Spectral Graph Theory	Instructor: Padraic Bartlett
Lecture 1: An Introduction to Spectral Theory	
Week 3	*Mathcamp 2011*

1 Motivation/Definitions

For those of you who've been in the graph theory courses for the last week, something you might have noticed about graph theory is that we really, um, don't understand very much. We don't have terribly good characterizations of graphs based on their chromatic number, we don't understand how k-flows work on graphs, we don't know the chromatic number of the unit distance graph, and we only just recently discovered a decent characterization of perfect graphs (which are honestly pretty simple, as far as graphs go!)

Conversely, for those of you who've been in the linear algebra sequence for the last week: something you might have noticed is that we're (honestly) pretty good at this stuff! If I give you a linear map $f : \mathbb{R}^n \to \mathbb{R}^n$, you can:

- Describe f as a $n \times n$ matrix, A_f!
- Find all of the eigenvectors and eigenvalues of A_f – i.e. find all of the subspaces E_i such that $f(E_i) = E_i$, and find the constants that f dilates these subspaces by!
- If f happens to be particularly nice and has n eigenvalues (counting multiplicity), then we can make a basis for $\mathbb{R}^n$ out of its eigenvectors!
- Furthermore, in such a nice case, we can write A_f as UDU^T, where U is a unitary matrix[1] and D is a diagonal matrix[2] with entries made out of A_f's eigenvalues!

So: Graphs are hard. Linear algebra is easy! How can we combine these?

One way (specifically, the way we're going to focus on in this course:) the adjacency matrix!

Definition. Given a graph G with vertex set $\{1, \ldots n\}$, we define its **adjacency matrix** A_G as the following $n \times n$ matrix:

$$A = \{a_{ij} : a_{ij} = 1 \text{ if } \{i,j\} \in E(G), \text{ and } 0 \text{ otherwise.}\}$$

There are a bunch of alternate/different ways to associate graphs to matrices (the Laplacian and incidence matrices, in particular, are fairly interesting things to study;) for now, however, we'll focus on the adjacency matrix, as it's easy to work with and fairly nice. Specifically, one property of these matrices that we'll use constantly is that these matrices are **symmetric** and **real-valued**! This allows us to apply the **spectral theorem**:

[1] A **unitary matrix** is a matrix whose columns form an orthonormal basis for $\mathbb{R}^n$.

[2] A **diagonal matrix** is a matrix where its only nonzero entries are on the diagonal.

Theorem 1 *Any real-valued symmetric matrix A has n eigenvalues (counting multiplicity) and n corresponding eigenvectors, which we can choose to all be orthogonal to each other. More explicitly, we can write $A = EDE^T$, where*

- *E is an orthonormal matrix whose columns are the eigenvectors of A, and*
- *D is a diagonal matrix whose entries are the corresponding eigenvalues of A.*

For reference, we calculate a few easy-to-find adjacency matrices:

Example.

1. The graph K_n has adjacency matrix with 0's on the diagonal and 1's everywhere else:

$$\begin{bmatrix} 0 & 1 & 1 & \dots & 1 \\ 1 & 0 & 1 & \dots & 1 \\ 1 & 1 & 0 & \dots & 1 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & 1 & \dots & 0 \end{bmatrix}$$

2. The empty graph $\overline{K_n}$'s adjacency matrix is identically 0:

$$\begin{bmatrix} 0 & 0 & \dots & 0 \\ 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 0 \end{bmatrix}$$

3. Enumerate the vertices of the cycle graph C_n as $\{1, 2, \dots n\}$ and its edges as $\{\{i, i+1\} : 1 \leq i \leq n\}$. Then, its adjacency matrix has ones as depicted below:

$$\begin{bmatrix} 0 & 1 & 0 & 0 & \dots & 1 \\ 1 & 0 & 1 & 0 & \dots & 0 \\ 0 & 1 & 0 & 1 & \dots & 0 \\ 0 & 0 & 1 & 0 & \ddots & 0 \\ \vdots & \vdots & \vdots & \ddots & \ddots & 1 \\ 1 & 0 & 0 & \dots & 1 & 0 \end{bmatrix}$$

2 Applications of the Adjacency Matrix: Counting Paths / Walks

So: we have these graphs, and we've turned them into matrices. How can we use these matrices to get back information about these graphs?

One quick application is to **counting walks on a graph**! Specifically: suppose we have a graph G on n vertices, and two nodes $i, j \in V(G)$. How do we count all of the walks[3] of length k from i to j?

[3] A **walk** of length n from vertex i to vertex j is a sequence $P = (i, \{i, x_1\}, x_1, \{x_1, x_2\}, \dots \{x_n, j\}, j)$.

