϕ_1, \dots, ϕ_n . Let $\mathbf{\Phi}$ be the $n \times n$ matrix whose k^{th} column is ϕ_k , and let $\mathbf{\Lambda}$ be the $n \times n$ diagonal matrix with $\mathbf{\Lambda}(k, k) = \lambda_k$. Then for any polynomial p ,

$$p(\mathbf{B}) = \mathbf{\Phi} p(\mathbf{\Lambda}) \mathbf{\Phi}^T.$$

Lemma 36. Let $G = (V, E)$ be a graph and let \mathbf{B} be a matrix such that

$$\forall a, b \in V, \quad a \neq b, \quad b \notin N(a) \implies \mathbf{B}(a, b) = 0.$$

Then,

$$\forall a, b \in V, \quad a \neq b, \quad b \notin N_m(a) \implies (\mathbf{B}^m)(a, b) = 0.$$

I leave the proof of these lemmas to you. Let us now prove Theorem 34.

Proof of Theorem 34. Recalling (39), we write the convolution of a signal \mathbf{x} against the filter \mathbf{h} in matrix form as

$$\mathbf{x} * \mathbf{h} = \mathbf{\Psi} p(\mathbf{\Lambda}) \mathbf{\Psi}^T \mathbf{x}.$$

Since $\mathbf{L} = \mathbf{\Psi} \mathbf{\Lambda} \mathbf{\Psi}^T$, using Lemma 35 we have:

$$\mathbf{\Psi} p(\mathbf{\Lambda}) \mathbf{\Psi}^T = p(\mathbf{L}).$$

Thus,

$$\mathbf{x} * \mathbf{h} = p(\mathbf{L}) \mathbf{x} = \sum_{j=0}^m c_j \mathbf{L}^j \mathbf{x}. \quad (41)$$

Now applying Lemma 36 to \mathbf{L} , we see that

$$\forall a, b \in V, \quad a \neq b, \quad b \notin N_j(a) \implies (\mathbf{L}^j)(a, b) = 0. \quad (42)$$

Thus we can write

$$\mathbf{L}^j \mathbf{x}(a) = \mathbf{L}^j(a, a) \mathbf{x}(a) + \sum_{b \in N_j(a)} \mathbf{L}^j(a, b) \mathbf{x}(b). \quad (43)$$

Combining equations (41), (42), and (43), we obtain:

$$\begin{aligned}
\mathbf{x} * \mathbf{h}(a) &= \sum_{j=0}^m c_j \mathbf{L}^j \mathbf{x}(a) = \sum_{j=0}^m c_j \left(\mathbf{L}^j(a, a) \mathbf{x}(a) + \sum_{b \in N_j(a)} \mathbf{L}^j(a, b) \mathbf{x}(b) \right) \\
&= \underbrace{\left(\sum_{j=0}^m c_j \mathbf{L}^j(a, a) \right)}_{\mathbf{A}_m(a, a)} \mathbf{x}(a) + \sum_{j=0}^m c_j \left(\sum_{b \in N_j(a)} \mathbf{L}^j(a, b) \mathbf{x}(b) \right) \\
&= \mathbf{A}_m(a, a) \mathbf{x}(a) + \sum_{b \in N_m(a)} \underbrace{\left(\sum_{j=0}^m c_j \mathbf{L}^j(a, b) \right)}_{\mathbf{A}_m(a, b)} \mathbf{x}(b),
\end{aligned}$$

and the proof is completed. \square

22.5 Translations of a graph filter

With the notion graph convolution now defined through the graph Fourier transform, we can reverse engineer a definition of translation on graphs as well. This will be useful for visualizing filters \mathbf{h} defined on a graph, as it will allow us to localize and center the filter at an arbitrary vertex of G , which we will see gives a better visualization of how \mathbf{h} acts as a filter compared with plotting \mathbf{h} directly. This is how Figure 11 was created.

In order to motivate matters, let us first consider a function $f : \mathbb{R} \rightarrow \mathbb{R}$ and define the translation of f to $u \in \mathbb{R}$ as

$$f_u(t) := f(t - u).$$

We now write $f_u(t)$ in frequency. To do so, we first note that its Fourier transform is:

$$\widehat{f}_u(\omega) = \int_{-\infty}^{+\infty} f(t - u) e^{-i\omega t} dt = \int_{-\infty}^{+\infty} f(\tilde{t}) e^{-i\omega(\tilde{t} + u)} d\tilde{t} = e^{-i\omega u} \int_{-\infty}^{+\infty} f(\tilde{t}) e^{-i\omega \tilde{t}} d\tilde{t} = e^{-i\omega u} \widehat{f}(\omega),$$

where we used the change of variables $\tilde{t} = t - u$. Now, notice that, using the notation $e_\omega(t) = e^{i\omega t}$, this means we have:

$$\widehat{f}_u(\omega) = \overline{e_\omega(u)} \widehat{f}(\omega),$$

and using the Fourier inversion formula:

$$f(t - u) = f_u(t) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \widehat{f}(\omega) e^{i\omega t} d\omega = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \widehat{f}(\omega) \overline{e_\omega(u)} e_\omega(t) d\omega. \quad (44)$$

Now let $G = (V, E, w)$ be a weighted graph. Motivated by (44), we define the *graph translation* of a graph signal $\mathbf{x} : V \rightarrow \mathbb{R}$ to the vertex $b \in V$ by:

$$\mathbf{x}_b(a) := \sqrt{n} \sum_{k=1}^n \widehat{\mathbf{x}}(k) \psi_k(b) \psi_k(a). \quad (45)$$

We use the normalization factor \sqrt{n} to preserve the mean of \mathbf{x} (see Theorem 37 below). On the other hand, traditional Euclidean translations also preserve the norm of the signal. However, graph translation does not, and in particular $\|\mathbf{x}_b\| \neq \|\mathbf{x}\|$ generically. Figure illustrates graph translation on the Minnesota graph.

Theorem 37. *Let $G = (V, E, w)$ be a connected graph, let $\mathbf{x} : V \rightarrow \mathbb{R}$ be a graph signal, and let $\mathbf{x}_b : V \rightarrow \mathbb{R}$ be the graph translation of \mathbf{x} to the vertex b , as defined in (45). Then*

$$\sum_{a \in V} \mathbf{x}(a) = \sum_{a \in V} \mathbf{x}_b(a).$$

I leave the proof to you in Homework 06. As an additional remark on graph translation, we note that we can write it as the graph convolution of \mathbf{x} against the Dirac delta graph signal $\delta_b : V \rightarrow \mathbb{R}$, which is defined as

$$\delta_b(a) := \begin{cases} 1 & a = b \\ 0 & a \neq b \end{cases}.$$

The graph Fourier transform of δ_b is

$$\widehat{\delta_b}(k) = \langle \delta_b, \psi_k \rangle = \sum_{a \in V} \delta_b(a) \psi_k(a) = \psi_k(b).$$

Therefore we have

$$\mathbf{x}_b = \sqrt{n} \sum_{k=1}^n \widehat{\mathbf{x}}(k) \psi_k(b) \psi_k = \sqrt{n} \sum_{k=1}^n \widehat{\mathbf{x}}(k) \widehat{\delta_b}(k) \psi_k = \sqrt{n} (\mathbf{x} * \delta_b). \quad (46)$$

Note that a similar formula holds for Euclidean signal translation as well (although one has to be careful in how δ_u is defined).

Now consider a graph filter \mathbf{h} . We want to center \mathbf{h} at the vertex $b \in V$ using \mathbf{h}_b , the translation of \mathbf{h} to the vertex b . Using the matrix representation (37) of \mathbf{h} and (46), we have:

$$\frac{1}{\sqrt{n}} \mathbf{h}_b = \mathbf{h} * \delta_b = \Psi \mathbf{H} \Psi^T \delta_b = \text{column } b \text{ of the filter matrix } \Psi \mathbf{H} \Psi^T.$$

Thus, the translation of the filter \mathbf{h} to the vertex b extracts the b^{th} column of the associated filter matrix $\Psi \mathbf{H} \Psi^T$. Since this matrix is symmetric, that is equivalent to extracting the b^{th} row of $\Psi \mathbf{H} \Psi^T$. Of course, it is exactly the inner product of the b^{th} row of $\Psi \mathbf{H} \Psi^T$ with \mathbf{x} that defines $\mathbf{x} * \mathbf{h}(b)$, and in particular we have

$$\mathbf{x} * \mathbf{h}(b) = \frac{1}{\sqrt{n}} \langle \mathbf{x}, \mathbf{h}_b \rangle.$$

Notice the previous equation mimics how one can write Euclidean convolutions. We see that \mathbf{h}_b represents how the filter \mathbf{h} aggregates information of a signal \mathbf{x} in the neighborhood of the vertex b .

Lecture 14: Graph Signal Processing, Part III

March 9, 2021

*Lecturer: Matthew Hirn***22.6 Dilations and graph wavelets**

In classical signal processing wavelet transforms are a common way of filtering a signal $f : \mathbb{R} \rightarrow \mathbb{R}$. A *wavelet* h is a filter that is localized in time/space, its Fourier transform $\hat{h} : \mathbb{R} \rightarrow \mathbb{R}$ is localized in frequency, and the filter has zero average:

$$\hat{h}(0) = \int_{\mathbb{R}} h(t) dt = 0. \quad (47)$$

A typical wavelet, along with its Fourier transform, is plotted in Figure 27. Notice that because it has zero average, it oscillates. But unlike $e_{\omega}(t) = e^{i\omega t}$, the support of this waveform is localized in space. This localization allows the wavelet, when used to filter f via $f * h$, to isolate localized phenomena in the signal f , such as sharp transitions (edges in images), which is not possible with standard Fourier analysis.




Figure 27: A typical wavelet (left), and its Fourier transform (right). Figure taken from [4].

Now, because we do not know exactly what scale we want to use for a given signal f , and because an individual signal f may exhibit different types of patterns at different scales, we dilate the wavelet to make it larger and smaller. The standard approach is to do this using dyadic scales:

$$h_j(t) := 2^{-j} h(2^{-j} t), \quad j \in \mathbb{Z}.$$

When $j > 0$, the filter h_j widens its support relative to h , and when $j < 0$, the filter h_j shrinks its support relative to h . We then filter information in f at different scales by computing $f * h_j$ for $j < J$ for some maximum scale 2^J . Notice that

$$\widehat{h_j}(\omega) = \widehat{h}(2^j \omega). \quad (48)$$

Thus, using the Fourier convolution theorem (35), we have:

$$\widehat{f * h_j}(\omega) = \widehat{f}(\omega) \widehat{h}(2^j \omega),$$

and so, from the frequency perspective, $f * h_j$ captures the high frequencies of f in proportion to 2^{-j} (hence the use of the letter h to denote this filter).

In order to capture large scale phenomena in the signal f at the scale larger than 2^J , we use a *low pass filter*. Let $\ell : \mathbb{R} \rightarrow \mathbb{R}$ denote a low pass filter, which we will assume means ℓ is localized in time/space, it has unit integral,

$$\widehat{\ell}(0) = \int_{\mathbb{R}} \ell(t) dt = 1,$$

and $\widehat{\ell}(\omega)$ is localized around $\omega = 0$. A typical low pass filter is plotted in Figure 28.





Figure 28: A typical low pass filter (left), and its Fourier transform (right). Figure taken from [4].

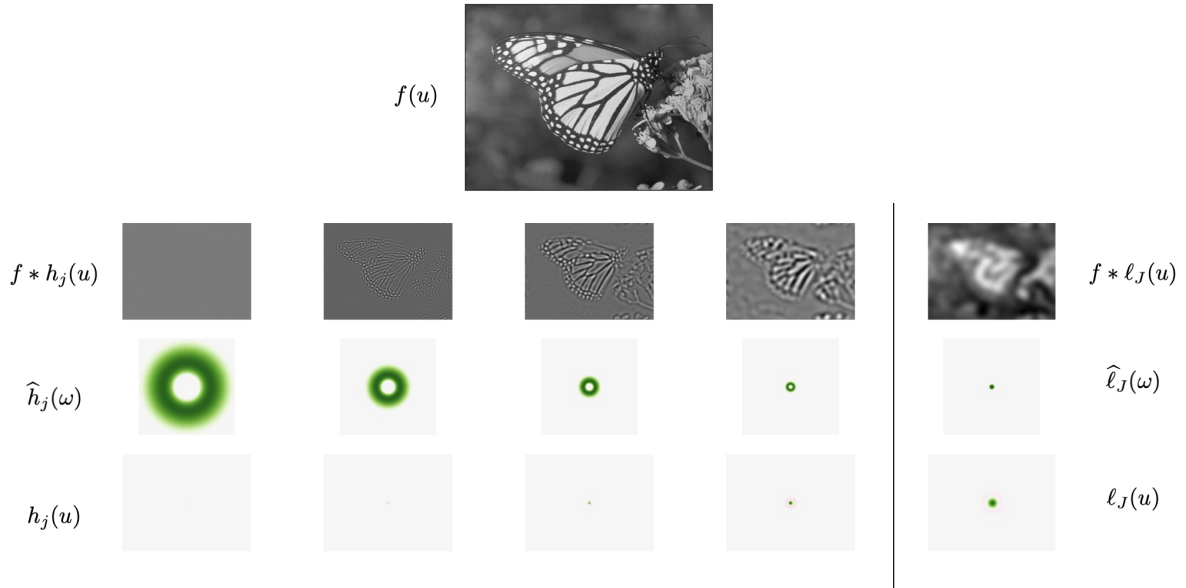
The *wavelet transform* of f computes:

$$W_J f := \{f * \ell_J, f * h_j : j < J\},$$

which captures the low frequencies of f via $f * \ell_J$ (hence the use of the letter ℓ for this filter), the high frequencies of f via $f * h_j$ for $j < J$, and thus all the information in the signal f (recall the Fourier inversion formula (31)). In practice, one places a minimum value on j in proportion to the sampling resolution of the signal f (i.e., how many samples of



(a) 1D wavelet transform (figure taken from [4])



(b) 2D wavelet transform

Figure 29: The 1D wavelet transform (a) and the 2D wavelet transform (b). In both cases, the low pass filtering $f * \ell_J$ smooths the original signal f . The high pass filtering coefficients, $f * h_j$, on the other hand, recover sharp transitions in the signal f , such as jump discontinuities in the 1D case and edges in the 2D case.

$f(t)$ you get on the computer). Figure 29 shows a one-dimension wavelet transform and a two-dimensional wavelet transform.

Our goal is to extend the wavelet transform to graphs. In order to do so, we need a notion of dilation. Similar to translations, it is difficult to define the dilation of a filter \mathbf{h} directly on the vertices of the graph. However we can define dilation in frequency. Let $\hat{\mathbf{h}}$ be a function of the eigenvalues of \mathbf{L} , i.e.,

$$\hat{\mathbf{h}}(k) = g(\lambda_k),$$

for some function $g : [0, \infty) \rightarrow \mathbb{R}$. Using (48) as our inspiration, we define the dilation of \mathbf{h} via:

$$\hat{\mathbf{h}}_j(k) := g(2^j \lambda_k).$$

We can define a graph low pass filter ℓ via its graph Fourier transform as

$$\hat{\ell}_J(k) = g_{\text{low}}(2^J \lambda_k),$$

where $g_{\text{low}} : [0, \infty) \rightarrow \mathbb{R}$ is designed so that $g_{\text{low}}(0) = 1$, $g_{\text{low}}(t) \geq 0$, and g_{low} is a (not necessarily monotonic) decreasing function. Similarly, we can define a high pass wavelet filter \mathbf{h} via its graph Fourier transform as

$$\hat{\mathbf{h}}_j(k) = g_{\text{high}}(2^j \lambda_k),$$

where $g_{\text{high}} : [0, \infty) \rightarrow \mathbb{R}$ is designed so that $g_{\text{high}}(0) = 0$, $g_{\text{high}}(t) \geq 0$, and the essential support of $g_{\text{high}}(t)$ is localized around a fixed value t_0 (so that $g(2^j t)$ is localized around $2^{-j} t_0$). One can verify that

$$\sum_{a \in V} \mathbf{h}_j(a) = 0,$$

which mimics the zero-average condition (47) on Euclidean wavelets. See Figure 30 for an illustration of how $\hat{\ell}_J$ and $\hat{\mathbf{h}}_j$ fit together. Figure 31, on the other hand, shows graph wavelets plotted on people manifolds that are approximated numerically as graphs.




Figure 30: Illustration of the graph low pass (blue) and graph wavelet (red) Fourier transforms, when taken as a function of the eigenvalue of the graph Laplacian.




Figure 31: Graph wavelets, $(\mathbf{h}_j)_b$, plotted in space using the graph translation operator (45), for increasing j from left to right. The first two rows are the same manifold, but the wavelet is translated to two different vertices b . The last two rows are different manifolds, but the wavelet is translated to the same vertex b . Figure taken from [2].

We can define g_{low} and g_{high} in such a way so that the graph wavelet transform runs over $0 \leq j < J$. In that case, the graph wavelet transform $\mathbf{W}_J : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times (J+1)}$ is defined as

$$\mathbf{W}_J \mathbf{x} := \{\mathbf{x} * \ell_J, \mathbf{x} * \mathbf{h}_j : 0 \leq j < J\}.$$

The graph wavelet transform provides a multiscale representation of the signal $\mathbf{x} : V \rightarrow \mathbb{R}$ with respect to the underlying graph structure G . It represents \mathbf{x} in terms of $J + 1$ graph signals that decompose \mathbf{x} at different scales according to the graph G . See Figure 32 for an example applied to the spherical MNIST data.




Figure 32: The graph wavelet transform applied to an MNIST handwritten digit projected onto the sphere. Numerically, the sphere is approximated by a graph and the digit's responses on the vertices of this graph form a graph signal \mathbf{x} .

We can define the norm of $\mathbf{W}_J \mathbf{x}$ as

$$\|\mathbf{W}_J \mathbf{x}\|^2 := \|\mathbf{x} * \ell_J\|^2 + \sum_{j=0}^{J-1} \|\mathbf{x} * \mathbf{h}_j\|^2.$$

The next theorem shows that under certain conditions on g_{low} and g_{high} , the graph wavelet transform will preserve the norm of \mathbf{x} and is invertible.

Theorem 38. *Let $G = (V, E, w)$ be a weighted graph with graph Laplacian eigenvalues $0 = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$. If*

$$|g_{\text{low}}(2^J t)|^2 + \sum_{j=0}^{J-1} |g_{\text{high}}(2^j t)|^2 = 1, \quad \forall t \in [0, \lambda_n], \quad (49)$$

then

$$\|\mathbf{W}_J \mathbf{x}\| = \|\mathbf{x}\| \quad \text{and} \quad \mathbf{x} = \mathbf{W}_J^T \mathbf{W}_J \mathbf{x} = \mathbf{x} * \ell_J * \ell_J + \sum_{j=0}^{J-1} \mathbf{x} * \mathbf{h}_j * \mathbf{h}_j.$$

I leave the proof of Theorem 38 for your homework. One possibility for satisfying the requirements of Theorem 38 is the following. Let $G = (V, E, w)$ be a connected graph and let g_{low} be any low pass filter function such that

$$g_{\text{low}}(t) = 1, \quad \forall t \in [0, \lambda_n],$$

with $g_{\text{low}}(t)$ decreasing for $t > \lambda_n$; see Figure 33 for an illustration. Then set

$$g_{\text{high}}(t) = (|g_{\text{low}}(t)|^2 - |g_{\text{low}}(2t)|^2)^{1/2}. \quad (50)$$

One can verify this combination of g_{low} and g_{high} will satisfy (49), but there are other ways as well.




Figure 33: One possibility for g_{low} that, when combined with g_{high} as defined in (50), will satisfy the conditions of Theorem 38.

23 The isoperimetric ratio

We now return to standard spectral graph theory and shift into graph partitioning and clustering. As we will see, these results will be intimately related to the second eigenvalue of the graph Laplacian and the to-be-defined normalized graph Laplacian. In short, these results will quantify the intuition we have sometimes already espoused, which is that even if G is connected, λ_2 quantifies how well G is connected. In light of our discussions on eigenvector frequency, this means that it is the low frequency eigenvectors that are good for clustering G .

To begin, let $G = (V, E)$ be an unweighted graph and let $S \subset V$. Let $V - S$ denote all the vertices in the graph that are not in S , i.e.,

$$V - S := \{a \in V : a \notin S\}.$$

One way to measure how well connected S is to the rest of the graph G is to measure the number of edges going from S to $V - S$. These edges are called the *boundary* of S and are collected in the set ∂S :

$$\partial S := \{(a, b) \in E : a \in S \text{ and } b \in V - S\}.$$

Instead of counting the number of edges in ∂S , it is more useful to measure the ratio of edges in ∂S to the size of S . Indeed, for example if ∂S is small but S is also small, then such a set is relatively well connected to the rest of G . We define this ratio as the *isoperimetric ratio* of S :

$$\theta(S) := \frac{|\partial S|}{|S|}.$$

The isoperimetric ratio of G is the minimum of $\theta(S)$ over all $S \subset V$ such that $|S| \leq n/2$, i.e.,

$$\theta_G := \min_{\substack{S \subset V \\ |S| \leq n/2}} \theta(S).$$

Intuitively, if the graph G has a “bottleneck” or two well connected parts that are only weakly connected to each other, then θ_G will be small, e.g., the barbell graph from Homework 03. In the more extreme case, if G is disconnected, then $\theta_G = 0$. On the other hand, θ_G will be large if any division of V into S and $V - S$ has many edges between S and $V - S$, e.g.,

the complete graph. Graphs with large θ_G are well connected in the sense that there is no $S \subset V$ such that $G(S)$ can be removed from G by only cutting a relatively small number of edges. The following theorem shows that λ_2 gives a lower bound for θ_G .

Theorem 39. *Let $G = (V, E)$, let $0 = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ be its graph Laplacian eigenvalues, and let $S \subset V$. Then*

$$\theta(S) \geq \lambda_2(1 - s), \quad s := \frac{|S|}{n}, \quad (51)$$

and as such,

$$\theta_G \geq \frac{\lambda_2}{2}. \quad (52)$$

Proof. Since θ_G minimizes over $S \subset V$ with $|S| \leq n/2$ we have that $s \leq 1/2$ and as such $1 - s \geq 1/2$. That proves (52) assuming we can prove (51).

To prove (51), recall from Theorem 13 that

$$\lambda_2 = \min_{\substack{\mathbf{x} \in \mathbb{R}^n \\ \langle \mathbf{x}, \mathbf{1} \rangle = 0}} \frac{\mathbf{x}^T \mathbf{L} \mathbf{x}}{\mathbf{x}^T \mathbf{x}}.$$

Thus,

$$\forall \mathbf{x} \in \mathbb{R}^n \text{ such that } \langle \mathbf{x}, \mathbf{1} \rangle = 0, \quad \frac{\mathbf{x}^T \mathbf{L} \mathbf{x}}{\mathbf{x}^T \mathbf{x}} \geq \lambda_2.$$

This is of course the test vector technique we discussed in Section 12. In this case, we need to find a test vector \mathbf{x} such that its Rayleigh quotient contains $\theta(S)$.

We would like to pick $\mathbf{x} = \mathbf{1}_S$, where

$$\mathbf{1}_S(a) := \begin{cases} 1 & a \in S \\ 0 & a \notin S \end{cases}.$$

Indeed, recall $E(S) := \{(a, b) \in E : a, b \in S\}$ and notice that

$$\begin{aligned} \mathbf{1}_S^T \mathbf{L} \mathbf{1}_S &= \sum_{(a,b) \in E} (\mathbf{1}_S(a) - \mathbf{1}_S(b))^2 \\ &= \sum_{(a,b) \in E(S)} (\mathbf{1}_S(a) - \mathbf{1}_S(b))^2 + \sum_{(a,b) \in \partial S} (\mathbf{1}_S(a) - \mathbf{1}_S(b))^2 + \sum_{(a,b) \in E(V-S)} (\mathbf{1}_S(a) - \mathbf{1}_S(b))^2 \\ &= 0 + |\partial S| + 0 \\ &= |\partial S|. \end{aligned}$$

However, $\langle \mathbf{1}_S, \mathbf{1} \rangle > 0$ and so we cannot use $\mathbf{1}_S$ as our test vector. So we use the next best thing, which is

$$\mathbf{x} = \mathbf{1}_S - s\mathbf{1},$$

and so

$$\mathbf{x}(a) = \begin{cases} 1 - s & a \in S \\ -s & a \notin S \end{cases}.$$

We have

$$\begin{aligned}
\langle \mathbf{x}, \mathbf{1} \rangle &= \sum_{a \in S} 1 - s + \sum_{a \in V-S} -s \\
&= |S|(1 - s) - |V - S|s \\
&= |S|(1 - s) - (n - |S|)s \\
&= |S| - s|S| - ns + s|S| \\
&= |S| - |S| \\
&= 0.
\end{aligned}$$

Thus \mathbf{x} is a test vector. Additionally,

$$\mathbf{x}^T \mathbf{L} \mathbf{x} = \sum_{(a,b) \in E} ((\mathbf{1}_S(a) - s) - (\mathbf{1}_S(b) - s))^2 = \sum_{(a,b) \in E} (\mathbf{1}_S(a) - \mathbf{1}_S(b))^2 = |\partial S|.$$

To complete the proof we need to compute $\mathbf{x}^T \mathbf{x}$:

$$\begin{aligned}
\mathbf{x}^T \mathbf{x} &= (\mathbf{1}_S - s\mathbf{1})^T (\mathbf{1}_S - s\mathbf{1}) \\
&= \mathbf{1}_S^T \mathbf{1}_S - s\mathbf{1}_S^T \mathbf{1} - s\mathbf{1}^T \mathbf{1}_S + s^2 \mathbf{1}^T \mathbf{1} \\
&= |S| - s|S| - s|S| + s^2 n \\
&= |S| - 2s|S| + s|S| \\
&= |S|(1 - s).
\end{aligned}$$

Therefore,

$$\frac{|\partial S|}{|S|(1 - s)} = \frac{\mathbf{x}^T \mathbf{L} \mathbf{x}}{\mathbf{x}^T \mathbf{x}} \geq \lambda_2 \implies \theta(S) \geq \lambda_2(1 - s).$$

□

Remark 21. Theorem 39 says that graphs G with large λ_2 are well connected.

