CS 270: Combinatorial Algorithms and Data Structures	Spring 2023

Lecture 24 — April 18th, 2023

Prof. Jelani Nelson *Scribe: Andrew Huang*

1 Overview

Today's and Thursday's lectures will be on Spectral Graph Theory. Today, we'll start it off with definitions, do examples, before proving basic properties of Laplacians. We'll also discuss the easier half of Cheeger's inequality, (the harder side will be discussed Thursday).

2 What is Spectral Graph Theory?

We know what graph theory means, and generally we can talk about its adjacency matrix. Spectral graph theory discusses the spectrum of the matrices associated with the graphs, particularly its eigenvalues and eigenvectors. One of these matrices that we generally look at is the adjacency matrix, another more complicated one is the Laplacian.

In other words, "How much can we learn from G by looking at the eigenvalues, eigenvectors, of associated matrices? (For example, adjacency matrix, or Laplacian)"

Nota Bene: The area of algorithmal graph theory is kind of this intersection between regular and spectral graph theory, kind of trying to use this information from the matrices and then use that to create algorithms doing something.

2.1 Intro Assumptions

In this lecture, we will be looking at undirected, weighted graphs G, where for every edge e, the corresponding weight $w_e \geq 0$.

Definition 2.1. The *adjacency matrix* $A(G)$ is defined by

$$A(G)_{u,v} = \begin{cases} w_e & \text{if } (u,v) = e \in E \\ 0 & \text{otherwise} \end{cases}$$

Note that $A(G) \in \mathbb{R}^{n\times n}$ is symetric and has nonnegative entries. So we can form this correspondence:

Symmetric matrix in $\mathbb{R}^{n\times n}$ with nonnegative entries	$\leftrightarrow$	unweighted, undirected graph G with all edge weights nonnegative

Definition 2.2. The *Laplacian matrix* $L(G)$ is defined as $D(G)-A(G)$, where $D(G)$ is the diagonal matrix of size $n \times n$ where $D(G)_{u,u} = \deg_G(u) = \sum_{e=(u,\cdot)\in E} w_e$.

Some examples:

Example 2.3. For the complete graph, K_n (all with weights of 1) we have:

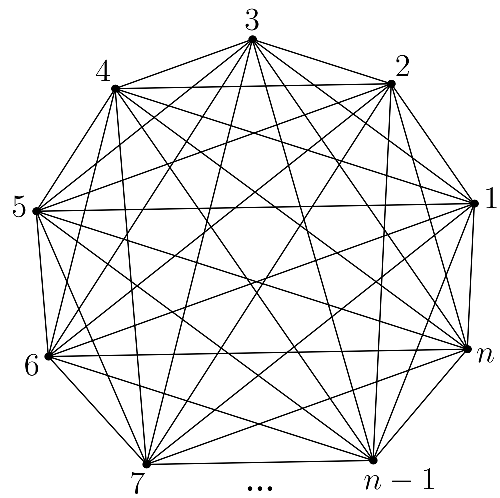

$$L(K_n) = \begin{bmatrix} n-1 & -1 & -1 & \cdots & -1 \\ -1 & n-1 & -1 & \cdots & -1 \\ -1 & -1 & n-1 & \cdots & -1 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ -1 & -1 & -1 & \cdots & n-1 \end{bmatrix}$$

Example 2.4. For a star graph S_n, if we define the middle vertex to be vertex 1 (one-indexing), we have:

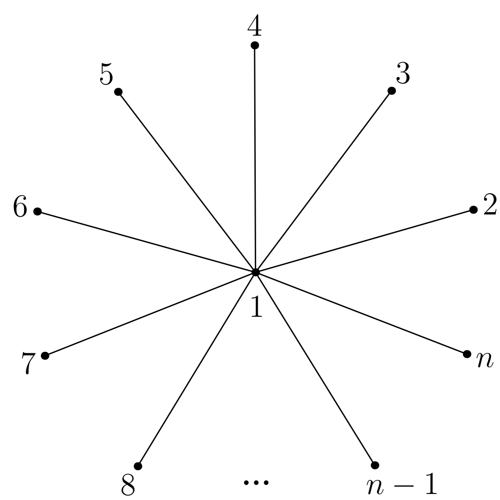

$$L(S_n) = \begin{bmatrix} n-1 & -1 & -1 & \cdots & -1 \\ -1 & 1 & 0 & \cdots & 0 \\ -1 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ -1 & 0 & 0 & \cdots & 1 \end{bmatrix}$$

2.2 Positive Semidefiniteness of Laplacians

It turns out that the Laplacian is PSD. Let's first recall what that means.

