

The Laplacian

Daniel A. Spielman

September 4, 2009

2.1 Eigenvectors and Eigenvalues

I'll begin this lecture by recalling some definitions of eigenvectors and eigenvalues, and some of their basic properties. First, recall that a vector \mathbf{v} is an eigenvector of a matrix M of eigenvalue λ if

$$M\mathbf{v} = \lambda\mathbf{v}.$$

As almost all the matrices we encounter in this class will be symmetric (or morally symmetric), I'll remind you of the special properties of the spectra of symmetric matrices. If \mathbf{v}_1 is an eigenvector of M of eigenvalue λ_1 , \mathbf{v}_2 is an eigenvector of M of eigenvalue $\lambda_2 \neq \lambda_1$, and M symmetric, then \mathbf{v}_1 is orthogonal to \mathbf{v}_2 . To see this, note that

$$\lambda_1 \mathbf{v}_1^T \mathbf{v}_2 = \mathbf{v}_1^T M \mathbf{v}_2 = \mathbf{v}_1^T \lambda_2 \mathbf{v}_2 = \lambda_2 \mathbf{v}_1^T \mathbf{v}_2$$

implies $\mathbf{v}_1^T \mathbf{v}_2 = 0$, assuming $\lambda_1 \neq \lambda_2$. On the other hand, if \mathbf{v}_1 and \mathbf{v}_2 are both eigenvectors of eigenvalue λ , then $\mathbf{v}_1 + \mathbf{v}_2$ is as well.

For a symmetric matrix M , the *multiplicity* of an eigenvalue λ is the dimension of the space of eigenvectors of eigenvalue λ . Also recall that every n -by- n symmetric matrix has n eigenvalues, counted with multiplicity. Thus, it has an orthonormal basis of eigenvectors, $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ with eigenvalues $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ so that

$$M\mathbf{v}_i = \lambda_i \mathbf{v}_i, \quad \text{for all } i.$$

If we let V be the matrix whose i th column is \mathbf{v}_i and Λ be the diagonal matrix whose i th diagonal is λ_i , we can write this more compactly as

$$MV = V\Lambda.$$

Multiplying by V^T on the right, we obtain the eigen-decomposition of M :

$$M = MVV^T = V\Lambda V^T = \sum_i \lambda_i \mathbf{v}_i \mathbf{v}_i^T.$$

2.2 The Laplacian Matrix

We will now recall the definition of the Laplacian matrix of a weighted graph, and present it in a more useful form. Recall that a weighted undirected graph $G = (V, E, w)$ is just an undirected

graph $G = (V, E)$ along with a function $w : E \rightarrow \mathbb{R}^+$, where \mathbb{R}^+ denotes the set of positive Real numbers. The adjacency matrix of a weighted graph G will be denoted A_G , and is given by

$$A_G(i, j) = \begin{cases} w(i, j) & \text{if } (i, j) \in E, \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

The degree matrix of a weighted graph G will be denoted D_G , and is the diagonal matrix such that

$$D_G(i, i) = \sum_j A_G(i, j).$$

The Laplacian matrix of a weighted graph G will be denoted L_G . Last class, we defined it by

$$L_G = D_G - A_G.$$

We will now see a more convenient definition of the Laplacian. To begin, let $G_{1,2}$ be the graph on two vertices with one edge¹ of weight 1. We define

$$L_{G_{1,2}} \stackrel{\text{def}}{=} \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}$$

Note that

$$\mathbf{x}^T L_{G_{1,2}} \mathbf{x} = (\mathbf{x}(1) - \mathbf{x}(2))^2. \quad (2.1)$$

For the graph with n vertices and just one edge between vertices u and v , we can define the Laplacian similarly. For concreteness, I'll call this graph $G_{u,v}$. It's Laplacian matrix is the n -by- n matrix whose only non-zero entries are in the intersections of rows and columns u and v . The two-by-two matrix at the intersections of these rows and columns is, of course,

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

For a weighted graph $G = (V, E, w)$, we now define

$$L_G \stackrel{\text{def}}{=} \sum_{(u,v) \in E} w(u, v) L_{G_{u,v}}.$$

You should verify for yourself that this is equivalent to the definition I gave before.