24 The normalized graph Laplacian

Remark 22. Whenever discussing the normalized graph Laplacian, we will assume G has no isolated vertices. That is, every vertex in G is connected to at least one other vertex and hence $\deg(a) > 0$ for all $a \in V$.

To take our analysis of graph partitioning and clustering further, it will be useful to introduce a new matrix called the *normalized graph Laplacian*. We will denote it by \mathbf{N} . For a diagonal matrix \mathbf{A} with positive entries on its diagonal, we define \mathbf{A}^α for any $\alpha \in \mathbb{R}$ as

$$\mathbf{A}^\alpha(a, b) := \begin{cases} \mathbf{A}(a, a)^\alpha & b = a \\ 0 & b \neq a \end{cases}.$$

Now define the normalized graph Laplacian as:

$$\mathbf{N} := \mathbf{D}^{-1/2} \mathbf{L} \mathbf{D}^{-1/2} = \mathbf{I} - \mathbf{D}^{-1/2} \mathbf{M} \mathbf{D}^{-1/2}.$$

Note this definition works for unweighted and weighted graphs. Recall that if $G = (V, E, w)$ is weighted, then the degree of $a \in V$ is defined as

$$\deg(a) := \sum_{b \in N(a)} w(a, b).$$

Recall also that we defined $\mathbf{d} : V \rightarrow \mathbb{R}$ as the degree vector G , i.e.,

$$\mathbf{d}(a) := \deg(a).$$

We also have:

$$d_{\min} := \min_{a \in V} \deg(a) \quad \text{and} \quad d_{\max} := \max_{a \in V} \deg(a).$$

The normalized graph Laplacian provides a degree independent representation of G . Indeed, notice that G and $c \cdot G = (V, E, cw)$ have the same normalized graph Laplacian. We will denote the eigenvalues of \mathbf{N} by $\nu_1 \leq \nu_2 \leq \dots \leq \nu_n$, and, when needed, (orthonormal) eigenvectors of \mathbf{N} by $\phi_1, \phi_2, \dots, \phi_n$:

$$\mathbf{N} \phi_k = \nu_k \phi_k, \quad 1 \leq k \leq n.$$

Before utilizing \mathbf{N} for graph partitioning and clustering, let us first discuss some basic properties of the normalized graph Laplacian. This first theorem will relate the eigenvalues of \mathbf{L} to the eigenvalues of \mathbf{N} .

Theorem 40. *Let $G = (V, E, w)$, let $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ be the eigenvalues of \mathbf{L}_G , and let $\nu_1 \leq \nu_2 \leq \dots \leq \nu_n$ be the eigenvalues of \mathbf{N}_G . Then*

$$\frac{\lambda_k}{d_{\max}} \leq \nu_k \leq \frac{\lambda_k}{d_{\min}}, \quad 1 \leq k \leq n.$$

Proof. By the Courant-Fischer theorem (Theorem 3) we know that

$$\nu_k = \min_{\substack{S \subseteq \mathbb{R}^n \\ \dim(S)=k}} \max_{\substack{\mathbf{x} \in S \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T \mathbf{N} \mathbf{x}}{\mathbf{x}^T \mathbf{x}} = \min_{\substack{S \subseteq \mathbb{R}^n \\ \dim(S)=k}} \max_{\substack{\mathbf{x} \in S \\ \mathbf{x} \neq \mathbf{0}}} \frac{(\mathbf{D}^{-1/2} \mathbf{x})^T \mathbf{L} (\mathbf{D}^{-1/2} \mathbf{x})}{\mathbf{x}^T \mathbf{x}}. \quad (53)$$

Now make the change of variables $\mathbf{y} = \mathbf{D}^{-1/2} \mathbf{x}$. Since $\mathbf{x} = \mathbf{D}^{1/2} \mathbf{y}$ implies that $\mathbf{x}^T \mathbf{x} = \mathbf{y}^T \mathbf{D} \mathbf{y}$ and since $\mathbf{D}^{-1/2}$ has full rank, we can rewrite (53) as:

$$\nu_k = \min_{\substack{S \subseteq \mathbb{R}^n \\ \dim(S)=k}} \max_{\substack{\mathbf{y} \in S \\ \mathbf{y} \neq \mathbf{0}}} \frac{\mathbf{y}^T \mathbf{L} \mathbf{y}}{\mathbf{y}^T \mathbf{D} \mathbf{y}}$$

Now observe that

$$\mathbf{y}^T \mathbf{D} \mathbf{y} = \sum_{a \in V} \deg(a) |\mathbf{y}(a)|^2 \leq d_{\max} \sum_{a \in V} |\mathbf{y}(a)|^2 = d_{\max} \mathbf{y}^T \mathbf{y}.$$

Therefore:

$$\nu_k = \min_{\substack{S \subseteq \mathbb{R}^n \\ \dim(S)=k}} \max_{\substack{\mathbf{y} \in S \\ \mathbf{y} \neq \mathbf{0}}} \frac{\mathbf{y}^T \mathbf{L} \mathbf{y}}{\mathbf{y}^T \mathbf{D} \mathbf{y}} \geq \min_{\substack{S \subseteq \mathbb{R}^n \\ \dim(S)=k}} \max_{\substack{\mathbf{y} \in S \\ \mathbf{y} \neq \mathbf{0}}} \frac{\mathbf{y}^T \mathbf{L} \mathbf{y}}{d_{\max} \mathbf{y}^T \mathbf{y}} = \frac{1}{d_{\max}} \min_{\substack{S \subseteq \mathbb{R}^n \\ \dim(S)=k}} \max_{\substack{\mathbf{y} \in S \\ \mathbf{y} \neq \mathbf{0}}} \frac{\mathbf{y}^T \mathbf{L} \mathbf{y}}{\mathbf{y}^T \mathbf{y}} = \frac{\lambda_k}{d_{\max}},$$

where we used the Courant-Fischer Theorem one more time in the last equality. The upper bound is proved in a similar fashion. \square

Remark 23. In the proof of Theorem 40 we showed that

$$\nu_k = \min_{\substack{S \subseteq \mathbb{R}^n \\ \dim(S)=k}} \max_{\substack{\mathbf{y} \in S \\ \mathbf{y} \neq \mathbf{0}}} \frac{\mathbf{y}^T \mathbf{L} \mathbf{y}}{\mathbf{y}^T \mathbf{D} \mathbf{y}} = \max_{\substack{T \subseteq \mathbb{R}^n \\ \dim(T)=n-k+1}} \min_{\substack{\mathbf{y} \in T \\ \mathbf{y} \neq \mathbf{0}}} \frac{\mathbf{y}^T \mathbf{L} \mathbf{y}}{\mathbf{y}^T \mathbf{D} \mathbf{y}}. \quad (54)$$

This formulation of ν_k is an important alternative formula for computing ν_k , so let us collect it separately here.

Theorem 40 almost immediately gives us the following corollary.

Corollary 41. *Let $G = (V, E, w)$. Then the normalized graph Laplacian, \mathbf{N}_G , is positive semidefinite. Furthermore, $\nu_1 = 0$ and one can take $\phi_1(a) = \mathbf{d}^{1/2}(a) := \deg(a)^{1/2}$. Finally, $\nu_2 > 0$ if and only if G is connected.*

Proof. Since \mathbf{L} is positive semidefinite and $\lambda_1 = 0$, it is immediate from Theorem 40 that \mathbf{N} is positive semidefinite and $\nu_1 = 0$. Additionally, $\nu_2 > 0 \iff G$ is connected, immediately follows by combining Theorem 40 with Theorem 5 (recall the latter says that $\lambda_2 > 0 \iff G$ is connected). Now let us compute $\mathbf{N} \mathbf{d}^{1/2}$:

$$\mathbf{N} \mathbf{d}^{1/2} = \mathbf{D}^{-1/2} \mathbf{L} \mathbf{D}^{-1/2} \mathbf{d}^{1/2} = \mathbf{D}^{-1/2} \mathbf{L} \mathbf{1} = \mathbf{D}^{-1/2} \mathbf{0} = \mathbf{0}.$$

\square

Thus the normalized graph Laplacian is in many ways similar to the graph Laplacian. Corollary 41 indicates one difference, which is that $\psi_1 = \mathbf{1}$ while $\phi_1 = \mathbf{d}^{1/2}$ (although, notice if G is d -regular then ϕ_1 is constant). The next theorem gives another difference between \mathbf{L} and \mathbf{N} .

Theorem 42. *Let $G = (V, E, w)$. The mean eigenvalue of \mathbf{L}_G is the average degree of G , whereas the mean eigenvalue of \mathbf{N}_G is 1.*

Proof. Recall from Homework 02, Exercise 02, that the sum of the eigenvalues of a matrix is equal to its trace. Therefore:

$$\frac{1}{n} \sum_{k=1}^n \lambda_k = \frac{1}{n} \text{Tr}(\mathbf{L}) = \frac{1}{n} \text{Tr}(\mathbf{D} - \mathbf{M}) = \frac{1}{n} \text{Tr}(\mathbf{D}) = \frac{1}{n} \sum_{a \in V} \deg(a),$$

On the other hand,

$$\frac{1}{n} \sum_{k=1}^n \nu_k = \frac{1}{n} \text{Tr}(\mathbf{N}) = \frac{1}{n} \text{Tr}(\mathbf{I} - \mathbf{D}^{-1/2} \mathbf{M} \mathbf{D}^{-1/2}) = \frac{1}{n} \text{Tr}(\mathbf{I}) = \frac{1}{n} \cdot n = 1.$$

□

Theorem 42 gives another indication that \mathbf{N} is a degree independent representation of G , whereas \mathbf{L} is not. Here is another fact along similar lines.

Theorem 43. *Let $G = (V, E, w)$. All eigenvalues of \mathbf{N}_G lie in the interval $[0, 2]$, i.e.,*

$$0 = \nu_1 \leq \nu_2 \leq \cdots \leq \nu_n \leq 2.$$

I leave the proof of Theorem 43 to you in your homework.

Proof. We know from (54) that

$$\nu_n = \max_{\substack{\mathbf{y} \in \mathbb{R}^n \\ \mathbf{y} \neq \mathbf{0}}} \frac{\mathbf{y}^T \mathbf{L} \mathbf{y}}{\mathbf{y}^T \mathbf{D} \mathbf{y}}.$$

We also know that

$$\mathbf{y}^T \mathbf{L} \mathbf{y} = \sum_{(a,b) \in E} w(a,b) (\mathbf{y}(a) - \mathbf{y}(b))^2 \quad \text{and} \quad \mathbf{y}^T \mathbf{D} \mathbf{y} = \sum_{a \in V} \deg(a) \mathbf{y}(a)^2.$$

Now, for any two real numbers $\alpha, \beta \in \mathbb{R}$, we have (if you have not seen this inequality before, you should verify it for yourself):

$$(\alpha + \beta)^2 \leq 2(\alpha^2 + \beta^2).$$

Therefore,

$$\begin{aligned}
\mathbf{y}^T \mathbf{L} \mathbf{y} &\leq 2 \sum_{(a,b) \in E} w(a,b) (\mathbf{y}(a)^2 + \mathbf{y}(b)^2) \\
&= 2 \sum_{(a,b) \in E} w(a,b) \mathbf{y}(a)^2 + 2 \sum_{(a,b) \in E} w(a,b) \mathbf{y}(b)^2 \\
&= 4 \sum_{(a,b) \in E} w(a,b) \mathbf{y}(a)^2 \\
&= 2 \sum_{a \in V} \sum_{b \in N(a)} w(a,b) \mathbf{y}(a)^2 \\
&= 2 \sum_{a \in V} \mathbf{y}(a)^2 \sum_{b \in N(a)} w(a,b) \\
&= 2 \sum_{a \in V} \deg(a) \mathbf{y}(a)^2
\end{aligned}$$

Thus:

$$\nu_n = \max_{\substack{\mathbf{y} \in \mathbb{R}^n \\ \mathbf{y} \neq \mathbf{0}}} \frac{\mathbf{y}^T \mathbf{L} \mathbf{y}}{\mathbf{y}^T \mathbf{D} \mathbf{y}} \leq \max_{\substack{\mathbf{y} \in \mathbb{R}^n \\ \mathbf{y} \neq \mathbf{0}}} \frac{2 \sum_{a \in V} \deg(a) \mathbf{y}(a)^2}{\sum_{a \in V} \deg(a) \mathbf{y}(a)^2} = 2.$$

□

As our last general fact about \mathbf{N} , we give an analogue of Theorem 13 (recall this theorem gave a formula for λ_2).

Theorem 44. *Let $G = (V, E, w)$ and let $0 = \nu_1 \leq \nu_2 \leq \dots \leq \nu_n \leq 2$ be the eigenvalues of \mathbf{N}_G . Then*

$$\nu_2 = \min_{\substack{\mathbf{x} \in \mathbb{R}^n \\ \langle \mathbf{x}, \mathbf{d}^{1/2} \rangle = 0}} \frac{\mathbf{x}^T \mathbf{N} \mathbf{x}}{\mathbf{x}^T \mathbf{x}} = \min_{\substack{\mathbf{y} \in \mathbb{R}^n \\ \langle \mathbf{y}, \mathbf{d} \rangle = 0}} \frac{\mathbf{y}^T \mathbf{L} \mathbf{y}}{\mathbf{y}^T \mathbf{D} \mathbf{y}}.$$

Proof. The proof of the first equality is nearly identical to the proof of Theorem 13, remembering that $\phi_1 = \mathbf{d}^{1/2} / \|\mathbf{d}^{1/2}\|$ and that we can take the eigenvectors $\phi_1, \phi_2, \dots, \phi_n$ of \mathbf{N} to be an ONB for \mathbb{R}^n .

For the proof of the second equality, make the change of variables $\mathbf{y} = \mathbf{D}^{-1/2} \mathbf{x}$; we know from the proof of Theorem 40 that we get the argument of the minimum. Note also that with this change of variables:

$$0 = \langle \mathbf{x}, \mathbf{d}^{1/2} \rangle = \langle \mathbf{D}^{1/2} \mathbf{y}, \mathbf{d}^{1/2} \rangle = \langle \mathbf{y}, \mathbf{D}^{1/2} \mathbf{d}^{1/2} \rangle = \langle \mathbf{y}, \mathbf{d} \rangle.$$

□

Lecture 16: Graph Conductance

March 16, 2021

Lecturer: Matthew Hirn

25 Conductance

The conductance of a weighted graph $G = (V, E, w)$ provides an alternative measure to the isoperimetric ratio by which to measure its connectivity. We set some preliminary notation before defining the conductance. First, recall that in a weighted graph the degree of a vertex $a \in V$ is the sum of the weights of the edges connected to a :

$$d(a) := \deg(a) := \sum_{b \in N(a)} w(a, b).$$

For a subset $S \subseteq V$, define $d(S)$ as the sum of the degrees of all vertices in S :

$$d(S) := \sum_{a \in S} \deg(a) = \sum_{a \in S} \sum_{b \in N(a)} w(a, b).$$

Notice that

$$d(V) = 2 \sum_{(a,b) \in E} w(a, b).$$

For a subset $F \subseteq E$, define $w(F)$ to be the sum of the weights of the edges in F :

$$w(F) := \sum_{(a,b) \in F} w(a, b).$$

Finally, recall that the boundary of S is:

$$\partial S := \{(a, b) \in E : a \in S \text{ and } b \in V - S\}.$$

We define the *conductance* of S as

$$\varphi(S) := \frac{w(\partial S)}{\min(d(S), d(V - S))}.$$

The conductance of the graph G minimizes $\varphi(S)$ over all subsets S :

$$\varphi_G := \min_{S \subseteq V} \varphi(S).$$

In order to get a better feel for the conductance of a graph, it may be useful to consider the following calculation for $d(S)$:

$$\begin{aligned}
d(S) &= \sum_{a \in S} \sum_{b \in N(a)} w(a, b) \\
&= \sum_{a \in S} \sum_{\substack{(a,b) \in E \\ b \in S}} w(a, b) + \sum_{a \in S} \sum_{\substack{(a,b) \in E \\ b \in V-S}} w(a, b) \\
&= 2w(E(S)) + w(\partial S).
\end{aligned} \tag{55}$$

Thus, if $\min(d(S), d(V-S)) = d(S)$, then

$$\varphi(S) = \frac{w(\partial S)}{2w(E(S)) + w(\partial S)}.$$

On the other hand, note that $w(\partial S) = w(\partial(V-S))$. Thus, if $\min(d(S), d(V-S)) = d(V-S)$, we have

$$\varphi(S) = \frac{w(\partial(V-S))}{2w(E(V-S)) + w(\partial(V-S))}.$$

These two facts also imply that

$$\varphi(S) \leq 1 \quad \text{and} \quad \varphi(S) = \varphi(V-S).$$

Remark 24. Let us compare the isoperimetric ratio of a set S to its conductance for an unweighted graph $G = (V, E)$. Let us assume that $\min(d(S), d(V-S)) = d(S)$ for a bit of added simplicity. In this case we have

$$\theta(S) = \frac{|\partial S|}{|S|} \quad \text{and} \quad \varphi(S) = \frac{|\partial S|}{d(S)}.$$

Thus in both cases the numerator is the same, but it is the denominator that changes. In particular, the isoperimetric ratio, $\theta(S)$, places more importance on the number of vertices being removed if one were to remove $G(S) = (S, E(S))$ from G , as indicated by having $|S|$ in the denominator. On the other hand, the conductance places more importance on the number of edges being removed, since $d(S) = 2|E(S)| + |\partial S|$.

Remark 25. If $G = (V, E)$ is d -regular and $|S| \leq n/2$, then $\theta(S)$ and $\varphi(S)$ differ by a factor of d :

$$\varphi(S) = \frac{|\partial S|}{d(S)} = \frac{|\partial S|}{d|S|} = \frac{\theta(S)}{d}.$$

The above considerations indicate the conductance is a degree-invariant measure of connectivity of a weighted graph $G = (V, E, w)$. It thus makes sense to try to relate it to ν_2 , the second eigenvalue of the normalized graph Laplacian \mathbf{N} , as opposed to λ_2 , the second eigenvalue of the graph Laplacian \mathbf{L} . Our goal is to prove that

$$\frac{\nu_2}{2} \leq \varphi_G \leq \sqrt{2\nu_2},$$

and to construct a set $S \subset V$ such that $\varphi(S) \leq \sqrt{2\nu_2}$. In doing so, we will have shown that ν_2 characterizes the connectivity of G , and furthermore, we will be able to turn the construction of S into an algorithm that allows us to compute a nearly optimal partition of G . The next theorem provides the lower bound for φ_G , which is an analogue of the similar result for the isoperimetric ratio given in Theorem 39.

Theorem 45. *Let $G = (V, E, w)$, let $0 = \nu_1 \leq \nu_2 \leq \dots \leq \nu_n \leq 2$ be its normalized graph Laplacian eigenvalues, and let $S \subset V$. Then*

$$\frac{d(V)w(\partial S)}{d(S)d(V-S)} \geq \nu_2,$$

and as such,

$$\varphi(S) \geq \nu_2/2 \implies \varphi_G \geq \nu_2/2.$$

Proof. The proof is pretty similar to the proof of Theorem 39. By Theorem 44 we know that

$$\nu_2 = \min_{\substack{\mathbf{y} \in \mathbb{R}^n \\ \langle \mathbf{y}, \mathbf{d} \rangle = 0}} \frac{\mathbf{y}^T \mathbf{L} \mathbf{y}}{\mathbf{y}^T \mathbf{D} \mathbf{y}},$$

and so we again want to use the test vector technique to bound ν_2 from above. We need a test vector that will give us the right quantities for the set S . As before, $\mathbf{y} = \mathbf{1}_S$ would be a nice choice, but it is not orthogonal to \mathbf{d} . So instead we select:

$$\mathbf{y} = \mathbf{1}_S - \sigma \mathbf{1}, \quad \sigma = \frac{d(S)}{d(V)}.$$

We check that $\langle \mathbf{y}, \mathbf{d} \rangle = 0$; indeed:

$$\langle \mathbf{y}, \mathbf{d} \rangle = (\mathbf{1}_S - \sigma \mathbf{1})^T \mathbf{d} = \mathbf{1}_S^T \mathbf{d} - \sigma \mathbf{1}^T \mathbf{d} = d(S) - \sigma d(V) = d(S) - \frac{d(S)}{d(V)} d(V) = 0.$$

Using a similar argument as in the proof of Theorem 39, let us compute:

$$\begin{aligned} \mathbf{y}^T \mathbf{L} \mathbf{y} &= \sum_{(a,b) \in E} w(a,b) ((\mathbf{1}_S(a) - \sigma) - (\mathbf{1}_S(b) - \sigma))^2 \\ &= \sum_{(a,b) \in E} w(a,b) (\mathbf{1}_S(a) - \mathbf{1}_S(b))^2 \\ &= \sum_{(a,b) \in \partial S} w(a,b) \\ &= w(\partial S). \end{aligned}$$

Now we compute the denominator:

$$\begin{aligned}
\mathbf{y}^T \mathbf{D} \mathbf{y} &= \sum_{a \in V} \mathbf{d}(a) \mathbf{y}(a)^2 = \sum_{a \in S} (1 - \sigma)^2 \mathbf{d}(a) + \sum_{a \in V-S} (-\sigma)^2 \mathbf{d}(a) \\
&= (1 - \sigma)^2 d(S) + \sigma^2 d(V - S) \\
&= d(S) - 2\sigma d(S) + \sigma^2 d(S) + \sigma^2 d(V - S) \\
&= d(S) - 2\sigma d(S) + \sigma^2 d(V) \\
&= d(S) - 2\sigma d(S) + \sigma \frac{d(S)}{d(V)} d(V) \\
&= d(S) - \sigma d(S) \\
&= \left(1 - \frac{d(S)}{d(V)}\right) d(S) \\
&= \frac{(d(V) - d(S))d(S)}{d(V)} \\
&= \frac{d(V - S)d(S)}{d(V)}.
\end{aligned}$$

Putting together our computations for $\mathbf{y}^T \mathbf{L} \mathbf{y}$ and $\mathbf{y}^T \mathbf{D} \mathbf{y}$ we have the result:

$$\nu_2 \leq \frac{\mathbf{y}^T \mathbf{L} \mathbf{y}}{\mathbf{y}^T \mathbf{D} \mathbf{y}} = \frac{d(V)w(\partial S)}{d(V - S)d(S)}.$$

To complete the proof, note that since $d(V) = d(S) + d(V - S)$, we have:

$$\max(d(S), d(V - S)) \geq d(V)/2 \implies \frac{d(V)}{\max(d(S), d(V - S))} \leq 2.$$

Therefore,

$$\begin{aligned}
\nu_2 &\leq \frac{d(V)w(\partial S)}{d(S)d(V - S)} \\
&= \frac{d(V)w(\partial S)}{\max(d(S), d(V - S)) \min(d(S), d(V - S))} \\
&\leq \frac{2w(\partial S)}{\min(d(S), d(V - S))} \\
&= 2\varphi(S).
\end{aligned}$$

□

Lecture 17: Preliminaries to Cheeger's Inequality

March 18, 2021

Lecturer: Matthew Hirn

26 Cheeger's inequality

Our goal now is to prove an upper bound for φ_G , the conductance of a weighted graph $G = (V, E, w)$, in terms of ν_2 , the second eigenvalue of \mathbf{N}_G . To that end, we will prove Cheeger's inequality:

$$\varphi_G \leq \sqrt{2\nu_2}. \quad (56)$$

In order to prove (56), we will consider vectors $\mathbf{y} \in \mathbb{R}^n$ such that $\langle \mathbf{y}, \mathbf{d} \rangle = 0$ and the sets $S_{\mathbf{y},t} \subset V$ defined as

$$S_{\mathbf{y},t} := \{a \in V : \mathbf{y}(a) \geq t\}.$$

We will prove that for any such \mathbf{y} , there is a $t \in \mathbb{R}$ such that $\varphi(S_{\mathbf{y},t})$ is small. Here small means we will show:

$$\frac{\varphi(S_{\mathbf{y},t})^2}{2} \leq \frac{\mathbf{y}^T \mathbf{L} \mathbf{y}}{\mathbf{y}^T \mathbf{D} \mathbf{y}}, \quad \text{for appropriately chosen } t.$$

Notice if we set $\mathbf{y} = \mathbf{D}^{-1/2} \phi_2$ then (56) follows from the previous equation and Theorem 44. Our proof for showing the existence of such a value t will be constructive and will allow us to develop an algorithm for finding the set $S_{\mathbf{y},t}$.

While proving a lower bound for φ_G was relatively straightforward and leveraged ideas we have seen before (namely the test vector technique), proving the upper bound (56) for φ_G is significantly harder. We will need several lemmas to prove our main result. Let us collect some of the general ones (i.e., not specific to graphs) now.

