An Introduction to Spectral Graph Theory

Jason Miller
Math 336

June 8, 2020

Abstract

Spectral graph theory is the study of the eigenvalues and eigenvectors of matrices associated with graphs. This paper is a review of Cvetković's *GRAPHS AND THEIR SPECTRA* [1], and builds up to a proof of Kirchhoff's Matrix Tree Theorem, which provides a formula for computing the number of spanning trees in a graph.

Contents

1 Introduction

Given a graph, we can construct the adjacency and Laplacian matrices, whose eigenvalues and eigenvectors convey surprisingly powerful information. In particular, Cvetković's paper, which is ultimately an overview of results in spectral graph theory, begins by introducing the fundamental properties of a spectrum of a graph, and proves combinatorial properties of graphs such as the number of walks on graphs and the characteristic polynomial of complete products of graphs. Then, the paper uses graph spectra to investigate the structure of special classes of graphs, and lastly proves results concerning the spanning trees of a graph. This review could be read as a presentation of the "best hits of (elementary) spectral graph theory," as it condenses the paper into several central results outlined in it while providing additional commentary, and applying a result in a way not done in the paper.

2 Some Introductory Graph-Theoretic Definitions

Broadly, graph theory is the study of graphs, which are networks of vertices connected by edges. The first results in spectral graph theory that this paper presents concerns the number of walks in an (undirected, unweighted) graph. In order to provide the graph-theoretic background for these results, we first present some definitions:

Definition 2.1. A *graph* G is an ordered pair (V, E), where V is any nonempty set and $E \subseteq \{\{v_1, v_2\} : (v_1, v_2) \in V \times V\}$. V is called the set of *vertices*, or nodes, of G, and elements of E are called *edges* of G.

If $\{a, b\} \in E$, then we say that a is *adjacent* to b. Additionally, a vertex v is *incident* to edge $\{a, b\}$ if $v = a$ or $v = b$. If there is more than one edge connecting two vertices, we call these edges *multiple edges*- note that our definition of graph does not allow for this. A *loop* is an edge that connects a node to itself, or a singleton set in E. We say that a graph is *simple* if it contains no loops or multiple edges. A graph is finite if $|V| < \infty$. Throughout this paper, we assume graphs are simple and finite.

Remark 1. If $E \subseteq V \times V$, so that the pairs of vertices are ordered, then we call G a directed graph. A graph is undirected if it is not directed. Throughout this paper, we discuss undirected graphs unless otherwise stated.

Definition 2.2. The *degree* of a vertex v, $deg(v)$, is equal to the number of edges adjacent to v.

Definition 2.3. A *walk* on a graph $G = (V, E)$ is a sequence of vertices $(v_1, \ldots v_k)$ such that $\{v_i, v_{i+1}\} \in E$ for all $1 \leq i \leq k - 1$. The length of the walk is $k-1$, i.e. the number of edges that were traversed along the walk.

If all vertices and edges along a walk are distinct, then we call it a *path*. We say that a path is *closed* if the first and last vertices in the seqence are the same (this is a slight abuse of notation, since this means that the first and last vertices are not distinct). A closed path is also called a *cycle*. Note that for each pair (v_i, v_{i+1}), where v_i and v_{i+1} are in the walk, the edge between them is unique due to the fact that there are no multiple edges. Additionally, a walk between vertex a and vertex b is called a *a-b walk*.

Definition 2.4. If v_1 and v_2 are vertices of G, and there exists a path from v_1 to v_2, then we say that v_1 and v_2 are *connected*. Furthermore, G is *connected* if each pair of vertices in G is connected.

Now, we may wish to store the information of a graph into a spreadsheet. The simplest, or at least most intuitive way to do this is with the *adjacency matrix*. Note that we can order the vertices in a graph G as $\{v_1, v_2 \dots v_n\}$.

Definition 2.5. For $G = (V, E)$, the adjacency matrix $A \in \mathbb{R}^{|V| \times |V|}$ is defined by $A_{ij} = \begin{cases} 1 & \{v_i, v_j\} \in E \\ 0 & \text{otherwise} \end{cases}$.

3 Motivating Question

Right away, we can notice some interesting properties of the adjacency matrix, namely that is is square and symmetric. Additionally, we recall that by the Spectral Theorem, any real symmetrix $n \times n$ matrix has all real eigenvalues and n orthonormal eigenvectors. Therefore the adjacency matrix has only real-valued eigenvalues and an orthonormal eigenbasis.

However, since we were motivated to define the adjacency matrix purely for the sake of storing data, it seems somewhat puzzling to treat it like actual linear transformation between vector spaces. Isn't this just a *spreadsheet*, after all? And yet, the question, "what if we treat this like any ordinary matrix?" is the fundamental question that birthed spectral graph theory!

