Lecture 17

Lecturer: Debmalya Panigrahi   *Scribe: Kevin Sun*

1 Overview

In this lecture, we look at the fundamental concepts of spectral graph theory.

2 Spectral Graph Theory

The basic premise of spectral graph theory is that we can study graphs by considering their matrix representations. Thus, we begin by briefly reviewing the basic properties of real square matrices, which is sufficient for our purposes.

2.1 Linear Algebra Review

Let A be an $n \times n$ matrix over the reals; we often write $A \in \mathbb{R}^{n \times n}$. The *column space* is the set of all linear combinations of the columns of A; its dimension is known as the *rank* of A. The null space of A is the vectors x such that $Ax = 0$, and its dimension is known as the *nullity* of A. The rank-nullity theorem tells us that the sum of the rank and nullity of A is equal to n.

The *determinant* of A is defined as follows:

$$\det(A) = \sum_{\sigma:[n] \to [n]} \text{sgn}(\sigma) \prod_{i=1}^{n} A_{i,\sigma(i)}$$

where $[n] = \{1, 2, \ldots, n\}$, $\text{sgn}(\sigma)$ denotes the *sign* of the permutation σ (1 if σ has an even number of inversions and -1 otherwise), and $A_{i,j}$ denotes the (i, j)-th element of A.

A non-zero vector x is an *eigenvector* of A associated with eigenvalue λ if $Ax = \lambda x$. Note that if x is an eigenvector of A, then so is any scaled version of x, so we often assume x has unit length. Also, $Ax = \lambda x$ occurs for non-zero x if and only if $\det(A - \lambda I) = 0$, where I denotes the $n \times n$ identity matrix. The polynomial $\det(A - \lambda I)$ is known as the *characteristic polynomial* of A, and the eigenvalues of A are its roots.

The *trace* of a square matrix, defined as the sum of its diagonal entries, is equal to the sum of its eigenvalues. Also, the determinant of a matrix is equal to the product of its eigenvalues.

Theorem 1 (The Spectral Theorem). *Let $A \in \mathbb{R}^{n \times n}$ be a symmetric matrix. Then there exist $\lambda_1, \ldots, \lambda_n \in \mathbb{R}$ and mutually orthonormal vectors $x_1, \ldots, x_n$ such that x_i is an eigenvector of A with eigenvalue x_i for every $i \in [n]$.*

If $A \in \mathbb{R}^{n \times n}$ is symmetric with rank r, then r of the vectors given by Theorem 1 span the column space of A, and the remaining $n - r$ vectors span the null space of A. Furthermore, the eigenvalue associated with each of these $n - r$ vectors is 0, because $Ax_i = \lambda_i x_i = 0$ implies $\lambda_i = 0$.

A square matrix is *diagonal* if its only non-zero entries are on its diagonal. It is well-known that any symmetric matrix $A \in \mathbb{R}^{n \times n}$ is *diagonalizable*, that is, there exist matrices $L, D \in \mathbb{R}^{n \times n}$ such that $A = LDL^{-1}$, and D is diagonal. Furthermore, the diagonal entries of D are the eigenvalues of A, and the columns of L are the corresponding eigenvectors. Since we can assume the eigenvectors are orthonormal, we have $L^{\top} L = I$, which implies $L^{\top} = L^{-1}$.

2.2 Matrices of Graphs

Now we consider an undirected graph $G = (V, E)$ and a few of its matrix representations. The *adjacency matrix* of G is defined as follows:

$$A_{ij} = \begin{cases} 1 & \text{if } (i, j) \in E \\ 0 & \text{otherwise.} \end{cases}$$

Example 1: Let A be the adjacency matrix of the complete graph on n vertices; we shall compute the eigenvalues of A. Notice that $A = \mathbb{1} - I$, where $\mathbb{1}$ denotes the all-ones $n \times n$ matrix. We see that the all-ones vector is an eigenvector of $\mathbb{1}$ with eigenvalue n. Since the rank of $\mathbb{1}$ is 1, by the discussion above, the remaining eigenvalues are all 0. Since $Ax = (\mathbb{1} - I)x = \mathbb{1}x - x$, the eigenvalues of A are simply the eigenvalues of $\mathbb{1}$ shifted down by one. Thus, the eigenvalues of A are $n - 1$ (with multiplicity 1) and -1 (with multiplicity $n - 1$).

