Spectral Graph Theory Lecture 5

The other eigenvectors of the Laplacian

Daniel A. Spielman September 16, 2009

5.1 Overview

We are now going to begin our study of the other eigenvalues and eigenvectors of the Laplacian. I will begin the lecture by showing how much of the theory we established can be preserved. We will then determine the eigenvalues of the hypercube, and begin to see why λ_2 is so important.

5.2 Remember the path

Recall that we showed that the kth eigenvector of a path graph crosses the origin at most $k-1$ times. For example, here are the first three non-constant eigenvectors of the path graph on 13 vertices, with a line drawn at the origin. Today, we will prove a result of Fiedler [Fie75] which says

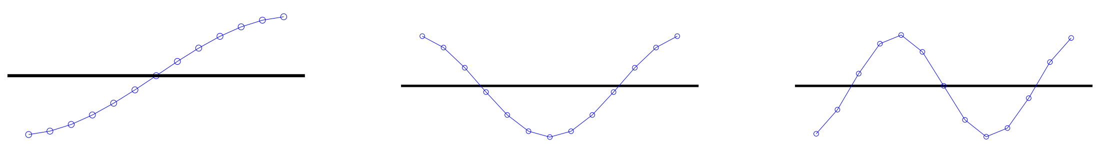

that for every G the graph induced on the vertices that are non-negative in the kth eigenvector has at most $k-1$ connected components. In particular, it says that the non-negative vertices in $\boldsymbol{v}_2$ are connected.

First, we need to recall a little linear algebra.

5.3 The Perron-Frobenius Theorem for Laplacians

In Lecture 3, we proved the Perron-Frobenius Theorem for non-negative matrices. I wish to quickly observe that this theory may also be applied to Laplacian matrices, to principal sub-matrices of Laplacian matrices, and to any matrix with non-positive off-diagonal entries. The difference is that it then involves the eigenvector of the smallest eigenvalue, rather than the largest eigenvalue.

Corollary 5.3.1. *Let M be a matrix with non-positive off-diagonal entries, such that the graph of the non-zero off-diagonally entries is connected. Let λ_1 be the smallest eigenvalue of M and let $\boldsymbol{v}_1$ be the corresponding eigenvector. Then $\boldsymbol{v}_1$ may be taken to be strictly positive, and λ_1 has multiplicity* 1.

Proof. Consider the matrix $A = \sigma I - M$, for some large σ. For σ sufficiently large, this matrix will be non-negative, and the graph of its non-zero entries is connected. So, we may apply the Perron-Frobenius theory to A to conclude that its largest eigenvalue α_1 has multiplicity 1, and the corresponding eigenvector $\boldsymbol{v}_1$ may be assumed to be strictly positive. We then have $\lambda_1 = \sigma - \alpha_1$, and $\boldsymbol{v}_1$ is an eigenvector of λ_1. □

5.4 Eigenvalue Interlacing

We will often use the following elementary consequence of the Courant-Fischer Theorem. I recommend deriving it for yourself.

Theorem 5.4.1 (Eigenvalue Interlacing)**.** *Let A be an n-by-n symmetric matrix and let B be a principal submatrix of A of dimension $n-1$ (that is, B is obtained by deleting the same row and column from A). Then,*

$$\alpha_1 \geq \beta_1 \geq \alpha_2 \geq \beta_2 \geq \cdots \geq \alpha_{n-1} \geq \beta_{n-1} \geq \alpha_n,$$

where $\alpha_1 \geq \alpha_2 \geq \cdots \geq \alpha_n$ and $\beta_1 \geq \beta_2 \geq \cdots \geq \beta_{n-1}$ are the eigenvalues of A and B, respectively.

Corollary 5.4.2 (Eigenvalue Interlacing)**.** *Let A be an n-by-n symmetric matrix and let B be a principal submatrix of A of dimension $n-k$ (that is, B is obtained by deleting the same set of k rows and columns from A). Then,*

$$\alpha_i \geq \beta_i \geq \alpha_{i+k},$$

for all $1 \leq i \leq n-k$, where $\alpha_1 \geq \alpha_2 \geq \cdots \geq \alpha_n$ and $\beta_1 \geq \beta_2 \geq \cdots \geq \beta_{n-k}$ are the eigenvalues of A and B, respectively.

5.5 Fiedler's Nodal Domain Theorem

Given a graph $G = (V, E)$ and a subset of vertices, $W \subseteq V$, recall that the *graph induced by G on W* is the graph with vertex set W and edge set

$$\{(i,j) \in E, i \in W \text{ and } j \in W\}.$$

This graph is sometimes denoted $G(W)$.

