MA431 Spectral Graph Theory: Lecture 5

Ahmad Abdi Neil Olver

Last time we stated the following theorem. Let us prove it.

Theorem 8.2 (Matrix-Tree Theorem)**.** *Let G be an n-vertex graph, and let L be its Laplacian matrix. Then $T(G)$ is equal to the determinant of any $(n-1) \times (n-1)$ principal submatrix of L.*

Proof. We shall proceed by induction on the number of edges of G. If G has no edge, then $T(G) = 0$, and since L is the zero matrix, the result follows. If $n = 2$, given that G has m parallel edges between the two vertices, we have $T(G) = m$, so the result follows since $L = \left(\begin{smallmatrix} m & -m \\ -m & m \end{smallmatrix}\right)$. For the induction step, assume that G has at least one edge, and $n \geq 3$. For $i \in [n]$, denote by $L[i]$ the principal submatrix of L obtained after removing row i and column i. It suffices to prove that $\det(L[n]) = T(G)$.

If G is not connected, then Proposition 7.5 implies that L has rank at most $n-2$, so $L[n]$ is a singular matrix, implying that $\det(L[n]) = 0 = T(G)$.

Otherwise, G is connected. Pick an edge e incident with n, say $e = \{n-1, n\}$. Consider the deletion $G \setminus e$ and the contraction G/e. For the latter, denote by $n-1$ the vertex obtained from identifying $n-1, n$. Denote by L^d, L^c the Laplacian matrices of $G \setminus e, G/e$, respectively. By the induction hypothesis, $T(G \setminus e) = \det(L^d[n])$ and $T(G/e) = \det(L^c[n-1])$. Let us recalculate the two determinants in terms of subdeterminants of L.

First, observe that $L^c[n-1] = L[n][n-1]$, so

$$\det(L^c[n-1]) = \det(L[n][n-1]).$$

Secondly, observe that L^d, L differ in only four entries, namely, $L^d_{n-1,n-1} = L_{n-1,n-1} - 1$, $L^d_{n,n} = L_{n,n} - 1$, $L^d_{n-1,n} = L_{n-1,n} + 1$ and $L^d_{n,n-1} = L_{n,n-1} + 1$. Subsequently, by a Laplace expansion along row $n-1$ of $L^d[n]$, we see that

$$\det(L^d[n]) = \det(L[n]) - \det(L[n][n-1]).$$

Consequently,

$$T(G \setminus e) + T(G/e) = \det(L^d[n]) + \det(L^c[n-1]) = \det(L[n]).$$

By Lemma 8.1, however, the LHS is equal to $T(G)$, so $T(G) = \det(L[n])$, thereby completing the induction step. □

As a consequence, we get a proof of *Cayley's formula*:

Corollary 8.3. *For every integer $n \geq 2$, the number of spanning trees K_n is n^{n-2}.*

Proof. Let L be the Laplacian matrix of K_n. Then $L = (n-1)I - (J - I) = nI - J$, where J is the all-ones matrix. As a result, any $(n-1) \times (n-1)$ principal submatrix of L is equal to $nI_{n-1} - J_{n-1}$. The spectrum of this submatrix is $1, n^{(n-2)}$, implying in turn that it has determinant n^{n-2}. The result now follows from the Matrix-Tree Theorem. □

We also have the following consequence of the Matrix-Tree Theorem:

Theorem 8.4. *Let G be a graph, and let $\lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_n$ be its Laplacian spectrum. Then*

$$T(G) = \frac{1}{n} \prod_{i=2}^{n} \lambda_i.$$

Proof. Exercise. □

For another application of the Matrix-Tree Theorem, see Exercise 6.

9 Extensions to weighted graphs

The Matrix-Tree Theorem has a useful extension to the weighted setting. Let $G = (V, E)$ be a graph. The *Kirchhoff polynomial of* G is the following polynomial with variables $(x_e : e \in E)$:

$$\operatorname{Kir}(G; x) := \sum_{T \text{ a spanning tree}} \prod_{e \in T} x_e.$$

Observe that if $w \in \mathbb{R}^E$ is a set of edge weights, then $\operatorname{Kir}(G; w)$ computes the sum of the "multiplicative weights" of the spanning trees. In particular, $\operatorname{Kir}(G; \mathbf{1})$ is nothing but the number of spanning trees of the graph G, i.e. $\operatorname{Kir}(G; \mathbf{1}) = T(G)$. Much like $T(G)$, the Kirchhoff polynomial has a powerful recursive formula.