Definition 2.5. A symmetric matrix $M \in \mathbb{R}^{n \times n}$ is *positive semidefinite* (PSD) if any of the following equivalent properties hold:

(1) $\forall x \in \mathbb{R}^n$, $x^\top M x \geq 0$,

(2) $\exists Q$ s.t. $M = Q^\top Q$.

(3) All eigenvalues of M are ≥ 0. (Recall that all eigenvalues of a symmetric matrix are already real)

Note: we can write $L(G) = \sum_{e \in E} w_e \cdot L(e)$, where $L(e)$ is the Laplacian matrix of the graph with the same vertices but removing all edges except e from E.

In other words, for example if $e = (u, v)$, then $L(e)$ is the sparse matrix (with unlabelled entries being zero):

		u		v	
u		1		-1	
v		-1		1	

We can intuitively see that this equation ($L(G) = \sum_{e \in E} w_e \cdot L(e)$) holds as the only $L(e)$ will affect the entry in the uth row and vth column, and for a diagonal entry, the sum $\sum$.

Now, let's go back to our claim that L is PSD. We have:

$$x^\top L x = \sum_{e \in E} w_e x^\top L(e) x.$$

Now, we observe that $L_e = (\mathbb{1}_u - \mathbb{1}_v)(\mathbb{1}_u - \mathbb{1}_v)^\top$, where $\mathbb{1}_u$ is the column vector where the only nonzero entry is a 1 in the uth entry. This is in part because $(ZZ^\top)_{ij} = Z_i Z_j$.

Thus, we have:

$$\begin{aligned} x^\top L x &= \sum_{e \in E} w_e \underbrace{(x^\top(\mathbb{1}_u - \mathbb{1}_v))^2}_{\langle X, \mathbb{1}_u - \mathbb{1}_v \rangle} \\ &= \sum_{e \in E} w_e (x_u - x_v)^2 \qquad (*) \end{aligned}$$

Claim 2.6. L is PSD.
Proof: We just showed the first condition for being PSD in Eq. $(*)$.

Now, this is sufficient for showing that L is PSD, but we can also directly show the second condition for being PSD. To do this, we define the edge incidence matrix.

Definition 2.7. The edge-vertex incidence matrix, $B(G)$, is a matrix in $\mathbb{R}^{m \times n}$, where the row corresponding to $e = (u, v)$ is $(\mathbb{1}_u - \mathbb{1}_v)^\top$.

For example, it may look like:

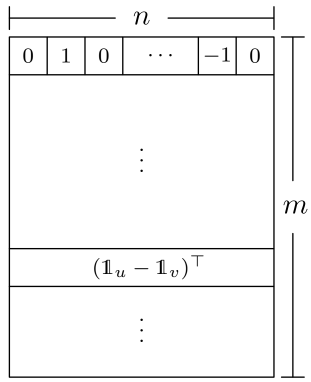

Now that we have the edge vertex incidence matrix, we can proceed. Let us quickly define $W \in \mathbb{R}^{m \times m}$ as

$$W = \begin{bmatrix} w_{e_1} & 0 & 0 & \cdots & 0 \\ 0 & w_{e_2} & 0 & \cdots & 0 \\ 0 & 0 & w_{e_3} & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & w_{e_m} \end{bmatrix}.$$

We have that

$$L = B^\top W B = B^\top W^{1/2} W^{1/2} B = (W^{1/2} B)^\top (W^{1/2} B).$$

Taking $Q = W^{1/2} B$, we have proved the second condition.