Theorem 5.5.1 ([Fie75])**.** *Let $G = (V, E, w)$ be a weighted connected graph, and let L_G be its Laplacian matrix. Let $0 = \lambda_1 < \lambda_2 \leq \cdots \leq \lambda_n$ be the eigenvalues of L_G and let $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_n$ be the corresponding eigenvectors. For any $k \geq 2$, let*

$$W_k = \{i \in V : \boldsymbol{v}_k(i) \geq 0\}.$$

Then, the graph induced by G on W_k has at most $k-1$ connected components.

Proof. To see that W_k is non-empty, recall that $\boldsymbol{v}_1 = \mathbf{1}$ and that $\boldsymbol{v}_k$ is orthogonal $\boldsymbol{v}_1$. So, $\boldsymbol{v}_k$ must have both positive and negative entries.

Assume that $G(W_k)$ has t connected components. After re-ordering the vertices so that the vertices in one connected component of $G(W_k)$ appear first, and so on, we may assume that L_G and $\boldsymbol{v}_k$ have the forms

$$L_G = \begin{bmatrix} B_1 & \mathbf{0} & \mathbf{0} & \cdots & C_1 \\ \mathbf{0} & B_2 & \mathbf{0} & \cdots & C_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & B_t & C_t \\ C_1^T & C_2^T & \cdots & C_t^T & D \end{bmatrix} \qquad \boldsymbol{v}_k = \begin{pmatrix} \boldsymbol{x}_1 \\ \boldsymbol{x}_2 \\ \vdots \\ \boldsymbol{x}_t \\ \boldsymbol{y} \end{pmatrix},$$

and

$$\begin{bmatrix} B_1 & \mathbf{0} & \mathbf{0} & \cdots & C_1 \\ \mathbf{0} & B_2 & \mathbf{0} & \cdots & C_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & B_t & C_t \\ C_1^T & C_2^T & \cdots & C_t^T & D \end{bmatrix} \begin{pmatrix} \boldsymbol{x}_1 \\ \boldsymbol{x}_2 \\ \vdots \\ \boldsymbol{x}_t \\ \boldsymbol{y} \end{pmatrix} = \lambda_k \begin{pmatrix} \boldsymbol{x}_1 \\ \boldsymbol{x}_2 \\ \vdots \\ \boldsymbol{x}_t \\ \boldsymbol{y} \end{pmatrix}.$$

The first t sets of rows and columns correspond to the t connected components. So, $\boldsymbol{x}_i \geq 0$ for $1 \leq i \leq t$ and $\boldsymbol{y} < 0$ (when I write this for a vector, I mean it holds for each entry). We also know that the graph of non-zero entries in each B_i is connected, and that each C_i is non-positive, and has at least one non-zero entry (otherwise the graph G would be disconnected).

We will now prove that the smallest eigenvalue of B_i is smaller than λ_k. We know that

$$B_i \boldsymbol{x}_i + C_i \boldsymbol{y} = \lambda_k \boldsymbol{x}_i.$$

As each entry in C_i is non-positive and $\boldsymbol{y}$ is strictly negative, each entry of $C_i\boldsymbol{y}$ is non-negative and some entry of $C_i\boldsymbol{y}$ is positive. Thus, $\boldsymbol{x}_i$ cannot be all zeros,

$$B_i \boldsymbol{x}_i = \lambda_k \boldsymbol{x}_i - C_i \boldsymbol{y} \leq \lambda_k \boldsymbol{x}_i$$

and

$$\boldsymbol{x}_i^T B_i \boldsymbol{x}_i \leq \lambda_k \boldsymbol{x}_i^T \boldsymbol{x}_i.$$

If $\boldsymbol{x}_i$ has *any* zero entries, then the Perron-Frobenius theorem tells us that $\boldsymbol{x}_i$ cannot be an eigenvector of smallest eigenvalue, and so the smallest eigenvalue of B_i is less than λ_k. On the other hand, if $\boldsymbol{x}_i$ is strictly positive, then $\boldsymbol{x}_i^T C_i \boldsymbol{y} > 0$, and

$$\boldsymbol{x}_i^T B_i \boldsymbol{x}_i = \lambda_k \boldsymbol{x}_i^T \boldsymbol{x}_i - \boldsymbol{x}_i^T C_i \boldsymbol{y} < \lambda_k \boldsymbol{x}_i^T \boldsymbol{x}_i.$$

Thus, the matrix

$$\begin{bmatrix} B_1 & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & B_2 & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & B_t \end{bmatrix}$$

has at least t eigenvalues less than λ_k. By the eigenvalue interlacing theorem, this implies that L_G has at least t eigenvalues less than λ_k. We may conclude that t, the number of connected components of $G(W_k)$, is at most $k - 1$. □

We remark that Fiedler actually proved a somewhat stronger theorem. He showed that the same holds for

$$W = \{i : \boldsymbol{v}_k(i) \geq t\},$$

for every $t \leq 0$.