Many elementary properties of the Laplacian follow from this definition. In particular, it is immediate that for all $\mathbf{x} \in \mathbb{R}^V$

$$\mathbf{x}^T L_G \mathbf{x} = \sum_{(u,v) \in E} w(u, v) (\mathbf{x}(u) - \mathbf{x}(v))^2. \quad (2.2)$$

¹Generally, we will view unweighted graphs as graphs in which all edges have weight 1. If I do not mention the weight of an edge, assume it is 1.

For an eigenvector \mathbf{v} of eigenvalue λ , this tells us that

$$\mathbf{v}^T L_G \mathbf{v} = \lambda \mathbf{v}^T \mathbf{v} \geq 0.$$

So, every eigenvalue of a Laplacian matrix is non-negative. That is, the matrix is *positive semi-definite*.

Remark Since the vertex set really doesn't matter, I actually prefer the notation $L(E)$ where E is a set of edges. Had I used this notation above, it would have eliminated some subscripts. For example, I could have written $L_{\{u,v\}}$ instead of $L_{G_{u,v}}$.

2.3 Connectivity

From (2.2), we see that if all entries of \mathbf{x} are the same, then $\mathbf{x}^T L_G \mathbf{x}$ equals zero. From the definition $L_G = D_G - A_G$, we can immediately see that $L_G \mathbf{x} = \mathbf{0}$, so the constant vectors are eigenvectors of eigenvalue 0.

Lemma 2.3.1. *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 Laplacian matrix. Then, $\lambda_2 > 0$ if and only if G is connected.*

Proof. We first show that $\lambda_2 = 0$ if G is disconnected. If G is disconnected, then it can be described as the union of two graphs, G_1 and G_2 . After suitably re-numbering the vertices, we can write

$$L_G = \begin{bmatrix} L_{G_1} & 0 \\ 0 & L_{G_2} \end{bmatrix}.$$

So, L_G has at least two orthogonal eigenvectors of eigenvalue zero:

$$\begin{bmatrix} \mathbf{0} \\ \mathbf{1} \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} \mathbf{1} \\ \mathbf{0} \end{bmatrix}.$$

where we have partitioned the vectors as we did the matrix L_G .

On the other hand, assume that G is connected and that \mathbf{x} is an eigenvector of L_G of eigenvalue 0. As

$$L_G \mathbf{x} = \mathbf{0},$$

we have

$$\mathbf{x}^T L_G \mathbf{x} = \sum_{(u,v) \in E} (\mathbf{x}(u) - \mathbf{x}(v))^2 = 0.$$

Thus, for each pair of vertices (u, v) connected by an edge, we have $\mathbf{x}(u) = \mathbf{x}(v)$. As every pair of vertices u and v are connected by a path, we may inductively apply this fact to show that $\mathbf{x}(u) = \mathbf{x}(v)$ for all vertices u and v . Thus, \mathbf{x} must be a constant vector. We conclude that the eigenspace of eigenvalue 0 has dimension 1. \square

Of course, the same holds for weighted graphs.

2.4 Some Fundamental Graphs

We now examine the eigenvalues and eigenvectors of the Laplacians of some fundamental graphs. In particular, we will examine

- The complete graph on n vertices, K_n , which has edge set $\{(u, v) : u \neq v\}$.
- The star graph on n vertices, S_n , which has edge set $\{(1, u) : 2 \leq u \leq n\}$.
- The path graph on n vertices, P_n , which has edge set $\{(u, u + 1) : 1 \leq u < n\}$.
- The ring graph on n vertices, R_n , which has all the edges of the path graph, plus the edge $(1, n)$.