4 Number of Walks and the Adjacency Matrix

Our first result concerns the number of walks in a graph. The statement is surprisingly simple, and we can obtain the result without considering the eigenvalues or eigenvectors of the adjacency matrix. It is stated as follows:

Theorem 4.1. *For a graph $G = (V, E)$ with adjacency matrix A and $V = \{v_1 \dots v_n\}$, and $m \in \mathbb{Z}$, the number of $v_i - v_j$ walks of length m is equal to $(A^m)_{ij}$.*

Proof. The proof is by induction.

Base case: when $m = 1$, then $A_{ij} = 1$ if and only if $\{v_i, v_j\} \in E$. This means that (v_i, v_j) is a $v_i - v_j$ walk of length 1, and this is the only such walk in G because all $v_i - v_j$ walks must visit both v_i and v_j, and if it visited any other vertex the length of the walk would be greater than 1.

Now, we assume that the statement holds for $m = k$ and prove it for $k + 1$.

First, since $A^{k+1} = (A^k)(A)$, we have that $(A^{k+1})_{ij} = \sum_{r=1}^{n} (A^k)_{ir} A_{rj}$. Then, since by the definition of A $A_{rj} = 0$ if $\{v_r, v_j\} \notin E$ and $A_{rj} = 1$ if $\{v_r, v_j\} \in E$, and $(A^k)_{ir}$ is the number of $v_i - v_r$ walks of length k by the inductive hypothesis, this means that $\sum_{r=1}^{n} (A^k)_{ir} A_{rj}$ is the number of $v_i - v_r$ walks of length k such that $\{v_r, v_j\} \in E$, and therefore the sequence $(v_i, \dots v_r, v_j)$ is a $v_i - v_j$ walk of length $k + 1$. Since all $v_i - v_j$ walks of length $k + 1$ can be written this way, this is the total number of $v_i - v_j$ walks of length $k + 1$. Therefore this property holds for all natural numbers by the induction principle, completing the proof. □

This is a highly encouraging result because it shows that multiplying out the adjacency matrix yields useful, nontrivial information about the graph, which means that "treating the adjacency matrix like a matrix" is indeed a meaningful action.

A very similar result concerns the number of *closed* walks of a given length between a pair of vertices, as well as an alternative way of writing this result that could be used to make clever combinatorial arguments. We present these results here as well. First, we need to define the spectrum of a graph and provide some properties of it:

Definition 4.1. The characteristic polynomial of a graph G with adjacency matrix A is the characteristic polynomial of A; that is, the function $P_G : \mathbb{C} \to \mathbb{C}$ defined by $P_G(\lambda) = \det(\lambda I - A)$, where I is the identity matrix with the same dimensions as A.

Definition 4.2. The *spectrum* of a graph G with adjacency matrix A is the set of eigenvalues of A, denoted $\text{Spec}(G)$.

Lemma 4.2. *If $M \in \mathbb{R}^{n\times n}$, and λ is an eigenvalue of M, then λ^n is an eigenvalue of M^n.*

Proof. The proof is by induction.

Base case: if $n = 1$, then of course λ^1 is an eigenvalue of M^1.

We assume the result holds for $n = k$, and prove it for $k + 1$.

Let v satisfy $Mv = \lambda v$, which exists because λ is an eigenvalue of M. Then $\lambda^{k+1}v = \lambda(\lambda^k v) = \lambda M^k v$ by the inductive hypothesis. Then,

$$\lambda M^k v = M^k \lambda v = M^k(Mv) = M^{k+1}v,$$

so λ^{k+1} is an eigenvalue of M^{k+1}. This proves the result for all natural numbers by the induction principle. □

Lemma 4.3. *For matrices $A \in \mathbb{R}^{n\times m}$ and $B \in \mathbb{R}^{m\times n}$,* $\text{Tr}(AB) = \text{Tr}(BA)$.

Proof. We multiply terms to obtain:

$$\text{Tr}(AB) = \sum_{i=1}^{n}(AB)_{ii} = \sum_{i=1}^{n}\sum_{j=1}^{m} A_{ij}B_{ji} = \sum_{j=1}^{m}\sum_{i=1}^{n} B_{ij}A_{ji} = \text{Tr}(BA).$$

□

Since we can apply the Spectral Theorem to A, let $\{u_1, \ldots u_n\}$ be an orthonormal eigenbasis for A, let $\lambda_1, \ldots \lambda_n$ be the eigenvalues that correspond to them, respectively, and let U be a matrix where the i^{th} column of U equals u_i. This gives us two interesting results: it allows us to rewrite Theorem 4.1:

Corollary 4.3.1. *Let u_{ij} be the $(i, j)^{th}$ entry of U. Then*

$$(A^m)_{ij} = \sum_{r=1}^{n} u_{ir}u_{jr}\lambda_r^m.$$

And it also allows us to calculate the number of closed walks along a graph:

Corollary 4.3.2. *The number of closed walks on G is* $\text{Tr}(A^m) = \sum_{k=1}^{n} \lambda_k^m$.