Example 2: Now let A be the adjacency matrix of a bipartite graph on n vertices, and let z be an eigenvector of A corresponding to eigenvalue λ. Notice that by relabeling the vertices, we can write the equation $Az = \lambda z$ as follows:

$$\begin{bmatrix} 0 & M^{\top} \\ M & 0 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} M^{\top} y \\ Mx \end{bmatrix} = \begin{bmatrix} \lambda x \\ \lambda y \end{bmatrix},$$

where $z = \begin{bmatrix} x & y \end{bmatrix}^{\top}$, which implies $M^{\top} y = \lambda x$ and $Mx = \lambda y$. Thus, we also have

$$\begin{bmatrix} 0 & M^{\top} \\ M & 0 \end{bmatrix} \begin{bmatrix} x \\ -y \end{bmatrix} = \begin{bmatrix} -M^{\top} y \\ Mx \end{bmatrix} = -\lambda \begin{bmatrix} x \\ -y \end{bmatrix},$$

so $-\lambda$ is also an eigenvalue of A. In fact, the converse is also true: if the eigenvalues of an adjacency matrix A can be paired up in this manner, then the underlying graph must be bipartite.

The Laplacian: The *Laplacian matrix* L of G is defined as $L = D - A$, where D is the diagonal matrix containing the degree values of G, and A is the adjacency matrix of G. It can be shown that L is *positive semidefinite* (PSD), that is, L is symmetric and $x^{\top} Lx \geq 0$ for every $x \in \mathbb{R}^n$.

Theorem 2. *Let $M \in \mathbb{R}^{n \times n}$ be a symmetric matrix. Then the following are equivalent:*

1. *For every $x \in \mathbb{R}^n$, $x^{\top} Mx \geq 0$.*

2. *Every eigenvalue of M is non-negative.*

3. *There exists a matrix B such that $M = B^{\top} B$.*

Proof. (1) implies (2): Let v be an eigenvector of M corresponding to eigenvalue λ. Then $Mv = \lambda v$, which implies $v^\top M v = \lambda v^\top v \geq 0$. Since $v^\top v$ is non-negative, λ must also be non-negative.

(2) implies (3): Since M is symmetric, we can write $M = QPQ^\top$ where the Q is orthogonal and P is the diagonal matrix whose entries are the eigenvalues of M. Every eigenvalue of M is non-negative, so $P = R^\top R$ for some diagonal matrix R. Thus, $M = QR^\top RQ^\top = (RQ^\top)^\top RQ$.

(3) implies (1): Notice $x^\top M x = x^\top B^\top B x = (Bx)^\top (Bx) \geq 0$. $\qquad\square$

The Laplacian of a matrix is associated with a *quadratic form*: if $x \in \mathbb{R}^n$ is a vector (i.e., a function from V to $\mathbb{R}$), then

$$x^\top L x = \sum_{(i,j) \in E} (x_i - x_j)^2.$$

(Notice that this equality immediately implies that L is positive semidefinite.) This notion is useful because it generalizes cuts in a graph: for any $S \subseteq V$, if we set $x_u = 1$ if $u \in S$ and 0 otherwise, then $x^\top L x$ precisely captures the number of edges with exactly one endpoint in S.

Recall that the goal of graph sparsification was the following: given a graph with Laplacian matrix L, find a sparser graph that preserves all cut values. By the discussion above, this is equivalent to preserving the value of $x^\top L x$ for every $x \in \{0, 1\}$. In spectral sparsification, we want to preserve this value for any $x \in \mathbb{R}^n$. For this problem, Batson, Spielman, and Srivastava [BSS12] gave a solution containing $O(n)$ edges. Notice that this result is stronger than the graph sparsification results we saw in Lecture 10 in a few ways: it preserves every $x \in \mathbb{R}^n$, there are no logarithmic factors, the algorithm is deterministic.

3   Summary

In this lecture, we introduced the foundations of spectral graph theory.

References

[BSS12] Joshua Batson, Daniel A Spielman, and Nikhil Srivastava. Twice-ramanujan sparsifiers. *SIAM Journal on Computing*, 41(6):1704–1721, 2012.