Lemma 46. *Let $A, B, C, D \geq 0$. Then*

$$\frac{A+B}{C+D} \geq \min \left(\frac{A}{C}, \frac{B}{D} \right).$$

Proof. Without loss of generality suppose

$$\frac{A}{C} = \min \left(\frac{A}{C}, \frac{B}{D} \right) \implies \frac{A}{C} \leq \frac{B}{D} \implies \frac{D}{C} \leq \frac{B}{A}.$$

Then:

$$\frac{A+B}{C+D} = \frac{A(1+B/A)}{C(1+D/C)} \geq \frac{A}{C}.$$

□

Lemma 47. *Let $\alpha, \beta \in \mathbb{R}$. Then for all $0 < \delta < 1$,*

$$(\alpha - \beta)^2 \geq \delta \alpha^2 - \frac{\delta}{1 - \delta} \beta^2.$$

Proof. If $\alpha = 0$ or $\beta = 0$ the result is clearly true. So suppose $\alpha, \beta \neq 0$. Define the function

$$f(\delta) = \delta \alpha^2 - \frac{\delta}{1 - \delta} \beta^2.$$

Calculating the derivative of f we obtain

$$f'(\delta) = \alpha^2 - (1 - \delta)^{-1} \beta^2 - \delta(1 - \delta)^{-2} \beta^2.$$

Setting $f'(\delta) = 0$ and solving we obtain:

$$f'(\delta^*) = 0 \implies \delta^* = 1 \pm \frac{\beta}{\alpha}.$$

Let us also compute the second derivative of f so that we can check whether a critical point is a maximum or a minimum. We obtain:

$$f''(\delta) = -2(1 - \delta)^{-2} \beta^2 - 2\delta(1 - \delta)^{-3} \beta^2.$$

Recall that a critical point δ^* is a maximum if $f''(\delta^*) < 0$. Since $\beta \neq 0$, we see that

$$f''(\delta) < 0 \iff (1 - \delta)^{-2}(1 + \delta(1 - \delta)^{-1}) > 0. \quad (57)$$

Now observe that

$$\begin{aligned} 1 - \beta/\alpha < 1 &\implies 1 + \beta/\alpha > 1, \\ 1 + \beta/\alpha < 1 &\implies 1 - \beta/\alpha > 1. \end{aligned}$$

We conclude that only one of the critical points can lie in the interval $(-\infty, 1)$. Let us examine the two possible cases.

Suppose that $1 - \beta/\alpha \in (-\infty, 1)$. Let us evaluate the right hand side of (57) at this point:

$$\text{RHS of (57) at } 1 - \beta/\alpha = \alpha^3/\beta^3.$$

Now, if $1 - \beta/\alpha < 1$, then $\beta/\alpha > 0$ and so $\alpha^3/\beta^3 > 0$ as well. Therefore $f(1 - \beta/\alpha)$ is the maximum value of $f(\delta)$ on the interval $(-\infty, 1)$. We compute:

$$f(1 - \beta/\alpha) = (\alpha - \beta)^2,$$

and so the lemma holds in this case.

Now suppose that $1 + \beta/\alpha \in (-\infty, 1)$. Evaluating the right hand side of (57) at this point we obtain:

$$\text{RHS of (57) at } 1 + \beta/\alpha = -\alpha^3/\beta^3.$$

Now, if $1 + \beta/\alpha < 1$, then $\beta/\alpha < 0$ and so $\alpha^3/\beta^3 < 0$. Therefore $-\alpha^3/\beta^3 > 0$ and $f(1 + \alpha/\beta)$ is the maximum value of $f(\delta)$ on the interval $(-\infty, 1)$. We compute:

$$f(1 + \beta/\alpha) = (\alpha + \beta)^3.$$

Notice, though, that $\beta/\alpha < 0$ implies that α and β must have different signs. Therefore $(\alpha - \beta)^2 \geq (\alpha + \beta)^2$ and the result holds in this case as well. \square

Lemma 48. *Let $z_1, \dots, z_k \in \mathbb{R}$ be such that*

$$z_1 \geq z_2 \geq \dots \geq z_k \geq 0.$$

If $\alpha_1, \dots, \alpha_k \in \mathbb{R}$ and $\beta_1, \dots, \beta_k \in \mathbb{R}$ satisfy

$$\sum_{i=1}^k \alpha_i z_i \geq \sum_{i=1}^k \beta_i z_i,$$

then there exists $j \in \{1, \dots, k\}$ such that

$$\sum_{i=1}^j \alpha_i \geq \sum_{i=1}^j \beta_i. \quad (58)$$

Additionally, we may take j such that either $z_j > z_{j+1}$ or $j = k$.

Proof. For the existence of at least one such j , we will do a proof by induction on k . The result is clearly true for $k = 1$. Now assume the result is true for all $1 \leq k \leq \ell$ ($\ell \geq 1$) and let us prove the result for $k = \ell + 1$. Define

$$\tilde{\alpha}_2 = \alpha_1 \frac{z_1}{z_2} + \alpha_2 \quad \text{and} \quad \tilde{\beta}_2 = \beta_1 \frac{z_1}{z_2} + \beta_2,$$

so that

$$\tilde{\alpha}_2 z_2 + \sum_{i=3}^{\ell+1} \alpha_i z_i = \sum_{i=1}^{\ell+1} \alpha_i z_i \geq \sum_{i=1}^{\ell+1} \beta_i z_i = \tilde{\beta}_2 z_2 + \sum_{i=3}^{\ell+1} \beta_i z_i.$$

By the inductive hypothesis, there exists $j \in \{2, \dots, \ell + 1\}$ so that

$$\tilde{\alpha}_2 + \sum_{i=3}^j \alpha_i \geq \tilde{\beta}_2 + \sum_{i=3}^j \beta_i \quad \implies \quad \alpha_1 \frac{z_1}{z_2} + \sum_{i=2}^j \alpha_i \geq \beta_1 \frac{z_1}{z_2} + \sum_{i=2}^j \beta_i.$$

If $\alpha_1 \geq \beta_1$ we are done, so let us consider the case when $\alpha_1 < \beta_1$. Since $z_1 \geq z_2$, we have

$$\left(\frac{z_1}{z_2} - 1 \right) \alpha_1 < \left(\frac{z_1}{z_2} - 1 \right) \beta_1 \quad \implies \quad - \left(\frac{z_1}{z_2} - 1 \right) \alpha_1 > - \left(\frac{z_1}{z_2} - 1 \right) \beta_1.$$

But we also have:

$$\sum_{i=1}^j \alpha_i = \left(\alpha_1 \frac{z_1}{z_2} + \sum_{i=2}^j \alpha_i \right) - \left(\frac{z_1}{z_2} - 1 \right) \alpha_1 > \left(\beta_1 \frac{z_1}{z_2} + \sum_{i=2}^j \beta_i \right) - \left(\frac{z_1}{z_2} - 1 \right) \beta_1 = \sum_{i=1}^j \beta_i.$$

For the second statement, let j_0 be the minimum possible index that satisfies (58). If $j_0 = k$ we are done, or if $z_{j_0} > z_{j_0+1}$ we are done. So let us suppose that $j_0 < k$ and $z_{j_0} = z_{j_0+1}$. Consider a new sequence of length $k-1$,

$$z_1 \geq \cdots \geq z_{j_0-1} \geq z_{j_0} \geq z_{j_0+2} \geq \cdots \geq z_k,$$

with corresponding weights

$$\alpha_1, \dots, \alpha_{j_0-1}, \alpha_{j_0} + \alpha_{j_0+1}, \alpha_{j_0+2}, \dots, \alpha_k,$$

and

$$\beta_1, \dots, \beta_{j_0-1}, \beta_{j_0} + \beta_{j_0+1}, \beta_{j_0+2}, \dots, \beta_k.$$

Apply the result we just proved to this setup to get an index j for this new sequence. We know that $j \geq j_0$ since otherwise (58) will hold for $j < j_0$ which contradicts the minimality of j_0 . But for $j \geq j_0$ we have

$$\sum_{i=1}^{j_1} \alpha_i \geq \sum_{i=1}^{j_1} \beta_i,$$

where $j_1 \geq j_0 + 1$. If $z_{j_1} > z_{j_1+1}$ or if $j_1 = k$, we are finished. If not, then $z_{j_1} = z_{j_1+1}$. If that is the case, then it must be that $z_{j_0} = z_{j_0+1} = z_{j_0+2}$. Consider the following sequence of length $k-2$,

$$z_1 \geq \cdots \geq z_{j_0-1} \geq z_{j_0} \geq z_{j_0+3} \geq \cdots \geq z_k,$$

with weights

$$\alpha_1, \dots, \alpha_{j_0-1}, \alpha_{j_0} + \alpha_{j_0+1} + \alpha_{j_0+2}, \alpha_{j_0+3}, \dots, \alpha_k,$$

and

$$\beta_1, \dots, \beta_{j_0-1}, \beta_{j_0} + \beta_{j_0+1} + \beta_{j_0+2}, \beta_{j_0+3}, \dots, \beta_k.$$

Now apply the same argument. We conclude this time that there is a $j_2 \geq j_0 + 2$ such that

$$\sum_{i=1}^{j_2} \alpha_i \geq \sum_{i=1}^{j_2} \beta_i.$$

Again, if $z_{j_2} > z_{j_2+1}$ or if $j_2 = k$, we are finished. If not, repeat the argument again. If we repeat the argument ℓ times, we get an index $j_\ell \geq j_0 + \ell$ such that

$$\sum_{i=1}^{j_\ell} \alpha_i \geq \sum_{i=1}^{j_\ell} \beta_i.$$

Since $j_\ell \rightarrow \infty$ as $\ell \rightarrow \infty$, at some point we must either have $z_{j_\ell} > z_{j_\ell+1}$ or $j_\ell = k$. □

Lecture 18: Centered Vectors

March 23, 2021

Lecturer: Matthew Hirn

Now we introduce the notion of a centered vector. Let $G = (V, E, w)$ and let \mathbf{d} be the degree vector of G . We say that $\mathbf{z} : V \rightarrow \mathbb{R}$ is *centered with respect to \mathbf{d}* if

$$\sum_{a: \mathbf{z}(a) > 0} \mathbf{d}(a) \leq d(V)/2 \quad \text{and} \quad \sum_{a: \mathbf{z}(a) < 0} \mathbf{d}(a) \leq d(V)/2.$$

It is not hard to center a function $\mathbf{x} : V \rightarrow \mathbb{R}$. Let $V = \{1, \dots, n\}$ and order the vertices of G so that

$$\mathbf{x}(1) \geq \mathbf{x}(2) \geq \dots \geq \mathbf{x}(n).$$

To center \mathbf{x} , let $k \in V$ be the least number for which

$$\sum_{a=1}^k \mathbf{d}(a) > d(V)/2. \tag{59}$$

Then set

$$\mathbf{z} = \mathbf{x} - \mathbf{x}(k)\mathbf{1}.$$

Clearly $\mathbf{z}(k) = 0$, and one can show that \mathbf{z} is centered with respect to \mathbf{d} . Indeed, since k is the least index such that (59) holds, we have that

$$\sum_{a: \mathbf{z}(a) > 0} \mathbf{d}(a) \leq \sum_{i=1}^{k-1} \mathbf{d}(a) \leq d(V)/2.$$

Additionally, again using (59),

$$\sum_{a: \mathbf{z}(a) < 0} \mathbf{d}(a) \leq \sum_{i=k+1}^n \mathbf{d}(a) = d(V) - \sum_{i=1}^k \mathbf{d}(a) < d(V)/2.$$

In the proof of Cheeger's inequality, it will be useful to replace \mathbf{y} with its centered version. The following lemma shows that in doing so we do not increase the generalized Rayleigh quotient.

Lemma 49. *Let $G = (V, E, w)$ and let $\mathbf{x} : V \rightarrow \mathbb{R}$. Define*

$$\mathbf{x}_s := \mathbf{x} - s\mathbf{1}.$$

Then

$$\min_s \mathbf{x}_s^T \mathbf{D} \mathbf{x}_s,$$

is achieved at the s such that $\langle \mathbf{x}_s, \mathbf{d} \rangle = 0$. As such, if $\langle \mathbf{y}, \mathbf{d} \rangle = 0$ and $\mathbf{z} = \mathbf{y} - \mathbf{y}(k)\mathbf{1}$ is the centered version of \mathbf{y} , then

$$\frac{\mathbf{z}^T \mathbf{L} \mathbf{z}}{\mathbf{z}^T \mathbf{D} \mathbf{z}} \leq \frac{\mathbf{y}^T \mathbf{L} \mathbf{y}}{\mathbf{y}^T \mathbf{D} \mathbf{y}}.$$

Proof. We first compute:

$$\mathbf{x}_s^T \mathbf{D} \mathbf{x}_s = \sum_{a \in V} \mathbf{d}(a) \mathbf{x}_s(a)^2 = \sum_{a \in V} \mathbf{d}(a) (\mathbf{x}(a) - s)^2 = \sum_{a \in V} \mathbf{d}(a) [\mathbf{x}(a)^2 - 2s\mathbf{x}(a) + s^2].$$

Therefore the derivative of the function $f(s) = \mathbf{x}_s^T \mathbf{D} \mathbf{x}_s$ is:

$$f'(s) = \sum_{a \in V} \mathbf{d}(a) (2s - 2\mathbf{x}(a)) = 2 \sum_{a \in V} \mathbf{d}(a) (s - \mathbf{x}(a)).$$

Setting $f'(s) = 0$, we conclude that

$$f'(s) = 0 \implies 0 = \sum_{a \in V} \mathbf{d}(a) (\mathbf{x}(a) - s) = \langle \mathbf{x} - s\mathbf{1}, \mathbf{d} \rangle = \langle \mathbf{x}_s, \mathbf{d} \rangle.$$

For the second statement, using Lemma 4 and the fact that $\langle \mathbf{1}, \boldsymbol{\psi}_i \rangle = 0$ for $i \geq 2$, we first note that

$$\mathbf{z}^T \mathbf{L} \mathbf{z} = \sum_{i=2}^n \lambda_i |\langle \mathbf{z}, \boldsymbol{\psi}_i \rangle|^2 = \sum_{i=2}^n \lambda_i |\langle \mathbf{y} - \mathbf{y}(k)\mathbf{1}, \boldsymbol{\psi}_i \rangle|^2 = \sum_{i=2}^n \lambda_i |\langle \mathbf{y}, \boldsymbol{\psi}_i \rangle|^2 = \mathbf{y}^T \mathbf{L} \mathbf{y}.$$

Additionally, since $\langle \mathbf{y}, \mathbf{d} \rangle = 0$, by the first part of this theorem we know that

$$\mathbf{y}^T \mathbf{D} \mathbf{y} \leq \mathbf{z}^T \mathbf{D} \mathbf{z}.$$

Combining these previous two statements we conclude:

$$\frac{\mathbf{z}^T \mathbf{L} \mathbf{z}}{\mathbf{z}^T \mathbf{D} \mathbf{z}} = \frac{\mathbf{y}^T \mathbf{L} \mathbf{y}}{\mathbf{z}^T \mathbf{D} \mathbf{z}} \leq \frac{\mathbf{y}^T \mathbf{L} \mathbf{y}}{\mathbf{y}^T \mathbf{D} \mathbf{y}}.$$

□

The next lemma shows that in fact we can replace the generalized Rayleigh quotient of a centered vector \mathbf{z} by the Rayleigh quotient of the vector that only retains either the positive entries of \mathbf{z} or the negative entries of \mathbf{z} .

Lemma 50. Let $G = (V, E, w)$ and let $\mathbf{x} : V \rightarrow \mathbb{R}$. Define

$$\mathbf{x}_+(a) := \begin{cases} \mathbf{x}(a) & \text{if } \mathbf{x}(a) > 0 \\ 0 & \text{if } \mathbf{x}(a) \leq 0 \end{cases} \quad \text{and} \quad \mathbf{x}_-(a) := \begin{cases} \mathbf{x}(a) & \text{if } \mathbf{x}(a) < 0 \\ 0 & \text{if } \mathbf{x}(a) \geq 0 \end{cases}.$$

Then

$$\min \left(\frac{\mathbf{x}_+^T \mathbf{L} \mathbf{x}_+}{\mathbf{x}_+^T \mathbf{D} \mathbf{x}_+}, \frac{\mathbf{x}_-^T \mathbf{L} \mathbf{x}_-}{\mathbf{x}_-^T \mathbf{D} \mathbf{x}_-} \right) \leq \frac{\mathbf{x}^T \mathbf{L} \mathbf{x}}{\mathbf{x}^T \mathbf{D} \mathbf{x}}.$$

Proof. Using Lemma 46 we know that

$$\min \left(\frac{\mathbf{x}_+^T \mathbf{L} \mathbf{x}_+}{\mathbf{x}_+^T \mathbf{D} \mathbf{x}_+}, \frac{\mathbf{x}_-^T \mathbf{L} \mathbf{x}_-}{\mathbf{x}_-^T \mathbf{D} \mathbf{x}_-} \right) \leq \frac{\mathbf{x}_+^T \mathbf{L} \mathbf{x}_+ + \mathbf{x}_-^T \mathbf{L} \mathbf{x}_-}{\mathbf{x}_+^T \mathbf{D} \mathbf{x}_+ + \mathbf{x}_-^T \mathbf{D} \mathbf{x}_-}.$$

Thus we need to show

$$\frac{\mathbf{x}_+^T \mathbf{L} \mathbf{x}_+ + \mathbf{x}_-^T \mathbf{L} \mathbf{x}_-}{\mathbf{x}_+^T \mathbf{D} \mathbf{x}_+ + \mathbf{x}_-^T \mathbf{D} \mathbf{x}_-} \leq \frac{\mathbf{x}^T \mathbf{L} \mathbf{x}}{\mathbf{x}^T \mathbf{D} \mathbf{x}}.$$

Noting that $\mathbf{x} = \mathbf{x}_+ + \mathbf{x}_-$ and $\mathbf{x}_+(a)\mathbf{x}_-(a) = 0$ for all $a \in V$, for the denominator we have:

$$\begin{aligned} \mathbf{x}^T \mathbf{D} \mathbf{x} &= (\mathbf{x}_+ + \mathbf{x}_-)^T \mathbf{D} (\mathbf{x}_+ + \mathbf{x}_-) \\ &= \mathbf{x}_+^T \mathbf{D} \mathbf{x}_+ + \mathbf{x}_-^T \mathbf{D} \mathbf{x}_- + 2 \sum_{a \in V} d(a) \mathbf{x}_+(a) \mathbf{x}_-(a) \\ &= \mathbf{x}_+^T \mathbf{D} \mathbf{x}_+ + \mathbf{x}_-^T \mathbf{D} \mathbf{x}_-. \end{aligned}$$

The numerator is a bit trickier. First, recall

$$\mathbf{x}^T \mathbf{L} \mathbf{x} = \sum_{(a,b) \in E} w(a,b) (\mathbf{x}(a) - \mathbf{x}(b))^2.$$

We can divide the edges E up into two types. Let the first type be those $(a,b) \in E$ such that $\mathbf{x}(a)$ and $\mathbf{x}(b)$ have the same sign. In that case we have $\mathbf{x}(a) - \mathbf{x}(b) = \mathbf{x}_+(a) - \mathbf{x}_+(b)$ and $\mathbf{x}_-(a) = \mathbf{x}_-(b) = 0$, or vice versa, and so

$$(\mathbf{x}(a) - \mathbf{x}(b))^2 = (\mathbf{x}_+(a) - \mathbf{x}_+(b))^2 + (\mathbf{x}_-(a) - \mathbf{x}_-(b))^2.$$

Let the other type of edge be when $\mathbf{x}(a)$ and $\mathbf{x}(b)$ have the opposite sign. Then:

$$\begin{aligned} (\mathbf{x}(a) - \mathbf{x}(b))^2 &= \mathbf{x}(a)^2 + \mathbf{x}(b)^2 - 2\mathbf{x}(a)\mathbf{x}(b) \\ &\geq \mathbf{x}(a)^2 + \mathbf{x}(b)^2 \\ &= (\mathbf{x}_+(a) - \mathbf{x}_+(b))^2 + (\mathbf{x}_-(a) - \mathbf{x}_-(b))^2. \end{aligned}$$

Either way we have $(\mathbf{x}(a) - \mathbf{x}(b))^2 \geq (\mathbf{x}_+(a) - \mathbf{x}_+(b))^2 + (\mathbf{x}_-(a) - \mathbf{x}_-(b))^2$ and so

$$\begin{aligned} \mathbf{x}^T \mathbf{L} \mathbf{x} &= \sum_{(a,b) \in E} w(a,b) (\mathbf{x}(a) - \mathbf{x}(b))^2 \\ &\geq \sum_{(a,b) \in E} [w(a,b) (\mathbf{x}_+(a) - \mathbf{x}_+(b))^2 + w(a,b) (\mathbf{x}_-(a) - \mathbf{x}_-(b))^2] \\ &= \mathbf{x}_+^T \mathbf{L} \mathbf{x}_+ + \mathbf{x}_-^T \mathbf{L} \mathbf{x}_-. \end{aligned}$$

That completes the proof. □

Lecture 19: Cheeger's Inequality and Spectral Clustering

March 25, 2021

Lecturer: Matthew Hirn

At last we are ready to formally state and prove our theorem on Cheeger's inequality.

Theorem 51 (Cheeger's inequality). *Let $G = (V, E, w)$ and let $\mathbf{y} : V \rightarrow \mathbb{R}$ be orthogonal to \mathbf{d} . Then there exists a $t \in \mathbb{R}$ for which the set of vertices*

$$S = \{a \in V : \mathbf{y}(a) \geq t\},$$

satisfies

$$\frac{\varphi(S)^2}{2} \leq \frac{\mathbf{y}^T \mathbf{L} \mathbf{y}}{\mathbf{y}^T \mathbf{D} \mathbf{y}}.$$

As such, if we set $\mathbf{y} = \mathbf{D}^{-1/2} \phi_2$, where ϕ_2 is the second eigenvector of \mathbf{N} , the normalized graph Laplacian, then we obtain a set $S \subset V$ such that

$$\frac{\varphi(S)^2}{2} \leq \nu_2,$$

where ν_2 is the eigenvalue of ϕ_2 . It follows that

$$\varphi_G \leq \sqrt{2\nu_2}.$$

Proof. Let $V = \{1, \dots, n\}$, order the vertices so that $\mathbf{y}(1) \geq \mathbf{y}(2) \geq \dots \geq \mathbf{y}(n)$, and let \mathbf{z} be the centered version of \mathbf{y} :

$$\mathbf{z} = \mathbf{y} - \mathbf{y}(k) \mathbf{1},$$

where $k \in V$ is the least number for which (59) holds. Applying Lemma 49 and Lemma 50, we have

$$\min \left(\frac{\mathbf{z}_+^T \mathbf{L} \mathbf{z}_+}{\mathbf{z}_+^T \mathbf{D} \mathbf{z}_+}, \frac{\mathbf{z}_-^T \mathbf{L} \mathbf{z}_-}{\mathbf{z}_-^T \mathbf{D} \mathbf{z}_-} \right) \leq \frac{\mathbf{z}^T \mathbf{L} \mathbf{z}}{\mathbf{z}^T \mathbf{D} \mathbf{z}} \leq \frac{\mathbf{y}^T \mathbf{L} \mathbf{y}}{\mathbf{y}^T \mathbf{D} \mathbf{y}}. \quad (60)$$

Without loss of generality let us assume that

$$\frac{\mathbf{z}_+^T \mathbf{L} \mathbf{z}_+}{\mathbf{z}_+^T \mathbf{D} \mathbf{z}_+} = \min \left(\frac{\mathbf{z}_+^T \mathbf{L} \mathbf{z}_+}{\mathbf{z}_+^T \mathbf{D} \mathbf{z}_+}, \frac{\mathbf{z}_-^T \mathbf{L} \mathbf{z}_-}{\mathbf{z}_-^T \mathbf{D} \mathbf{z}_-} \right).$$

We have

$$\mathbf{z}(1) \geq \mathbf{z}(2) \geq \dots \geq \mathbf{z}(k-1) \geq \mathbf{z}(k) = 0 \geq \mathbf{z}(k+1) \geq \dots \geq \mathbf{z}(n).$$

For each $j \in V$ define the set $S_j \subseteq V$ as:

$$S_j := \{a \in V : 1 \leq a \leq j\}.$$

Now let us collect the vertices on which \mathbf{z} takes the same value, so that:

$$\begin{aligned} \mathbf{z}(1) &= \cdots = \mathbf{z}(j_1) > \mathbf{z}(j_1 + 1), \quad j_1 \geq 1, \\ \mathbf{z}(j_1 + 1) &= \cdots = \mathbf{z}(j_2) > \mathbf{z}(j_2 + 1), \quad j_2 \geq j_1 + 1 \\ &\vdots \\ \mathbf{z}(j_{m-1} + 1) &= \cdots = \mathbf{z}(j_m), \quad j_m = n. \end{aligned}$$

Notice we have

$$\mathbf{z}(j_1) > \mathbf{z}(j_2) > \cdots > \mathbf{z}(j_{\ell-1}) > \mathbf{z}(j_\ell) = 0 > \mathbf{z}(j_{\ell+1}) > \cdots > \mathbf{z}(j_m),$$

where ℓ is some index between 1 and m . Note that $k \in S_{j_\ell}$ but $k \notin S_{j_{\ell-1}}$. It follows that:

$$\begin{aligned} S_{j_i} &= \{a \in V : 1 \leq a \leq j_i\} = \{a \in V : \mathbf{z}(a) \geq \mathbf{z}(j_i)\} \\ &= \{a \in V : \mathbf{y}(a) - \mathbf{y}(k) \geq \mathbf{y}(j_i) - \mathbf{y}(k)\} \\ &= \{a \in V : \mathbf{y}(a) \geq \mathbf{y}(j_i)\}. \end{aligned}$$

The remainder of the proof will be devoted to showing that one of the sets S_{j_i} satisfies the theorem. We will have two cases:

1. $\ell \geq 2$
2. $\ell = 1$

For the first case in which $\ell \geq 2$, we will prove there exists an $i^* \in \{1, \dots, \ell - 1\}$ such that

$$\frac{\varphi(S_{j_{i^*}})^2}{2} \leq \frac{\mathbf{z}_+ \mathbf{L} \mathbf{z}_+}{\mathbf{z}_+ \mathbf{D} \mathbf{z}_+}. \quad (61)$$

That will prove the result with $t = \mathbf{y}(j_{i^*})$. To facilitate the proof of (61), set:

$$\sigma = \min_{j \in \{j_1, \dots, j_{\ell-1}\}} \varphi(S_j).$$

By the definition of the index k ,

$$\sum_{a=1}^{k-1} \mathbf{d}(a) \leq \frac{d(V)}{2},$$

and so $d(S_j) \leq d(V - S_j)$ for all $j \in \{j_1, \dots, j_{\ell-1}\}$. It follows that

$$\sigma \leq \varphi(S_j) = \frac{w(\partial S_j)}{\min(d(S_j), d(V - S_j))} = \frac{w(\partial S_j)}{d(S_j)}, \quad j \in \{j_1, \dots, j_{\ell-1}\}. \quad (62)$$

Recall that $\varphi(S) \leq 1$. Now we have two sub-cases: (i) $\sigma = 1$; and (ii) $0 < \sigma < 1$ (note if $\sigma = 0$ we immediately satisfy (61)). Let us tackle the $\sigma = 1$ case first. If $\sigma = 1$ then $\varphi(S_j) = 1$ for all $j \in \{j_1, \dots, j_{\ell-1}\}$, and so showing (61) means we need to show

$$\frac{1}{2} \leq \frac{\mathbf{z}_+ \mathbf{L} \mathbf{z}_+}{\mathbf{z}_+ \mathbf{D} \mathbf{z}_+}. \quad (63)$$

In (55) we showed previously that $d(S) = 2w(E(S)) + w(\partial S)$, and so $\varphi(S_j) = 1$ if and only if $w(E(S_j)) = 0$. Thus, if $\sigma = 1$, it must be that $w(E(S_{j_{\ell-1}})) = 0$, which means that $E(S_{j_{\ell-1}}) = \emptyset$. Recall that $z_+(a) = 0$ for all $a \in V - S_{j_{\ell-1}}$. We compute:

$$\begin{aligned}
z_+^T L z_+ &= \sum_{(a,b) \in E} w(a,b)(z_+(a) - z_+(b))^2 \\
&= \left[\sum_{(a,b) \in E(S_{j_{\ell-1}})} + \sum_{(a,b) \in \partial S_{j_{\ell-1}}} + \sum_{(a,b) \in E(V - S_{j_{\ell-1}})} \right] w(a,b)(z_+(a) - z_+(b))^2 \\
&= \sum_{(a,b) \in \partial S_{j_{\ell-1}}} w(a,b)(z_+(a) - z_+(b))^2 \\
&= \sum_{a=1}^{j_{\ell-1}} \sum_{b \in N(a)} w(a,b) z_+(a)^2 \\
&= \sum_{a=1}^{j_{\ell-1}} d(a) z_+(a)^2 \\
&= z_+^T D z_+.
\end{aligned}$$

Therefore the right hand side of (63) is equal to one, and so (63) holds.

