Spectral Graph Theory	Instructor: Padraic Bartlett

Lecture 4: Spectra and χ

Week 3	*Mathcamp 2011*

This lecture is going to be **awesome.**

Theorem 1 *For any graph G on n vertices, we have $\chi(G) \geq 1 - \frac{\lambda_{max}(G)}{\lambda_{min}(G)}$.*

Proof. This proof is easily the hardest/most conceptually difficult thing we're going to do in this class, and involves some rather strange/mysterious steps. To make this less mysterious, we're going to begin this proof with a "roadmap:" i.e. before we start, I want to talk about how the proof is going to go, and what tricks we're going to use later (so that they're not so baffling when they do show up!)

So: roadmap. We're studying the object A_G, G's adjacency matrix; specifically, we want a way to think about $\chi(G)$ while working with A_G. How can we do this? Well: one way is the following:

- Take G, and turn it into a n-dimensional vector space, by associating to each vertex v_j the basis vector $\mathbf{e}_j$ of $\mathbb{R}^n$ that's got a 1 in the i-th coördinate and 0's everywhere else.

- Once you've done this, take any $\chi(G) = k$-coloring of G, and let $C_1, \ldots C_k$ be the k distinct color classes of the vertices in G.

- So we've taken our graph G, turned it into a vector space, and used this abstraction to give us a way to "condense" G along its color classes $C_1, \ldots C_k$. How can we use these color classes to talk about A_G, and specifically about its eigenvalues?

- Well: let $\lambda_{\max}$ be the largest eigenvalue of A_G, and $\mathbf{v}$ be a corresponding eigenvector for λ. Because the basis vectors $\{\mathbf{e}_1, \ldots \mathbf{e}_n\}$ for $\mathbb{R}^n$ are each in one of the U_i's, we can find a way of writing $\mathbf{v}$ as a sum of elements

$$\sum_{i=1}^{k} c_i \cdot \mathbf{u}_i,$$

  where each $\mathbf{u}_i$ is in U_i, and they're all of length 1. Notice that all of the $\mathbf{u}_i$ are orthogonal, as each $\mathbf{u}_i$ only has nonzero coördinates at the locations where U_i contains the appropriate basis vector $\mathbf{e}_i$.

- Now, let U be the vector space generated by taking the vectors $\{\mathbf{u}_1, \mathbf{u}_2, \ldots \mathbf{u}_k\}$ as a **basis**: i.e. look at the space formed by taking linear combinations made precisely out of these $\mathbf{u}_i$'s. This is a k-dimensional space, and can be thought of as a way to "collapse" our original vector space along the k color classes we have, in a way that preserves the largest eigenvector of A_G, as it's in this space!

- Let S denote the $n \times k$ matrix that sends a vector in U (written in the k-dimensional form $(a_1\mathbf{u}_1, a_2\mathbf{u}_2, \ldots a_k\mathbf{u}_k)$ to the same vector as expressed in $\mathbb{R}^n$ (i.e as the n-dimensional vector $\sum a_i\mathbf{u}_i$). This matrix is specifically given as

$$S = \begin{bmatrix} \vdots & \vdots & & \vdots \\ \mathbf{u}_1 & \mathbf{u}_2 & \ldots & \mathbf{u}_n \\ \vdots & \vdots & & \vdots \end{bmatrix}.$$

- Finally, examine the linear map $B = S^T \cdot A_G \cdot S$, which takes in k-dimensional things (i.e. elements in U) and spits out other k-dimensional things (because the dimension of this matrix is $(k \times n) \cdot (n \times n) \cdot (n \times k) = (k \times k)$.) In essence, this map takes in elements in our condensed space U, interprets them as vectors in R^n, acts on them by A_G, and then takes them back into U. This certainly seems like a promising object to study! – it seems to be designed to preserve the largest eigenvalue of A_G, and yet only deal with a $\chi(G)$-sized subspace.

- Explicitly, we claim that this "condensing map" B has the following properties:
  - It has $\chi(G)$-many eigenvalues.
  - Its maximal eigenvalue is the same as the eigenvalue of A_G.
  - Its minimal eigenvalue is bounded below by the minimal eigenvalue of A_G.
  - The sum of all of the eigenvalues for this graph is 0.

- Notice that if we can prove these observations, we are done! I.e. if you plug in these three observations together, bounding all of the non-maximal eigenvalues below by $\lambda_{\min}$ and noting that the maximal one is $\lambda_{\max}$, we will have shown that $\lambda_{\max}(G) + (\chi(G) - 1)\lambda_{\min}(G) \leq 0$, which after some rearranging[1] gives our inequality above.

