Spectral Graph Theory Lecture 7

Fiedler's Theorems on Nodal Domains

Daniel A. Spielman September 19, 2012

7.1 About these notes

These notes are not necessarily an accurate representation of what happened in class. The notes written before class say what I think I should say. The notes written after class way what I wish I said.

I think that these notes are mostly correct.

7.2 Overview

Nodal domains are the connected parts of a graph on which an eigenvector is negative or positive. In this lecture, we will cover some of the fundamental theorems of Fiedler on nodal domains of the Laplacian. The first theorem that I will present says that the kth eigenvector of a weighted path graph changes sign $k-1$ times. So, the alternation that we observed when we derived the eigenvectors of the path holds even if the path is weighted. In fact, the theorem is stronger. A similar fact holds for trees. Next lecture, I hope to use this theorem to show how eigenvectors can be used to reconstruct tree metrics.

The second theorem in this lecture will be one of our best extensions of this fact to general graphs.

7.3 Sylverter's Law of Interia

When I introduced the normalized Laplacian last lecture, I assumed but did not prove that it is positive semidefinite. Since no one complained, I assumed that it was obvious. But, just to be safe I will tell you why.

Claim 7.3.1. *If A is positive semidefinite, then so is $B^T AB$ for any matrix B.*

Proof. For any x,

$$x^T B^T ABx = (Bx)^T A(Bx) \geq 0,$$

since A is positive semidefinite. □

In this lecture, we will make use of Sylvester's law of intertia, which is a powerful generalization of this fact. I will state and prove it now.

Theorem 7.3.2 (Sylvester's Law of Intertia)**.** *Let A be any symmetric matrix and let B be any non-singular matrix. Then, the matrix BAB^T has the same number of positive, negative and zero eigenvalues as A.*

Proof. It is clear that A and BAB^T have the same rank, and thus the same number of zero eigenvalues.

We will prove that A has at least as many positive eigenvalues as BAB^T. One can similarly prove that that A has at least as many negative eigenvalues, which proves the theorem.

Let $\gamma_1, \ldots, \gamma_k$ be the positive eigenvalues of BAB^T and let Y_k be the span of the corresponding eigenvectors. Now, let S_k be the span of the vectors $B^T \boldsymbol{y}$, for $\boldsymbol{y} \in Y_k$. As B is non-singluar, S_k has dimension k. By the Courant-Fischer Theorem, we have

$$\alpha_k = \max_{\substack{S \subseteq \mathbb{R}^n \\ \dim(S)=k}} \min_{x \in S} \frac{\boldsymbol{x}^T A \boldsymbol{x}}{\boldsymbol{x}^T \boldsymbol{x}} \geq \min_{x \in S_k} \frac{\boldsymbol{x}^T A \boldsymbol{x}}{\boldsymbol{x}^T \boldsymbol{x}} = \min_{y \in Y_k} \frac{\boldsymbol{y}^T BAB^T \boldsymbol{y}}{\boldsymbol{y}^T BB^T \boldsymbol{y}} > 0.$$

So, A has at least k positive eigenvalues (The point here is that the denominators are always positive, so we only need to think about the numerators.) $\square$

7.4 Weighted Trees

We will now prove a theorem of Fiedler [Fie75].

Theorem 7.4.1. *Let T be a weighted tree graph on n vertices, let $\boldsymbol{L}_T$ have eigenvalues $0 = \lambda_1 < \lambda_2 \cdots \leq \lambda_n$, and let $\boldsymbol{\psi}_k$ be an eigenvector of λ_k. If there is no vertex u for which $\boldsymbol{\psi}_k(u) = 0$, then there are exactly $k-1$ edges for which $\boldsymbol{\psi}_k(u)\boldsymbol{\psi}_k(v) < 0$.*

In the case of a path graph, this means that the eigenvector changes sign $k-1$ times along the path. We will consider eigenvectors with zero entries in the next problem set.

Our analysis will rest on an understanding of Laplacians of trees that are allowed to have negative edges weights.

Lemma 7.4.2. *Let $T = (V, E)$ be a tree, and let*

$$\boldsymbol{M} = \sum_{(u,v) \in E} w_{u,v} \boldsymbol{L}_{u,v},$$

where the weights $w_{u,v}$ are non-zero and we recall that $\boldsymbol{L}_{u,v}$ is the Laplacian of the edge (u, v). The number of negative eigenvalues of $\boldsymbol{M}$ equals the number of negative edge weights.

Proof. Note that

$$\boldsymbol{x}^T \boldsymbol{M} \boldsymbol{x} = \sum_{(u,v) \in E} w_{u,v} (\boldsymbol{x}(u) - \boldsymbol{x}(v))^2.$$

We now perform a change of variables that will diagonalize the matrix $\boldsymbol{M}$. Begin by re-numbering the vertices so that for every vertex v there is a path from vertex 1 to vertex v in which the numbers of vertices encountered are increasing. Under such and ordering, for every vertex v there will be exactly one edge (u, v) with $u < v$. Let $\boldsymbol{\delta}(1) = \boldsymbol{x}(1)$, and for every node v let $\boldsymbol{\delta}(v) = \boldsymbol{x}(v) - \boldsymbol{x}(u)$ where (u, v) is the edge with $u < v$.