Nota Bene: The theorem we'll prove actually has roots considering manifolds in differential geometry, which is why it will have terminology coming from it. It was then ported over to spectral graph theory, which in a sense means that we'll have to take analogous definitions.

3 Quick Linear Algebra Review

Recall that "v is an eigenvector of $M \in \mathbb{R}^{n \times n}$ with eigenvalue $\lambda \in \mathbb{C}$ if $Mv = \lambda v$."

Claim 3.1. Suppose v_1, v_2 are eigenvectors with eigenvalues λ_1, λ_2 respectively. Furthermore, suppose M is symmetric. Then:

(1) $\lambda_1 \neq \lambda_2 \Rightarrow \langle v_1, v_2 \rangle = 0$ (in orther words, eigenvectors with different eigenvalues are pairwise orthogonal)

(2) $\lambda_1 = \lambda_2 \Rightarrow \forall \alpha, \beta \in \mathbb{R}$, $\alpha v_1 + \beta v_2$ is an eigenvector also with eigenvalue λ_1 (In other words, the set of all eigenvectors with fixed eigenvalue λ forms a vector space).

Proof:

(1) We have that:

$$\lambda_1 v_1^\top v_2 = (Mv_1)^\top v_2 = v_1^\top M v_2 = \lambda_2 \vec{v}_1^\top v_2.$$

Therefore, $0 = \lambda_1 v_1^\top v_2 - \lambda_2 v_1^\top v_2 = (\underbrace{\lambda_1 - \lambda_2}_{\neq 0}) v_1^\top v_2$, and it must be the case that $v_1^\top v_2 = 0$.

(2) $M(\alpha v_1 + \beta v_2) = \alpha M v_1 + \beta M v_2 = \alpha \lambda v_1 + \beta \lambda v_2 = \lambda(\alpha v_1 + \beta v_2)$.

Now, for our last bit of Linear Algebra review, we state (but don't prove) the Spectral Theorem:

Theorem 3.2 (Spectral Theorem)**.** *Let $M \in \mathbb{R}^{n\times n}$, symmetric.*

Then, $\exists$ an orthonormal basis, $v_1, v_2, \cdots, v_m$ of eigenvectors with corresponding real eigenvalues $\lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_n$,

Then, we have that $M = V\Lambda V^\top = \sum_{i=1}^n \lambda_i v_i v_i^\top$, where

$$V = \begin{bmatrix} | & | & & | \\ v_1 & v_2 & \cdots & v_n \\ | & | & & | \end{bmatrix}, \Lambda = \begin{bmatrix} \lambda_1 & 0 & 0 & \cdots & 0 \\ 0 & \lambda_2 & 0 & \cdots & 0 \\ 0 & 0 & \lambda_3 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & \lambda_n \end{bmatrix}$$

Nota Bene: note that in general, $AB^\top = \sum_i a_i b_i^\top$.

4 Laplacians and Eigenvectors

Now, note that if $\lambda_1 = 0$, if we define $\mathbb{1}$ to be the all ones vector we have

$$\begin{aligned} L\mathbb{1} &= \sum_e w_e \underbrace{L(e)}_{(\mathbb{1}_u - \mathbb{1}_v)(\mathbb{1}_u - \mathbb{1}_v)^\top} \mathbb{1} \\ &= \sum_e w_e (\mathbb{1}_u - \mathbb{1}_v)\langle \mathbb{1}_u - \mathbb{1}_v, \mathbb{1}\rangle \\ &= 0 \\ &= L\mathbb{1}. \end{aligned}$$

Now, note that all other eigenvectors with other eigenvalues must therefore be orthogonal to the all-ones vector.

Now, let's take our same two examples again.

Example 4.1. Let's first return to the complete graph. For K_n, we have:

$$\begin{aligned} L(K_n) &= \begin{bmatrix} n-1 & -1 & -1 & \cdots & -1 \\ -1 & n-1 & -1 & \cdots & -1 \\ -1 & -1 & n-1 & \cdots & -1 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ -1 & -1 & -1 & \cdots & n-1 \end{bmatrix} \\ &= n \cdot I - \mathbb{1}\mathbb{1}^\top. \end{aligned}$$

Next, for any arbitrary v, let's suppose that $\langle v, \mathbb{1} \rangle = 0$. Then,

$$\begin{aligned} L(K_n)v &= n \cdot Iv - \mathbb{1}\underbrace{\mathbb{1}^\top v}_{=0} \\ &= nv. \end{aligned}$$

Therefore, v must be an eigenvector with eigenvalue of n. Clearly, the dimension of the space of vectors orthogonal to v is $n-1$, so we can create a linearly independent basis of eigenvectors.