Now let us move on to the second sub-case, $0 < \sigma < 1$. Define:

$$\begin{aligned}
\alpha_a &:= \sum_{\substack{b \in N(a) \\ b > a}} w(a,b), \\
\beta_a &:= \sum_{\substack{b \in N(a) \\ b < a}} w(a,b).
\end{aligned}$$

Notice that

$$d(a) = \alpha_a + \beta_a,$$

which in turn implies

$$d(S_j) = \sum_{a=1}^j d(a) = \sum_{a=1}^j (\alpha_a + \beta_a), \quad 1 \leq j \leq n. \quad (64)$$

We also have the following claim, which we will state as a lemma:

Lemma 52. *The following holds:*

$$w(\partial S_j) = \sum_{a=1}^j (\alpha_a - \beta_a), \quad 1 \leq j \leq n. \quad (65)$$

Proof of Lemma. We will prove the result by induction on j . The base case is $j = 1$. We have

$$w(\partial S_1) = \mathbf{d}(1) = \sum_{\substack{b \in N(1) \\ b > 1}} = \alpha_1,$$

and

$$\beta_1 = \sum_{\substack{b \in N(1) \\ b < 1}} w(a, b) = 0.$$

Thus the base case is proven.

For the induction step, suppose the result is true for S_{j-1} (with $j \geq 2$), and let us prove it for S_j . Using this induction hypothesis, we write:

$$\begin{aligned} w(\partial S_j) &= \sum_{a=1}^j \sum_{\substack{b \in N(a) \\ b > j}} w(a, b) \\ &= \sum_{a=1}^{j-1} \sum_{\substack{b \in N(a) \\ b > j-1}} w(a, b) + \sum_{\substack{b \in N(j) \\ b > j}} w(j, b) - \sum_{\substack{a \in N(j) \\ a < j}} w(a, j) \\ &= w(\partial S_{j-1}) + \alpha_j - \beta_j \\ &= \sum_{a=1}^{j-1} (\alpha_a - \beta_b) + \alpha_j - \beta_j, \end{aligned}$$

and so the lemma is proved. □

Subtract (65) from (64) and divide both sides by two. We obtain:

$$\sum_{a=1}^j \beta_a = \frac{d(S_j) - w(\partial S_j)}{2}, \quad 1 \leq j \leq n. \quad (66)$$

Now let us see how we can use α_a and β_a . Using Lemma 47,

$$\begin{aligned}
\mathbf{z}_+^T \mathbf{L} \mathbf{z}_+ &= \sum_{(a,b) \in E} (\mathbf{z}_+(a) - \mathbf{z}_+(b))^2 \\
&= \sum_{a=1}^n \sum_{\substack{b \in N(a) \\ b > a}} w(a,b) (\mathbf{z}_+(a) - \mathbf{z}_+(b))^2 \\
&\geq \sum_{a=1}^n \sum_{\substack{b \in N(a) \\ b > a}} w(a,b) \left[\sigma \mathbf{z}_+(a)^2 - \frac{\sigma}{1-\sigma} \mathbf{z}_+(b)^2 \right] \\
&= \sigma \sum_{a=1}^n \mathbf{z}_+(a)^2 \sum_{\substack{b \in N(a) \\ b > a}} w(a,b) - \frac{\sigma}{1-\sigma} \sum_{a=1}^n \sum_{\substack{b \in N(a) \\ b > a}} w(a,b) \mathbf{z}_+(b)^2 \\
&= \sigma \sum_{a=1}^n \alpha_a \mathbf{z}_+(a)^2 - \frac{\sigma}{1-\sigma} \sum_{b=1}^n \mathbf{z}_+(b)^2 \sum_{\substack{a \in N(b) \\ a < b}} w(a,b) \\
&= \sigma \sum_{a=1}^n \alpha_a \mathbf{z}_+(a)^2 - \frac{\sigma}{1-\sigma} \sum_{b=1}^n \beta_b \mathbf{z}_+(b)^2.
\end{aligned} \tag{67}$$

Since $\mathbf{z}_+(a) = 0$ for all $a > j_{\ell-1}$, we can rewrite (67) as:

$$\mathbf{z}_+^T \mathbf{L} \mathbf{z}_+ \geq \sum_{a=1}^{j_{\ell-1}} \left[\sigma \alpha_a - \frac{\sigma}{1-\sigma} \beta_a \right] \mathbf{z}_+(a)^2. \tag{68}$$

Set:

$$R(\mathbf{z}_+) := \frac{\mathbf{z}_+^T \mathbf{L} \mathbf{z}_+}{\mathbf{z}_+^T \mathbf{D} \mathbf{z}_+}.$$

Using (68) we have

$$R(\mathbf{z}_+) \sum_{a=1}^{j_{\ell-1}} \mathbf{d}(a) \mathbf{z}_+(a)^2 = R(\mathbf{z}_+) \mathbf{z}_+^T \mathbf{D} \mathbf{z}_+ = \mathbf{z}_+^T \mathbf{L} \mathbf{z}_+ \geq \sum_{a=1}^{j_{\ell-1}} \left[\sigma \alpha_a - \frac{\sigma}{1-\sigma} \beta_a \right] \mathbf{z}_+(a)^2.$$

Since $\mathbf{z}_+(1)^2 \geq \mathbf{z}_+(2)^2 \geq \dots \geq \mathbf{z}_+(j_{\ell-1})^2 > 0$, Lemma 48 proves there exists a $1 \leq j \leq j_{\ell-1}$ such that

$$R(\mathbf{z}_+) \sum_{a=1}^j \mathbf{d}(a) \geq \sum_{a=1}^j \left[\sigma \alpha_a - \frac{\sigma}{1-\sigma} \beta_a \right]. \tag{69}$$

Furthermore, by the second part of Lemma 48, we know that we can take j such that $\mathbf{z}_+(j)^2 > \mathbf{z}_+(j+1)^2$ or $j = j_{\ell-1}$; that means $j \in \{j_1, \dots, j_{\ell-1}\}$. The left hand side of (69) is:

$$R(\mathbf{z}_+) \sum_{a=1}^j \mathbf{d}(a) = R(\mathbf{z}_+) d(S_j).$$

For the right hand side, we use (62), (65), and (66) to make the following calculation:

$$\begin{aligned}
\sum_{a=1}^j \left[\sigma \alpha_a - \frac{\sigma}{1-\sigma} \beta_a \right] &= \sigma \sum_{a=1}^j (\alpha_a - \beta_a) - \frac{\sigma^2}{1-\sigma} \sum_{a=1}^j \beta_a \\
&= \sigma w(\partial S_j) - \frac{\sigma^2}{1-\sigma} \left[\frac{d(S_j) - w(\partial S_j)}{2} \right] \\
&\geq \sigma^2 d(S_j) + \frac{\sigma^2 w(\partial S_j) - \sigma^2 d(S_j)}{2(1-\sigma)} \\
&\geq \frac{2\sigma^2(1-\sigma)d(S_j) + \sigma^3 d(S_j) - \sigma^2 d(S_j)}{2(1-\sigma)} \\
&= \frac{2\sigma^2 d(S_j) - 2\sigma^3 d(S_j) + \sigma^3 d(S_j) - \sigma^2 d(S_j)}{2(1-\sigma)} \\
&= \frac{(\sigma^2 - \sigma^3)d(S_j)}{2(1-\sigma)} \\
&= \frac{\sigma^2}{2} d(S_j).
\end{aligned}$$

Therefore we have

$$R(\mathbf{z}_+)d(S_j) \geq \frac{\sigma^2}{2} d(S_j) \implies R(\mathbf{z}_+) \geq \frac{\sigma^2}{2}.$$

Recalling that $\sigma = \min_{j \in \{j_1, \dots, j_\ell\}} \varphi(S_j)$ completes the proof of this first case.

Recall the second case is when $\ell = 1$. If $\ell = 1$, then $\mathbf{z}(a) \leq 0$ for all $a \in V$ which means that $\mathbf{z}_+(a) = 0$ for all $a \in V$. Thus \mathbf{z}_+ is not very useful, but on the other hand we have $\mathbf{z}_- = \mathbf{z}$. Set

$$\tilde{\mathbf{z}} := -\mathbf{z},$$

so that

$$\tilde{\mathbf{z}}(n) \geq \tilde{\mathbf{z}}(n-1) \geq \dots \geq \tilde{\mathbf{z}}(1) = 0,$$

and

$$\begin{aligned}
\tilde{\mathbf{z}}(j_m) &= \dots = \tilde{\mathbf{z}}(j_{m-1} + 1) > \tilde{\mathbf{z}}(j_{m-1}) \\
\tilde{\mathbf{z}}(j_{m-1}) &= \dots = \tilde{\mathbf{z}}(j_{m-2} + 1) > \tilde{\mathbf{z}}(j_{m-2}) \\
&\vdots \\
\tilde{\mathbf{z}}(j_2) &= \dots = \tilde{\mathbf{z}}(j_1 + 1) > \tilde{\mathbf{z}}(j_1) \\
\tilde{\mathbf{z}}(j_1) &= \dots = \tilde{\mathbf{z}}(1) = 0.
\end{aligned}$$

Additionally, note that $m \geq 2$. Indeed, $\langle \mathbf{y}, \mathbf{d} \rangle = 0$ means that \mathbf{y} must take positive and negative values, which means that \mathbf{y} takes at least two distinct values. Thus \mathbf{z} takes at least two distinct values too. Since $m \geq 2$ and since $1 \leq k \leq j_1$, we can apply the same proof as in case 1 to $\tilde{\mathbf{z}}$. We conclude there exists a $j_i \in \{j_2, \dots, j_m\}$ such that

$$\tilde{S} := \{a \in V : \tilde{\mathbf{z}}(a) \geq \tilde{\mathbf{z}}(j_i)\} = \{a \in V : -\mathbf{z}(a) \geq -\mathbf{z}(j_i)\} = \{a \in V : \mathbf{z}(a) \leq \mathbf{z}(j_i)\},$$

satisfies

$$\frac{\varphi(\tilde{S})^2}{2} \leq \frac{\tilde{\mathbf{z}}^T \mathbf{L} \tilde{\mathbf{z}}}{\tilde{\mathbf{z}}^T \mathbf{D} \tilde{\mathbf{z}}} = \frac{\mathbf{z}^T \mathbf{L} \mathbf{z}}{\mathbf{z}^T \mathbf{D} \mathbf{z}}.$$

That's pretty good except that \tilde{S} does not quite have the right form. We can fix that pretty easily. First, recall that $\varphi(V - \tilde{S}) = \varphi(\tilde{S})$. Furthermore,

$$V - \tilde{S} = \{a \in V : \mathbf{z}(a) > \mathbf{z}(j_i)\}.$$

By the way j_i was defined, and since $i \geq 2$, we know that all the vertices $a \in V$ satisfying $\mathbf{z}(a) > \mathbf{z}(j_i)$ in fact satisfy $\mathbf{z}(a) \geq \mathbf{z}(j_{i-1})$. Thus:

$$V - \tilde{S} = S_{j_{i-1}} = \{a \in V : \mathbf{z}(a) \geq \mathbf{z}(j_{i-1})\}.$$

That completes the proof. □

27 Spectral clustering

Given a weighted graph $G = (V, E, w)$ we may be interested in dividing G into two (or more) clusters. For example, we could have a data set $A = \{a_1, \dots, a_n\} \subset \mathbb{R}^d$ that we represent as a weighted graph $G = (V, E, w)$ in which

- $V = A$
- For the edges E , we give each vertex k neighbors (this is called a k nearest neighbor graph, or k -NN graph) according to the following rule. For each $a_i \in A$ we re-order the other vertices so that

$$\|a_i - a_{j_1}\| \leq \|a_i - a_{j_2}\| \leq \dots \leq \|a_i - a_{j_{n-1}}\|, \quad j_\ell \neq i.$$

We then set the initial neighborhood of a_i to be:

$$\tilde{N}(a_i) := \{a_{j_\ell} : 1 \leq \ell \leq k\}.$$

However, this results in asymmetric neighborhoods, i.e., just because $a_j \in \tilde{N}(a_i)$ it does not mean that $a_i \in \tilde{N}(a_j)$. To fix this, define the symmetric neighborhood of $a_i \in V$ as:

$$N(a_i) := \{a_j \in V : a_j \in \tilde{N}(a_i) \text{ or } a_i \in \tilde{N}(a_j)\}.$$

- We give the edges weights according to the inverse of the distance between data points, for example:

$$w(a_i, a_j) = \begin{cases} e^{-\|a_i - a_j\|^2 / 2\sigma^2} & (a_i, a_j) \in E \\ 0 & (a_i, a_j) \notin E \end{cases},$$

where σ is a parameter that we must set.

We can use the conductance to measure the optimal division of G into two clusters. In particular we can set:

$$S^* = \arg \min_{S \subset V} \varphi(S),$$

which means that S^* is the subset of vertices such that $\varphi(S^*) = \varphi_G$. We then divide the vertices of G into two clusters, S^* and $V - S^*$ (with associated subgraphs $G(S^*) = (S^*, E(S^*), w|_{E(S^*)})$ and $G(V - S^*)$ if we like). However, finding S^* on a computer is very expensive. Indeed, we would need to search over all subsets $S \subset V$ in order to find it. Thus the time to compute it is $O(2^n)$ since there are 2^n subsets of a set with n elements.

We can use Cheeger's inequality (Theorem 51) to get a clustering that is nearly as good as the clustering given by S^* , but which has a much better run-time. Here is the algorithm:

1. Compute the normalized graph Laplacian, \mathbf{N} , of the graph G . Worst case is this step takes $O(n^2)$ time, but it can be reduced if $|E| \ll n(n-1)/2$ and hence \mathbf{N} is sparse.
2. Compute the second eigenvalue ν_2 and the second eigenvector ϕ_2 of \mathbf{N} . Worst case is this step takes $O(n^3)$ time, but that is the time to compute all eigenvalues and eigenvectors of \mathbf{N} . One can take advantage of the fact that we just want ν_2 and ϕ_2 , and if \mathbf{N} is sparse (see the previous step), we can get additional time savings.
3. If $\nu_2 = 0$ then G is disconnected and we can cluster G by connected components, similarly to Remark 8. Otherwise continue:
4. Set $\mathbf{y} = \mathbf{D}^{-1/2} \phi_2$. This step requires $O(n)$ time.
5. Sort the values of \mathbf{y} in descending order and order the vertices of G accordingly, so that

$$\mathbf{y}(1) \geq \mathbf{y}(2) \geq \dots \geq \mathbf{y}(n).$$

This step requires $O(n \log n)$ times using, for example, merge sort.

6. For each set

$$S_j = \{a \in V : 1 \leq a \leq j\},$$

compute its conductance and set

$$j^* = \arg \min_j \varphi(S_j).$$

This takes $O(n^2)$ time.

7. Cluster the vertices of G using S_{j^*} and $V - S_{j^*}$.

Thus the time to compute S_{j^*} is significantly less than the time to compute S^* ; indeed, it is naively no worse than $O(n^3)$ and can be improved in certain cases. Furthermore, while

S_{j^*} is not necessarily the optimal partition, it is not too far from the optimal partition since from Theorem 45 and Theorem 51 we know that

$$\nu_2/2 \leq \varphi(S^*) \leq \sqrt{2\nu_2},$$

and we know from the proof of Theorem 51 that

$$\nu_2/2 \leq \varphi(S_{j^*}) \leq \sqrt{2\nu_2}.$$

Lecture 20: Random Walks on Graphs, Part I

March 30, 2021

Lecturer: Matthew Hirn

28 Random walks on graphs

Now we study random walks on graphs. We will see they are closely related to the normalized graph Laplacian. Among other things, they will give us another perspective, in addition to Cheeger's inequality, on the connectivity of a graph. Throughout let $G = (V, E, w)$ be a weighted, connected graph.

28.1 Definitions

A *random walk* on a graph G follows the following rules:

- Currently at vertex $a \in V$.
- Move to vertex $b \in N(a)$ with probability $p(b|a) = w(a, b)/d(a)$.
- $a \leftarrow b$ and repeat.

As such, a random walk generates a sequence of vertices a_1, a_2, a_3, \dots such that $(a_i, a_{i+1}) \in E$ for all $i \geq 1$. While we are sometimes interested in realizations of random walks, it is often the case that we want to analyze the expected behavior of the random walk. That is what we will do here.

Let $\mathbf{p}_t : V \rightarrow \mathbb{R}$ denote the probability distribution after $t \geq 0$ steps. In other words, $\mathbf{p}_t(a)$ denotes the probability we are at vertex a after t steps of the random walk. Since \mathbf{p}_t is a probability distribution that means that

$$\mathbf{p}_t(a) \geq 0 \quad \text{and} \quad \sum_{a \in V} \mathbf{p}_t(a) = 1.$$

The distribution \mathbf{p}_0 is the initial probability distribution. Often we will take $\mathbf{p}_0 = \delta_b$ for some vertex $b \in V$; i.e., we start the random walk at vertex $b \in V$. However, we can take \mathbf{p}_0 as any initial distribution.

Our first task is to derive $\mathbf{p}_{t+1}(a)$ from $\mathbf{p}_t(a)$. In order to get to a at step $t+1$, we must be at some $b \in N(a)$ at step t ; the probability of being at b at step t is $\mathbf{p}_t(b)$. The probability of walking from b to a in one step is $p(a|b) = w(a, b)/d(b)$. Thus we have:

$$\mathbf{p}_{t+1}(a) = \sum_{b \in N(a)} p(a|b) \mathbf{p}_t(b) = \sum_{b \in N(a)} \frac{w(a, b)}{d(b)} \mathbf{p}_t(b). \quad (70)$$

Define the $n \times n$ random walk matrix \mathbf{W} as

$$\mathbf{W} := \mathbf{M}\mathbf{D}^{-1}.$$

Since

$$\mathbf{W}(a, b) = \frac{w(a, b)}{d(b)} = p(a|b),$$

it is clear from (70) that

$$\mathbf{p}_{t+1} = \mathbf{W}\mathbf{p}_t.$$

It follows that

$$\mathbf{p}_t = \mathbf{W}^t \mathbf{p}_0.$$

Thus \mathbf{W} propagates the walk one step and \mathbf{W}^t propagates it t steps.

Notice that the column sums of \mathbf{W} are all equal to one. It follows that \mathbf{p}_t is a probability distribution so long as \mathbf{p}_0 is a probability distribution. Once we verify this for \mathbf{p}_1 the result will follow by induction. Clearly if $\mathbf{p}_0(a) \geq 0$ for all $a \in V$ then $\mathbf{p}_1(a) \geq 0$ for all a since $\mathbf{W}(a, b) \geq 0$ for all $a, b \in V$. Now we need to check that \mathbf{p}_1 sums to one:

$$\begin{aligned} \sum_{a \in V} \mathbf{p}_1(a) &= \sum_{a \in V} \sum_{b \in N(a)} \frac{w(a, b)}{d(b)} \mathbf{p}_0(b) \\ &= \sum_{a \in V} \sum_{b \in V} \frac{M(a, b)}{d(b)} \mathbf{p}_0(b) \\ &= \sum_{b \in V} \frac{\mathbf{p}_0(b)}{d(b)} \sum_{a \in V} M(a, b) \\ &= \sum_{b \in V} \frac{\mathbf{p}_0(b)}{d(b)} d(b) \\ &= \sum_{b \in V} \mathbf{p}_0(b) \\ &= 1. \end{aligned}$$

Now, it is often the case that we will use a *lazy random walk* instead of a random walk. Notice that in the random walk, if you are at vertex $a \in V$ then you can walk to any $b \in N(a)$, but you cannot stay at the vertex a . A lazy random walk changes this by allowing the walk to stay at a with probability $1/2$, and to move to one of the neighbors of a with probability $1/2$:

$$\mathbf{p}_{t+1}(a) = \frac{1}{2} \mathbf{p}_t(a) + \frac{1}{2} \sum_{b \in N(a)} \frac{w(a, b)}{d(b)} \mathbf{p}_t(b).$$

The walk matrix associated to the lazy random walk is:

$$\widetilde{\mathbf{W}} := \frac{1}{2}(\mathbf{I} + \mathbf{W}),$$

and so

$$\mathbf{p}_{t+1} = \widetilde{\mathbf{W}} \mathbf{p}_t.$$

28.2 Spectra of random walks

In this section we compute the eigenvectors and eigenvalues of \mathbf{W} and $\widetilde{\mathbf{W}}$. Even though they are not symmetric matrices, they are similar to symmetric matrices, and thus we can guarantee they will have n eigenvalues and n eigenvectors, although their eigenvectors will not be orthogonal (recall Homework assignments 01 and 02). Define the *normalized adjacency matrix* of $G = (V, E, w)$ as

$$\mathbf{A} := \mathbf{D}^{-1/2} \mathbf{M} \mathbf{D}^{-1/2} = \mathbf{D}^{-1/2} \mathbf{W} \mathbf{D}^{1/2}.$$

Notice that

$$\mathbf{N} = \mathbf{I} - \mathbf{D}^{-1/2} \mathbf{M} \mathbf{D}^{-1/2} = \mathbf{I} - \mathbf{A}.$$

Thus the eigenvectors of \mathbf{N} , ϕ_1, \dots, ϕ_n , are the eigenvectors of \mathbf{A} and the eigenvalues of \mathbf{A} are

$$\mu_i = 1 - \nu_i,$$

where we recall $0 = \nu_1 < \nu_2 \leq \dots \leq \nu_n \leq 2$ are the eigenvalues of \mathbf{N} .

Theorem 53. *The vector ϕ is an eigenvector of \mathbf{A} with eigenvalue μ if and only if $\mathbf{D}^{1/2}\phi$ is an eigenvector of \mathbf{W} with eigenvalue μ if and only if $\mathbf{D}^{1/2}\phi$ is an eigenvector of $\widetilde{\mathbf{W}}$ with eigenvalue $\frac{1}{2}(1 + \mu)$.*

Proof. Let $\mathbf{A}\phi = \mu\phi$. Then

$$\mathbf{W} \mathbf{D}^{1/2} \phi = \mathbf{D}^{1/2} \mathbf{A} \mathbf{D}^{-1/2} \mathbf{D}^{1/2} \phi = \mathbf{D}^{1/2} \mathbf{A} \phi = \mu \mathbf{D}^{1/2} \phi.$$

Additionally,

$$\widetilde{\mathbf{W}} \mathbf{D}^{1/2} \phi = \frac{1}{2}(\mathbf{I} + \mathbf{W}) \mathbf{D}^{1/2} \phi = \frac{1}{2}(1 + \mu) \mathbf{D}^{1/2} \phi.$$

The reverse directions are similar. □

Since any eigenvalue ν of \mathbf{N} lies in $[0, 2]$, we see that any eigenvalue μ of \mathbf{W} lies in $[-1, 1]$. In turn, any eigenvalue ω of $\widetilde{\mathbf{W}}$ lies in $[0, 1]$. Let

$$1 = \omega_1 > \omega_2 \geq \dots \geq \omega_n \geq 0,$$

be the eigenvalues of $\widetilde{\mathbf{W}}$. The eigenvalues of $\widetilde{\mathbf{W}}$ are related to the eigenvalues of \mathbf{N} via

$$\omega_i = \frac{1}{2}(1 + \mu_i) = \frac{1}{2}(1 + 1 - \nu_i) = 1 - \frac{\nu_i}{2}.$$

Additionally, we have

$$\begin{aligned}
\widetilde{\boldsymbol{W}} &= \frac{1}{2}(\boldsymbol{I} + \boldsymbol{W}) \\
&= \frac{1}{2}(\boldsymbol{I} + \boldsymbol{D}^{1/2} \boldsymbol{A} \boldsymbol{D}^{-1/2}) \\
&= \frac{1}{2}(\boldsymbol{I} + \boldsymbol{D}^{1/2}(\boldsymbol{I} - \boldsymbol{N}) \boldsymbol{D}^{-1/2}) \\
&= \frac{1}{2}(2\boldsymbol{I} - \boldsymbol{D}^{1/2} \boldsymbol{N} \boldsymbol{D}^{-1/2}) \\
&= \boldsymbol{I} - \frac{1}{2} \boldsymbol{D}^{1/2} \boldsymbol{N} \boldsymbol{D}^{-1/2} .
\end{aligned}$$

Lecture 21: Random Walks on Graphs, Part II

April 1, 2021

Lecturer: Matthew Hirn

28.3 Stationary distribution and mixing time

It turns out that no matter what distribution \mathbf{p}_0 we start with, the lazy random walk on a connected graph $G = (V, E, w)$ will always converge to the *stationary distribution* $\boldsymbol{\pi} : V \rightarrow \mathbb{R}$, which is defined as

$$\boldsymbol{\pi} = \frac{\mathbf{d}}{\|\mathbf{d}\|_1},$$

where

$$\|\mathbf{d}\|_1 = \sum_{a \in V} |\mathbf{d}(a)| = \sum_{a \in V} \mathbf{d}(a),$$

since $\mathbf{d}(a) \geq 0$ for all $a \in V$. First, observe that \mathbf{d} is an eigenvector of $\widetilde{\mathbf{W}}$ with eigenvalue $\omega_1 = 1$:

$$\widetilde{\mathbf{W}}\mathbf{d} = \left(\mathbf{I} - \frac{1}{2}\mathbf{D}^{1/2}\mathbf{N}\mathbf{D}^{-1/2} \right) \mathbf{d} = \mathbf{d} - \frac{1}{2}\mathbf{D}^{1/2}\mathbf{N}\mathbf{d}^{1/2} = \mathbf{d},$$

since $\mathbf{d}^{1/2}$ is an eigenvector of \mathbf{N} with eigenvalue $\nu_1 = 0$. Thus $\boldsymbol{\pi}$ is an eigenvector of $\widetilde{\mathbf{W}}$ with eigenvalue $\omega_1 = 1$. Since all the other eigenvalues of $\widetilde{\mathbf{W}}$ are non-negative and strictly less than one, we will see that $\mathbf{p}_t = \widetilde{\mathbf{W}}^t \mathbf{p}_0$ will converge to $\boldsymbol{\pi}$ as $t \rightarrow \infty$. The following theorem gives a precise statement.