Proof. Here we prove both corollaries. By Lemma 4.2, $\lambda_1^k, \dots \lambda_n^k$ are eigenvalues of A^k. Let Λ be a diagonal matrix where $\Lambda_{ii} = \lambda_i$. Then $A = U\Lambda U^{\mathrm{T}}$, which proves Corollary 4.3.1. Then, by Theorem 4.1, $(A_{ii})^m$ is the total number of closed walks of length m from v_i to itself on G, so $\mathrm{Tr}(A^m)$ is the total number of closed walks on G. and by Lemma 4.3,

$$\mathrm{Tr}(A^m) = \mathrm{Tr}((U\Lambda U^{\mathrm{T}})^m) = \mathrm{Tr}(U\Lambda^m U^{\mathrm{T}}) = \mathrm{Tr}(UU^{\mathrm{T}}\Lambda^m) = \mathrm{Tr}(\Lambda^m) = \sum_{k=1}^{n} \lambda_k^m.$$

Therefore the total number of closed walks on G equals $\mathrm{Tr}(A^m) = \sum_{k=1}^{n} \lambda_k^m$, completing the proof. □

Remark 2. The $\lambda_i's$ are not necessarily distinct.

Now, we extend the results in Cvetković's paper by applying Corollary 4.3.1 to complete graphs. This is not a novel result, but it can illustrate how it can be used to derive closed-form expressions for combinatorial properties of graphs. First, we define what a complete graph is.

Definition 4.3. A *complete graph* K_n is a graph with n vertices such that every pair of distinct vertices is connected by an edge.

Proposition 4.1. *Let K_n be a complete graph with adjacency matrix A. Then*

$$(A^m)_{ii} = \frac{1}{n}((n-1)^m + (n-1)(-1)^m).$$

Proof. Let J be a $n \times n$ matrix with all $1's$. It is immediately apparent that $A = J - I$. The eigenvalues of J are 0 with multiplicity $n-1$, and n with multiplicity 1. This means that $J - I$ has eigenvalues -1 with multiplicity $n-1$ and $n-1$ with multiplicity 1. Therefore by Corollary 4.3.2, we have that $\mathrm{Tr}(A^m) = ((-1)^m(n-1) + (n-1)^m(1))$, and every $(A^m)_{ii}$ is equal by symmetry, so $(A^m)_{ii} = \frac{1}{n}((n-1)^m + (n-1)(-1)^m)$. □

Now, we transition to an overview of some powerful properties of the Laplacian matrix of a graph.

5 The Laplacian Matrix and Connectedness

So far, we've examined combinatorial properties of the adjacency matrix of a graph. But there's another way to represent a graph as a matrix than the adjacency matrix- the Laplacian matrix! First, let's define it in two ways.

Definition 5.1. For a graph $G = (V, E)$, where $V = \{v_1 \dots v_n\}$, the degree matrix of G, called D, is an $n \times n$ diagonal matrix where $D_{ii} = deg(v_i)$.

Definition 5.2. The Laplacian matrix, L, equals $D - A$, where A is the adjacency matrix of G.

From this construction, we see that the row entries of L sum to zero by the definition of degree.

This definition is straightforward, but there is an equivelant formulation of the Laplacian from which many fundamental properties follow easily.

Definition 5.3. For $a, b \in V$, define the $n \times n$ matrix $L_{\{a,b\}}$ by

$$(L_{\{a,b\}})_{ij} = \begin{cases} 1 & i = j \text{ and } (v_i = a \text{ or } v_i = b) \\ -1 & (v_i = a \text{ and } v_j = b) \text{ or vice versa} \\ 0 & \text{otherwise} \end{cases}$$

Note that for any $x \in \mathbb{R}^n$, we have that $x^{\mathrm{T}} L_{\{v_i,v_j\}} x = (x_i - x_j)^2 \geq 0$. This means that $L_{\{a,b\}}$ is positive-semidefinite.

Definition 5.4. The Laplacian matrix is defined as

$$L = \sum_{e \in E} L_e$$

It immediately follows that these two definitions of L are equivelent. Additionally, just like the adjacency matrix, L is real-valued and symmetric, so all eigenvalues are real by the Spectral Theorem. Unlike A, however, L is positive semidefinite.