Well: let's limit ourselves to just walks of length 1. Then, it's trivially just 1 if there is an edge connecting i and j, and 0 otherwise. What about walks of length 2? Well: any walk of length two will have to connect i to some vertex v, and then connect v to j: i.e. it's the sum

$$\sum_{v=1}^{n} \text{isEdge}(i, v) \cdot \text{isEdge}(v, j).$$

But wait! We've defined these isEdge functions earlier – specifically, we defined the adjacency matrix A_G of G in such a way that $a_{ij} = 1$ whenever there is an edge from i to j, and 0 otherwise. So, in this notation, we have that the number of walks from i to j is just

$$\sum_{v=1}^{n} a_{iv} \cdot a_{vk},$$

which we can recognize as the dot product

$$\begin{bmatrix} a_{i1} & a_{i2} & \dots & a_{in} \end{bmatrix} \cdot \begin{bmatrix} a_{1j} \\ a_{2j} \\ \dots \\ a_{nj} \end{bmatrix}.$$

But this is just the dot product of the i-th row and the j-th row of A_G! So, we've just proven the following:

Proposition 2 *Suppose G is a graph with vertex set $\{1, \dots n\}$ with adjacency matrix A. Then the (i, j)-th entry of A^2 denotes the number of walks of length 2 from i to j.*

We can easily generalize this to walks of length k:

Theorem 3 *Suppose G is a graph with vertex set $\{1, \dots n\}$ with adjacency matrix A. Then the (i, j)-th entry of A^k denotes the number of distinct walks of length k from i to j.*

Proof. As discussed above, this is trivially obvious for $k = 1$.

We proceed by induction on k. Suppose that we know that the entries of A^k correspond to the number of walks of length k from i to j. Given i and j, how can we find all of the walks of length $k + 1$ from i and j? Well: any walk of length $k + 1$ from i to j can be described as a walk from i to some vertex v of length k, and then a walk of length 1 from v to j itself! So, if we just simply use the summation trick we used before, we can see that

$$\begin{aligned} \text{numberOfWalks}_{k+1}(i, j) &= \sum_{v=1}^{n} \text{numberOfWalks}_k(i, v) \cdot \text{isEdge}(v, j) \\ &= (i, j) - \text{th entry of } A_G^k \cdot A_G \\ &= (i, j) - \text{th entry of } A_G^{k+1}. \end{aligned}$$

As a quick corollary, we have the following:

Corollary 4 *Suppose G is a graph with vertex set $\{1, \ldots n\}$ with adjacency matrix A. The number of distinct triangles*[4] *(v_1, v_2, v_3) contained within G is $tr(A^3)/6$.*

Proof. A triangle with a fixed starting point and order in which to visit its vertices is precisely a closed walk of length 3. There are three possible starting points (v_1, v_2, v_3) and two possible orientations (clockwise, counterclockwise) in which to traverse any such closed walk; therefore, the number of triangles is just 1/6-th of the number of closed walks on a graph of length 3.

But the number of closed walks on a graph of length 3 is just the sum over all $v \in V(G)$ of the closed length-3 walks starting at v: i.e. the sum of the diagonal entries in A^3, which is (by definition) $tr(A^3)$.

Question 5 *Can you derive a similar formula for 4-gons?*

3 Adjacency Matrices and Isomorphism

As we saw in the introduction to graph theory class, to really work with graphs we need to consider them only up to isomorphism – i.e. we want to think of the Petersen graph as just anything with the same vertex-edge relations as the Petersen graph, and not care so much about the labeling of its vertices. However, adjacency matrices care **very much** about the labeling of our vertices: i.e. for the two graphs below,

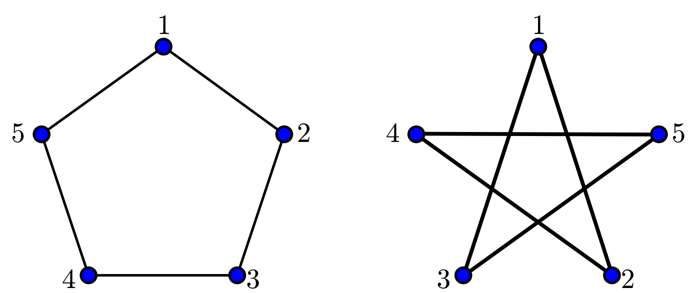

despite the fact that they're both "pentagons," their adjacency matrices are quite different:

$$A_{G_1} = \begin{bmatrix} 0 & 1 & 0 & 0 & 1 \\ 1 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 \\ 1 & 0 & 0 & 1 & 0 \end{bmatrix}, \qquad A_{G_2} = \begin{bmatrix} 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 1 & 0 \end{bmatrix}$$

This is ... troublesome. If we're going to use linear algebra to study our graphs, getting different results whenever we label our graph differently is going to give us no end of trouble. So: can we say anything about the relation between these matrices at all?