Lemma 9.1. *For every edge,*

$$\operatorname{Kir}(G; x) = x_e \cdot \operatorname{Kir}(G/e; x^e) + \operatorname{Kir}(G \setminus e; x^e),$$

where x^e denotes the vector obtained from x after dropping the coordinate corresponding to e.

Proof. Exercise. □

Define $L(G, x)$ to be the matrix whose rows and columns are indexed by the vertices, and whose entries are defined as follows:

1. for each vertex u, the uu-entry is $\sum(x_e : e$ is an edge incident with $u)$,
2. for adjacent vertices u, v, the uv-entry is $-\sum(x_e : e$ has ends $u, v)$,
3. for non-adjacent vertices u, v, the uv-entry is 0.

Theorem 9.2. *Let G be an n-vertex graph, let $w \in \mathbb{R}^E$, and let $L_w := L(G, w)$. Then $\operatorname{Kir}(G; w)$ is equal to the determinant of any $(n-1) \times (n-1)$ principal submatrix of L_w.*

Proof. Exercise. □

We are now ready to define the Laplacian matrix in a particular weighted setting:

Definition 9.3. Let $G = (V, E)$ be a graph, and let $w \in \mathbb{R}_+^E$. The *Laplacian matrix* of the weighted graph (G, w) is the matrix $L(G, w)$.

Observe that the Laplacian matrix of the weighted graph $(G, \mathbf{1})$ is just the Laplacian matrix of the graph G. The nonnegativity of the edge weights is needed in order to guarantee the positive semidefinite-ness of the Laplacian matrix. More generally, we have the following:

Proposition 9.4. *Let $G = (V, E)$ be a graph, let $w \in \mathbb{R}_+^E$, and let L_w be the Laplacian of the weighted graph (G, w). Then the following statements hold:*

1. $L_w = \sum_{e=\{u,v\}\in E} w_e \cdot (e_u - e_v)(e_u - e_v)^\top$,
2. *for each* $x \in \mathbb{R}^V$, $x^\top L_w x = \sum_{e=\{u,v\}\in E} w_e (x_u - x_v)^2$,
3. L_w *is a positive semidefinite matrix,*
4. $\mathbf{1}$ *is an eigenvector of* L_w *with eigenvalue* 0,
5. *if every edge has nonzero weight, then* L_w *has rank* $n - c$, *and* 0 *as an eigenvalue has multiplicity* c, *where* c *is the number of connected components of* G

Proof. Exercise. □

In the weighted setting, for all intents and purposes, we may assume that the graph G is simple (i.e. it has no loops or parallel edges), and every edge has a strictly positive weight. These two assumptions can be made after deleting all edges of weight zero, and after collapsing all parallel edges to a single edge whose weight is the sum of the previous weights.

Exercises

1. Prove Proposition 7.2.
2. Recall the cospectral pair of graphs from Lecture 1, displayed in Figure 1. Find the Laplacian spectrum of each graph. Then conclude that cospectral graphs may not necessarily have the same Laplacian spectra.
3. Let G be an n-vertex simple graph, and let $\overline{G}$ be its complement. Prove the following statements:
   (a) $\lambda_i(\overline{G}) = n - \lambda_{n-i+2}(G)$ for $2 \leq i \leq n$,
   (b) $\lambda_n(G) \leq n$,
   (c) if $\overline{G}$ has $\overline{c}$ connected components, and $\overline{c} \geq 2$, then $\lambda_n(G) = n$ and its multiplicity is $\overline{c} - 1$.

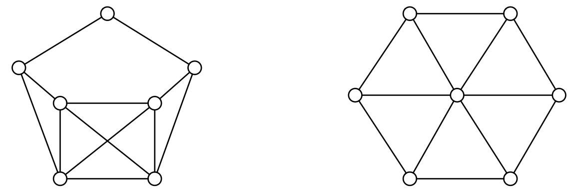

Figure 1: Two cospectral graphs with different Laplacian spectra.

4. Let $G = (V, E)$ be a connected graph, let $\lambda_2 := \lambda_2(G)$, and let $f \in \mathbb{R}^V$ be a corresponding eigenvector.