Proof. Let $x \in \mathbb{R}^n$. We can write

$$x^{\mathrm{T}} L x = x^{\mathrm{T}} \sum_{\{v_i,v_j\} \in E} L_{\{v_i,v_j\}} x = \sum_{\{v_i,v_j\} \in E} x^{\mathrm{T}} L_{\{v_i,v_j\}} x = \sum_{\{v_i,v_j\} \in E} (x_i - x_j)^2 \geq 0,$$

with equality when $x = \mathbf{0}$. □

Because of this, the eigenvalues of L aren't just real; they're nonnegative.

Proof. Suppose λ is an eigenvalue of L and v is an eigenvector. Then,

$$v^{\mathrm{T}} L v = v^{\mathrm{T}} (\lambda v) = \lambda v^{\mathrm{T}} v.$$

Since both $v^{\mathrm{T}} L v \geq 0$ (by positive-semidefiniteness) and $v^{\mathrm{T}} v \geq 0$, this means that $\lambda \geq 0$. □

Another consequence of the Spectral Theorem is that there is an orthonormal basis of eigenvectors of L. This means that we can pick an orthonormal basis $\{u_1 \dots u_n\}$, and then take the corresponding eigenvalues $\lambda_1, \dots \lambda_n$ (as with the adjacency matrix, these aren't necessarily distinct). Without loss of generality, assume the eigenvalues are ordered such that $\lambda_n \geq \lambda_{n-1} \dots \lambda_1 \geq 0$.

We will now see that the eigenvalues of the Laplacian matrix are intricately tied to the connectedness of the graph. Immediately, there is an important result concerning the first eigenvalue for all graphs, and the second-smallest eigenvalues for connected graphs:

Proposition 5.1. *For all graphs,* $\lambda_1 = 0$.

Proof. Let $\mathbf{1} \in \mathbb{R}^n$ be a vector with all $1's$. Then, since the rows of L sum to zero, $L\mathbf{1} = \mathbf{0}$. Therefore 0 is an eigenvalue of L with eigenvector $\mathbf{1}$. □

Lemma 5.1. *If G is connected, then* $\lambda_2 > 0$.

Proof. Take w to be a nontrivial eigenvector of 0. Then,

$$\sum_{\{v_i, v_j\} \in E} (w_i - w_j)^2 = w^{\mathrm{T}} L w = w^{\mathrm{T}}(0) = 0.$$

Therefore, for all $\{v_i, v_j\} \in E$, $w_i = w_j$. Since G is connected, this means $w_i = w_j$ for all $v_i, v_j \in V$. Therefore $w \in \operatorname{span}\{\mathbf{1}\}$. Therefore the dimension of the eigenspace of $\lambda_1 = 1$. Therefore the multiplicity is 1, so $\lambda_2 \neq \lambda_1$, so $\lambda_2 > 0$. □

There is, however, a far more general result about the number of connected components of a graph and the multiplicity of the 0 eigenvalue. First, we need more terminology:

Definition 5.5. A *subgraph* G' of G is a graph (V', E') such that $V' \subseteq V$ and $E' \subseteq E$.

Definition 5.6. A *connected component* G' of G is a connected subgraph such that any pair of vertices $a, b \in V, a \in V', b \in V \setminus V'$ is not connected.

Theorem 5.2. *The multiplicity of the* 0 *eigenvalue equals the number of connected components of* G.

Proof. Suppose G has k connected components $(V_1, E_1), \dots (V_k, E_k)$. Similar to the proof of Lemma 5.1, let $W \in \mathbb{R}^{n \times k}$ be defined by $W_{ij} = \begin{cases} 1 & j \in V_i \\ 0 & \text{otherwise} \end{cases}$. Let x be a nontrivial eigenvector with eigenvalue 0.

Then by Lemma 5.1, $x_i = x_j$ for any $v_i, v_j \in V$ that are in the same connected component. Therefore the eigenspace of 0 is simply the span of the columns of W. Then, since the columns of W are clearly linearly independent, the multiplicity of 0 equals k. □

Next, we prove the Matrix-Tree Theorem.

6 The Matrix-Tree Theorem

The Matrix-Tree Theorem, first proven by Kirchhoff in 1847, provides a formula for the number of spanning trees in a graph that can be computed in polynomial time. To proceed, we will establish more terminology.

Definition 6.1. A *tree* is a connected graph with no cycles.

Definition 6.2. A connected subgraph $G' = (V', E')$ of G is a *spanning subgraph* of G if $V' = V$. It is a *spanning tree* if it is a tree and a spanning subgraph of G.