Every variable $\boldsymbol{x}(1), \dots, \boldsymbol{x}(n)$ can be expressed as a linear combination of the variables $\boldsymbol{\delta}(1), \dots, \boldsymbol{\delta}(n)$. To see this, let v be any vertex and let $1, u_1, \dots, u_k, v$ be the vertices on the path from 1 to v, in order. Then,

$$\boldsymbol{x}(v) = \boldsymbol{\delta}(1) + \boldsymbol{\delta}(u_1) + \cdots + \boldsymbol{\delta}(u_k) + \boldsymbol{\delta}(v).$$

So, there is a square matrix $\boldsymbol{L}$ of full rank such that

$$\boldsymbol{x} = \boldsymbol{L}\boldsymbol{\delta}.$$

By Sylvester's law of intertia, we know that

$$\boldsymbol{L}^T \boldsymbol{M} \boldsymbol{L}$$

has the same number of positive, negative, and zero eigenvalues as $\boldsymbol{M}$. On the other hand,

$$\boldsymbol{\delta}^T \boldsymbol{L}^T \boldsymbol{M} \boldsymbol{L} \boldsymbol{\delta} = \sum_{(u,v) \in E, u < v} w_{u,v} (\boldsymbol{\delta}(v))^2.$$

So, this matrix clearly has one zero eigenvalue, and as many negative eigenvalues as there are negative $w_{u,v}$. $\square$

Proof of Theorem 7.4.1. Let $\boldsymbol{\Psi}_k$ denote the diagonal matrix with $\boldsymbol{\psi}_k$ on the diagonal, and let λ_k be the corresponding eigenvalue. Consider the matrix

$$\boldsymbol{M} = \boldsymbol{\Psi}_k (\boldsymbol{L}_P - \lambda_k \boldsymbol{I}) \boldsymbol{\Psi}_k.$$

The matrix $\boldsymbol{L}_P - \lambda_k \boldsymbol{I}$ has one zero eigenvalue and $k - 1$ negative eigenvalues. As we have assumed that $\boldsymbol{\psi}_k$ has no zero entries, $\boldsymbol{\Psi}_k$ is non-singular, and so we may apply Sylvester's Law of Intertia to show that the same is true of $\boldsymbol{M}$.

I claim that

$$\boldsymbol{M} = \sum_{(u,v) \in E} w_{u,v} \boldsymbol{\psi}_k(u) \boldsymbol{\psi}_k(v) \boldsymbol{L}_{u,v}.$$

To see this, first check that this agrees with the previous definition on the off-diagonal entries. To verify that these expression agree on the diagonal entries, we will show that the sum of the entries in each row of both expressions agree. As we know that all the off-diagonal entries agree, this implies that the diagonal entries agree. We compute

$$\boldsymbol{\Psi}_k (\boldsymbol{L}_P - \lambda_k \boldsymbol{I}) \boldsymbol{\Psi}_k \boldsymbol{1} = \boldsymbol{\Psi}_k (\boldsymbol{L}_P - \lambda_k \boldsymbol{I}) \boldsymbol{\psi}_k = \boldsymbol{\Psi}_k (\lambda_k \boldsymbol{\psi}_k - \lambda_k \boldsymbol{\psi}_k) = \boldsymbol{0}.$$

As $\boldsymbol{L}_{u,v} \boldsymbol{1} = \boldsymbol{0}$, the row sums agree. Lemma 7.4.2 now tells us that the matrix $\boldsymbol{M}$ has as many negative eigenvalues as there are edges (u, v) for which $\boldsymbol{\psi}_k(u) \boldsymbol{\psi}_k(v) < 0$. $\square$

7.5 More linear algebra

There are a few more facts from linear algebra that we will need for the rest of this lecture. We stop to prove them now.

7.5.1 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 7.5.1. *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. $\square$

7.5.2 Eigenvalue Interlacing

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

Theorem 7.5.2 (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.

7.6 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 7.6.1 ([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.

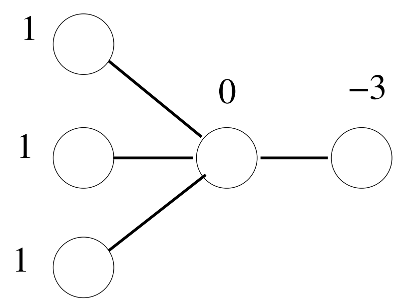

Figure 7.1: The star graph on 5 vertices, with an eigenvector of $\lambda_2 = 1$.

References

[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.