More Spectral Graph Theory

Rasmus Kyng

Lecture 5 — Wednesday, March 18th

We continue our introduction to spectral graph theory that we started in Lecture 4. We're getting used to recalling the Courant-Fischer Theorem, and in these tough times, it's nice to see something familiar, so let's do that again. We'll include the eigenbasis version of the theorem, which makes explicit what subspace obtains the extreme value.

Theorem 0.1 (The Courant-Fischer Theorem)**.** *Let $\boldsymbol{A}$ be a symmetric matrix in $\mathbb{R}^{n\times n}$, with eigenvalues $\lambda_1 \leq \lambda_2 \leq \ldots \leq \lambda_n$. Then*

1.
$$\lambda_i = \min_{\substack{\text{subspace } W \subseteq \mathbb{R}^n \\ \dim(W)=i}} \max_{\boldsymbol{x}\in W, \boldsymbol{x}\neq \boldsymbol{0}} \frac{\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x}}{\boldsymbol{x}^\top \boldsymbol{x}} = \max_{\substack{\boldsymbol{x}\perp \boldsymbol{x}_{i+1},\ldots \boldsymbol{x}_n \\ \boldsymbol{x}\neq \boldsymbol{0}}} \frac{\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x}}{\boldsymbol{x}^\top \boldsymbol{x}}$$

2.
$$\lambda_i = \max_{\substack{\text{subspace } W \subseteq \mathbb{R}^n \\ \dim(W)=n+1-i}} \min_{\boldsymbol{x}\in W, \boldsymbol{x}\neq \boldsymbol{0}} \frac{\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x}}{\boldsymbol{x}^\top \boldsymbol{x}} = \min_{\substack{\boldsymbol{x}\perp \boldsymbol{x}_{1},\ldots \boldsymbol{x}_{i-1} \\ \boldsymbol{x}\neq \boldsymbol{0}}} \frac{\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x}}{\boldsymbol{x}^\top \boldsymbol{x}}$$

1 Understanding Eigenvalues of the Laplacian, Continued

We'll continue our efforts to understand the eigevalues of the Laplacian of weighted undirected graphs.

In Lecture 4, we first saw a complete characterization of the eigenvalues and eigenvectors of the unit weight complete graph on n vertices, K_n. Namely, $\boldsymbol{L}_{K_N} = n\boldsymbol{I} - \boldsymbol{1}\boldsymbol{1}^\top$, and this means that *every* vector $\boldsymbol{y} \perp \boldsymbol{1}$ is an eigenvector of eigenvalue 1.

We then looked at eigenvalues of P_n, the unit weight path on n vertices, and we showed using *test vector* bounds that

$$\lambda_2(\boldsymbol{L}_{P_n}) \leq \frac{12}{n^2} \text{ and } 1 \leq \lambda_n(\boldsymbol{L}_{P_n}). \tag{1}$$

Ideally we would like to prove an almost matching upper bound on λ_2 and an almost matching lower bound on λ_n, but it is not clear how to get that from the Courant-Fischer theorem.

To get there, we start we need to introduce some more tools.

1.1 The Loewner Order, aka. the Positive Semi-Definite Order

We'll now introduce an ordering on symmetric matrices called the *Loewner order*, which I also like to just call the positive semi-definite order. As we will see in a moment, it is a partial order on symmetric matrices, we denote it by "$\preceq$". For conveniece, we allow ourselves to both write $\boldsymbol{A} \preceq \boldsymbol{B}$ and equivalently $\boldsymbol{B} \succeq \boldsymbol{A}$.

For a symmetric matrix $\boldsymbol{A} \in \mathbb{R}^{n \times n}$ we define that

$$\boldsymbol{A} \succeq \boldsymbol{0}$$

if and only if $\boldsymbol{A}$ is positive semi-definite.

More generally, when we have two symmetric matrices $\boldsymbol{A}, \boldsymbol{B} \in \mathbb{R}^{n \times n}$, we will write

$$\boldsymbol{A} \preceq \boldsymbol{B} \text{ if and only if for all } \boldsymbol{x} \in \mathbb{R}^n \text{ we have } \boldsymbol{x}^\top \boldsymbol{A} \boldsymbol{x} \leq \boldsymbol{x}^\top \boldsymbol{B} \boldsymbol{x} \tag{2}$$

This is a partial order, because it satisfies the three requirements of

1. Reflexivity: $\boldsymbol{A} \preceq \boldsymbol{A}$.
2. Anti-symmetry: $\boldsymbol{A} \preceq \boldsymbol{B}$ and $\boldsymbol{B} \preceq \boldsymbol{A}$ implies $\boldsymbol{A} = \boldsymbol{B}$
3. Transitivity: $\boldsymbol{A} \preceq \boldsymbol{B}$ and $\boldsymbol{B} \preceq \boldsymbol{C}$ implies $\boldsymbol{A} \preceq \boldsymbol{C}$

Check for yourself that these properties hold!

