Lecture 2

In which we introduce the Laplacian matrix and we prove our first results in spectral graph theory.

1   The Basics of Spectral Graph Theory

Given an undirected graph $G = (V, E)$, the approach of spectral graph theory is to associate a symmetric real-valued matrix to G, and to related the eigenvalues of the matrix to combinatorial properties of G.

For the sake of this lecture, we will restrict ourselves to the case in which G is a d-regular graph, and we will then see how to extend our results to apply to irregular graphs as well.

The most natural matrix to associate to G is the adjacency matrix A such that $A_{i,j} = 1$ if $\{i, j\} \in E$ and $A_{i,j} = 0$ otherwise. In the second part of the course, in which we will study expander graphs, the adjacency matrix will indeed be the most convenient matrix to work with. For the sake of the algorithms that we will analyze in the first part of the course, however, a slight variation called the *normalized Laplacian* is more convenient.

There are a few ways to motivate the definition of the Laplacian. One way is the following: the variational characterization of the eigenvalues of real symmetric matrices tells us that we can think of the eigenvalues of M as optima of min-max optimization problems in which vectors $\mathbf{x} \in \mathbb{R}^V$ are feasible solutions and the cost function is the Rayleigh quotient

$$R_M(x) = \frac{\mathbf{x}^T M \mathbf{x}}{\mathbf{x}^T \mathbf{x}}$$

We know that every homogeneous polynomial of degree 2 can be realized as $x^T M x$ for some matrix M, and,if we want to study cuts in a graph $G = (V, E)$, it makes sense to choose a matrix M such that

$$\mathbf{x}^T M \mathbf{x} = \sum_{\{u,v\} \in E} (x_u - x_v)^2$$

because, if $\mathbf{x} \in \{0, 1\}^V$ is a Boolean vector, representing a cut in the graph, then the right-hand-side expression above is counting the number of edges that cross the cut, and so optimization problems with the above cost functions will be relaxations of cut problems.

Some calculations show that the matrix having such a property is $dI - A$, which is called the Laplacian matrix of G. Indeed, we can verify that

$$\mathbf{x}^T(dI - A)\mathbf{x} = \sum_{\{u,v\} \in E} (x_u - x_v)^2$$

because both expressions are easily seen to be equal to

$$\sum_v dx_v^2 - 2 \sum_{\{u,v\} \in E} x_u x_v$$

As we will see in a moment, the eigenvalues of $dI - A$ are in the range $[0, 2d]$, and it is not hard to see that their sum is dn, so it is convenient to divide the Laplacian matrix by d so that the range and the average values of the eigenvalues of the resulting matrix are independent of the degree. (This degree independence will make it possible to generalize results to the irregular case.)

We have thus reached the following definition.

Definition 1 (Normalized Laplacian) *The normalized Laplacian matrix of an undirected d-regular graph $G = (V, E)$ is $L := I - \frac{1}{d}A$.*

We shall now prove the following relations between the eigenvalues of L and certain purely combinatorial properties of G.

Theorem 2 *Let G be a d-regular undirected graph, let A be the adjacency matrix of G, and $L = I - \frac{1}{d} \cdot A$ be the normalized Laplacian matrix of G. Let $\lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_n$ be the real eigenvalues of L with multiplicities, in nondecreasing order. Then*

1. *$\lambda_1 = 0$ and $\lambda_n \leq 2$.*

2. *$\lambda_k = 0$ if and only if G has at least k connected components.*

3. *$\lambda_n = 2$ if and only if at least one of the connected components of G is bipartite.*

Note that the first two properties imply that the multiplicity of 0 as an eigenvalue is precisely the number of connected components of G.

PROOF: By the characterization of the Rayleigh quotient of L that we established above, and from the variational characterization of eigenvalues, we have

$$\lambda_1 = \min_{\mathbf{x} \in \mathbb{R}^n - \{\mathbf{0}\}} \frac{\sum_{\{u,v\} \in E} (x_u - x_v)^2}{d \sum_v x_v^2}$$

and so $\lambda_1 \geq 0$ because the Rayleigh quotient, being a ratio of sums of squares, is always non-negative.

If we take $\mathbf{1} = (1, \ldots, 1)$ to be the all-one vector, we see that its Rayleigh quotient is 0, and so 0 is the smallest eigenvalue of L, with $\mathbf{1}$ being one of the vectors in the eigenspace of 1.

We also have the following formula for λ_k:

$$\lambda_k = \min_{S \; k-\text{dimensional subspace of } \mathbb{R}^n} \; \max_{\mathbf{x} \in S - \{\mathbf{0}\}} \frac{\sum_{\{u,v\} \in E} (x_u - x_v)^2}{d \sum_v x_v^2}$$

So, if $\lambda_k = 0$, there must exist a k-dimensional space S such that for every $\mathbf{x} \in S$ and every $\{u, v\} \in E$, we have $x_u = x_v$, and so $x_u = x_v$ for every u, v which are in the same connected component. This means that each $\mathbf{x} \in S$ must be constant within each connected component of G, and so the dimension of S can be at most the number of connected components of G, meaning that G has at least k connected components.

Conversely, if G has at least k connected components, we can let S be the space of vectors that are constant within each component, and S is a space of dimension at least k such that for every element $\mathbf{x}$ of S we have

$$\sum_{\{u,v\} \in E} (x_u - x_v)^2 = 0$$

meaning that S is a witness of the fact that $\lambda_k = 0$.

Finally, to study λ_n, we first note that we have the formula

$$\lambda_n = \max_{\mathbf{x} \in \mathbb{R}^n - \{\mathbf{0}\}} \frac{\mathbf{x}^T L \mathbf{x}}{\mathbf{x}^T \mathbf{x}}$$

from the variational characterization of eigenvalues (see Handout 0).

We also observe that for every vector $\mathbf{x} \in \mathbb{R}^n$ we have

$$2\mathbf{x}^T \mathbf{x} - \mathbf{x}^T L \mathbf{x} = \frac{1}{d} \sum_{\{u,v\} \in E} (x_u + x_v)^2$$

and so

$$\lambda_n = 2 - \min_{\mathbf{x} \in \mathbb{R}^n - \{\mathbf{0}\}} \frac{\sum_{\{u,v\} \in E} (x_u + x_v)^2}{d \sum_v x_v^2}$$

from which it follows that

$$\lambda_n \leq 2$$

and if $\lambda_n = 2$ then there must be a non-zero vector $\mathbf{x}$ such that

$$\sum_{\{u,v\} \in E} (x_u + x_v)^2 = 0$$

which means that $x_u = -x_v$ for every edge $(u,v) \in E$.

Let us now define $A := \{v : x_v > 0\}$ and $B := \{v : x_v < 0\}$. The set $A \cup B$ is non-empty (otherwise we would have $\mathbf{x} = \mathbf{0}$) and is either the entire graph, or else it is disconnected from the rest of the graph, because otherwise an edge with an endpoint in $A \cup B$ and an endpoint in $V - (A \cup B)$ would give a positive contribution to $\sum_{\{u,v\} \in E} (x_u - x_v)^2$; furthermore, every edge incident on a vertex on A must have the other endpoint in B, and vice versa. Thus, $A \cup B$ is a connected component, or a collection of connected components, of G which is bipartite, with the bipartition A, B. $\square$