This theorem breaks down if we instead consider the set

$$W = \{i : \boldsymbol{v}_k(i) > 0\}.$$

The star graphs provide counter-examples.

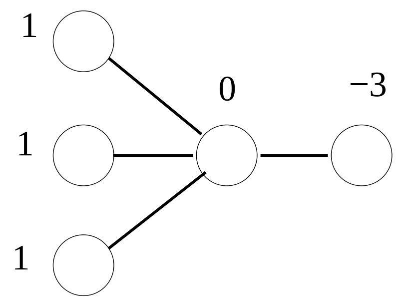

Figure 5.1: The star graph on 5 vertices, with an eigenvector of $\lambda_2 = 1$.

5.6 The second Laplacian eigenvalue

The most important eigenvalue will be λ_2. It is the answer to many questions about graphs, and will entertain us for a few weeks of this course. Let me begin to tell you why.

Recall that $\lambda_2 = 0$ if and only if a graph is disconnected. Fiedler [Fie73] observed that as λ_2 becomes further from 0, a graph becomes better connected. We will see many versions of this statement, capturing varying ways of measuring connectivity. For the first, we consider the boundary of a set of vertices,

$$\delta(S) = \{(i,j) \in E : i \in S \text{ and } j \notin S\}.$$

Theorem 5.6.1. *Let $G = (V, E)$ be a graph and let L_G be its Laplacian matrix. Let $S \subset V$ and set $\sigma = |S| / |V|$. Then,*

$$|\delta(S)| \geq \lambda_2 |S| (1 - \sigma).$$

Proof. Recall that

$$\lambda_2 = \min_{\boldsymbol{v}:\boldsymbol{v}^T\mathbf{1}=0} \frac{\boldsymbol{v}^T L_G \boldsymbol{v}}{\boldsymbol{v}^T \boldsymbol{v}},$$

for for every non-zero $\boldsymbol{v}$ orthogonal to $\mathbf{1}$,

$$\boldsymbol{v}^T L_G \boldsymbol{v} \geq \lambda_2 \boldsymbol{v}^T \boldsymbol{v}.$$

To apply this bound, we need to construct a vector related to the set S. A natural choice is χ_S. In fact, we have

$$\chi_S^T L_G \chi_S = \sum_{(i,j)\in E} (\chi_S(i) - \chi_S(j))^2 = |\delta(S)| .$$

However, χ_S is not orthogonal to $\mathbf{1}$. To fix this, use

$$\boldsymbol{v} = \chi_S - \sigma \mathbf{1}.$$

We have $\boldsymbol{v}^T \mathbf{1} = 0$, and as $L_G \mathbf{1} = \mathbf{0}$

$$\boldsymbol{v}^T L_G \boldsymbol{v} = \chi_S^T L_G \chi_S = |\delta(S)| .$$

To finish the proof, we compute

$$\boldsymbol{v}^T \boldsymbol{v} = |S| (1-\sigma)^2 + (|V| - |S|)\sigma^2 = |S| (1 - 2\sigma + \sigma^2) + |S| \sigma - |S| \sigma^2 = |S| (1-\sigma).$$

□

This theorem says that if λ_2 is big, then G is very well connected: the boundary of every small set of vertices is at least λ_2 times something just slightly smaller than the number of vertices in the set.

This lemma also provides an easy technique for proving upper bounds on λ_2. Next lecture, we will see techniques for proving lower bonds on λ_2.

For now, I'd like to show you that there are interesting graphs with large λ_2.

5.7 The Hypercube

The hypercube graph is the graph with vertex set $\{0,1\}^d$, with edges between vertices whose names differ in exactly one bit. The hypercube may also be expressed as the product of the one-edge graph with itself $d-1$ times, with the proper definition of graph product.