Lemma 2.4.1. *The Laplacian of K_n has eigenvalue 0 with multiplicity 1 and n with multiplicity $n - 1$.*

Proof. The multiplicity of the zero eigenvalue follows from Lemma 2.3.1.

To compute the non-zero eigenvalues, let \mathbf{v} be any non-zero vector orthogonal to the all-1s vector, so

$$\sum_i \mathbf{v}(i) = 0. \quad (2.3)$$

Assume, without loss of generality, that $\mathbf{v}(1) \neq 0$. We may now compute the first coordinate of $L_{K_n} \mathbf{v}$, and then divide by $\mathbf{v}(1)$ to compute λ . We find

$$(L_{K_n} \mathbf{v})(1) = (n - 1)\mathbf{v}(1) - \sum_{j=2}^n \mathbf{v}(j) = n\mathbf{v}(1), \quad \text{by (2.3).}$$

So, every vector orthogonal to the all-1s vector is an eigenvector of eigenvalue n . □

To determine the eigenvalues of S_n , we first observe that each vertex $i \geq 2$ has degree 1, and that each of these degree-one vertices has the same neighbor. Whenever two degree-one vertices share the same neighbor, they provide an eigenvector of eigenvalue 1.

Lemma 2.4.2. *Let $G = (V, E)$ be a graph, and let i and j be vertices of degree one that are both connected to another vertex k . Then, the vector \mathbf{v} given by*

$$\mathbf{v}(u) = \begin{cases} 1 & u = i \\ -1 & u = j \\ 0 & \text{otherwise,} \end{cases}$$

is an eigenvector of the Laplacian of G of eigenvalue 1.

Proof. One can immediately verify that $L_G \mathbf{v} = \mathbf{v}$. □

The existence of this eigenvector implies that $\mathbf{v}(i) = \mathbf{v}(j)$ for every eigenvector \mathbf{v} of a different eigenvalue.

Lemma 2.4.3. *The graph S_n has eigenvalue 0 with multiplicity 1, eigenvalue 1 with multiplicity $n - 2$, and eigenvalue n with multiplicity 1.*

Proof. The multiplicity of the eigenvalue 0 follows from Lemma 2.3.1. Applying Lemma 2.4.2 to vertices i and $i + 1$ for $2 \leq i < n$, we find $n - 2$ linearly independent eigenvectors of eigenvalue 1. To determine the last eigenvalue, recall that the trace² of a matrix equals the sum of its eigenvalues. We know that the trace of L_{S_n} is $2n - 2$, and we have identified $n - 1$ eigenvalues that sum to $n - 2$. So, the remaining eigenvalue must be n . Knowing this, and the fact that the corresponding eigenvector must be constant across vertices 2 through n , make it an easy exercise to compute the last eigenvector. \square

Lemma 2.4.4. *The Laplacian of R_n has eigenvectors*

$$\begin{aligned}\mathbf{x}_k(u) &= \sin(2\pi ku/n), \text{ and} \\ \mathbf{y}_k(u) &= \cos(2\pi ku/n),\end{aligned}$$

for $1 \leq k \leq n/2$. When n is even, $\mathbf{x}_{n/2}$ is the all-zero vector, so we only have $\mathbf{y}_{n/2}$. Eigenvectors \mathbf{x}_k and \mathbf{y}_k have eigenvalue $2 - 2\cos(2\pi k/n)$.