Example 4.2. Now, let's return to the star graph with center vertex of 1. Recall the Laplacian:

$$L(S_n) = \begin{bmatrix} n-1 & -1 & -1 & \cdots & -1 \\ -1 & 1 & 0 & \cdots & 0 \\ -1 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ -1 & 0 & 0 & \cdots & 1 \end{bmatrix}$$

Claim 4.3. $\mathbb{1}_u - \mathbb{1}_v$ is an eigenvector with eigenvalue 1 for any $u, v \neq 1$ (Note that this is essentially taking two spokes of the wheel/star and subtracting one out from the other).
Proof: We have that:

$$\begin{aligned} L(S_n)(\mathbb{1}_u - \mathbb{1}_v) &= \sum_{e=(a,b)} w_e(\mathbb{1}_a - \mathbb{1}_b)(\mathbb{1}_a - \mathbb{1}_b)^\top(\mathbb{1}_u - \mathbb{1}_v) \\ &= (\mathbb{1}_1 - \mathbb{1}_u)(\mathbb{1}_1 - \mathbb{1}_u)^\top(\mathbb{1}_u - \mathbb{1}_v) + (\mathbb{1}_1 - \mathbb{1}_v)(\mathbb{1}_1 - \mathbb{1}_v)^\top(\mathbb{1}_u - \mathbb{1}_v) \\ &\quad + \sum_{\substack{e=(1,w) \\ w \neq u,v}} (\mathbb{1}_1 - \mathbb{1}_w)(\mathbb{1}_1 - \mathbb{1}_w)^\top(\mathbb{1}_u - \mathbb{1}_v) \\ &= (\mathbb{1}_1 - \mathbb{1}_u)(-1) + (\mathbb{1}_1 - \mathbb{1}_v)(1) + \sum_{\substack{e=(1,w) \\ w \neq u,v}} 0 \\ &= 1(\mathbb{1}_u - \mathbb{1}_v). \end{aligned}$$

So, now it'd be nice to create a linearly independent and orthogonal basis of eigenvectors. Let's do so:

- We start of course with $\mathbb{1}$.
- Next, for every new vertex, we can simply add $\mathbb{1}_i - \mathbb{1}_{i+1}$. Clearly, since we haven't used vertex $i+1$ before, it'll be linearly independent to all previous vectors. We can do this for $i = 2, \cdots, n-1$ (recall that we cannot consider the center vertex 1)
- Finally, we can find something that is orthogonal to all of the previous vectors. In particular, in order to have it be orthogonal to all previous vectors, we want the 2nd to nth entries all be equal. WLOG make them all 1. Moreover, we want the sum of the entries to be 0 in order for it to be orthogonal to $\mathbb{1}$. Therefore, we find that we must choose: $(-(n-1), 1, 1, \cdots, 1)$.

5 Second Smallest Eigenvalue of Laplacian

Now, we know that $v_1 = \mathbb{1}$, and $\lambda_1 = 0$ is a smallest eigenvector-eigenvalue pair always. Now, what about λ_2, v_2, our eigenvector-eigenvalue pair with second smallest eigenvalue?