Theorem 54. *Let $G = (V, E, w)$ be connected, let $\mathbf{p}_0 : V \rightarrow \mathbb{R}$ be a probability distribution, and set*

$$\mathbf{p}_t = \widetilde{\mathbf{W}}^t \mathbf{p}_0, \quad t \geq 0.$$

Then,

$$\|\mathbf{p}_t - \boldsymbol{\pi}\| \leq \omega_2^t \left(\frac{d_{\max}}{d_{\min}} \right)^{1/2} \|\mathbf{p}_0\|.$$

Proof. The main difficulty of this proof is that the eigenvectors of $\widetilde{\mathbf{W}}$ are not orthogonal. So let $\boldsymbol{\phi}_1, \dots, \boldsymbol{\phi}_n$ be an orthonormal basis of eigenvectors of \mathbf{N} , which means they are an orthonormal basis of eigenvectors of \mathbf{A} . By Theorem 53, we know that $\mathbf{D}^{1/2}\boldsymbol{\phi}_1, \dots, \mathbf{D}^{1/2}\boldsymbol{\phi}_n$ are eigenvectors of $\widetilde{\mathbf{W}}$. While they are not orthonormal, they do form a basis for \mathbb{R}^n . Therefore we can write

$$\mathbf{p}_0 = \sum_{i=1}^n \alpha_i \mathbf{D}^{1/2} \boldsymbol{\phi}_i,$$

for some unique coefficients α_i . This in turn implies,

$$\mathbf{D}^{-1/2}\mathbf{p}_0 = \sum_{i=1}^n \alpha_i \phi_i.$$

Now since ϕ_1, \dots, ϕ_n does form an orthonormal basis, we know that

$$\alpha_i = \langle \mathbf{D}^{-1/2}\mathbf{p}_0, \phi_i \rangle.$$

In particular,

$$\alpha_1 = \langle \mathbf{D}^{-1/2}\mathbf{p}_0, \phi_1 \rangle = (\mathbf{D}^{-1/2}\mathbf{p}_0)^T \frac{\mathbf{d}^{1/2}}{\|\mathbf{d}^{1/2}\|} = \frac{\mathbf{p}_0^T \mathbf{D}^{-1/2} \mathbf{d}^{1/2}}{\|\mathbf{d}^{1/2}\|} = \frac{\mathbf{p}_0^T \mathbf{1}}{\|\mathbf{d}^{1/2}\|} = \frac{1}{\|\mathbf{d}^{1/2}\|}.$$

Now we compute

$$\begin{aligned} \mathbf{p}_t &= \widetilde{\mathbf{W}}^t \mathbf{p}_0 = \widetilde{\mathbf{W}}^t \sum_{i=1}^n \alpha_i \mathbf{D}^{1/2} \phi_i \\ &= \sum_{i=1}^n \alpha_i \widetilde{\mathbf{W}}^t \mathbf{D}^{1/2} \phi_i \\ &= \sum_{i=1}^n \omega_i^t \alpha_i \mathbf{D}^{1/2} \phi_i \\ &= \alpha_1 \mathbf{D}^{1/2} \phi_1 + \sum_{i=2}^n \omega_i^t \alpha_i \mathbf{D}^{1/2} \phi_i \\ &= \frac{\mathbf{D}^{1/2} \phi_1}{\|\mathbf{d}^{1/2}\|} + \sum_{i=2}^n \omega_i^t \alpha_i \mathbf{D}^{1/2} \phi_i \\ &= \frac{1}{\|\mathbf{d}^{1/2}\|} \mathbf{D}^{1/2} \frac{\mathbf{d}^{1/2}}{\|\mathbf{d}^{1/2}\|} + \sum_{i=2}^n \omega_i^t \alpha_i \mathbf{D}^{1/2} \phi_i \\ &= \frac{\mathbf{d}}{\|\mathbf{d}\|_1} + \sum_{i=2}^n \omega_i^t \alpha_i \mathbf{D}^{1/2} \phi_i \\ &= \boldsymbol{\pi} + \sum_{i=2}^n \omega_i^t \alpha_i \mathbf{D}^{1/2} \phi_i. \end{aligned}$$

It follows that

$$\mathbf{p}_t - \boldsymbol{\pi} = \sum_{i=2}^n \omega_i^t \alpha_i \mathbf{D}^{1/2} \phi_i \implies \mathbf{D}^{-1/2}(\mathbf{p}_t - \boldsymbol{\pi}) = \sum_{i=2}^n \omega_i^t \alpha_i \phi_i.$$

Thus,

$$\begin{aligned}
\|\mathbf{D}^{-1/2}(\mathbf{p}_t - \boldsymbol{\pi})\|^2 &= \sum_{i=2}^n \omega_i^{2t} |\alpha_i|^2 \\
&\leq \omega_2^{2t} \sum_{i=2}^n |\alpha_i|^2 \\
&\leq \omega_2^{2t} \sum_{i=1}^n |\alpha_i|^2 \\
&= \omega_2^{2t} \|\mathbf{D}^{-1/2} \mathbf{p}_0\|^2.
\end{aligned}$$

To finish the proof, we observe that for any $\mathbf{x} : V \rightarrow \mathbb{R}$,

$$\|\mathbf{D}^{-1/2} \mathbf{x}\|^2 = \sum_{a \in V} \frac{1}{d(a)} \mathbf{x}(a)^2 \geq \frac{1}{d_{\max}} \sum_{a \in V} \mathbf{x}(a)^2 = \frac{\|\mathbf{x}\|^2}{d_{\max}},$$

and

$$\|\mathbf{D}^{-1/2} \mathbf{x}\|^2 = \sum_{a \in V} \frac{1}{d(a)} \mathbf{x}(a)^2 \leq \frac{1}{d_{\min}} \sum_{a \in V} \mathbf{x}(a)^2 = \frac{\|\mathbf{x}\|^2}{d_{\min}}.$$

Therefore,

$$\frac{\|\mathbf{p}_t - \boldsymbol{\pi}\|^2}{d_{\max}} \leq \|\mathbf{D}^{-1/2}(\mathbf{p}_t - \boldsymbol{\pi})\|^2 \leq \omega_2^{2t} \|\mathbf{D}^{-1/2} \mathbf{p}_0\|^2 \leq \omega_2^{2t} \frac{\|\mathbf{p}_0\|^2}{d_{\min}},$$

and so:

$$\|\mathbf{p}_t - \boldsymbol{\pi}\| \leq \omega_2^t \left(\frac{d_{\max}}{d_{\min}} \right)^{1/2} \|\mathbf{p}_0\|.$$

□

It follows that the closer ω_2 is to 0, the faster \mathbf{p}_t converges to $\boldsymbol{\pi}$. Recall that $\omega_2 = 1 - \nu_2/2$. Theorem 54 says:

$$\|\mathbf{p}_t - \boldsymbol{\pi}\| \leq \left(1 - \frac{\nu_2}{2}\right)^t \left(\frac{d_{\max}}{d_{\min}}\right)^{1/2} \|\mathbf{p}_0\|.$$

Now remember that ν_2 measures the connectivity of G since, by Theorems 45 and 51, we have

$$\frac{\nu_2}{2} \leq \varphi_G \leq \sqrt{2\nu_2},$$

where φ_G is the conductance of G . The upper bound on φ_G (Cheeger's inequality, Theorem 51) can be used to rewrite Theorem 54 one more time as:

$$\|\mathbf{p}_t - \boldsymbol{\pi}\| \leq \left(1 - \frac{\varphi_G^2}{4}\right)^t \left(\frac{d_{\max}}{d_{\min}}\right)^{1/2} \|\mathbf{p}_0\|.$$

Unfortunately, since $\varphi_G \leq 1$, this bound is a little too loose in this approach and it is better to consider ν_2 (equivalently ω_2) directly as the measure of connectivity of G .

Thus, for graphs that are well connected (meaning that ν_2 will be large, or equivalently ω_2 will be small) will have \mathbf{p}_t converge fast. Indeed, if G is well connected then there will be many paths from any vertex a to any vertex b and many realizations of the lazy random walk will get from a to b in a small number of steps. Thus the lazy random walk will traverse the graph fast (on average) and in turn converge to $\boldsymbol{\pi}$ quickly. On the other hand, graphs that are not well connected will have small ν_2 (equivalently, large ω_2) and the bound for the convergence of \mathbf{p}_t to $\boldsymbol{\pi}$ will be weak. In fact, for many such graphs convergence will indeed be slow. For example, if there exist pairs of vertices $a, b \in V$ for which there are only a few paths from a to b , then only a few realizations of the lazy random walk will get from a to b in a small number of steps, and most realizations of the lazy random walk will take many steps to get from a to b . As such, the convergence of \mathbf{p}_t to $\boldsymbol{\pi}$ will be slow.

To quantify the number of steps it takes for a walk to mix through G , let us define the *mixing time* of the lazy random walk as the minimum value of t for which

$$\|\mathbf{p}_t - \boldsymbol{\pi}\| \leq \frac{\|\boldsymbol{\pi}\|}{2}.$$

Applying Theorem 54, we see the mixing time t is no more than the t for which

$$\begin{aligned} \|\mathbf{p}_t - \boldsymbol{\pi}\| &\leq \omega_2^t \left(\frac{d_{\max}}{d_{\min}} \right)^{1/2} \|\mathbf{p}_0\| \leq \frac{\|\boldsymbol{\pi}\|}{2} \\ \iff \left(1 - \frac{\nu_2}{2} \right)^t \left(\frac{d_{\max}}{d_{\min}} \right)^{1/2} \|\mathbf{p}_0\| &\leq \frac{\|\boldsymbol{\pi}\|}{2} \\ \iff \left(1 - \frac{\nu_2}{2} \right)^t &\leq \left(\frac{d_{\min}}{d_{\max}} \right)^{1/2} \frac{\|\boldsymbol{\pi}\|}{2\|\mathbf{p}_0\|} \end{aligned}$$

Since $1 - x \leq e^{-x}$, we have $(1 - x)^t \leq e^{-xt}$. As such, we can again upper bound the mixing time t via:

$$\begin{aligned} \left(1 - \frac{\nu_2}{2} \right)^t \leq e^{-\nu_2 t/2} &\leq \left(\frac{d_{\min}}{d_{\max}} \right)^{1/2} \frac{\|\boldsymbol{\pi}\|}{2\|\mathbf{p}_0\|} \iff \frac{-\nu_2 t}{2} \leq \log \left[\left(\frac{d_{\min}}{d_{\max}} \right)^{1/2} \frac{\|\boldsymbol{\pi}\|}{2\|\mathbf{p}_0\|} \right] \\ \iff t &\geq \frac{-2}{\nu_2} \log \left[\left(\frac{d_{\min}}{d_{\max}} \right)^{1/2} \frac{\|\boldsymbol{\pi}\|}{2\|\mathbf{p}_0\|} \right] \\ \iff t &\geq \frac{2}{\nu_2} \log \left[\left(\frac{d_{\max}}{d_{\min}} \right)^{1/2} \frac{2\|\mathbf{p}_0\|}{\|\boldsymbol{\pi}\|} \right]. \end{aligned}$$

We thus obtain an upper bound for the mixing time that is proportional to $1/\nu_2$. Thus if G is well connected, and hence ν_2 is large, then the mixing time will be small. On the other hand, if G is not well connected (meaning ν_2 is small), the mixing time could be large.

In the case of when G is d -regular, we can simplify the log term. Indeed, $d_{\max} = d_{\min} = d$ in this case and

$$\|\boldsymbol{\pi}\| = \frac{\|\mathbf{d}\|}{\|\mathbf{d}\|_1} = \frac{\sqrt{nd}}{nd} = \frac{1}{\sqrt{n}}.$$

If $\mathbf{p}_0 = \boldsymbol{\delta}_b$ for some b as well, then $\|\mathbf{p}_0\| = 1$ and we obtain:

$$t \geq \frac{2}{\nu_2} \log(2\sqrt{n}) = \frac{2}{\nu_2} \log((4n)^{1/2}) = \frac{\log(4n)}{\nu_2}.$$

In some graphs the $\log(4n)$ factor is too pessimistic and in fact the mixing time is $O(1/\nu_2)$.

28.4 Diffusion

In this section we consider the operators

$$\mathbf{P} = \mathbf{W}^T \quad \text{and} \quad \tilde{\mathbf{P}} := \tilde{\mathbf{W}}^T.$$

Recall that $\tilde{\mathbf{W}}$ maps probability distributions to probability distributions, which is sometimes referred to as a Markov operator. We also have $\tilde{\mathbf{W}}(a, b) = p(a|b)$, i.e., the probability of walking to a given that you are located at b . On the other hand,

$$\tilde{\mathbf{P}}(a, b) = \tilde{\mathbf{W}}^T = p(b|a),$$

that is, the probability of walking to b given that you are a . Since the column sums of $\tilde{\mathbf{W}}$ are equal to one, the row sums of $\tilde{\mathbf{P}}$ are equal to one. This means that $\tilde{\mathbf{P}}$ acts as a diffusion, or averaging operator, since:

$$\begin{aligned} \tilde{\mathbf{P}}\mathbf{x}(a) &= \tilde{\mathbf{P}}(a, a)\mathbf{x}(a) + \sum_{b \in N(a)} \tilde{\mathbf{P}}(a, b)\mathbf{x}(b) \\ &= \frac{1}{2}\mathbf{x}(a) + \frac{1}{2} \sum_{b \in N(a)} \mathbf{P}(a, b)\mathbf{x}(b) \\ &= \frac{1}{2}\mathbf{x}(a) + \frac{1}{2} \sum_{b \in N(a)} \frac{w(a, b)}{\mathbf{d}(a)}\mathbf{x}(b), \end{aligned}$$

and in particular if $G = (V, E)$ is not weighted, then

$$\tilde{\mathbf{P}}\mathbf{x}(a) = \frac{1}{2} \left[\mathbf{x}(a) + \frac{1}{\mathbf{d}(a)} \sum_{b \in N(a)} \mathbf{x}(b) \right] = \frac{1}{2} \left[\mathbf{x}(a) + \frac{1}{|N(a)|} \sum_{b \in N(a)} \mathbf{x}(b) \right].$$

Thus $\tilde{\mathbf{P}}\mathbf{x}(a)$ replaces the value $\mathbf{x}(a)$ at $a \in V$ with a weighted average of $\mathbf{x}(a)$ and the values $\mathbf{x}(b)$ for $b \in N(a)$.

Notice that $\tilde{\mathbf{P}}^t(a, b)$ is the probability of walking from a to b in exactly t steps. We can use the family of matrices $\tilde{\mathbf{P}}^t$ to define the family of *diffusion distances* between a and b :

$$d_t(a, b) := \left(\sum_{c \in V} \left[\tilde{\mathbf{P}}^t(a, c) - \tilde{\mathbf{P}}^t(b, c) \right]^2 \frac{1}{d(c)} \right)^{1/2}.$$

Since

$$\text{Row } a \text{ of } \tilde{\mathbf{P}}^t = (\tilde{\mathbf{P}}^t(a, c))_{c \in V} = \boldsymbol{\delta}_a^T \tilde{\mathbf{P}}^t = \boldsymbol{\delta}_a^T (\widetilde{\mathbf{W}}^T)^t = \boldsymbol{\delta}_a^T (\widetilde{\mathbf{W}}^t)^T = (\widetilde{\mathbf{W}}^t \boldsymbol{\delta}_a)^T,$$

we see that

$$d_t(a, b) = \left(\sum_{c \in V} \left[\widetilde{\mathbf{W}}^t \boldsymbol{\delta}_a(c) - \widetilde{\mathbf{W}}^t \boldsymbol{\delta}_b(c) \right]^2 \frac{1}{d(c)} \right)^{1/2}.$$

In other words, the diffusion distance measures the distance between a and b by measuring the overlap of the t -step random walk distribution for walks started at a with the t -step random walk distribution for walks started at b . It provides an alternate distance to the shortest path (geodesic) distance between a and b . In some cases the diffusion distance between a and b better reflects the clustering structure of the graph G ; see for example Figure 34. It is also more robust to noise than the shortest path distance if the graph G is generated from noisy data.

We can rewrite the diffusion distance in terms of the eigenvectors and eigenvalues of $\tilde{\mathbf{P}}$. Let us first observe that

$$\begin{aligned} \tilde{\mathbf{P}} &= \widetilde{\mathbf{W}}^T = \frac{1}{2}(\mathbf{I} + \mathbf{W}^T) = \frac{1}{2}(\mathbf{I} + \mathbf{D}^{-1}\mathbf{M}) \\ &= \frac{1}{2}\mathbf{D}^{-1/2}(\mathbf{I} + \mathbf{D}^{-1/2}\mathbf{M}\mathbf{D}^{-1/2})\mathbf{D}^{1/2} = \mathbf{D}^{-1/2}(\mathbf{I}/2 + \mathbf{A}/2)\mathbf{D}^{1/2}. \end{aligned}$$

On the other hand, since \mathbf{I} and \mathbf{A} are symmetric,

$$\widetilde{\mathbf{W}} = \mathbf{D}^{1/2}(\mathbf{I}/2 + \mathbf{A}/2)\mathbf{D}^{-1/2}.$$

Thus the eigenvalues of $\tilde{\mathbf{P}}$ and $\widetilde{\mathbf{W}}$ are the same, given by $1 = \omega_1 \geq \omega_2 \geq \dots \geq \omega_n \geq 0$. Furthermore, if the eigenvectors of \mathbf{A} are $\boldsymbol{\phi}_1, \dots, \boldsymbol{\phi}_n$ (reminder, these are also the eigenvectors of \mathbf{N}), then as we saw earlier the eigenvectors of $\widetilde{\mathbf{W}}$ are $\mathbf{D}^{1/2}\boldsymbol{\phi}_1, \dots, \mathbf{D}^{1/2}\boldsymbol{\phi}_n$, but the eigenvectors of $\tilde{\mathbf{P}}$ are $\mathbf{D}^{-1/2}\boldsymbol{\phi}_1, \dots, \mathbf{D}^{-1/2}\boldsymbol{\phi}_n$. Since $\boldsymbol{\phi}_1 = \mathbf{d}^{1/2}$, it follows that the first eigenvector of $\tilde{\mathbf{P}}$ is $\mathbf{1}$ (although we could have verified this directly using that the row sums of $\tilde{\mathbf{P}}$ are equal to one):

$$\tilde{\mathbf{P}}\mathbf{1} = \mathbf{1}.$$

Denote the eigenvectors of $\tilde{\mathbf{P}}$ by

$$\tilde{\boldsymbol{\phi}}_i := \mathbf{D}^{-1/2}\boldsymbol{\phi}_i, \quad 1 \leq i \leq n.$$




Figure 34: Illustration of the shortest path distance versus the diffusion distance. Here G is a barbell type graph, with $a, b, c \in V$. The shortest path distance, as indicated by the red lines, is approximately the same between a and b as it is between b and c . However, their diffusion distances, which are inversely proportional to the overlap of the shaded disks centered at each vertex, are very different. Indeed, due to the bottleneck in the barbell graph, the t -step distribution started from a overlaps very little with the t -step distribution started at b . On the other hand, since b and c lie in the same cluster, their t -step distributions overlap significantly.

Define the *diffusion map* $\tilde{\Phi}_t : V \rightarrow \mathbb{R}^{n-1}$ as

$$\tilde{\Phi}_t(a) := (\omega_2^t \tilde{\phi}_2(a), \dots, \omega_n^t \tilde{\phi}_n(a)).$$

The diffusion map should remind you of the eigenvector graph embeddings we studied in Section 9, except that we use the eigenvectors of $\tilde{\mathbf{P}}$ instead of \mathbf{L} and we re-scale the eigenvector coordinates by ω_i^t . The following theorem shows the diffusion distance between a and b is equal to the Euclidean distance between $\tilde{\Phi}_t(a)$ and $\tilde{\Phi}_t(b)$.

Theorem 55. *Let $G = (V, E, w)$ be a connected graph. Then*

$$d_t(a, b) = \|\tilde{\Phi}_t(a) - \tilde{\Phi}_t(b)\|, \quad \forall a, b \in V, \quad \forall t \geq 0.$$

I leave the proof to you as an upcoming homework exercise. Whereas the eigenvector embedding of the graph using the eigenvectors of \mathbf{L} preserved local relations between vertices in the graph, the diffusion map embedding preserves the diffusion distance. Additionally, since for a connected graph we have $1 > \omega_2 \geq \dots \geq \omega_n \geq 0$, if the eigenvalues of $\tilde{\mathbf{P}}$ decay

fast and/or if t is large, we can approximate $d_t(a, b)$ by truncating $\tilde{\Phi}_t(a)$ to only include the first k entries (where k depends on the desired error upper bound and the decay of the eigenvalues of the size of t).

The diffusion distance and the diffusion map were introduced in [5]. In that paper they explore several different aspects of the diffusion map. We highlight here one part of that paper, related to clustering. The idea is that the diffusion distances provides a family of multiscale distances that reflect hierarchical clusters in a graph (e.g., a graph derived from data). Small times t distinguish many of the clusters, while medium times t merge together nearby clusters and large times t collapse nearly all clusters together. See Figure 35 for an illustration.




Figure 35: Left: A 3-cluster data set in which two clusters are closer together than their distance from the third. Right: The matrix $\tilde{\mathbf{P}}^t$ for a graph $G = (V, E, w)$ derived from the data. (a) $t = 8$: The rows of $\tilde{\mathbf{P}}^t$, and hence the diffusion distance, distinguish the three clusters. (b) $t = 64$: The rows of $\tilde{\mathbf{P}}^t$ are the nearly identical for the two closer clusters, but are still different than the rows of the vertices in the third cluster. Hence the diffusion distance effectively merges the closer two clusters together but keeps them separate from the third cluster. (c) $t = 1024$: All rows of $\tilde{\mathbf{P}}^t$ have nearly converged to the stationary distribution π (see Theorem 54) and hence do not distinguish any of the clusters. Figure taken from [5].

Lecture 22: Expander Graphs, Part I

April 6, 2021

Lecturer: Matthew Hirn

29 Expander graphs

In the next few lectures we will explore how to approximate a graph with a different graph that has far fewer edges. We will begin our studies with expander graphs. A graph $G = (V, E)$ is a d -regular, ϵ -expander graph ((d, ϵ) -expander) if

$$|\mu_i| \leq \epsilon d, \quad 2 \leq i \leq n,$$

where we let $d = \mu_1 \geq \mu_2 \geq \dots \geq \mu_n$ denote the eigenvalues of the adjacency matrix, \mathbf{M} , of G . Since

$$\mathbf{L} = d\mathbf{I} - \mathbf{M}, \quad \text{if } G \text{ is } d\text{-regular},$$

this definition is equivalent to

$$|\lambda_i - d| \leq \epsilon d, \quad 2 \leq i \leq n,$$

where $0 = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ are the eigenvalues of \mathbf{L} . We will show expander graphs have some very interesting properties, including:

1. Expander graphs approximate the complete graph.
2. The number of edges between subsets of vertices in an expander graph is approximately the same as the number of edges between subsets of vertices in a random d -regular graph.
3. Small subsets of vertices have many neighbors.
4. We will prove that as $n \rightarrow \infty$, one must have $\epsilon \geq 2\sqrt{d-1}/d$.