Proof. The best way to see that x_k and y_k are eigenvectors is to plot the graph on the circle using these vectors as coordinates. That they are eigenvectors is geometrically obvious. To compute the eigenvalue, just consider vertex 1, and use the double-angle formula to compute:

$$\begin{aligned}(L_{R_n} \mathbf{x}_k)(1) &= 2\mathbf{x}_k(1) - \mathbf{x}_k(0) - \mathbf{x}_k(2) \\ &= 2\sin(2\pi k/n) - \sin(2\pi k \cdot 0/n) - \sin(2\pi k \cdot 2/n) \\ &= 2\sin(2\pi k/n) - 2\sin(2\pi k/n)\cos(2\pi k/n) \\ &= (2 - \cos(2\pi k/n))\mathbf{x}_k(1).\end{aligned}$$

The computation for \cos follows similarly. \square

Lemma 2.4.5. *The Laplacian of P_n has the same eigenvalues as R_{2n} , and eigenvectors*

$$\mathbf{v}_k(u) = \sin(\pi ku/n + \pi/2n).$$

for $0 \leq k < n$

Proof. This is our first interesting example. We derive the eigenvectors and eigenvalues by treating P_n as a quotient of R_{2n} : we will identify vertex u of P_n with both vertices u and $2n + 1 - u$ of R_{2n} . Let \mathbf{z} be an eigenvector of R_{2n} in which $\mathbf{z}(u) = \mathbf{z}(2n + 1 - u)$ for all u . I then claim that the first n components of \mathbf{z} give an eigenvector of P_n .

²the sum of its diagonal entries

To obtain such an eigenvector \mathbf{z} , take

$$\begin{aligned}\mathbf{z}_k(u) &= \sin(2\pi ku/(2n) + \pi/(2n)) \\ &= \sin(2\pi ku/(2n)) \cos(\pi/(2n)) \\ &\quad + \cos(2\pi ku/(2n)) \sin(\pi/(2n)) \\ &= \mathbf{x}_k(u) \cos(\pi/(2n)) + \mathbf{y}_k(u) \sin(\pi/(2n)).\end{aligned}$$

So, \mathbf{z}_k is an eigenvector of R_{2n} of eigenvalue $\lambda_k \stackrel{\text{def}}{=} 2 - \cos(2\pi k/2n)$.

We now set $\mathbf{v}_k(i) = \mathbf{z}_k(i)$ for $1 \leq i \leq n$. To see why \mathbf{v}_k is an eigenvector of L_{P_n} of eigenvalue λ_k , note that for $1 < i < n$,

$$\begin{aligned}(L_{P_n} \mathbf{v}_k)(i) &= 2\mathbf{v}_k(i) - \mathbf{v}_k(i-1) - \mathbf{v}_k(i+1) \\ &= \frac{1}{2} \left(2\mathbf{z}_k(i) - \mathbf{z}_k(i-1) - \mathbf{z}_k(i+1) \right. \\ &\quad \left. + 2\mathbf{z}_k(2n+1-i) - \mathbf{z}_k(2n+1-(i-1)) - \mathbf{z}_k(2n+1-(i+1)) \right) \\ &= \frac{1}{2} (\lambda_k \mathbf{z}_k(i) + \lambda_k \mathbf{z}_k(2n+1-i)) \\ &= \lambda_k \mathbf{v}_k(i).\end{aligned}$$

For $i = 1$, we have

$$\begin{aligned}(L_{P_n} \mathbf{v}_k)(1) &= \mathbf{v}_k(1) - \mathbf{v}_k(2) \\ &= 2\mathbf{v}_k(1) - \mathbf{v}_k(2) - \mathbf{v}_k(1) \\ &= 2\mathbf{z}_k(1) - \mathbf{z}_k(2) - \mathbf{z}_k(2n) \\ &= \lambda_k \mathbf{z}_k(1) \\ &= \lambda_k \mathbf{v}_k(1).\end{aligned}$$

Of course, the other end is similar. □

The quotient construction used in this proof is an example of a generally applicable technique.

We have now seen that the k th eigenvector of the path graph alternates in sign $k-1$ times. This is consistent with our intuition that the Laplacian of the path graph is a discretization of a continuous string, and that its eigenvectors are approximations of its fundamental modes of vibration when its ends are free.

If this intuition is correct, then it should continue to be true even if we discretize a string whose material changes along its length. This corresponds to a weighted path graph.