So, it suffices to prove these observations. We do this in the following series of lemmas and definitions:

Definition. Given two vectors $\mathbf{u}, \mathbf{v}$ in the same vector space, their **inner product** $\langle \mathbf{u}, \mathbf{v} \rangle$ is just their dot product:

$$\langle \mathbf{u}, \mathbf{v} \rangle = \mathbf{u}^T \cdot \mathbf{v} = \sum_{i=1}^{n} u_i v_i.$$

It's just notation, but it helps clean up a lot of things. We'll use it heavily throughout the following proofs.

Lemma 2 *B is a symmetric matrix.*

Proof. To see this, just take its transpose, remembering that A_G is itself a symmetric matrix: $B^T = (S^T \cdot A_G \cdot S)^T = (S^T)^T \cdot A_G^T \cdot S^T = S \cdot A_G \cdot A^T$.

[1] and noticing that the smallest eigenvalue $\lambda_{\min}$ is always negative for adjacency matrices of loopless graphs!

Corollary 3 *B has $\chi(G)$-many eigenvalues.*

Proof. Just use the spectral theorem!

Furthermore, by applying the spectral theorem on B again, you can prove (on the HW!) the following proposition:

Proposition 4 *For any real symmetric matrix A and for any vector $\mathbf{v}$ with $||v|| = 1$, we have*

$$\mu_{min} \leq \langle A\mathbf{v}, \mathbf{v} \rangle \leq \mu_{max},$$

where μ_{min}, μ_{max} are the smallest and largest eigenvectors of A, respectively. Furthermore, these bounds are always attained (specifically, by taking $\mathbf{v}$ to be an eigenvector corresponding to either the smallest or largest eigenvalue.)

The reason we care about the above proposition is the following lemma:

Lemma 5 *For any two vectors $\mathbf{v}, \mathbf{u} \in U$, we have*

$$\langle B\mathbf{u}, \mathbf{v} \rangle = \langle (S^T \cdot A_G \cdot S)\mathbf{u}, \mathbf{v} \rangle = \langle (A_G \cdot S)\mathbf{u}, S\mathbf{v} \rangle$$

Proof. By definition, we have

$$\begin{aligned}
\langle (S^T \cdot A_G \cdot S)\mathbf{u}, \mathbf{v} \rangle &= (S^T \cdot A_G \cdot S\mathbf{u})^T \cdot \mathbf{v} \\
&= (\mathbf{u}^T S^T \cdot A_G^T \cdot S) \cdot \mathbf{v} \\
&= (\mathbf{u}^T S^T \cdot A_G^T) \cdot (S\mathbf{v}) \\
&= (A_G \cdot S\mathbf{u})^T \cdot (S\mathbf{v}) \\
&= \langle (A_G \cdot S)\mathbf{u}, S\mathbf{v} \rangle.
\end{aligned}$$

Why do we mention this lemma? Well, it allows us to prove another one of B's claimed properties:

Corollary 6 *The eigenvalues of B are all bounded above by A_G's maximum eigenvalue λ_{max}, and below by A_G's minimum eigenvalue λ_{min}.*

Proof. In lemma 4, we proved that every inner product $\langle B\mathbf{v}), \mathbf{v} \rangle$ can be written in the form $\langle (A_G \cdot S)\mathbf{u}, S\mathbf{u} \rangle$. So, if we apply our proposition above to the real symmetric matrix A_G, we have just shown that these values $\langle (A_G \cdot S)\mathbf{u}, S\mathbf{u} \rangle$ are bounded above by $\lambda_{\max}$ and below by $\lambda_{\min}$, where these are A_G's maximum and minimum eigenvalues.

Therefore, we know that we must have $\mu_{\max} \leq \lambda_{\max}$ and $\mu_{\min} \leq \lambda_{\min}$, as these values are bounding all of the possible results for $\langle B\mathbf{v}, \mathbf{v} \rangle$, and therefore in specific are bounding B's maximum and minimum eigenvalues $\mu_{\min}, \mu_{\max}$.

This is another one of the properties we wanted to prove: i.e. that all of B's eigenvalues are bounded below by $\lambda_{\min}$!