29.1 Expanders approximate the complete graph

Let G, H be two graphs with n vertices. Recall that

$$H \preceq G \iff \mathbf{x}^T \mathbf{L}_H \mathbf{x} \leq \mathbf{x}^T \mathbf{L}_G \mathbf{x}, \quad \forall \mathbf{x} \in \mathbb{R}^n.$$

Recall as well that $c \cdot G$ is the same graph as G but in which every edge weight is multiplied by c . For unweighted graphs this means that $c \cdot G$ replaces all the edge weights equal to 1 with edge weights equal to c . For this part of the course, we will say that G is an ϵ -approximation of H if

$$(1 - \epsilon)H \preceq G \preceq (1 + \epsilon)H.$$

Theorem 56. *Let G be (d, ϵ) -expander. Then G is an ϵ approximation of $H = (d/n)K_n$.*

Proof. Suppose G is a (d, ϵ) -expander so that $|d - \lambda_i| \leq \epsilon d$ for all $i \geq 2$. Let $\mathbf{x} \in \mathbb{R}^n$ be orthogonal to $\mathbf{1}$, i.e., $\langle \mathbf{x}, \mathbf{1} \rangle = 0$, so that, using Lemma 4, we have

$$\begin{aligned} \mathbf{x}^T \mathbf{L}_G \mathbf{x} &= \sum_{i=2}^n \lambda_i |\langle \mathbf{x}, \boldsymbol{\psi}_i \rangle|^2 \\ &= \sum_{i=2}^n (\lambda_i - d + d) |\langle \mathbf{x}, \boldsymbol{\psi}_i \rangle|^2 \\ &= \sum_{i=2}^n (\lambda_i - d) |\langle \mathbf{x}, \boldsymbol{\psi}_i \rangle|^2 + d \sum_{i=2}^n |\langle \mathbf{x}, \boldsymbol{\psi}_i \rangle|^2 \\ &= \sum_{i=2}^n (\lambda_i - d) |\langle \mathbf{x}, \boldsymbol{\psi}_i \rangle|^2 + d \|\mathbf{x}\|^2, \end{aligned}$$

and

$$\sum_{i=2}^n (\lambda_i - d) |\langle \mathbf{x}, \boldsymbol{\psi}_i \rangle|^2 \begin{cases} \leq \sum_{i=2}^n |\lambda_i - d| |\langle \mathbf{x}, \boldsymbol{\psi}_i \rangle|^2 \leq \epsilon d \|\mathbf{x}\|^2 \\ \geq -\sum_{i=2}^n |\lambda_i - d| |\langle \mathbf{x}, \boldsymbol{\psi}_i \rangle|^2 \geq -\epsilon d \|\mathbf{x}\|^2 \end{cases}.$$

It follows that for all \mathbf{x} with $\langle \mathbf{x}, \mathbf{1} \rangle = 0$ we have

$$(1 - \epsilon) d \mathbf{x}^T \mathbf{x} \leq \mathbf{x}^T \mathbf{L}_G \mathbf{x} \leq (1 + \epsilon) d \mathbf{x}^T \mathbf{x}.$$

On the other hand, we know from Theorem 7 that all non-zero eigenvalues of the complete graph, K_n , are equal to n . In other words, all $\mathbf{x} \in \mathbb{R}^n$ with $\langle \mathbf{x}, \mathbf{1} \rangle = 0$ satisfy

$$\mathbf{x}^T \mathbf{L}_{K_n} \mathbf{x} = n \mathbf{x}^T \mathbf{x}.$$

Set

$$H = \frac{d}{n} K_n,$$

so that

$$\mathbf{x}^T \mathbf{L}_H \mathbf{x} = d \mathbf{x}^T \mathbf{x}, \quad \forall \mathbf{x} \in \mathbb{R}^n \text{ with } \langle \mathbf{x}, \mathbf{1} \rangle = 0.$$

Therefore we have shown that G is an ϵ -approximation of $H = (d/n)K_n$. □

For a matrix \mathbf{B} define the norm of \mathbf{B} as

$$\|\mathbf{B}\| := \max_{\substack{\mathbf{x} \in \mathbb{R}^n \\ \mathbf{x} \neq \mathbf{0}}} \frac{\|\mathbf{B}\mathbf{x}\|}{\|\mathbf{x}\|}.$$

If \mathbf{B} is symmetric, then

$$\|\mathbf{B}\| = \max_{1 \leq i \leq n} |\mu_i(\mathbf{B})|,$$

where $\mu_1(\mathbf{B}), \dots, \mu_n(\mathbf{B})$ are the eigenvalues of \mathbf{B} . Now, observe that

$$(1 - \epsilon)\mathbf{L}_H \preceq \mathbf{L}_G \preceq (1 + \epsilon)\mathbf{L}_H \implies -\epsilon\mathbf{L}_H \preceq \mathbf{L}_G - \mathbf{L}_H \preceq \epsilon\mathbf{L}_H.$$

Using Theorem 15 we have for all $i \geq 2$,

$$-\epsilon d = -\epsilon \lambda_i(\mathbf{L}_H) \leq \lambda_i(\mathbf{L}_G - \mathbf{L}_H) \leq \epsilon \lambda_i(\mathbf{L}_H) = \epsilon d,$$

from which we conclude

$$\|\mathbf{L}_G - \mathbf{L}_H\| \leq \epsilon d. \quad (71)$$

This inequality will be useful in a bit; it also gives another way by which to interpret G as an ϵ -approximation of H .

Remark 26. Note that K_n has $n(n-1)/2$ edges but a (d, ϵ) -expander has $dn/2$ edges. Thus, if $d \ll n$ and if $\epsilon \ll 1$, we will have obtained a good approximation of K_n , but using a graph G that has far fewer edges. Later on we will see how small ϵ can be as a function of d .

29.2 Quasi-random properties of expanders

Let $G = (V, E)$ and consider generating a subset $S \subseteq V$ by including each vertex in S independently with probability α , and generate another subset $T \subseteq V$ by including each vertex in T independently with probability β . On average, the number of ordered pairs $(a, b) \in S \times T$ such that $(a, b) \in E$ will be $(\alpha\beta) \cdot (2|E|)$, since the total number of ordered edges is $2|E|$.

Now, let $S \subset V$ and $T \subset V$ such that $S \cap T = \emptyset$, and define

$$E(S, T) := \{(a, b) \in E : a \in S \text{ and } b \in T\}.$$

$E(S, T)$ is the set of edges going between S and T , and $|E(S, T)|$ is the number of such edges. Since a d -regular graph has $dn/2$ edges, the following theorem proves that $|E(S, T)|$ in (d, ϵ) -expander graphs is approximately equal to the expected number of edges between a randomly selected S and T .

Theorem 57. *Let $G = (V, E)$ be a d -regular graph that ϵ -approximates $H = (d/n)K_n$. Then, for every $S \subset V$ and $T \subset V$ with $S \cap T = \emptyset$, $|S| = \alpha n$, and $|T| = \beta n$, we have*

$$||E(S, T)| - \alpha\beta dn| \leq \epsilon dn \sqrt{(\alpha - \alpha^2)(\beta - \beta^2)}.$$

Proof. First observe that

$$\begin{aligned} \mathbf{1}_S^T \mathbf{L}_G \mathbf{1}_T &= \mathbf{1}_S^T \mathbf{D}_G \mathbf{1}_T - \mathbf{1}_S^T \mathbf{M}_G \mathbf{1}_T = d|S \cap T| - \sum_{a \in V} \sum_{b \in V} M_G(a, b) \mathbf{1}_S(a) \mathbf{1}_T(b) \\ &= - \sum_{a \in S} \sum_{b \in T} M(a, b) = -|E(S, T)|. \end{aligned}$$

Now define \mathbf{J} as the matrix of all ones so that $\mathbf{L}_{K_n} = n\mathbf{I} - \mathbf{J}$. We have

$$\begin{aligned}
\mathbf{1}_S^T \mathbf{L}_H \mathbf{1}_T &= \mathbf{1}_S^T (d\mathbf{I} - (d/n)\mathbf{J}) \mathbf{1}_T = d|S \cap T| - \frac{d}{n} \sum_{a \in V} \sum_{b \in V} \mathbf{J}(a, b) \mathbf{1}_S(a) \mathbf{1}_T(b) \\
&= -\frac{d}{n} \sum_{a \in S} \sum_{b \in T} 1 \\
&= -\frac{d}{n} |S| |T| \\
&= -\alpha \beta d n.
\end{aligned}$$

Therefore

$$|E(S, T)| - \alpha \beta d n = |\mathbf{1}_S^T \mathbf{L}_G \mathbf{1}_T - \mathbf{1}_S^T \mathbf{L}_H \mathbf{1}_T| = |\mathbf{1}_S^T (\mathbf{L}_G - \mathbf{L}_H) \mathbf{1}_T|.$$

Now remember from the proof of Theorem 39 (recall this was the lower bound of the isoperimetric ratio) that

$$\mathbf{x}_S := \mathbf{1}_S - s\mathbf{1}, \quad s := \frac{|S|}{n} = \alpha,$$

satisfies $\langle \mathbf{x}_S, \mathbf{1} \rangle = 0$, and more importantly here,

$$\mathbf{x}_S^T \mathbf{x}_S = |S|(1 - s) = \alpha n(1 - \alpha) = n(\alpha - \alpha^2).$$

Define \mathbf{x}_T analogously, so that

$$\mathbf{x}_T := \mathbf{1}_T - t\mathbf{1}, \quad t := \frac{|T|}{n} = \beta \implies \mathbf{x}_T^T \mathbf{x}_T = n(\beta - \beta^2).$$

Since $\mathbf{L}_G \mathbf{1} = \mathbf{L}_H \mathbf{1} = \mathbf{0}$, and since by (71) we have $\|\mathbf{L}_G - \mathbf{L}_H\| \leq \epsilon d$, using the Cauchy-Schwarz inequality we compute:

$$\begin{aligned}
|\mathbf{1}_S^T (\mathbf{L}_G - \mathbf{L}_H) \mathbf{1}_T| &= |\mathbf{x}_S^T (\mathbf{L}_G - \mathbf{L}_H) \mathbf{x}_T| \\
&\leq \|\mathbf{x}_S\| \|(\mathbf{L}_G - \mathbf{L}_H) \mathbf{x}_T\| \\
&\leq \|\mathbf{x}_S\| \|\mathbf{L}_G - \mathbf{L}_H\| \|\mathbf{x}_T\| \\
&\leq \epsilon d n \sqrt{(\alpha - \alpha^2)(\beta - \beta^2)}.
\end{aligned}$$

That completes the proof. □

Remark 27. One can allow S and T to overlap (i.e., $S \cap T \neq \emptyset$) if one replaces $E(S, T)$ with the set of ordered pairs $(a, b) \in S \times T$.

Lecture 23: Expander Graphs, Part II

April 8, 2021

Lecturer: Matthew Hirn

29.3 Vertex expansion

Now let us explain the name expander graph. For $S \subseteq V$, let $N(S)$ denote the collection of neighbors of the vertices in S ,

$$N(S) := \{b \in V : b \in N(a) \text{ for some } a \in S\}.$$

Tanner's theorem provides a lower bound on $|N(S) \cup S|$, and in particular shows that the number of such vertices can be much larger than $|S|$.

Theorem 58 (Tanner's Theorem). *Let $G = (V, E)$ be a d -regular graph that ϵ -approximates $(d/n)K_n$. Then, for all $S \subseteq V$ with $|S| = \alpha n$,*

$$|N(S) \cup S| \geq \frac{|S|}{\epsilon^2(1 - \alpha) + \alpha}.$$

Proof. Let $R = N(S) \cup S$ and set $T = V - R$. Then

$$T \cap S = \emptyset \quad \text{and} \quad E(S, T) = \emptyset.$$

Also set

$$|T| = \beta n \text{ and } |R| = \gamma n \implies \gamma = 1 - \beta.$$

Applying Theorem 57 we obtain

$$\alpha\beta dn = ||E(S, T)| - \alpha\beta dn| \leq \epsilon dn \sqrt{(\alpha - \alpha^2)(\beta - \beta^2)}.$$

Therefore, we have:

$$\begin{aligned}
\alpha\beta &\leq \epsilon\sqrt{(\alpha - \alpha^2)(\beta - \beta^2)} \\
\alpha^2\beta^2 &\leq \epsilon^2(\alpha - \alpha^2)(\beta - \beta^2) \\
\alpha\beta &\leq \epsilon^2(1 - \alpha)(1 - \beta) \\
\frac{\beta}{1 - \beta} &\leq \frac{\epsilon^2(1 - \alpha)}{\alpha} \\
\frac{1 - \gamma}{\gamma} &\leq \frac{\epsilon^2(1 - \alpha)}{\alpha} \\
\frac{1}{\gamma} &\leq \frac{\epsilon^2(1 - \alpha)}{\alpha} + 1 \\
\frac{1}{\gamma} &\leq \frac{\epsilon^2(1 - \alpha) + \alpha}{\alpha} \\
\gamma &\geq \frac{\alpha}{\epsilon^2(1 - \alpha) + \alpha}.
\end{aligned}$$

Since $\gamma = |R|/n$, we conclude that

$$|N(S) \cup S| = |R| \geq \frac{\alpha n}{\epsilon^2(1 - \alpha) + \alpha} = \frac{|S|}{\epsilon^2(1 - \alpha) + \alpha}.$$

□

Remark 28. We note, in particular, that if $\alpha \ll \epsilon^2$ then the right hand side of Tanner's Theorem is approximately $|S|/\epsilon^2$, which can be much larger than $|S|$.

29.4 How well can a graph approximate the complete graph?

We know that a (d, ϵ) -expander graph G is an ϵ -approximation of $H = (d/n)K_n$, but how small can we make ϵ ? It stands to reason that the lower bound of ϵ should depend upon d , since as d gets smaller the number of edges in G also gets smaller, and our ability to approximate $(d/n)K_n$ is reduced.

We can use Tanner's Theorem (Theorem 58) to get a first estimate for the lower bound of ϵ . To do so, apply Tanner's Theorem to $S = \{a\}$ for any vertex $a \in V$. Since G is d -regular we conclude:

$$\begin{aligned}
d + 1 &\geq \frac{1}{\epsilon^2(1 - 1/n) + 1/n} \\
\epsilon^2(1 - 1/n) + 1/n &\geq \frac{1}{d + 1} \\
\epsilon^2 \cdot \frac{n - 1}{n} &\geq \frac{1}{d + 1} - \frac{1}{n} \\
\epsilon^2 &\geq \frac{n}{(n - 1)d + 1} - \frac{1}{n - 1} \geq \frac{1}{d + 1} - \frac{1}{n - 1}.
\end{aligned}$$

Thus:

$$\epsilon \geq \sqrt{\frac{1}{d+1} - \frac{1}{n-1}} \approx \frac{C}{\sqrt{d}}.$$

The next theorem, which is referred to as the Alon–Boppana Bound, implies a more precise estimate for ϵ . For any two edges $(a_0, a_1), (b_0, b_1) \in E$ we define the distance between them as

$$\text{dist}((a_0, a_1), (b_0, b_1)) = \min\{\text{dist}(a_i, b_j) : 0 \leq i, j \leq 1\},$$

where we recall that $\text{dist}(a, b)$ is the length of the shortest path connecting a to b .

Theorem 59 (Alon–Boppana Bound). *Let G be a d -regular graph containing two edges (a_0, a_1) and (b_0, b_1) such that $\text{dist}((a_0, a_1), (b_0, b_1)) \geq 2k + 2$ for some integer $k \geq 0$. Then:*

$$\lambda_2 \leq d - 2\sqrt{d-1} + \frac{2\sqrt{d-1} - 1}{k+1}.$$

Proof. Define the following sets of vertices:

$$\begin{aligned} U_0 &= \{a_0, a_1\} \\ U_j &= N(U_{j-1}) - \bigcup_{i=0}^{j-1} U_i, \quad 1 \leq j \leq k \\ W_0 &= \{b_0, b_1\} \\ W_j &= N(W_{j-1}) - \bigcup_{i=0}^{j-1} W_i, \quad 1 \leq j \leq k. \end{aligned}$$

Put another way, U_j is exactly the set of vertices that are at distance j from U_0 , and W_j is exactly the set of vertices that are at distance j from W_0 . Note that since $\text{dist}((a_0, a_1), (b_0, b_1)) \geq 2k + 2$, there can be no vertices U_j and $W_{j'}$ for any $0 \leq j, j' \leq k$, i.e.,

$$U_j \cap W_{j'} = \emptyset, \quad \forall 0 \leq j, j' \leq k.$$

Furthermore, no vertex in U_j , for any $0 \leq j \leq k$, can be a neighbor of a vertex $W_{j'}$, for any $0 \leq j' \leq k$, since otherwise $\text{dist}((a_0, a_1), (b_0, b_1)) \leq 2k + 1$, which is a contradiction. See Figure 36 for an illustration.

As we learned in Section 12 and in particular using Theorem 13, we can upper bound λ_2 with a test vector $\mathbf{x} \in \mathbb{R}^n$ such that $\langle \mathbf{x}, \mathbf{1} \rangle = 0$. We need a good one to obtain the result, and so we pick:

$$\mathbf{x}(a) := \begin{cases} (d-1)^{-j/2} & a \in U_j \\ \beta(d-1)^{-j/2} & a \in W_j \\ 0 & \text{otherwise} \end{cases},$$

where $\beta \in \mathbb{R}$ is chosen so that $\langle \mathbf{x}, \mathbf{1} \rangle = 0$.




Figure 36: Illustration of the sets U_j and W_j .

Now we need to estimate $\mathbf{x}^T \mathbf{L} \mathbf{x}$ and $\mathbf{x}^T \mathbf{x}$ so we can estimate the Rayleigh quotient of \mathbf{x} . Let us begin with $\mathbf{x}^T \mathbf{x}$:

$$\begin{aligned}
 \mathbf{x}^T \mathbf{x} &= \sum_{a \in V} \mathbf{x}(a)^2 \\
 &= \sum_{j=0}^k \sum_{a \in U_j} \mathbf{x}(a)^2 + \sum_{j=0}^k \sum_{a \in W_j} \mathbf{x}(a)^2 \\
 &= \sum_{j=0}^k |U_j| (d-1)^{-j} + \beta^2 \sum_{j=0}^k |W_j| (d-1)^{-j} \\
 &= A + \beta^2 B,
 \end{aligned}$$

where we set

$$A := \sum_{j=0}^k |U_j| (d-1)^{-j} \quad \text{and} \quad B := \sum_{j=0}^k |W_j| (d-1)^{-j}.$$

To upper bound $\mathbf{x}^T \mathbf{L} \mathbf{x}$ we recall that

$$\mathbf{x}^T \mathbf{L} \mathbf{x} = \sum_{(a,b) \in E} (\mathbf{x}(a) - \mathbf{x}(b))^2.$$

We need to partition the edges in E , and in particular focus on those edges (a, b) for which $\mathbf{x}(a) \neq \mathbf{x}(b)$. First, notice any $a \in U_0$ only has neighbors in U_0 and U_1 . Similarly, for $1 \leq j \leq k-1$, any $a \in U_j$ only has neighbors in U_{j-1} , U_j , and U_{j+1} . Denote by $U \subset V$ all the vertices collectively in $\{U_j : 0 \leq j \leq k\}$, i.e.,

$$U := \bigcup_{j=0}^k U_j.$$

Define W analogously. Since U_k cannot be connected to any vertex in W , the neighbors of any vertex in U_k must be in U_{k-1} , U_k , or $V - (U \cup W)$. The same analysis applies to the sets W_j for $0 \leq j \leq k$. Therefore, the edges E are:

$$\begin{aligned} E = & \left(\bigcup_{j=0}^k E(U_j) \right) \cup \left(\bigcup_{j=0}^k E(W_j) \right) \cup E(V - (U \cup W)) \cup \dots \\ & \dots \cup \left(\bigcup_{j=0}^{k-1} E(U_j, U_{j+1}) \right) \cup E(U_k, V - (U \cup W)) \cup \dots \\ & \dots \cup \left(\bigcup_{j=0}^{k-1} E(W_j, W_{j+1}) \right) \cup E(W_k, V - (U \cup W)). \end{aligned}$$

Now, $\mathbf{x}(a) = \mathbf{x}(b)$ on all the edges in the first row of the decomposition of E , so we can ignore those edges. That leaves us with the second and third rows. We will need to bound the number of edges in $E(U_j, U_{j+1})$ and $E(U_k, V - (U \cup W))$, and likewise for the third row. Notice that each $a \in U_j$, for $0 \leq j \leq k-1$, must have at least one neighbor not in U_{j+1} . Indeed, $a_0, a_1 \in U_0$ are neighbors, and any $a \in U_j$ for $1 \leq j \leq k-1$ must be connected to at least one vertex in U_{j-1} . Therefore, since G is d -regular, each vertex in U_j has at most $d-1$ edges with vertices in U_{j+1} . Thus:

$$|E(U_j, U_{j+1})| \leq (d-1)|U_j|, \quad 0 \leq j \leq k-1.$$

A similar analysis shows that

$$|E(U_k, V - (U \cup W))| \leq (d-1)|U_k|.$$

Of course the same holds for the W_j edges in the third row.

Let us use $(\mathbf{x}^T \mathbf{L} \mathbf{x})|_U$ to denote the part of $\mathbf{x}^T \mathbf{L} \mathbf{x}$ that includes edges from the second row, and $(\mathbf{x}^T \mathbf{L} \mathbf{x})|_W$ the part of $\mathbf{x}^T \mathbf{L} \mathbf{x}$ that includes edges from the third row, so that $\mathbf{x}^T \mathbf{L} \mathbf{x} = (\mathbf{x}^T \mathbf{L} \mathbf{x})|_U + (\mathbf{x}^T \mathbf{L} \mathbf{x})|_W$. The analysis of these two components of $\mathbf{x}^T \mathbf{L} \mathbf{x}$ is similar,

so let us focus on $(\mathbf{x}^T \mathbf{L} \mathbf{x})|_U$:

$$\begin{aligned}
(\mathbf{x}^T \mathbf{L} \mathbf{x})|_U &= \sum_{j=0}^{k-1} \sum_{(a,b) \in E(U_j, U_{j+1})} (\mathbf{x}(a) - \mathbf{x}(b))^2 + \sum_{(a,b) \in E(U_k, V - (U \cup W))} (\mathbf{x}(a) - \mathbf{x}(b))^2 \\
&= \sum_{j=0}^{k-1} \sum_{(a,b) \in E(U_j, U_{j+1})} [(d-1)^{-j/2} - (d-1)^{-(j+1)/2}]^2 + \sum_{(a,b) \in E(U_k, V - (U \cup W))} (d-1)^{-k} \\
&\leq \sum_{j=0}^{k-1} (d-1)|U_j| [(d-1)^{-j/2} - (d-1)^{-(j+1)/2}]^2 + (d-1)|U_k|(d-1)^{-k} \\
&= \sum_{j=0}^{k-1} (d-1)|U_j| [(d-1)^{-(j+1)/2}((d-1)^{1/2} - 1)]^2 + |U_k|(d-1)^{-k+1} \\
&= \sum_{j=0}^{k-1} |U_j|(d-1)^{-j}(\sqrt{d-1} - 1)^2 + |U_k|(d-1)^{-k+1} \\
&= \sum_{j=0}^{k-1} |U_j|(d-1)^{-j}(d - 2\sqrt{d-1}) + |U_k|(d-1)^{-k}(d - 2\sqrt{d-1} + 2\sqrt{d-1} - 1) \\
&= (d - 2\sqrt{d-1}) \sum_{j=0}^k |U_j|(d-1)^{-j} + |U_k|(d-1)^{-k}(2\sqrt{d-1} - 1) \\
&= (d - 2\sqrt{d-1})A + (k+1)|U_k|(d-1)^{-k} \left(\frac{2\sqrt{d-1} - 1}{k+1} \right). \tag{72}
\end{aligned}$$

Now, the inequality $|E(U_j, U_{j+1})| \leq (d-1)|U_j|$ implies that

$$|U_{j+1}| \leq (d-1)|U_j| \implies |U_{j+1}|(d-1)^{-(j+1)} \leq |U_j|(d-1)^{-j}, \quad 0 \leq j \leq k-1.$$

In particular,

$$|U_k|(d-1)^{-k} \leq |U_j|(d-1)^{-j}, \quad 0 \leq j \leq k.$$

So continuing from (72) we have:

$$\begin{aligned}
(\mathbf{x}^T \mathbf{L} \mathbf{x})|_U &\leq (72) = (d - 2\sqrt{d-1})A + (k+1)|U_k|(d-1)^{-k} \left(\frac{2\sqrt{d-1} - 1}{k+1} \right) \\
&\leq (d - 2\sqrt{d-1})A + \left(\frac{2\sqrt{d-1} - 1}{k+1} \right) \sum_{j=0}^k |U_j|(d-1)^{-j} \\
&= (d - 2\sqrt{d-1})A + \left(\frac{2\sqrt{d-1} - 1}{k+1} \right) A \\
&= \left(d - 2\sqrt{d-1} + \frac{2\sqrt{d-1} - 1}{k+1} \right) A.
\end{aligned}$$

Since β factors out of all the terms in $(\mathbf{x}^T \mathbf{L} \mathbf{x})|_W$, we can apply the same approach to conclude that

$$(\mathbf{x}^T \mathbf{L} \mathbf{x})|_W \leq \left(d - 2\sqrt{d-1} + \frac{2\sqrt{d-1}-1}{k+1} \right) \beta^2 B.$$

Putting everything together we obtain:

$$\begin{aligned} \lambda_2 &\leq \frac{\mathbf{x}^T \mathbf{L} \mathbf{x}}{\mathbf{x}^T \mathbf{x}} = \frac{(\mathbf{x}^T \mathbf{L} \mathbf{x})|_U + (\mathbf{x}^T \mathbf{L} \mathbf{x})|_W}{\mathbf{x}^T \mathbf{x}} \\ &\leq \left(d - 2\sqrt{d-1} + \frac{2\sqrt{d-1}-1}{k+1} \right) \left(\frac{A + \beta^2 B}{A + \beta^2 B} \right) \\ &= d - 2\sqrt{d-1} + \frac{2\sqrt{d-1}-1}{k+1}. \end{aligned}$$

□

Recall that $\lambda_i = d - \mu_i$ for a d -regular graph, and so rearranging terms in Theorem 59 we obtain:

$$\mu_2 \geq 2\sqrt{d-1} - \frac{2\sqrt{d-1}-1}{k+1}.$$

In particular,

$$|\mu_2| \geq d \left(\frac{2\sqrt{d-1}}{d} - \frac{2\sqrt{d-1}-1}{d(k+1)} \right),$$

which implies that ϵ in a (d, ϵ) -expander graph must satisfy

$$\epsilon \geq \frac{2\sqrt{d-1}}{d} - \frac{2\sqrt{d-1}-1}{d(k+1)},$$

since by definition $|\mu_2| \leq \epsilon d$ in a (d, ϵ) -expander graph.