The PSD order has other very useful properties: $\boldsymbol{A} \preceq \boldsymbol{B}$ implies $\boldsymbol{A} + \boldsymbol{C} \preceq \boldsymbol{B} + \boldsymbol{C}$ for any symmetric matrix $\boldsymbol{C}$. Convince yourself of this too!

And, combining this observation with transitivity, we can see that $\boldsymbol{A} \preceq \boldsymbol{B}$ and $\boldsymbol{C} \preceq \boldsymbol{D}$ implies $\boldsymbol{A} + \boldsymbol{C} \preceq \boldsymbol{B} + \boldsymbol{D}$.

Here is another useful property: If $\boldsymbol{0} \preceq \boldsymbol{A}$ then for all $\alpha \geq 1$

$$\frac{1}{\alpha} \boldsymbol{A} \preceq \boldsymbol{A} \preceq \alpha \boldsymbol{A}.$$

Here is another one:

Claim 1.1. *If $\boldsymbol{A} \preceq \boldsymbol{B}$, then for all i*

$$\lambda_i(\boldsymbol{A}) \leq \lambda_i(\boldsymbol{B}).$$

Proof. We can prove this Claim by applying the subspace version of the Courant-Fischer theorem.

$$\lambda_i(\boldsymbol{A}) = \min_{\substack{\text{subspace } W \subseteq \mathbb{R}^n \\ \dim(W) = i}} \max_{\boldsymbol{x} \in W, \boldsymbol{x} \neq \boldsymbol{0}} \frac{\boldsymbol{x}^\top \boldsymbol{A} \boldsymbol{x}}{\boldsymbol{x}^\top \boldsymbol{x}} \leq \min_{\substack{\text{subspace } W \subseteq \mathbb{R}^n \\ \dim(W) = i}} \max_{\boldsymbol{x} \in W, \boldsymbol{x} \neq \boldsymbol{0}} \frac{\boldsymbol{x}^\top \boldsymbol{B} \boldsymbol{x}}{\boldsymbol{x}^\top \boldsymbol{x}} = \lambda_i(\boldsymbol{B}).$$

$\square$

Note that the converse of Clam 1.1 is very much false, for example the matrices $\boldsymbol{A} = \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix}$ and $\boldsymbol{B} = \begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix}$ have equal eigenvalues, but both $\boldsymbol{A} \not\preceq \boldsymbol{B}$ and $\boldsymbol{B} \not\preceq \boldsymbol{A}$.

Remark 1.2. It's useful to get used to and remember some of the properties of the Loewner order, but all the things we have established so far are almost immediate from the basic characterization in Equation (2). So, ideally, don't memorize all these facts, instead, try to see that they are simple consequences of the definition.

1.2 Upper Bounding a Laplacian's λ_n Using Degrees

In an earlier lecture, we observed that for any graph $G = (V, E, \boldsymbol{w})$, $\boldsymbol{L} = \boldsymbol{D} - \boldsymbol{A} \succeq \boldsymbol{0}$. We can see this from $\boldsymbol{x}^\top(\boldsymbol{D} - \boldsymbol{A})\boldsymbol{x} = \sum_{(u,v)\in E} \boldsymbol{w}(u,v)(\boldsymbol{x}(u) - \boldsymbol{x}(v))^2 \geq 0$. Similarly $\boldsymbol{D} + \boldsymbol{A} \succeq \boldsymbol{0}$. because $\boldsymbol{x}^\top(\boldsymbol{D} + \boldsymbol{A})\boldsymbol{x} = \sum_{(u,v)\in E} \boldsymbol{w}(u,v)(\boldsymbol{x}(u) + \boldsymbol{x}(v))^2 \geq 0$. But this means that $-\boldsymbol{A} \preceq \boldsymbol{D}$ and hence $\boldsymbol{L} = \boldsymbol{D} - \boldsymbol{A} \preceq 2\boldsymbol{D}$.

So, for the path graph P_n, we have $\boldsymbol{L}_{P_n} \preceq \boldsymbol{D} - \boldsymbol{A} \preceq 2\boldsymbol{D} \preceq 4\boldsymbol{I}$. So by Claim 1.1

$$\lambda_n(\boldsymbol{L}_{P_n}) \leq 4. \tag{3}$$

We can see that our test vector-based lower bound on $\lambda_n(\boldsymbol{L}_{P_n})$ from Equation (1) is tight up to a factor 4.

Since this type of argument works for any unit weight graph, it proves the following claim.

Claim 1.3. *For any unit weight graph G, $\lambda_n(\boldsymbol{L}_G) \leq 2\max_{v\in V} degree(v)$.*

This is tight on a graph consisting of a single edge.