5.1 Connected Components and the Smallest Eigenvalues

Claim 5.1. Suppose the connected components of G are $C_1, C_2, \cdots, C_k$. Then, $\left\{\mathbb{1}_{C_j}\right\}_{j=1}^{k}$ forms an (orthogonal) basis for $\mathrm{Ker}(L)$, and therefore, dim(eigenspace with $\lambda = 0$) $= k$.
Proof: First, we show that $\mathbb{1}_{C_j}$ is an eigenvector of L.

$$L\mathbb{1}_{C_j} = \sum_{e} w_e(\mathbb{1}_u - \mathbb{1}_v)(\mathbb{1}_u - \mathbb{1}_v)^\top \mathbb{1}_{C_j}.$$

First, we note that, for all edges not in C_j, the only nonnegative entries in $(\mathbb{1}_u - \mathbb{1}_v)(\mathbb{1}_u - \mathbb{1}_v)^\top$ are in columns u and v, which are not in C_j, which means that $(\mathbb{1}_u - \mathbb{1}_v)(\mathbb{1}_u - \mathbb{1}_v)^\top \mathbb{1}_{C_j} = 0$. Then, we can use the property that $\mathbb{1}$ is an eigenvector with eigenvalue 0 on the subgraph of G consisting of only edges and vertices in C_j to see that $\sum_{e \in C_j} w_e(\mathbb{1}_u - \mathbb{1}_v)(\mathbb{1}_u - \mathbb{1}_v)^\top \mathbb{1}_{C_j} = 0$.

Thus, we have that $(\mathbb{1}_u - \mathbb{1}_v)^\top \mathbb{1}_{C_j} = (\mathbb{1}_{C_j})_u - (\mathbb{1}_{C_j})_v$.

Now, we show that the kernel is no larger than this. Suppose that $x \in \mathrm{Ker}(L)$.

Then,

$$x^\top L x = 0$$
$$\Rightarrow \sum_{e=(u,v)\in E} w_e(x_u - x_v)^2 = 0$$
$$\Rightarrow \forall j = 1, \cdots, k, \ x \text{ is constant on } C_j$$
$$\text{(so, say } x_i = \alpha_j \ \forall i \in C_j)$$
$$\Rightarrow x = \sum_{j=1}^{k} \alpha_j \mathbb{1}_{C_j}.$$

5.2 A More Robust Version

Now, we've seen a necessary and sufficient condition for $\lambda_2 = 0$ (in particular, the graph being disconnected). But what happens if the graph is not disconnected? What properties does a graph has a small (but strictly positive) value of λ_2? For example, what if $\lambda_2 = 0.00001$? Does this mean that G is "almost" disconnected?

Now, intuitively, maybe something that we'd guess is that, say, the min cut over the graph is very small. Unfortunately, this is not exactly true, but it's relatively close. This is somewhat because we can kind of choose a single vertex, and that'll likely be a smaller cut than a more interesting one, like a cut of $n/2$ and $n/2$ vertices, or even $n - \sqrt{n}$ and $\sqrt{n}$ vertices. It turns out that the right idea is called the "conductance".

Definition 5.2. The *conductance* of a nonempty (and nontrivial) cut, $S \subsetneq V_G$, in a graph, $\Phi_G(S)$ is defined to be:

$$\Phi_G(S) = \frac{w(\partial S)}{\min\{\mathrm{vol}(S), \mathrm{vol}(V \setminus S)\}}$$

where $w(E)$ is the sum of the weights of the edges in E, where ∂S is the set of all edges "on the boundary of S," or namely all edges incident to exactly one vertex in S (and exactly one vertex in $V \setminus S$), and $\text{vol}(S) = \sum_{u \in S} \deg(u)$ is the "volume of S."

Then, we define the conductance of a graph, $\Phi(G)$, to be

$$\min_{\substack{S \subset V_G \\ S \neq \varnothing, V_G}} \Phi_G(S).$$

Nota bene: As noted earlier, a lot of this comes directly from differential geometry. In particular, these are analogs of the area of the boundary and the volumes of set.

Theorem 5.3 (Cheeger [1])**.** *We have the following bounds on* $\Phi(G)$*:*

$$\frac{\lambda_2}{2} \leq \Phi(G) \leq \sqrt{2\lambda_2}.$$

Interestingly, we actually prove the right-hand (upper bound) inequality algorithmically. In particular, in the next lecture, we will devise an algorithm using the Laplacian to partition the graph in a strong enough way.

References

[1] Cheeger Jeff. A lower bound for the smallest eigenvalue of the Laplacian. *Problems in Analysis.*, Princeton University Press, pp. 195-199, 1970.