The quantity k is related to the graph diameter. Recall that $\text{diam}(G)$ is the maximum length of the shortest path between any two vertices in G . Let $a_0, b_0 \in G$ be chosen such that

$$\text{dist}(a_0, b_0) = \text{diam}(G).$$

Let $a_1 \in N(a_0)$ and $b_1 \in N(b_0)$. Then

$$\text{dist}((a_0, a_1), (b_0, b_1)) \geq 2k+2 \geq \text{diam}(G)-2 \implies k+1 \geq \frac{\text{diam}(G)}{2}-1 \geq \left\lfloor \frac{\text{diam}(G)}{2} \right\rfloor - 1,$$

and so

$$\frac{-1}{k+1} \geq \frac{-1}{\lfloor \text{diam}(G)/2 \rfloor - 1}.$$

Therefore,

$$\mu_2 \geq 2\sqrt{d-1} - \frac{2\sqrt{d-1}-1}{k+1} \geq 2\sqrt{d-1} - \frac{2\sqrt{d-1}-1}{\lfloor \text{diam}(G)/2 \rfloor - 1}.$$

Now, if one fixes d and lets $|V| = n \rightarrow \infty$, then $\text{diam}(G) \rightarrow \infty$, and we get that $\mu_2 \geq 2\sqrt{d-1}$. A *Ramanujan graph* is a connected d -regular graph in which

$$|\mu_i| \leq 2\sqrt{d-1}, \quad i \geq 2.$$

As such, they are very good (d, ϵ) -expander graphs for when n is large. Figure 37, taken from these [notes](#) by Matija Bucic, gives some examples of circular Ramanujan graphs. Over the years various folks have been interested in coming up with infinite families of Ramanujan graphs; see for example the [Wikipedia article](#) on Ramanujan graphs, which lists several types of Ramanujan graphs.




Figure 37: Examples of Ramanujan graphs

Remark 29. In [1, Chapter 30], Spielman proves that for every $\epsilon > 0$ there is a d for which there is an infinite family of (d, ϵ) -expander graphs. The family is constructed explicitly, so check it out if you want to see more examples of expander graphs!

30 Graph sparsification via random sampling

30.1 Introduction

Recall that $H = (V, E_H, w_H)$ is an ϵ -approximation of a graph $G = (V, E, w)$ if

$$(1 - \epsilon)\mathbf{L}_G \preceq \mathbf{L}_H \preceq (1 + \epsilon)\mathbf{L}_G.$$

In Section 29, Theorem 56, we learned that (d, ϵ) -expander graphs are ϵ -approximations of the re-scaled complete graph $(d/n)K_n$. We also proved in Theorem 59 that

$$\epsilon \geq \frac{2\sqrt{d-1}}{d} - \frac{2\sqrt{d-1}-1}{d(\text{diam}(G)/2-1)},$$

and we observed that the Ramanujan graphs, which are a special type of expander graph, satisfy

$$\epsilon \leq \frac{2\sqrt{d-1}}{d},$$

meaning they are very good approximations of K_n , particularly when n is large. Recall that a (d, ϵ) -expander graph is a d -regular graph and so it has $dn/2$ edges. We can calculate the number of edges in a Ramanujan graph in terms of its approximation error ϵ . Indeed, by the above two inequalities we can take:

$$\epsilon = \frac{2\sqrt{d-1}}{d} \implies \epsilon^2 = \frac{4(d-1)}{d^2} = \frac{4}{d} - \frac{4}{d^2} \leq \frac{4}{d} \implies \epsilon^{-2} \geq \frac{d}{4} \implies d = O(\epsilon^{-2}).$$

Therefore, the number of edges in a Ramanujan graph that ϵ approximates $(d/n)K_n$ is equal to

$$\frac{dn}{2} = O(\epsilon^{-2}n).$$

In this section we will prove that any connected graph G , not just the complete graph, can be well approximated by a sparse graph. We will see that by using a careful random sampling of the edges in G , we can obtain an ϵ approximation of G with only $O(\epsilon^{-2}n \log n)$ edges, which is nearly as sparse as the Ramanujan graph (but, again, here we approximate any G , not just K_n). This is a very powerful statement, as it allows us to approximate any

graph with a sparse graph through a relatively simple algorithm. As an example of this power, remember that if H is an ϵ -approximation of G , then by Theorem 15 we have

$$(1 - \epsilon)\lambda_i(G) \leq \lambda_i(H) \leq (1 + \epsilon)\lambda_i(G), \quad 1 \leq i \leq n.$$

Additionally, for all sets $S \subset V$, the number of edges in the boundary of S with respect to H is similar to the number of edges in the boundary of S with respect to G since, recalling a calculation we performed in the proof of Theorem 39, we have:

$$(1 - \epsilon)|\partial_G S| = (1 - \epsilon)\mathbf{1}_S^T \mathbf{L}_G \mathbf{1}_S \leq |\partial_H S| = \mathbf{1}_S^T \mathbf{L}_H \mathbf{1}_S \leq (1 + \epsilon)\mathbf{1}_S^T \mathbf{L}_G \mathbf{1}_S = (1 + \epsilon)|\partial_G S|.$$

There are other properties for which G and H will be similar as well.

30.2 Overview of the algorithm

Now for the random sampling algorithm that will generate the graph $H = (V, E_H, w_H)$ from the graph $G = (V, E, w)$:

1. Carefully choose a probability $p_{a,b}$ for each edge $(a, b) \in E$, where $p_{a,b}$ gives the probability that we keep the edge in H .
2. Randomly sample the edges from G to include in H using the probabilities $p_{a,b}$.
3. If (a, b) is included in H , we give it weight

$$w_H(a, b) := \frac{w(a, b)}{p_{a,b}}.$$

We will show that with properly chosen probabilities $p_{a,b}$, the resulting graph H will have $|E_H| = C\epsilon^{-2}n \log n$ edges with high probability, and H will be an ϵ -approximation of G also with high probability.

As some initial intuition, the reason for re-weighting the edges is that it will imply the expected matrix of \mathbf{L}_H will be equal to \mathbf{L}_G . To see this, recall that

$$\mathbf{L}_G = \sum_{(a,b) \in E} w(a, b) \mathbf{L}_{a,b}.$$

On the other hand,

$$\mathbf{L}_H = \sum_{(a,b) \in E} W_H(a, b) \mathbf{L}_{a,b}, \tag{73}$$

where $W_H(a, b)$ is the random variable defined as

$$\mathbb{P}\left(W_H(a, b) = \frac{w(a, b)}{p_{a,b}}\right) = p_{a,b} \quad \text{and} \quad \mathbb{P}(W_H(a, b) = 0) = 1 - p_{a,b}. \tag{74}$$

It follows that the expected value of $W_H(a, b)$ is:

$$\begin{aligned}\mathbb{E}[W_H(a, b)] &= \mathbb{P}\left(W_H(a, b) = \frac{w(a, b)}{p_{a,b}}\right) \cdot \frac{w(a, b)}{p_{a,b}} + \mathbb{P}(W_H(a, b) = 0) \cdot 0 \\ &= p_{a,b} \cdot \frac{w(a, b)}{p_{a,b}} + (1 - p_{a,b}) \cdot 0 \\ &= w(a, b).\end{aligned}$$

Thus, on average, the weight of each edge is preserved and thus the overall weight of the graph on average is preserved. Furthermore,

$$\mathbb{E}[\mathbf{L}_H] = \sum_{(a,b) \in E} \mathbb{E}[W_H(a, b)] \mathbf{L}_{a,b} = \sum_{(a,b) \in E} w(a, b) \mathbf{L}_{a,b} = \mathbf{L}_G,$$

indicating that the average graph Laplacian of H over an infinite number of random draws of H will be equal to \mathbf{L}_G . However, for most random draws of H , \mathbf{L}_H will be much sparser than \mathbf{L}_G (if we pick the probabilities $p_{a,b}$ correctly).

Remark 30. Since on average the overall weight of G is preserved by H , one can think of the new weights $w_H(a, b) = w(a, b)/p_{a,b}$, which increase the weight of the edge, as aggregating the weights of the edges that are removed from G into the fewer remaining edges left in H . As an additional reference point, if $G = K_n$ with all edge weights equal to one, and if \tilde{H} is a (d, ϵ) -expander graph also with all edge weights equal to one, then it is $H = (n/d)\tilde{H}$ that ϵ -approximates K_n . In particular, we see that all the edge weights are blown up in this case too.

30.3 Chernoff bounds

To prove that the algorithm from Section 30.2 works, we will need two types of Chernoff bounds, one for random variables and one for random matrices. All of these results are types of concentration inequalities.

Theorem 60 (Random variable Chernoff bound). *Let X_1, \dots, X_m be independent random variables taking values in $\{0, 1\}$. Define*

$$X := \sum_{i=1}^m X_i,$$

and let $\mu := \mathbb{E}[X]$. Then for any $0 \leq \delta \leq 1$,

$$\begin{aligned}\mathbb{P}(X \leq (1 - \delta)\mu) &\leq e^{-\delta^2 \mu / 2}, \\ \mathbb{P}(X \geq (1 + \delta)\mu) &\leq e^{-\delta^2 \mu / 3}.\end{aligned}$$

Theorem 61 (Random matrix Chernoff bound). *Let $\mathbf{X}_1, \dots, \mathbf{X}_m$ be independent random $n \times n$, symmetric, positive semi-definite matrices so that $\|\mathbf{X}_i\| \leq R$ almost surely. Define*

$$\mathbf{X} := \sum_{i=1}^m \mathbf{X}_i,$$

and let μ_{\min} and μ_{\max} be the minimum and maximum eigenvalues of

$$\mathbb{E}[\mathbf{X}] = \sum_{i=1}^m \mathbb{E}[\mathbf{X}_i],$$

respectively. Then for any $0 \leq \epsilon \leq 1$,

$$\begin{aligned} \mathbb{P}(\lambda_{\min}(\mathbf{X}) \leq (1 - \epsilon)\mu_{\min}) &\leq ne^{-\epsilon^2\mu_{\min}/2R}, \\ \mathbb{P}(\lambda_{\max}(\mathbf{X}) \geq (1 + \epsilon)\mu_{\max}) &\leq ne^{-\epsilon^2\mu_{\max}/3R}. \end{aligned}$$

Remark 31. It is important to note that the random variables X_1, \dots, X_m in Theorem 60, and the random matrices $\mathbf{X}_1, \dots, \mathbf{X}_m$ in Theorem 61, can have different distributions.

Remark 32. Looking ahead, we will use Theorem 60 to prove that the number of edges in H is $O(\epsilon^{-2}n \log n)$ with high probability, and we will use Theorem 61 to prove that H is an ϵ -approximation of G with high probability.

Lecture 25: Graph Sparsification via Random Sampling, Part II

April 15, 2021

Lecturer: Matthew Hirn

30.4 The key transformation

From Section 30.2 we know that $\mathbb{E}[\mathbf{L}_H] = \mathbf{L}_G$. We could try to write \mathbf{L}_H as the sum of random matrices and apply Theorem 61 with $\mathbf{X} = \mathbf{L}_H$ and $\mathbb{E}[\mathbf{X}] = \mathbb{E}[\mathbf{L}_H] = \mathbf{L}_G$, but to make our life easier we will first make an important transformation into an equivalent problem for which $\mu_{\min} = \mu_{\max} = 1$.

For positive definite matrices \mathbf{A} and \mathbf{B} we have

$$\mathbf{A} \preceq c\mathbf{B} \iff \mathbf{B}^{-1/2}\mathbf{A}\mathbf{B}^{-1/2} \preceq c\mathbf{I}.$$

In the above, recall that

$$\mathbf{B} = \Phi\mathbf{\Lambda}\Phi^T,$$

where $\mathbf{\Lambda}$ is the diagonal matrix with the eigenvalues $\lambda_1, \dots, \lambda_n$ of \mathbf{B} on its diagonal. Since \mathbf{B} is positive definite, $\lambda_i > 0$ for all $1 \leq i \leq n$, and we define

$$\mathbf{B}^{-1/2} := \Phi\mathbf{\Lambda}^{-1/2}\Phi^T, \quad \mathbf{\Lambda}^{-1/2}(i, i) := 1/\sqrt{\lambda_i}.$$

The same property holds for positive semi-definite matrices so long as they have the same null space, and if we replace the inverse of \mathbf{B} with the pseudo-inverse of \mathbf{B} . More precisely, let \mathbf{B} be a positive semi-definite matrix. Then the pseudo-inverse of \mathbf{B} is defined as:

$$\mathbf{B}^+ := \Phi\mathbf{\Lambda}^+\Phi^T, \quad \mathbf{\Lambda}^+(i, i) := \begin{cases} 0 & \lambda_i = 0 \\ \lambda_i^{-1} & \lambda_i > 0 \end{cases}.$$

Similarly, we define $\mathbf{B}^{+/2}$ as:

$$\mathbf{B}^{+/2} := \Phi\mathbf{\Lambda}^{+/2}\Phi^T, \quad \mathbf{\Lambda}^{+/2}(i, i) := \begin{cases} 0 & \lambda_i = 0 \\ \lambda_i^{-1/2} & \lambda_i > 0 \end{cases}.$$

Now, if \mathbf{A} and \mathbf{B} are positive semi-definite matrices with the same null space, then

$$\mathbf{A} \preceq c\mathbf{B} \iff \mathbf{B}^{+/2}\mathbf{A}\mathbf{B}^{+/2} \preceq c\mathbf{B}^{+/2}\mathbf{B}\mathbf{B}^{+/2}.$$

In particular, if H and G are both connected graphs, then both \mathbf{L}_H and \mathbf{L}_G have the same null space given by $\{\alpha\mathbf{1} : \alpha \in \mathbb{R}\}$. As such:

$$\mathbf{L}_H \preceq (1 + \epsilon)\mathbf{L}_G \iff \mathbf{L}_G^{+/2}\mathbf{L}_H\mathbf{L}_G^{+/2} \preceq (1 + \epsilon)\mathbf{L}_G^{+/2}\mathbf{L}_G\mathbf{L}_G^{+/2}.$$

Set

$$\mathbf{\Pi} := \mathbf{L}_G^{+/2} \mathbf{L}_G \mathbf{L}_G^{+/2}.$$

We conclude that

$$(1 - \epsilon) \mathbf{L}_G \preceq \mathbf{L}_H \preceq (1 + \epsilon) \mathbf{L}_G \iff (1 - \epsilon) \mathbf{\Pi} \preceq \mathbf{L}_G^{+/2} \mathbf{L}_H \mathbf{L}_G^{+/2} \preceq (1 + \epsilon) \mathbf{\Pi}.$$

In other words, \mathbf{L}_H is an ϵ -approximation of \mathbf{L}_G if and only if $\mathbf{L}_G^{+/2} \mathbf{L}_H \mathbf{L}_G^{+/2}$ is an ϵ -approximation of $\mathbf{\Pi}$. So rather than prove that \mathbf{L}_H is an ϵ -approximation of G with high probability, we will instead prove that $\mathbf{L}_G^{+/2} \mathbf{L}_H \mathbf{L}_G^{+/2}$ is an ϵ -approximation of $\mathbf{\Pi}$ with high probability.

At this point, it is instructive to get a feel for what the operator $\mathbf{\Pi}$ does. In words, it projects a vector \mathbf{x} onto the range \mathbf{L}_G . Since G is connected, the range of \mathbf{L}_G is given by $\text{span}\{\psi_2, \dots, \psi_n\}$. To see that $\mathbf{\Pi}$ projects \mathbf{x} onto $\text{span}\{\psi_2, \dots, \psi_n\}$, we compute

$$\mathbf{\Pi} = \mathbf{\Psi} \mathbf{\Lambda}^{+/2} \mathbf{\Psi}^T \mathbf{\Psi} \mathbf{\Lambda} \mathbf{\Psi}^T \mathbf{\Psi} \mathbf{\Lambda}^{+/2} \mathbf{\Psi}^T = \mathbf{\Psi} \mathbf{\Lambda}^{+/2} \mathbf{\Lambda} \mathbf{\Lambda}^{+/2} \mathbf{\Psi}^T = \mathbf{\Psi} \begin{pmatrix} 0 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{pmatrix} \mathbf{\Psi}^T.$$

As such,

$$\mathbf{\Pi} \mathbf{x} = \sum_{i=2}^n \langle \mathbf{x}, \psi_i \rangle \psi_i,$$

and we note that $\lambda_1(\mathbf{\Pi}) = 0$ and $\lambda_2(\mathbf{\Pi}) = \dots = \lambda_n(\mathbf{\Pi}) = 1$.

30.5 The edge probabilities and graph approximation

In this section we prove that $\mathbf{L}_G^{+/2} \mathbf{L}_H \mathbf{L}_G^{+/2}$ is an ϵ -approximation of $\mathbf{\Pi}$ with high probability. In other words, we want to show:

$$\mathbb{P} \left(\mathbf{L}_G^{+/2} \mathbf{L}_H \mathbf{L}_G^{+/2} \preceq (1 - \epsilon) \mathbf{\Pi} \right) \ll 1, \quad (75)$$

$$\mathbb{P} \left(\mathbf{L}_G^{+/2} \mathbf{L}_H \mathbf{L}_G^{+/2} \succeq (1 + \epsilon) \mathbf{\Pi} \right) \ll 1. \quad (76)$$

Note, if $\mathbf{L}_G^{+/2} \mathbf{L}_H \mathbf{L}_G^{+/2} \succeq (1 + \epsilon) \mathbf{\Pi}$, then by Theorem 15 we have

$$\lambda_n(\mathbf{L}_G^{+/2} \mathbf{L}_H \mathbf{L}_G^{+/2}) \geq (1 + \epsilon) \lambda_n(\mathbf{\Pi}) = 1 + \epsilon.$$

Thus,

$$\mathbb{P} \left(\mathbf{L}_G^{+/2} \mathbf{L}_H \mathbf{L}_G^{+/2} \succeq (1 + \epsilon) \mathbf{\Pi} \right) \leq \mathbb{P} \left(\lambda_n(\mathbf{L}_G^{+/2} \mathbf{L}_H \mathbf{L}_G^{+/2}) \geq 1 + \epsilon \right).$$

As such, to show (76) it is sufficient to show

$$\mathbb{P} \left(\lambda_n(\mathbf{L}_G^{+/2} \mathbf{L}_H \mathbf{L}_G^{+/2}) \geq 1 + \epsilon \right) \ll 1.$$

Similarly, to show (75) it is sufficient to show

$$\mathbb{P}\left(\lambda_2(\mathbf{L}_G^{+/2}\mathbf{L}_H\mathbf{L}_G^{+/2}) \leq 1 - \epsilon\right) \ll 1.$$

Note that we use λ_2 instead of λ_1 since we are guaranteed to have $\lambda_1(\mathbf{L}_G^{+/2}\mathbf{L}_H\mathbf{L}_G^{+/2}) = \lambda_1(\mathbf{\Pi}) = 0$.

At this point you should be thinking that we will want to use Theorem 61. To do so we need to do two things. First, we need to compute $\mathbb{E}[\mathbf{L}_G^{+/2}\mathbf{L}_H\mathbf{L}_G^{+/2}]$. We have:

$$\mathbb{E}[\mathbf{L}_G^{+/2}\mathbf{L}_H\mathbf{L}_G^{+/2}] = \mathbf{L}_G^{+/2}\mathbb{E}[\mathbf{L}_H]\mathbf{L}_G^{+/2} = \mathbf{L}_G^{+/2}\mathbf{L}_G\mathbf{L}_G^{+/2} = \mathbf{\Pi}.$$

That is great because it means that on average $\mathbf{L}_G^{+/2}\mathbf{L}_H\mathbf{L}_G^{+/2}$ is equal to $\mathbf{\Pi}$, and we will be able to use Theorem 61 to prove that any individual realization of $\mathbf{L}_G^{+/2}\mathbf{L}_H\mathbf{L}_G^{+/2}$ should not deviate from $\mathbf{\Pi}$ by too much, with high probability. We note that,

$$1 + \epsilon = (1 + \epsilon)\lambda_n(\mathbf{\Pi}) = (1 + \epsilon)\lambda_n\left(\mathbb{E}[\mathbf{L}_G^{+/2}\mathbf{L}_H\mathbf{L}_G^{+/2}]\right),$$

and similarly,

$$1 - \epsilon = (1 - \epsilon)\lambda_2(\mathbf{\Pi}) = (1 - \epsilon)\lambda_2\left(\mathbb{E}[\mathbf{L}_G^{+/2}\mathbf{L}_H\mathbf{L}_G^{+/2}]\right).$$

The second thing we need to do is write $\mathbf{L}_G^{+/2}\mathbf{L}_H\mathbf{L}_G^{+/2}$ as the sum of random matrices. To that end, for each $(a, b) \in E$ define $\mathbf{X}_{a,b}$ as the random matrix such that

$$\mathbb{P}\left(\mathbf{X}_{a,b} = \frac{w(a,b)}{p_{a,b}}\mathbf{L}_G^{+/2}\mathbf{L}_{a,b}\mathbf{L}_G^{+/2}\right) = p_{a,b} \quad \text{and} \quad \mathbb{P}(\mathbf{X}_{a,b} = \mathbf{0}) = 1 - p_{a,b}.$$

Recalling how we defined H in (73) and (74), we see that

$$\mathbf{L}_G^{+/2}\mathbf{L}_H\mathbf{L}_G^{+/2} = \sum_{(a,b) \in E} \mathbf{X}_{a,b}.$$

Now we need to specify the probabilities $p_{a,b}$

$$p_{a,b} := \frac{1}{R}w(a,b)\|\mathbf{L}_G^{+/2}\mathbf{L}_{a,b}\mathbf{L}_G^{+/2}\|,$$

where R is a parameter that we will specify shortly. Since

$$\mathbf{X}_{a,b} = \frac{w(a,b)}{p_{a,b}}\mathbf{L}_G^{+/2}\mathbf{L}_{a,b}\mathbf{L}_G^{+/2},$$

when $(a, b) \in E$ is chosen, we see that

$$\|\mathbf{X}_{a,b}\| \leq R.$$

Now, there is a chance that $p_{a,b} > 1$. We will address that scenario at the end of our discussion. For now let us assume that $p_{a,b} \leq 1$.

Set

$$R := \frac{\epsilon^2}{\tau \log n}, \quad \tau > 3.$$

Applying Theorem 61 we have:

$$\begin{aligned} \mathbb{P}\left(\mathbf{L}_G^{+/2} \mathbf{L}_H \mathbf{L}_G^{+/2} \succcurlyeq (1 + \epsilon) \mathbf{\Pi}\right) &\leq \mathbb{P}\left(\lambda_n(\mathbf{L}_G^{+/2} \mathbf{L}_H \mathbf{L}_G^{+/2}) \geq 1 + \epsilon\right) \\ &= \mathbb{P}\left(\lambda_n(\mathbf{L}_G^{+/2} \mathbf{L}_H \mathbf{L}_G^{+/2}) \geq (1 + \epsilon) \lambda_n(\mathbb{E}[\mathbf{L}_G^{+/2} \mathbf{L}_H \mathbf{L}_G^{+/2}])\right) \\ &\leq n \exp\left(-\frac{\epsilon^2}{3R}\right) \\ &= n \exp\left(-\frac{\tau \log n}{3}\right) \\ &= n \cdot n^{-\tau/3} \\ &= n^{-(\tau-3)/3}. \end{aligned}$$

Since $\lambda_1(\mathbf{X}_{a,b}) = 0$ for all $(a, b) \in E$, we can apply Theorem 61 to λ_2 instead of λ_1 to obtain:

$$\begin{aligned} \mathbb{P}\left(\mathbf{L}_G^{+/2} \mathbf{L}_H \mathbf{L}_G^{+/2} \preccurlyeq (1 - \epsilon) \mathbf{\Pi}\right) &\leq \mathbb{P}\left(\lambda_2(\mathbf{L}_G^{+/2} \mathbf{L}_H \mathbf{L}_G^{+/2}) \leq 1 - \epsilon\right) \\ &= \mathbb{P}\left(\lambda_2(\mathbf{L}_G^{+/2} \mathbf{L}_H \mathbf{L}_G^{+/2}) \leq (1 - \epsilon) \lambda_2(\mathbb{E}[\mathbf{L}_G^{+/2} \mathbf{L}_H \mathbf{L}_G^{+/2}])\right) \\ &\leq n \exp\left(-\frac{\epsilon^2}{2R}\right) \\ &= n \exp\left(-\frac{\tau \log n}{2}\right) \\ &= n \cdot n^{-\tau/2} \\ &= n^{-(\tau-2)/2}. \end{aligned}$$

Remark 33. Based on the above inequalities, one might be tempted to set τ to be very large. However, we will see in Section 30.6 that $\mathbb{E}[|E_H|]$ increases with τ .