1.3 The Loewner Order and Laplacians of Graphs.

It's sometimes convenient to overload the for the PSD order to also apply to graphs. We will write

$$G \preceq H$$

if $\boldsymbol{L}_G \preceq \boldsymbol{L}_H$.

For example, given two unit weight graphs $G = (V, E)$ and $H = (V, F)$, if $H = (V, F)$ is a subgraph of G, then

$$\boldsymbol{L}_H \preceq \boldsymbol{L}_G.$$

We can see this from the Laplacian quadratic form:

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

Dropping edges will only decrease the value of the quadratic form. The same is for decreasing the weights of edges. The graph order notation is especially useful when we allow for scaling a graph by a constant, say $c > 0$,

$$c \cdot H \preceq G$$

What is $c \cdot H$? It is the same graph as H, but the weight of every edge is multiplied by c. Now we can make statements like $\frac{1}{2}H \preceq G \preceq 2H$, which turn out to be useful notion of the two graphs approximating each other.

1.4 The Path Inequality

Now, we'll see a general tool for comparing two graphs G and H to prove an inequalities like $cH \preceq G$ for some constant c. Our tools won't necessarily work well for all cases, but we'll see some examples where they do.

In the rest of the lecture, we will often need to compare two graphs define on the same vertex set $V = \{1, \ldots, n\} = [n]$.

We use $G_{i,j}$ to denote the unit weight graph on vertex set $[n]$ consisting of a single edge between vertices i and j.

Lemma 1.4 (The Path Inequality)**.**

$$(n-1) \cdot P_n \succeq G_{1,n},$$

Proof. We want to show that for every $\boldsymbol{x} \in\in \mathbb{R}^n$,

$$(n-1) \cdot \sum_{i=1}^{n-1} (\boldsymbol{x}(i+1) - \boldsymbol{x}(i))^2 \geq (\boldsymbol{x}(n) - \boldsymbol{x}(1))^2.$$

For $i \in [n-1]$, set

$$\boldsymbol{\Delta}(i) = \boldsymbol{x}(i+1) - \boldsymbol{x}(i).$$

The inequality we want to prove then becomes

$$(n-1) \sum_{i=1}^{n-1} (\boldsymbol{\Delta}(i))^2 \geq \left(\sum_{i=1}^{n-1} \boldsymbol{\Delta}(i) \right)^2.$$

But, this is immediate from the Caucy-Schwarz inequality $\boldsymbol{a}^\top \boldsymbol{b} \leq \|\boldsymbol{a}\|_2 \|\boldsymbol{b}\|_2$:

$$\begin{aligned}
(n-1) \sum_{i=1}^{n-1} (\boldsymbol{\Delta}(i))^2 &= \|\mathbf{1}_{n-1}\|^2 \cdot \|\boldsymbol{\Delta}\|^2 \\
&= (\|\mathbf{1}_{n-1}\| \cdot \|\boldsymbol{\Delta}\|^2)^2 \\
&\geq (\mathbf{1}_{n-1}^\top \boldsymbol{\Delta})^2 \\
&= (\sum_{i=1}^{n-1} \boldsymbol{\Delta}(i))^2
\end{aligned}$$

□

1.5 Bounding λ_2 of a Path Graph

We will now use Lemma 1.4 to prove a lower bound on $\lambda_2(\boldsymbol{L}_{P_n})$. Our strategy will be to prove that the path P_n is at least some multiple of the complete graph K_n, measured by the Loewner order, i.e. $K_n \preceq f(n) P_n$ for some function $f : \mathbb{N} \to \mathbb{R}$. We can combine this with our observation from the previous lecture that $\lambda_2(\boldsymbol{L}_{K_n}) = n$ to show that

$$f(n) \lambda_2(\boldsymbol{L}_{P_n}) \geq \lambda_2(\boldsymbol{L}_{K_n}) = n, \tag{4}$$

and this will give our lower bound on $\lambda_2(\boldsymbol{L}_{P_n})$. When establishing the inequality between P_n and K_n, we can treat each edge of the complete graph separately, by first noting that

$$\boldsymbol{L}_{K_n} = \sum_{i<j} \boldsymbol{L}_{G_{i,j}}$$

For every edge (i, j) in the complete graph, we apply the Path Inequality, Lemma 1.4:

$$\begin{aligned} G_{i,j} &\preceq (j-i) \sum_{k=i}^{j-1} G_{k,k+1} \\ &\preceq (j-i) P_n \end{aligned}$$

This inequality says that $G_{i,j}$ is at most $(j-i)$ times the part of the path connecting i to j, and that this part of the path is less than the whole.