   (a) Call a path $(u_1, \ldots, u_r)$ *strictly decreasing* if $f_{u_1} > \cdots > f_{u_r}$. Prove that if $f_u > 0$, then it is joined by a strictly decreasing path to some vertex v such that $f_v \leq 0$.

   (b) Prove that for any $c \leq 0$, the graph induced on the vertex set $\{v \in V : f_v \geq c\}$ is connected.

5. Prove Theorem 8.4.

6. Let M be an $n \times n$, and let C be its cofactor matrix. Recall that C is an $n \times n$ matrix whose ij-entry is $(-1)^{i+j}$ times the determinant of the submatrix of M obtained after removing row i and column j. By a Laplace expansion along any row of M, we get the matrix equation $C^\top M = \det(M)I$. The matrix $C^\top$ is called the *adjugate of* M, and denoted $\operatorname{adj}(M)$.

   Let G be a graph, and let L be its Laplacian matrix. Prove that every entry of $\operatorname{adj}(L)$ is equal to $T(G)$.

7. Prove Lemma 9.1.

8. Prove Theorem 9.2.

9. Prove Proposition 9.4.

10. Let G be a connected graph on n vertices. Prove that

$$\lambda_2(G) = \min_x \frac{n \sum_{ij \in E}(x_i - x_j)^2}{\sum_{i<j}(x_i - x_j)^2}$$

   where the minimum is taken over all non-constant vectors x.

11. Let T be a tree. Prove that $\lambda_2(T) \leq 1$, and equality holds if and only if T is a star.

12. The *Cartesian product* of two graphs G, H, denoted $G \square H$, is the graph over vertex set $V(G) \times V(H)$, where (u_1, u_2) and (v_1, v_2) are adjacent if $u_1 = v_1$ and u_2, v_2 are adjacent in H, or $u_2 = v_2$ and u_1, v_1 are adjacent in G.

    Prove that $\lambda_2(G \square H) = \min\{\lambda_2(G), \lambda_2(H)\}$.

13. Denote by Q_n the skeleton graph of the n-dimensional unit hypercube. Prove that $\lambda_2(Q_n) = 2$.

14. Let G be a connected graph on n vertices and with diameter d. Prove that $\lambda_2(G) \geq \frac{1}{nd}$.

15. Let L_1, L_2 be positive semidefinite matrices of the same dimensions such that $L_1 \succcurlyeq L_2$. Prove that the k^{th} largest eigenvalue of L_1 in its spectrum is greater than or equal to the k^{th} largest eigenvalue of L_2 in its spectrum.

16. Let $(G, w), (H, w')$ be weighted graphs on the same number of vertices and with positive edge weights. Let L_1, L_2 be the Laplacian matrices of $(G, w), (H, w')$, respectively. We write $(G, w) \succcurlyeq (H, w')$ if $L_1 \succcurlyeq L_2$.

    Prove that
    $$(P_n, (n-1)\mathbf{1}) \succcurlyeq (A_n, \mathbf{1})$$
    where the weighted graph on the left is the path on vertices $\{1, 2, \ldots, n\}$ with edges $\{1, 2\}, \{2, 3\}, \ldots, \{n-1, n\}$, whose weights are equal to $n-1$, and the weighted graph on the right is the graph on the same vertex set with just one edge, $\{1, n\}$, whose weight is 1.

17. (a) Prove that $\lambda_2(K_n) = n$.

    (b) Prove that $\lambda_2(P_n) \geq \frac{6}{(n+1)(n-1)}$.

10 The cut and cycle spaces

Let $G = (V, E)$ be a **connected** undirected graph with n vertices and m edges, and let $\vec{G} = (V, \vec{E})$ be any orientation of G. The arbitrariness of the orientation may be a bit disconcerting, but everything we discuss will in fact be insensitive to the choice of orientation (just as was the case when we used it to give one definition of the Laplacian). We will often make use of the natural bijection between $\vec{E}$ and E without comment; so for example, we will generally treat $\mathbb{R}^E$ and $\mathbb{R}^{\vec{E}}$ as the same space. We will also define $\overleftrightarrow{G} = (V, \overleftrightarrow{E})$ to be the *bidirection* of G, where each edge is replaced by both orientations.