We only have two more things to show, then: that $\lambda_{\max}$ is an eigenvalue of B (by the above lemma, we know that it would be a maximal eigenvalue if it is one), and that the sum of the eigenvalues of B is 0. We do this in two more lemmas:

Lemma 7 *B has λ_{max} as an eigenvalue.*

Proof. Specifically, notice that $\mathbf{v} = \sum_{i=1}^{k} c_i \cdot \mathbf{u}_i$ from earlier is an eigenvector for $\lambda_{\max}$. This is because for any of the basis vectors $\mathbf{u}_i$, we have

$$\begin{aligned}\langle B\mathbf{v}, \mathbf{u}_i\rangle &= \langle A_G\mathbf{v}, \mathbf{u}_i\rangle \\ &= (A_G\mathbf{v})^T \cdot \mathbf{u}_i \\ &= \lambda_{\max}(\mathbf{v})^T \cdot \mathbf{u}_i \\ &= \lambda_{\max}\left(\sum_{j=1}^{k} c_j \cdot \mathbf{u}_j^T\right) \cdot \mathbf{u}_i \\ &= \lambda_{\max}\left(\sum_{j=1}^{k} c_j \cdot (\mathbf{u}_j^T \cdot \mathbf{u}_i)\right) \\ &= \lambda_{\max} \cdot c_i \cdot ||\mathbf{u}_i|| \\ &= \lambda_{\max} \cdot c_i,\end{aligned}$$

where we justify those last two steps because the $\mathbf{u}_i$'s are all orthogonal and have norm 1. But this means that the i-th coördinate of $(S^T \cdot A_G \cdot S\mathbf{v}) \cdot \mathbf{v}$ is precisely $\lambda_{\max} \cdot c_i$, for every i: i.e. that

$$S^T \cdot A_G \cdot S\mathbf{v} = \lambda_{\max} \cdot (c_1\mathbf{u}_1, \ldots c_n\mathbf{u}_n) = \lambda_{\max} \cdot \mathbf{v},$$

and thus that $\lambda_{\max}$ is an eigenvalue, as claimed.

Lemma 8 *The trace*[2] *of B is 0.*

Proof. To see why, simply notice that we have

$$\langle B\mathbf{u}_i, \mathbf{u}_i\rangle = \langle A_G\mathbf{u}_i, \mathbf{u}_i\rangle.$$

However, notice that for any two basis vectors e_x, e_y of $\mathbb{R}^n$ that lie in the same color class C_i , we have

$$\langle A_G\mathbf{e}_x, \mathbf{e}_y\rangle = \langle (a_{1,x}, a_{2,x}, \ldots a_{n,x}), \mathbf{e}_y\rangle = a_{y,x} = 0,$$

as there are no edges between two vertices x, y with the same color i.

Because we can write $\mathbf{u}_i$ as the linear combination of several orthogonal elements all from the same color class, we know that in fact we have

$$\langle A_G\mathbf{u}_i, \mathbf{u}_i\rangle = 0,$$

and therefore that $\langle B\mathbf{u}_i, \mathbf{u}_i\rangle$ is also 0. Why do we care? Well, $B\mathbf{u}_i$ is the i-th column of B: taking its dot product with $\mathbf{u}_i$ then gives you the i-th element of that i-th column, i.e. the entry in (i, i). We've just proven that all of these entries are 0; therefore, the trace of B is trivially 0 as well.

[2] The **trace** of a matrix is the sum $\sum_{i=1}^{n} a_{i,i}$ of its diagonal elements.

Proposition 9 *The trace of a matrix is equal to the sum of its eigenvalues (counted with respect to their algebraic multiplicity.)*

Proof. The proof of this is simple: consider the characteristic polynomial! We defer the details to the HW.

So: let's combine these observations! We know that

- $\lambda_{\max}$ is an eigenvalue of this matrix.
- All of the other eigenvalues range from $\lambda_{\max}$ to $\lambda_{\min}$.
- There are k such eigenvalues counting multiplicity, by the spectral theorem.
- We know that the sum of all of these eigenvalues is 0.

So: if we bound the sum of all of the eigenvalues below by $\lambda_{\max}+(k-1)\lambda_{\min}$ by replacing all of the other nonmaximal eigenvalues with $\lambda_{\min}$ s, we get that

$$\begin{aligned}
&\lambda_{\max} + (k-1)\lambda_{\min} \leq 0 \\
\Rightarrow &(k-1)\lambda_{\min} \leq -\lambda_{\max} \\
\Rightarrow &k - 1 \geq -\frac{\lambda_{\max}}{\lambda_{\min}} \\
\Rightarrow &k \geq 1 - \frac{\lambda_{\max}}{\lambda_{\min}},
\end{aligned}$$

(where we switched the direction on our inequality above because $\lambda_{\min}$ is negative!)

This is what we sought to prove. Win.