Summing inequality (4.3) over all edges $(i, j) \in K_n$ gives

$$K_n = \sum_{i<j} G_{i,j} \preceq \sum_{i<j} (j-i) P_n.$$

To finish the proof, we compute

$$\sum_{i<j} (j-i) \leq \sum_{i<j} n \leq n^3$$

So

$$L_{K_n} \preceq n^3 \cdot L_{P_n}.$$

Plugging this into Equation (4) we obtain

$$\frac{1}{n^2} \leq \lambda_2(P_n).$$

This only differs from our test vector-based upper bound in Equation (1) by a factor 12.

We could make this consirably tigher by being more careful about the sums.

1.6 Laplacian Eigenvalues of the Complete Binary Tree

Let's do the same analysis with the complete binary tree with unit weight edges on $n = 2^{d+1} - 1$ vertices, which we denote by T_d.

T_d is the balanced binary tree on this many vertices, i.e. it consists of a root node, which has two children, each of those children have two children and so on until we reach a depth of d from the root, at which point the child vertices have no more children. A simple induction shows that indeed $n = 2^{d+1} - 1$.

We can also describe the edge set by saying that each node i has edges to its children $2i$ and $2i + 1$ whenever the node labels do not exceed n. We emphasize that we still think of the graph as undirected.

The largest eigenvalue. We'll start by above bounding $\lambda_n(\boldsymbol{L}_{T_d})$ using a test vector.

We let $\boldsymbol{x}(i) = 0$ for all nodes that have a child node, and $\boldsymbol{x}(i) = -1$ for even-numbered leaf nodes and $\boldsymbol{x}(i) = +1$ for odd-numbered leaf nodes. Note that there are $(n+1)/2$ leaf nodes, and every leaf node has a single edge, connecting it to a parent with value 0. Thus

$$\lambda_n(\boldsymbol{L}) = \max_{\boldsymbol{v} \neq \boldsymbol{0}} \frac{\boldsymbol{v}^\top \boldsymbol{L} \boldsymbol{v}}{\boldsymbol{v}^\top \boldsymbol{v}} \geq \frac{\boldsymbol{x}^\top \boldsymbol{L} \boldsymbol{x}}{\boldsymbol{x}^\top \boldsymbol{x}} = \frac{(n+1)/2}{(n+1)/2} = 1. \tag{5}$$

Meanwhile, every vertex has degree at most 3, so by Claim 1.3, $\lambda_n(\boldsymbol{L}) \leq 6$. So we can bound the largest eigenvalue above and below by constant.

λ_2 and diameter in any graph. The following lemma gives a simple lower bound on λ_2 for any graph.

Lemma 1.5. *For any unit weight graph G with diameter D,*

$$\lambda_2(\boldsymbol{L}_G) \geq \frac{1}{nD}.$$

Proof. We will again prove a lower bound comparing G to the complete graph. For each edge $(i,j) \in K_n$, let $G^{i,j}$ denote a shortest path in G from i to j. This path will have length at most D. So, we have

$$\begin{aligned} K_n &= \sum_{i<j} G_{i,j} \\ &\preceq \sum_{i<j} D G^{i,j} \\ &\preceq \sum_{i<j} D G \\ &\preceq n^2 D G. \end{aligned}$$

So, we obtain the bound

$$n^2 D \lambda_2(G) \geq n,$$

which implies our desired statement. $\square$

λ_2 in a tree. Since a complete binary tree T_d has diameter $2d \leq 2\log_2(n)$, by Lemma 1.5, $\lambda_2(\boldsymbol{L}_{T_d}) \geq \frac{1}{2n\log_2(n)}$.

Let us give an upper bound on λ_2 of the tree using a test vector. Let $\boldsymbol{x} \in \mathbb{R}^v$ have $\boldsymbol{x}(1) = 0$ and $\boldsymbol{x}(i) = -1$ for i in the left subtree and $\boldsymbol{x}(i) = +1$ in the right subtree. Then

$$\lambda_2(\boldsymbol{L}_{T_d}) = \min_{\substack{\boldsymbol{v} \neq \boldsymbol{0} \\ \boldsymbol{v}^\top \boldsymbol{1} = 0}} \frac{\boldsymbol{v}^\top \boldsymbol{L} \boldsymbol{v}}{\boldsymbol{v}^\top \boldsymbol{v}} \leq \frac{\boldsymbol{x}^\top \boldsymbol{L} \boldsymbol{x}}{\boldsymbol{x}^\top \boldsymbol{x}} = \frac{2}{n-1}.$$

So, we have shown $\frac{1}{2n\log_2(n)} \leq \lambda_2(\boldsymbol{L}_{T_d}) \leq \frac{2}{n-1}$, and unlike the previous examples, the gap is more than a constant.

In the graded homework, I will ask you to improve the lower bound to $1/(cn)$ for some constant c.