Throughout, we will use B to denote the vertex-edge incidence matrix associated with $\vec{G}$. That is, $B \in \mathbb{R}^{V \times \vec{E}} \simeq \mathbb{R}^{V \times E}$, with
$$B_{v,e} = \begin{cases} 1 & \text{if } v \text{ is the head of } e, \\ -1 & \text{if } v \text{ is the tail of } e, \\ 0 & \text{otherwise.} \end{cases}$$

A *flow* on G (or on $\vec{G}$; we won't distinguish) is simply any vector $f \in \mathbb{R}^E$ (or in $\mathbb{R}^{\vec{E}}$—again, we won't distinguish). Note that we do *not* require that $f_e \geq 0$ in a flow. While a positive flow on an edge $(u, v) \in \vec{E}$ should be interpreted as flow from u to v, a negative flow on this edge should be interpreted as flow in the reverse direction, from v to u. The purpose of an orientation is purely to indicate in which direction on an undirected edge a positive flow traverses.

The *net flow* into a node $v \in V$ induced by a flow $f \in \mathbb{R}^E$ is simply the total entering flow less the total leaving flow. We define ∇ to be the operator from $\mathbb{R}^E$ to $\mathbb{R}^V$ which maps a flow f to the vector b, where b_v is the net flow into v for each $v \in V$. That is,

$$(\nabla f)_v = \sum_{e=(u,v)\in \vec{E}} f_e - \sum_{e=(v,w)\in \vec{E}} f_e.$$

In the standard basis, the matrix representing ∇ is simply B. We will often write ∇f_v in place of $(\nabla f)_v$, as there can be no confusion.

A *circulation* is simply a flow f with $\nabla f_v = 0$ for all $v \in V$. More generally, given a *demand vector* $b \in \mathbb{R}^V$ with $\sum_{v\in V} b_v = 0$, we might be interested in flows which correctly match this demand vector, i.e., which satisfy $\nabla f = b$.

We will need some notation. Define, for any edge $e \in \overleftrightarrow{E}$, the vector $\chi^e \in \mathbb{R}^E$ by

$$(\chi^e)_a = \begin{cases} 1 & \text{if } a = e, \\ -1 & \text{if the reverse of } a \text{ is } e, \\ 0 & \text{otherwise.} \end{cases}$$

This is a *signed* characteristic vector of the edge e. For any $F \subseteq \overleftrightarrow{E}$, define $\chi(F) := \sum_{e\in F} \chi^e$. Given a set $S \subseteq V$, with $S \neq \emptyset$ and $S \neq V$, the *cut* associated with S, denoted by $\delta^+(S)$, is the set of arcs $(v, w) \in \overleftrightarrow{E}$ with $v \in S$ and $w \notin S$.

We now define two subspaces of $\mathbb{R}^E$. The subspaces do depend on the choice of orientation, but again, not in any important way.

The *cycle space* $W^\diamond$ is the set of all circulations in G, that is,

$$W^\diamond := \{f \in \mathbb{R}^E : \nabla f = \mathbf{0}\}.$$

We define the *cut space* (sometimes called the *star space*) as simply the orthogonal complement of $W^\diamond$, and denote it by $W^\star$: so $W^\star = (W^\diamond)^\perp$, and $W^\diamond \oplus W^\star = \mathbb{R}^E$. The reason for the names will become clear soon.

First, let's see an alternative description of $W^\diamond$, as well as a description of one possible basis. A reminder that given a spanning tree T[1] of G, every edge e not in T has an associated *fundamental cycle*, the cycle consisting of e along with the path in T between the endpoints of e. We will consider this as a directed cycle in $\overleftrightarrow{G}$, oriented so that e is included in the forward direction. There is also the *fundamental cut* associated with any edge $e \in T$: removing e from T splits it into two connected components, partitioning V into $S_e, V \setminus S_e$, where S_e contains the tail of e; the fundamental cut is $\delta^+(S_e)$.

Lemma 10.1. *The following statements about $W^\diamond$ hold.*

1. $W^\diamond = \mathrm{span}(\{\chi(C) : C$ *is a directed cycle in* $\overleftrightarrow{G}\})$.

2. *Given any spanning tree T, and taking C_e to be the fundamental cycle associated with e for each $e \notin T$, $\{\chi(C_e) : e \in E \setminus T\}$ is a basis for $W^\diamond$.*

[1] We will not distinguish between a spanning tree and its set of edges.

3. $\dim(W^\diamond) = m - n + 1$.