Thankfully, there is! To say precisely what it is, we need the following definition:

[4] A **triangle** in G is a triple (v_1, v_2, v_3) where all of the edges $\{v_1, v_2\}, \{v_2, v_3\}, \{v_3, v_1\}$ are contained within G.

Definition. A $n \times n$ matrix P whose entries are all either 0 or 1 is called a **permutation matrix** if P has exactly one 1 in each of its rows and columns. For example, the following matrix is a permutation matrix:

$$\begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \end{bmatrix}$$

The reason we call this a permutation matrix is because multiplying a vector $\mathbf{v}$ on the left by P "permutes" $\mathbf{v}$' entries! For example

$$\begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \end{bmatrix} \cdot \begin{bmatrix} v_1 \\ v_2 \\ v_3 \\ v_4 \\ v_5 \\ v_6 \end{bmatrix} = \begin{bmatrix} v_2 \\ v_6 \\ v_1 \\ v_5 \\ v_4 \\ v_3 \end{bmatrix}$$

Given a permutation $\sigma : \{1, \ldots n\} \to \{1, \ldots n\}$, we will sometimes write P_σ to denote the permutation matrix such that $P(v_1, \ldots v_n) = (v_{\sigma(1)}, \ldots v_{\sigma(n)})$. It bears noting that every permutation matrix P can be expressed as P_σ for some permutation σ, by just tracking where it sends a generic vector $(v_1, \ldots v_n)$.

We first note the following property of permutation matrices:

Proposition 6 *If P is a $n \times n$ permutation matrix with associated permutation σ, then $(v_1, \ldots v_n) \cdot P = (v_{\sigma^{-1}(1)}, \ldots v_{\sigma^{-1}(n)})$.*

Proof. On the HW!

Proposition 7 *If P is a $n \times n$ permutation matrix with associated permutation σ, then P^{-1} is also a permutation matrix with associated permutation σ^{-1} (and furthermore is equal to P^T.)*

Proof. On the HW!

Proposition 8 *If P is a $n \times n$ permutation matrix, P is unitary; furthermore, if E is any other unitary matrix, $P \cdot E$ is still unitary.*

Proof. On the HW!

Given this, we can prove the following remarkably useful fact about adjacency matrices of isomorphic graphs:

Proposition 9 *If G_1 and G_2 are a pair of isomorphic graphs with adjacency matrices A_1, A_2, then A_1 and A_2 are conjugate via a permutation matrix P: i.e.*

$$A_2 = PA_1P^{-1}.$$

Proof. Suppose that G_1 and G_2 are isomorphic graphs, both with vertex set $\{1, \ldots n\}$. Then there is some permutation σ of $\{1, \ldots n\}$ that realizes this isomorphism (i.e. such that (i, j) is an edge in A_1 iff $(\sigma(i), \sigma(j))$ is an edge in A_2.)

Let P be the associated permutation to this map σ; then, we have that

- PA_1 is the matrix where we've taken each column of A_1 and permuted its entries according to σ : in other words, PA_1 is A_1 if we permute its rows by σ.
- Similarly, A_1P^{-1} is the matrix where we permute A_1's columns by $(\sigma^{-1})^{-1} = \sigma$, by our earlier two lemmas.
- By combining these two results, PA_1P^{-1} is the matrix where we permute A_1's rows by σ, and then permute the resulting matrices' columns by σ again!

What does this mean? Well: we've started by taking any point (i, j) in A_1, and have sent it to $(\sigma(i), \sigma(j))$. But this means that we've sent the indicator function for the edge (i, j) to the location $(\sigma(i), \sigma(j))$! In other words, we've sent A_1 to A_2: i.e. we've proven $A_2 = PA_1P^{-1}$, as claimed.

One remarkable consequence of this is the following corollary:

Corollary 10 *If G_1 and G_2 are isomorphic graphs, their adjacency matrices A_1 and A_2 have the same set of eigenvalues (counting multiplicity.)*

Proof. Pick any permutation matrix P such that $A_2 = PA_1P^{-1}$. We know that A_1 is real-valued and symmetric; therefore, we can write it in the form EDE^T, for some unitary matrix E and diagonal matrix of eigenvalues D. But this means that we've written

$$\begin{aligned} A_2 &= (P \cdot E) \cdot D \cdot (E^T \cdot P^T) \\ &= (P \cdot E) \cdot D \cdot (P \cdot E)^T. \end{aligned}$$

But what have we done here? We've expressed A_2 as D conjugated by a unitary matrix. But this means that[5] A_2's eigenvalues are precisely the entries on the diagonal of D! So we've proven our claim.

This motivates us to make the following definition:

Definition. The **spectrum** of a graph G is the set of all of A_G's eigenvalues, counted with multiplicity. For example, we say that the empty graph on 3 vertices, $\overline{K_3}$, has spectrum $\{0^3\}$, where by 0^3 we mean that it has the eigenvalue 0 repeated three times.

[5]By a HW exercise!