Definition 5.7.1. *Let $G = (V, E)$ and $H = (W, F)$ be graphs. Then $G \times H$ is the graph with vertex set $V \times W$ and edge set*

$$\Big((v_1, w), (v_2, w)\Big) \text{ where } (v_1, v_2) \in E \text{ and}$$
$$\Big((v, w_1), (v, w_2)\Big) \text{ where } (w_1, w_2) \in F.$$

Let $G = (\{0,1\}, \{(0,1)\})$, and let H_d be the d-dimensional hypercube graph. You should check that $H_1 = G$ and that $H_d = H_{d-1} \times G$.

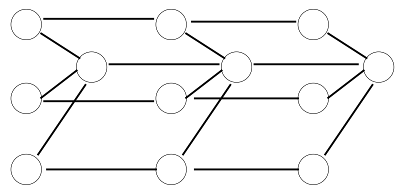

Figure 5.2: The product of a star graph on 4 vertices with a path on 3.

Theorem 5.7.2. *Let $G = (V, E)$ and $H = (W, F)$ be graphs with Laplacian eigenvalues $\lambda_1, \ldots, \lambda_n$ and $\mu_1, \ldots, \mu_m$, and eigenvectors $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n$ and $\boldsymbol{y}_1, \ldots, \boldsymbol{y}_m$, respectively. Then, for each $1 \leq i \leq n$ and $1 \leq j \leq m$, $G \times H$ has an eigenvector $\boldsymbol{z}$ of eigenvalue $\lambda_i + \mu_j$ such that*

$$\boldsymbol{z}(v, w) = \boldsymbol{x}_i(v)\boldsymbol{y}_j(w).$$

Proof. To see that this eigenvector has the proper eigenvalue, let L denote the Laplacian of $G \times H$, d_v the degree of node v in G, and e_w the degree of node w in H. We can then verify that

$$\begin{aligned}
(L\boldsymbol{z})(v, w) &= (d_v + e_w)\boldsymbol{x}_i(v)\boldsymbol{y}_j(w) - \sum_{(v,v_2)\in E} \boldsymbol{x}_i(v_2)\boldsymbol{y}_j(w) - \sum_{(w,w_2)\in F} \boldsymbol{x}_i(v)\boldsymbol{y}_j(w_2) \\
&= \left[(d_v)\boldsymbol{x}_i(v)\boldsymbol{y}_j(w) - \sum_{(v,v_2)\in E} \boldsymbol{x}_i(v_2)\boldsymbol{y}_j(w)\right] + \left[(e_w)\boldsymbol{x}_i(v)\boldsymbol{y}_j(w) - \sum_{(w,w_2)\in F} \boldsymbol{x}_i(v)\boldsymbol{y}_j(w_2)\right] \\
&= \boldsymbol{y}_j(w)\left(d_v\boldsymbol{x}_i(v) - \sum_{(v,v_2)\in E} \boldsymbol{x}_i(v_2)\right) + \boldsymbol{x}_i(v)\left(e_w\boldsymbol{y}_j(w) - \sum_{(w,w_2)\in F} \boldsymbol{y}_j(w_2)\right) \\
&= \boldsymbol{y}_j(w)\lambda_i\boldsymbol{x}_i(v) + \boldsymbol{x}_i(v)\mu_j\boldsymbol{y}_j(w) \\
&= (\lambda_i + \mu_j)(\boldsymbol{x}_i(v)\boldsymbol{y}_j(w)).
\end{aligned}$$

□

As the non-zero eigenvector of G is $(1, -1)$ and has eigenvalue 2, we see that H_d has eigenvalue $2k$ with multiplicity $\binom{d}{k}$, for $0 \leq k \leq d$. Using the above theorem, you should also confirm that the eigenvectors of H_d are given by the functions

$$\boldsymbol{v}_a(b) = (-1)^{a^T b},$$

where $a \in \{0, 1\}^d$, and we view vertices b as length-d vectors of zeros and ones. The eigenvalue of which $\boldsymbol{v}_a$ is an eigenvector is the number of ones in a.

Using Theorem 5.6.1 and our knowledge of the eigenvalues of the hypercube, we can immediately prove the following isoperimetric theorem for the hypercube.

Corollary 5.7.3. *Let S be a subset of $\{0, 1\}^d$ of size at most 2^{d-1}. Then,*

$$|\delta(S)| \geq |S|.$$

It is possible to prove this by more concrete combinatorial means. But, this proof is simpler.

References

[Fie73] M. Fiedler. Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal*, 23(98):298–305, 1973.

[Fie75] M. Fiedler. A property of eigenvectors of nonnegative symmetric matrices and its applications to graph theory. *Czechoslovak Mathematical Journal*, 25(100):618–633, 1975.