Proof. For any directed cycle C in $\overleftrightarrow{G}$, $\nabla(\chi(C)) = 0$, and hence $\chi(C) \in W^\diamond$. Since C_e does not contain e' for any $e \neq e' \in E \setminus T$, $\{\chi(C_e) : e \in E \setminus T\}$ is certainly linearly independent. It remains to show that $\dim(W^\diamond) = m - n + 1$; since a spanning tree has $n - 1$ edges, it then follows that $\{\chi(C_e) : e \in E \setminus T\}$ is indeed a basis, and so all parts of the lemma follow.

$W^\diamond$ is the kernel of the ∇ operator, and so it suffices to show that the rank of this operator, or equivalently the rank of B, is $n - 1$. We show this by demonstrating that the kernel of $B^\top$ is 1-dimensional. If $B^\top \alpha = 0$, we must have that $\alpha_u = \alpha_v$ for every $\{u, v\} \in E$, and hence by the connectivity of G, α is a multiple of $\mathbf{1}$. Further, clearly $B^\top \mathbf{1} = 0$. □

Now we move on to the cut space, which will in fact be the more important subspace for us. Both names (cut space as well star space) should become clear after this lemma.

Lemma 10.2. *The following statements about $W^\star$ hold.*

1. $W^\star = \operatorname{span}(\{\chi(\delta^+(S)) : \emptyset \subsetneq S \subsetneq V\})$.
2. $W^\star = \operatorname{span}(\{\chi(\delta^+(\{r\})) : r \in V\})$.
3. *Given any spanning tree T, and taking $\delta^+(S_e)$ to be the fundamental cut associated with e for each $e \in T$, $\{\chi(\delta^+(S_e)) : e \in T\}$ is a basis for $W^\star$.*
4. $\{\chi(\delta^+(\{r\})) : r \in V \setminus \{t\}\}$ *is a basis, for any choice of $t \in V$.*
5. $\dim(W^\star) = n - 1$.

Proof. Exercise. □

The grad operator (the *gradient*) is a linear map from $\mathbb{R}^V$ to $\mathbb{R}^E$; for any $\pi \in \mathbb{R}^V$, $f = \operatorname{grad}\pi$ is defined by $f_e = \pi_w - \pi_v$ for every $e = (v, w) \in \vec{E}$. The matrix of this linear operator in the standard basis is simply $B^\top$. As such, ∇ and grad are adjoint operators: $\langle \nabla f, \pi \rangle = \langle f, \operatorname{grad}\pi \rangle$.

Lemma 10.3. *The range of* grad *is precisely $W^\star$.*

Proof. $W^\diamond$ is the kernel of the ∇ operator. Thus $W^\star = (W^\diamond)^\perp$ is the range of the adjoint operator grad. □

We have the following relation to the Laplacian of G, which we will denote by L throughout.

Lemma 10.4. *The linear operator which is represented by L in the standard basis is precisely* $\nabla \operatorname{grad}$.

Proof. In the standard basis, $\nabla \operatorname{grad} = BB^\top$, which we already saw was one possible definition of the Laplacian. □

Remark 10.5. The Laplacian is also used to describe a differential operator. Given an appropriately smooth real-valued function $g : \mathbb{R}^k \to \mathbb{R}$, the Laplacian of g is written as $\nabla^2 g$; it is a map from $\mathbb{R}^k$ to $\mathbb{R}$ defined by

$$\nabla^2 g = \sum_{i=1}^{k} \frac{\partial^2 g}{\partial x_i^2}.$$

It can be viewed as first computing the gradient of g (often denoted by ∇g), which is

$$h = \nabla g = \left(\frac{\partial g}{\partial x_1}, \frac{\partial g}{\partial x_2}, \dots, \frac{\partial g}{\partial x_k} \right),$$

and then taking the divergence of h (often denoted by $\nabla \cdot h$), which is

$$\nabla \cdot h = \sum_{i=1}^{k} \frac{\partial h_i}{\partial x_i}.$$

In our discrete setting, we have exactly the same description of the Laplacian, as the divergence of the gradient[2]. The analogy is particularly clear if you consider the graph G describing a grid in k dimensions.

[2]We cannot get away with overloading ∇ for both the gradient and the divergence in the discrete setting without causing unnecessary confusion, so we use it only for the divergence.