Note that since if G is not connected, there are clearly no spanning trees- graphs that are not connected are irrelevant to this section and so we assume G is connected unless otherwise specified.

Definition 6.3. The *complexity* of G, denoted by $\kappa(G)$, is the number of spanning trees of G.

Definition 6.4. An *orientation* of G can be defined as follows: for each $\{a, b\} \in E$, choose an ordered pair (a, b) or (b, a). The first vertex in the list is called the *initial vertex* and the second is called the *final vertex.*

Similar to how we constructed $L_{\{a,b\}}$, we can also construct another matrix called the *incidence matrix*:

Definition 6.5. Suppose G has an orientation, and $E = \{e_1 \dots e_m\}$. Then the *incidence matrix* $M \in \mathbb{R}^{n \times m}$ is defined by

$$M_{ij} = \begin{cases} -1 & e_j \text{ has initial vertex } v_i \\ 1 & e_j \text{ has final vertex } v_i \\ 0 & \text{otherwise} \end{cases}$$

Definition 6.6. For a set S, denote $\binom{S}{k}$ as the set of k-element subsets of S. For $A \in \mathbb{R}^{n \times m}$, denote the $m \times m$ submatrix of A formed by taking the columns of A indexed by $\binom{S}{k}$ as $A[S]$.

Before we prove the Matrix-Tree Theorem, we must establish a lemma relating the incidence matrix to the Laplacian:

Lemma 6.1.

$$L = MM^T$$

Proof. Since $M_{ij} = \sum_{k=1}^{m} M_{ik}M_{jk}$, if $i \neq j$, for each k, if e_k connects v_i and v_j then exactly one of M_{ik} and M_{jk} equals 1, and the other equals -1. If $i = j$, then both equal zero if e_k isn't incident to v_i, or both equal 1 or -1 otherwise. □

This provides another interpretation for the Laplacian as a discretized version of the Laplacian operator (∇^2) for twice-differentiable functions on graphs.

Definition 6.7. M_0 is the matrix obtained by deleting the last row from M. For $S \in \binom{m}{n-1}$, denote the subset of edges indexed by S as $S(E)$. The graph formed by $S(E)$ (which we will abuse notation and also call $S(E)$) is the graph with edges equal to $S(E)$ and vertices incident to the edges in $S(E)$.

Now, we are ready to prove the Matrix-Tree Theorem:

Theorem 6.2.

$$\kappa(G) = \frac{1}{n}\lambda_2\lambda_3 \ldots \lambda_n$$

Proof. Recall the Cauchy-Binet formula, which states that if $A \in \mathbb{R}^{m\times n}, B \in \mathbb{R}^{n\times m}$, then

$$\det(AB) = \sum_S \det(A[S]) \det(B[S]).$$

Note that $M_0[S]^{\mathrm{T}} = M_0^{\mathrm{T}}[S]$. Since $L = MM^{\mathrm{T}}$, we have that $L_0 = M_0M_0^{\mathrm{T}}$. Then, by the Cauchy-Binet formula, we have that

$$\det(L_0) = \sum_S \det(M_0[S]) \det(M_0^{\mathrm{T}}[S]) = \sum_S \det(M_0[S])^2$$

And since we can see that $\det(M_0[S]) = 1$ if $M_0[S]$ forms a spanning tree, and equals zero otherwise, and a tree with n vertices has $n-1$ edges, the right-hand-side gives us the number of spanning trees. Now we compute that

$$\det(L - \lambda I) = \prod_{i=1}^{n}(\lambda_i - \lambda),$$

and so the first-order coefficient of this expression is $-n\det(L_0)$, which is also equal to $-\lambda(\prod_{i=1}^{n}\lambda_i)$, and the result follows. □

7 Conclusion

In this review, we highlighted several important results in spectral graph theory, which are representative of Cvetković's paper. Namely, we proved formulas for the number of walks of a given length on a graph using the adjacency matrix, results about the connectedness of a graph using the Laplacian matrix, and lastly provided a proof of the Matrix-Tree Theorem.

References

[1] Cvetković, Dragoš M. "GRAPHS AND THEIR SPECTRA." Publikacije Elektrotehničkog Fakulteta. Serija Matematika i Fizika, no. 354/356, 1971, pp. 1–50. JSTOR, www.jstor.org/stable/43667526. Accessed 7 June 2020.

[2] Harary, Frank. "Graphs and Matrices." SIAM Review, vol. 9, no. 1, 1967, pp. 83–90. JSTOR, www.jstor.org/stable/2027411. Accessed 8 June 2020