30.6 The number of edges

Now we know that with high probability H will be an ϵ -approximation of G , but we still need to compute the expected number of edges in H . To that end define the random variables $X_{a,b}$ as

$$\mathbb{P}(X_{a,b} = 1) = p_{a,b} \quad \text{and} \quad \mathbb{P}(X_{a,b} = 0) = 1 - p_{a,b}.$$

It follows that $\mathbb{E}[X_{a,b}] = p_{a,b}$ and

$$|E_H| = X := \sum_{(a,b) \in E} X_{a,b}.$$

Thus,

$$\mathbb{E}[|E_H|] = \mathbb{E}[X] = \sum_{(a,b) \in E} \mathbb{E}[X_{a,b}] = \sum_{(a,b) \in E} p_{a,b}.$$

So we need to compute the sum of the probabilities $p_{a,b}$. Let us first compute $\|\mathbf{L}_G^{+/2} \mathbf{L}_{a,b} \mathbf{L}_G^{+/2}\|$. Since $\mathbf{L}_{a,b}$ has rank one, the matrix $\mathbf{L}_G^{+/2} \mathbf{L}_{a,b} \mathbf{L}_G^{+/2}$ also has rank one. Thus it has one non-zero eigenvalue $\lambda_{a,b} > 0$. Therefore:

$$\begin{aligned} \|\mathbf{L}_G^{+/2} \mathbf{L}_{a,b} \mathbf{L}_G^{+/2}\| &= \lambda_{a,b} \\ &= \text{Tr}(\mathbf{L}_G^{+/2} \mathbf{L}_{a,b} \mathbf{L}_G^{+/2}) \\ &= \text{Tr}(\mathbf{L}_G^+ \mathbf{L}_{a,b}). \end{aligned}$$

As such:

$$\begin{aligned} \sum_{(a,b) \in E} p_{a,b} &= \frac{1}{R} \sum_{(a,b) \in E} w(a,b) \|\mathbf{L}_G^{+/2} \mathbf{L}_{a,b} \mathbf{L}_G^{+/2}\| \\ &= \frac{1}{R} \sum_{(a,b) \in E} w(a,b) \text{Tr}(\mathbf{L}_G^+ \mathbf{L}_{a,b}) \\ &= \frac{1}{R} \text{Tr} \left(\sum_{(a,b) \in E} w(a,b) \mathbf{L}_G^+ \mathbf{L}_{a,b} \right) \\ &= \frac{1}{R} \text{Tr} \left(\mathbf{L}_G^+ \sum_{(a,b) \in E} w(a,b) \mathbf{L}_{a,b} \right) \\ &= \frac{1}{R} \text{Tr}(\mathbf{L}_G^+ \mathbf{L}_G) \\ &= \frac{1}{R} \text{Tr}(\mathbf{\Pi}) \\ &= \frac{n-1}{R}. \end{aligned}$$

We conclude that

$$\mathbb{E}[|E_H|] = \frac{n-1}{R}.$$

Recall,

$$R = \frac{\epsilon^2}{\tau \log n}, \quad \tau > 3,$$

so that

$$\mathbb{E}[|E_H|] = \tau \epsilon^{-2} (n-1) \log n.$$

Thus the expected number of edges is $O(\epsilon^{-2}n \log n)$, which is good. But what about any single particular realization of H ? Let $0 \leq \delta \leq 1$ and apply Theorem 60:

$$\begin{aligned}
\mathbb{P}(|E_H| \geq (1 + \delta)\tau\epsilon^{-2}(n - 1) \log n) &= \mathbb{P}(|E_H| \geq (1 + \delta)\mathbb{E}[|E_H|]) \\
&\leq \exp\left(-\frac{\delta^2\mathbb{E}[|E_H|]}{3}\right) \\
&= \exp\left(-\frac{\delta^2(n - 1)}{3R}\right) \\
&= \exp\left(-\frac{\tau\delta^2(n - 1) \log n}{3\epsilon^2}\right) \\
&\leq \exp\left(-\frac{\delta^2(n - 1) \log n}{\epsilon^2}\right) \\
&\leq \exp\left(-\frac{\delta^2(n - 1)}{\epsilon^2}\right).
\end{aligned}$$

We can simplify even further by setting $\delta = \epsilon$, in which case we get:

$$\mathbb{P}(|E_H| \geq (1 + \epsilon)\tau\epsilon^{-2}(n - 1) \log n) \leq e^{-(n-1)} \leq 3e^{-n}.$$

In other words, it is exceedingly unlikely that $|E_H|$ will have more than $O(\epsilon^{-2}n \log n)$ edges, even for small n .

30.7 Collecting everything

In this section we collect all the main results that we proved in the previous sections. Here is our algorithm:

- **Input:** A connected graph $G = (V, E, w)$, a tolerance $0 < \epsilon < 1$, and the parameter $\tau > 3$.

1. For each edge $(a, b) \in E$ compute the probabilities

$$p_{a,b} := \min\left(\frac{1}{R}w(a, b)\|\mathbf{L}_G^{+/2}\mathbf{L}_{a,b}\mathbf{L}_G^{+/2}\|, 1\right), \quad \text{with } R := \frac{\epsilon^2}{\tau \log n}.$$

2. Randomly sample the edges from G for inclusion in $H = (V, E_H, w_H)$ according to the probabilities:

$$\mathbb{P}((a, b) \in E_H) = p_{a,b} \quad \text{and} \quad \mathbb{P}((a, b) \notin E_H) = 1 - p_{a,b}.$$

3. For each edge $(a, b) \in E_H$, assign it the weight

$$w_H(a, b) = \frac{w(a, b)}{p_{a,b}}.$$

- **Output:** The graph $H = (V, E_H, w_H)$.

The following theorem collects what we can say about the graph H .

Theorem 62. *In the algorithm above suppose that*

$$p_{a,b} = \frac{1}{R} w(a,b) \| \mathbf{L}_G^{+/2} \mathbf{L}_{a,b} \mathbf{L}_G^{+/2} \| \leq 1, \quad \forall (a,b) \in E.$$

Then the expected number of edges in H is

$$\mathbb{E}[|E_H|] = \tau \epsilon^{-2} (n-1) \log n,$$

and furthermore

$$|E_H| < 2\mathbb{E}[|E_H|] = 2\tau \epsilon^{-2} (n-1) \log n \quad \text{with probability } 1 - 3e^{-n}.$$

Additionally, H is an ϵ -approximation of G with probability $1 - 2n^{-(\tau-3)/3}$.

Remark 34. As we mentioned earlier, what if $R^{-1}w(a,b) \| \mathbf{L}_G^{+/2} \mathbf{L}_{a,b} \mathbf{L}_G^{+/2} \| > 1$? We have two options. One is use the above algorithm, which accounts for this possibility by setting $p_{a,b} = 1$ in this case. Theorem 62 does not apply, but one can adapt the Chernoff bounds to get a similar result.

The other option is to adapt the algorithm. For any $p_{a,b} = R^{-1}w(a,b) \| \mathbf{L}_G^{+/2} \mathbf{L}_{a,b} \mathbf{L}_G^{+/2} \| > 1$, set $k = \lfloor p_{a,b} \rfloor \geq 1$. Create k copies of the edge $(a,b) \in E$ such that each of these copies is put into E_H with probability 1; weight these edges as $w_H(a,b) = w(a,b)/p_{a,b}$. Associated to these edges, create k random matrices $\mathbf{X}_{a,b}^{(j)}$, $1 \leq j \leq k$, such that

$$\mathbb{P} \left(\mathbf{X}_{a,b}^{(j)} = \frac{w(a,b)}{p_{a,b}} \mathbf{L}_G^{+/2} \mathbf{L}_{a,b} \mathbf{L}_G^{+/2} \right) = 1, \quad 1 \leq j \leq k.$$

Create one additional copy of the edge (a,b) such that this copy is placed in E_H with probability $p_{a,b} - k$ and give it the same edge weight $w_H(a,b) = w(a,b)/p_{a,b}$. Associate to this edge the random matrix $\mathbf{X}_{a,b}^{(k+1)}$ such that

$$\mathbb{P} \left(\mathbf{X}_{a,b}^{(k+1)} = \frac{w(a,b)}{p_{a,b}} \mathbf{L}_G^{+/2} \mathbf{L}_{a,b} \mathbf{L}_G^{+/2} \right) = p_{a,b} - k \quad \text{and} \quad \mathbb{P} \left(\mathbf{X}_{a,b}^{(k+1)} = \mathbf{0} \right) = 1 - (p_{a,b} - k).$$

Algorithmically, combine all the copies of (a,b) that make it into E_H into a single edge, and assign this single edge the weight equal to the sum of the edge weights of the copies. The proof of Theorem 62 will go through with minor modifications; in particular, one will be able to use the same Chernoff bounds.

Lecture 26: Graph Convolution Networks

April 20, 2021

Lecturer: Matthew Hirn

31 Graph convolution networks

The field of graph neural networks is now a vast a rich subfield of deep learning, but its origins are in spectral graph theory and graph signal processing. We describe those origins in this section. The references for this material are [6, 7, 8].

31.1 Convolution neural networks

To begin, let us give a quick primer on how standard convolution neural networks (CNNs) work. Let x be an $N \times N$ image and let h be a small filter, such as a 3×3 filter. The convolution of x with h computes:

$$x * h(u) = \sum_v x(v)h(u - v),$$

where $u \in \{-N/2, \dots, N/2 - 1\} \times \{-N/2, \dots, N/2 - 1\}$ and the sum over v can be restricted to the support of h in the neighborhood of u . Note that one has to be careful at the boundary, either using a zero boundary condition or a periodic boundary condition, but we ignore those details here.

The main component of a CNN is the transformation that takes the channels at layer ℓ to the channels at layer $\ell + 1$; Figure 38 gives an illustration of the VGG network. Suppose that layer ℓ , for $\ell \geq 0$, has d_ℓ channels given by:

$$\text{Layer } \ell \text{ channels} = \left(x_i^{(\ell)} \right)_{i=1}^{d_\ell}, \quad x_i^{(\ell)} \in \mathbb{R}^{N \times N}.$$

Layer $\ell + 1$ has $d_{\ell+1}$ channels defined as:

$$\text{Layer } \ell + 1 \text{ channels} = x_j^{(\ell+1)} := \sigma \left[\sum_{i=1}^{d_\ell} x_i^{(\ell)} * h_{ij}^{(\ell)} \right], \quad 1 \leq j \leq d_{\ell+1},$$

where $(h_{ij}^{(\ell)} : 1 \leq i \leq d_\ell, 1 \leq j \leq d_{\ell+1})$ are learned small filters and σ is a pointwise nonlinearity,

$$\sigma(x)(u) := \sigma(x(u)).$$




Figure 38: The VGG convolution neural network.




Figure 39: Pixel locations as a graph.

For example, $\sigma(t) = \max(0, t)$ is the commonly used rectified linear unit (ReLU) nonlinearity. In image classification, the input layer of the network is either a single channel ($d_0 = 1$) if the image is grayscale, or it has three channels ($d_0 = 3$) if the image is a color image (red, green, and blue channels). At some layers there is also a pooling operation (these are illustrated by the narrowing of the network in Figure 38), but we ignore this operation in our analysis.

We can think of the underlying pixel locations as lying on the grid graph $P_N \times P_N$ that also includes diagonal edges; see Figure 39. It follows that if h is a 3×3 filter, then $h(a - b)$ (as a function of b) is supported on $\{a\} \cup N(a)$. In other words, h is a 1-hop filter.

31.2 Graph convolution networks

Let us now extend CNNs to general connected graphs $G = (V, E, w)$. Recall \mathbf{L} is the graph Laplacian of G , ψ_1, \dots, ψ_n are an orthonormal set of eigenvectors of \mathbf{L} with eigenvalues

$0 = \lambda_1 < \lambda_2 \leq \dots \leq \lambda_n$. Let $\mathbf{x} : V \rightarrow \mathbb{R}$ be a graph signal and let $\mathbf{h} : V \rightarrow \mathbb{R}$ be a graph filter. Recall the graph Fourier transform of \mathbf{x} is

$$\widehat{\mathbf{x}}(k) := \langle \mathbf{x}, \boldsymbol{\psi}_k \rangle,$$

and graph convolution is defined as

$$\mathbf{x} * \mathbf{h}(a) := \sum_{k=1}^n \widehat{\mathbf{x}}(k) \widehat{\mathbf{h}}(k) \boldsymbol{\psi}_k(a) = \boldsymbol{\Psi} \mathbf{H} \boldsymbol{\Psi}^T \mathbf{x},$$

where $\boldsymbol{\Psi}$ is the $n \times n$ for which $\boldsymbol{\psi}_k$ is column k , and \mathbf{H} is the $n \times n$ diagonal matrix with $\widehat{\mathbf{h}}$ down its diagonal:

$$\mathbf{H} := \begin{pmatrix} \widehat{\mathbf{h}}(1) & 0 & \cdots & 0 \\ 0 & \widehat{\mathbf{h}}(2) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \widehat{\mathbf{h}}(n) \end{pmatrix}.$$

In a graph convolution network (GCN), the transform from layer ℓ to layer $\ell + 1$ is:

$$\mathbf{x}_j^{(\ell+1)} := \sigma \left[\sum_{i=1}^{d_\ell} \mathbf{x}_i^{(\ell)} * \mathbf{h}_{ij}^{(\ell)} \right] = \sigma \left[\sum_{i=1}^{d_\ell} \boldsymbol{\Psi} \mathbf{H}_{ij}^{(\ell)} \boldsymbol{\Psi}^T \mathbf{x}_i^{(\ell)} \right], \quad 1 \leq j \leq d_{\ell+1}, \quad \mathbf{x}_j^{(\ell+1)} \in \mathbb{R}^n,$$

where the diagonals of the matrices $\mathbf{H}_{ij}^{(\ell)}$ are learned for each layer $0 \leq \ell < L =$ number of layers and all $1 \leq i \leq d_\ell$ and $1 \leq j \leq d_{\ell+1}$.

In a CNN, the convolution $\mathbf{x} * \mathbf{h}(a)$ replaces $\mathbf{x}(a)$ with information aggregated from $\mathbf{x}(b)$ for $b = a$ and pixels b close to a . Similarly, in a GCN, if $\widehat{\mathbf{h}}(k) = p(\lambda_k)$ for an m^{th} order polynomial p , then the convolution $\mathbf{x} * \mathbf{h}(a)$ replaces $\mathbf{x}(a)$ with information aggregated from $\mathbf{x}(a)$ and $\mathbf{x}(b)$ for b in the m -hop neighborhood of a (recall Theorem 34).

31.3 Graph learning tasks

In regular CNNs the graph (i.e., the pixel grid) is fixed and the task is to classify different signals (images) defined on the graph. We can extend this paradigm to the general graph setting as well. In this case we are given a fixed graph $G = (V, E, w)$ and many signals defined on the vertices of G , some labeled (the training set) and some not labeled (the test set). The task is to classify the unlabeled signals. Here are two examples:

- Spherical MNIST: The MNIST handwritten digits are projected onto the sphere. The task is to label the digits. Numerically, the sphere is approximated with a graph.
- Text document classification: The graph $G = (V, E)$ consists of words (vertices) and edges between pairs of words that are similar. To each text document, we associate a graph signal such that $\mathbf{x}(a)$ counts the number of times the word a appears in the document (i.e, \mathbf{x} is a word histogram). The task is to classify the text documents into different categories.

In the graph setting, though, there are other learning tasks in addition to signal classification. These other tasks include:

1. Node classification: We have a fixed graph $G = (V, E, w)$. Some of the vertices are labeled, others are not labeled; the task is to label the unlabeled vertices. For example, G is a gene-gene interaction network (vertices are genes), and some genes we know whether or not they are correlated with a certain disease. The task is to predict whether the other genes are correlated with this disease.
2. Link prediction: We have a fixed graph $G = (V, E)$. Based on existing edges, the task is to predict if there should exist an edge between two vertices $a, b \in V$ for which $(a, b) \notin E$. For example, friendship recommendations in social networks.
3. Graph classification: We have many graphs, some of them are labeled, others are not labeled. The task is to label the unlabeled graphs. For example, each graph represents a molecule (atoms are vertices, edges are bonds), and we want to predict whether the molecule is toxic or not.

It turns out that all of these tasks can be approached using a GCN. The case of signal classification over a fixed graph is clear. In the other three cases, even though they are not about signal classification, we can process a signal (or signals) over the graph(s) in order to extract information from G . These signals are sometimes given to us in the form of “side information” about the vertices of the graph, often referred to as vertex features. On the other hand, if such information is not given to us, we can generate surrogate signals on the graph as a means by which to extract information from G , similar to how we use test vectors in spectral graph theory.

Thus, in all four of the above graph learning tasks, our input to the GCN is a set of input channels $(\mathbf{x}_i^0)_{i=1}^{d_0}$, and the output of the GCN is a set of output channels $(\mathbf{x}_i^{(L)})_{i=1}^{d_L}$. What we do with the output channels depends on the task. In the case of signal classification, there is a single input channel (the signal to be classified) and the output channels are the representation of this signal. These output channels are passed to a classifier (such as an artificial neural network with a few fully connected layers) that assigns a label to the signal. For the other tasks, we proceed in the following manner:

1. Node classification: Each node (vertex in the graph) is assigned the representation

$$a \mapsto (\mathbf{x}_i^{(L)}(a))_{i=1}^{d_L} \in \mathbb{R}^{d_L}.$$

This representation is passed to a classifier that assigns a label to a .

2. Link prediction: Let $a, b \in V$ with $(a, b) \notin E$. We assign the candidate edge (a, b) a representation derived from the vertex representations of a and b , for example:

$$(a, b) \mapsto ((\mathbf{x}_i^{(L)}(a) - \mathbf{x}_i^{(L)}(b))^2)_{i=1}^{d_L} \in \mathbb{R}^{d_L}.$$

3. Graph classification: In this case the representation of the graph G must be independent of the number of vertices in G and must be invariant to re-indexation of the vertices/edges. We obtain such a representation by mapping G to:

$$G \mapsto \sum_{a \in V_G} \mathbf{x}_i^{(L)}(a) \in \mathbb{R}^{d_L}.$$

31.4 Parameterized graph filters

While the setup in the previous section has potential, it also has issues that need to be fixed. Here are two that we will address:

1. Each filter is defined by $\hat{\mathbf{h}} \in \mathbb{R}^n$, which means that we need to learn n values for each filter. But this is significantly more learned values per filter than in a standard CNN. Indeed, in a CNN, $n = N^2$ = the number of pixels, but most CNNs use small filters, e.g., 3×3 filters with 9 learned parameters.
2. For graph classification, the number of vertices n depends on the graph G , as do the eigenvectors and eigenvalues of \mathbf{L}_G that are used to define graph convolution. How do we transfer a graph convolution network across multiple graphs?

It turns out that we can solve both of these issues by parameterizing the filters. That is, we replace \mathbf{h} with \mathbf{h}_θ , where $\theta \in \mathbb{R}^m$ is a set of parameters that defines \mathbf{h}_θ , and m is independent of G . For example, $\hat{\mathbf{h}}_\theta(k)$ can be taken to be an $(m-1)^{\text{st}}$ order polynomial of λ_k :

$$\hat{\mathbf{h}}_\theta(k) = p_\theta(\lambda_k) = \sum_{t=0}^{m-1} \theta^{(t)} \lambda_k^t.$$

In this case, each filter in the GCN is parameterized as an $(m-1)^{\text{st}}$ degree polynomial and the network learns, for each filter, its respective m polynomial coefficients $\theta \in \mathbb{R}^m$. Therefore, instead of needing to learn n values for each filter, the network learns m values for each filter, and we can take $m \ll n$. Additionally, since each filter is a parameterized function of the eigenvalues of the graph, it can be transferred to a new graph \tilde{G} with Laplacian eigenvalues $0 = \tilde{\lambda}_1 < \tilde{\lambda}_2 \leq \dots \leq \tilde{\lambda}_n$ by defining the analogous filter on \tilde{G} as:

$$\hat{\mathbf{h}}_\theta(k) = p_\theta(\tilde{\lambda}_k).$$

Furthermore, from Theorem 34 we know that such a filter \mathbf{h}_θ corresponds to an $(m-1)$ -hop filter, thus giving us control over the support of the filter with respect to the graph.

Even so, the way in which we defined graph convolution requires computing the eigenvalues and eigenvectors of the graph Laplacian. For large graphs this may entail a prohibitive computational cost, and for the graph classification problem it requires diagonalizing the graph Laplacian for every graph we consider. However, if we take $\hat{\mathbf{h}}_\theta$ as a polynomial, then we have:

$$\mathbf{x} * \mathbf{h}_\theta = \Psi \mathbf{H}_\theta \Psi^T \mathbf{x} = \Psi p_\theta(\Lambda) \Psi^T \mathbf{x} = p_\theta(\Psi \Lambda \Psi^T) \mathbf{x} = p_\theta(\mathbf{L}) \mathbf{x}.$$

As such, the transform from layer ℓ to layer $\ell + 1$ can be written as:

$$\mathbf{x}_j^{(\ell+1)} = \sigma \left[\sum_{i=1}^{d_\ell} \mathbf{x}_i^{(\ell)} * \mathbf{h}_{\theta_{ij}^{(\ell)}} \right] = \sigma \left[\sum_{i=1}^{d_\ell} p_{\theta_{ij}^{(\ell)}}(\mathbf{L}) \mathbf{x}_i^{(\ell)} \right], \quad 1 \leq j \leq d_{\ell+1}.$$

In particular, we see there is no need to diagonalize the graph Laplacian \mathbf{L} . Furthermore, if G is sparse, then \mathbf{L} will be sparse, and thus $p_\theta(\mathbf{L})$ will be sparse if m is small, e.g., $m = 2$.

Notice that choosing $m = 2$ corresponds to a small 3×3 CNN filter since if $m = 2$ then p_θ is a 1st order polynomial, which means the resulting filter \mathbf{h}_θ is a 1-hop filter. In other words, $\mathbf{x} * \mathbf{h}_\theta(a)$ will replace $\mathbf{x}(a)$ with information aggregated from $\mathbf{x}(a)$ and $\mathbf{x}(b)$ for $b \in N(a)$, which is a direct analogue of how a 3×3 CNN filter aggregates information from the central pixel and its neighboring pixels. In this case we have:

$$p_\theta(\mathbf{L}) = \theta^{(0)} \mathbf{I} + \theta^{(1)} \mathbf{L} \implies \mathbf{x} * \mathbf{h}_\theta = p_\theta(\mathbf{L})\mathbf{x} = \theta^{(0)}\mathbf{x} + \theta^{(1)}\mathbf{L}\mathbf{x}.$$

Now, since re-scaling the weights of a graph $G = (V, E, w)$ is likely to not change the outcome of the learning task, we replace the graph Laplacian \mathbf{L} with the normalized graph Laplacian \mathbf{N} :

$$p_\theta(\mathbf{N}) = \theta^{(0)} \mathbf{I} + \theta^{(1)} \mathbf{N} = \theta^{(0)} \mathbf{I} + \theta^{(1)} (\mathbf{I} - \mathbf{D}^{-1/2} \mathbf{M} \mathbf{D}^{-1/2}) = (\theta^{(0)} + \theta^{(1)}) \mathbf{I} - \theta^{(1)} \mathbf{D}^{-1/2} \mathbf{M} \mathbf{D}^{-1/2}.$$

Even though this filter only has two learnable parameters, we can reduce this to one learnable parameter by replacing $\theta = (\theta^{(0)}, \theta^{(1)})$ with a scalar parameter θ satisfying:

$$\theta := -\theta^{(1)} = \theta^{(0)} + \theta^{(1)}.$$

We then have:

$$p_\theta(\mathbf{N}) = \theta(\mathbf{I} + \mathbf{D}^{-1/2} \mathbf{M} \mathbf{D}^{-1/2}) = \theta(\mathbf{I} + \mathbf{A}),$$

where we recall that $\mathbf{A} = \mathbf{D}^{-1/2} \mathbf{M} \mathbf{D}^{-1/2}$ is the normalized adjacency matrix. Recall that $\Lambda(\mathbf{A}) \subset [-1, 1]$, which implies that $\Lambda(\mathbf{I} + \mathbf{A}) \subset [0, 2]$. In [8], Kipf and Welling argue that because the spectrum lies in $[0, 2]$, it leads to vanishing and exploding gradients during training. In order to alleviate this phenomenon, they replace $\mathbf{I} + \mathbf{A}$ with

$$\tilde{\mathbf{A}} := \tilde{\mathbf{D}}^{-1/2} \tilde{\mathbf{M}} \tilde{\mathbf{D}}^{-1/2}, \quad \text{where} \quad \tilde{\mathbf{M}} := \mathbf{I} + \mathbf{M}.$$

While there is not much explanation in [8] for why this substitution improves the GCN training, some later papers attempted to explain it more. The final filter is:

$$\mathbf{x} * \mathbf{h}_\theta = \theta \tilde{\mathbf{A}} \mathbf{x}. \tag{77}$$

With these filters, we can now write the layer ℓ to layer $\ell + 1$ transform in the following way. Let $\mathbf{X}^{(\ell)}$ denote the $n \times d_\ell$ matrix with $\mathbf{x}_i^{(\ell)}$ as its i^{th} column. Using (77) and the definition of the layer ℓ to layer $\ell + 1$ transform, one can show:

$$\mathbf{X}^{(\ell+1)} = \sigma \left(\tilde{\mathbf{A}} \mathbf{X}^{(\ell)} \boldsymbol{\Theta}^{(\ell)} \right),$$

where $\boldsymbol{\Theta}^{(\ell)}$ is a $d_\ell \times d_{\ell+1}$ matrix of learned parameters. This is the basic form of many modern GCNs.

References

- [1] Daniel A. Spielman. Spectral and algebraic graph theory. Book draft, available at: <http://cs-www.cs.yale.edu/homes/spielman/sagt/>, 2019.
- [2] Michael Perlmutter, Feng Gao, Guy Wolf, and Matthew Hirn. Geometric scattering networks on compact Riemannian manifolds. In *Proceedings of The First Mathematical and Scientific Machine Learning Conference, Proceedings of Machine Learning Research*, volume 107, pages 570–604, 2020.
- [3] David I. Shuman, Sunil K. Narang, Pascal Frossard, Antonio Ortega, and Pierre Vandergheynst. The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains. *IEEE Signal Processing Magazine*, 30(3):83–98, 2013.
- [4] Stéphane Mallat. *A Wavelet Tour of Signal Processing, Third Edition: The Sparse Way*. Academic Press, 3rd edition, 2008.
- [5] Ronald R. Coifman and Stéphane Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21:5–30, 2006.
- [6] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. Spectral networks and deep locally connected networks on graphs. In *International Conference on Learning Representations (ICLR)*, 2014.
- [7] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in Neural Information Processing Systems 29*, pages 3844–3852, 2016.
- [8] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*, 2017.