Chapter 1

Matrices and Graphs

1.1 The Adjacency Matrix

This section is an introduction to the basic themes of the course.

Definition 1.1.1. *A simple undirected* graph $G = (V, E)$ *consists of a non-empty set* V *of vertices and a set* E *of unordered pairs of distinct elements of* V*, called edges.*

It is useful, and usual, to think a graph as a picture, in which the vertices are depicted with dots and the edges are represented by lines between the relevant pairs of dots. For example

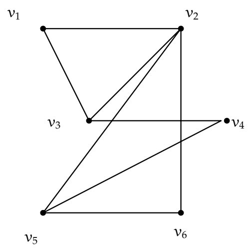

A *directed graph* is similar, except that edges are ordered instead of unordered pairs of vertices. In pictures, the ordering is indicated by an arrow pointing from the initial vertex of the edge to the terminal vertex. Other variants on the definition allow loops (edges from a vertex to itself) or multiple edges between the same pair of vertices. Graph Theory is the mathematical study of graphs and their variants. The subject has lots of applications to the analysis of situations in which members or subgroups of some population are interacting with each other in different ways, for example to the study of (e.g. electrical, traffic, social) networks.

Graphs can be infinite or finite, but in this course we will only consider finite graphs. An undirected graph is *connected* if it is all in one piece. In general the connected pieces of a graph are called *components*. Given a graph G, the numerical parameters describing G that you might care about include things like

- the *order* (the number of vertices);
- the number of edges (anything from zero to $\binom{n}{2}$ for a simple graph of order n);
- the number of connected components;
- the maximum (or minimum, or average) *vertex degree* - the degree of a vertex is the number of edges incident with that vertex;
- if G is connected, its *diameter* - this is the distance between a pair of vertices that are furthest apart in G;

- the length of the longest *cycle*;
- the size of the largest *clique*;
- the size of the largest *independent set*;
- if G is connected, its *vertex-connectivity* - the minimum number of vertices that must be deleted to disconect the graph;
- if G is connected, its *edge-connectivity* - the minimum number of edges that must be deleted to disconnect the graph;
- the list goes on ...

Thinking about graphs as pictures is definitely a very useful conceptual device, but it can be a bit misleading too. If you are presented with a picture of a graph with 100 vertices and lots of edges, and it is not obvious from the picture that the graph is disconnected, then deciding by looking at the picture whether the graph is connected is not at all easy (for example). We need some systematic ways of organising the information encoded in graphs so that we can interpret it. Luckily the machinery of linear algebra turns out to be extremely useful.

Definition 1.1.2. *Let* G *be a graph with vertex set* $\{v_1, \dots, v_n\}$*. The* adjacency matrix *of* G *is the* $n \times n$ *matrix that has a 1 in the* (i, j)*-position if there is an edge from* v_i *to* v_j *in* G *and a 0 in the* (i, j)*-position otherwise.*

Examples

1. An undirected graph and its adjacency matrix.

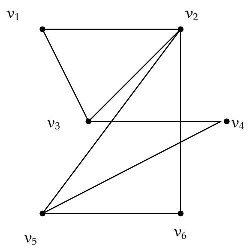

$$A = \begin{pmatrix} 0 & 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & 0 & 1 & 0 \end{pmatrix}$$

2. A directed graph and its adjacency matrix.

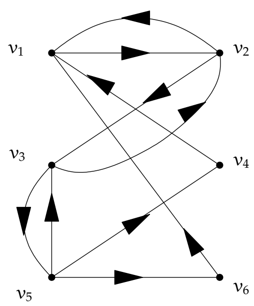

$$A = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

Notes

1. The adjacency matrix is *symmetric* (i.e. equal to its transpose) if the graph is undirected.
2. The adjacency matrix has zeros on its main diagonal (unless the graph has loops).
3. A graph can easily be reconstructed from its adjacency matrix.

4. The adjacency matrix of a graph G depends on a choice of ordering of the vertices of G (so technically we should talk about the adjacency matrix with respect to a particular ordering). The adjacency matrices A and A' of the same graph G with respect to different orderings are related by *permutation similarity*, i.e.

$$A' = P^{-1}AP,$$

where P is a permutation matrix - i.e. a $(0,1)$-matrix with exactly one entry in each row and column equal to 1. Note that a permutation matrix is *orthogonal*, its inverse is equal to its transpose (more on that later).
Exercise: Prove the above assertion about the connection between adjacency matrices corresponding to different orderings.

Given a graph G, its adjacency matrix is nothing more than a table that records where the edges are in the graph. It happens to be a matrix, but its definition does not involve anything to do with matrix algebra. So there is no good reason to expect that applying the usual considerations of matrix algebra (matrix multiplication, diagonalization, eigenvalues, rank etc) to A would give us anything meaningful in terms of the graph G. However it does. The first reason for that is the following theorem, which describes what the entries of the positive integer powers of A tell us about the graph G.

Theorem 1.1.3. *Let* A *be the adjacency matrix of a simple graph* G *on vertices* $v_1, v_2, \dots, v_n$*. Let* k *be a positive integer. Then the entry in the* (i,j)*-position of the matrix* A^k *is the number of walks of length* k *from* v_i *to* v_j *in* G.

Proof. We use induction on k. The theorem is clearly true in the case $k = 1$, since the (i,j)-entry is 1 if there is a walk of length 1 from v_i to v_j (i.e. an edge), and 0 otherwise.

Assume that the theorem holds for all positive integers up to $k-1$. Then

$$(A^k)_{ij} = \sum_{r=1}^{n} (A^{k-1})_{ir} A_{rj}.$$

We need to show that this is the number of walks of length k from v_i to v_j in G. By the induction hypothesis, $(A^{k-1})_{ir}$ is the number of walks of length $k-1$ from v_i to v_r. For a vertex v_r of G, think of the number of walks of length k from v_i to v_j that have v_r as their second-last vertex. If v_r is adjacent to v_j, this is the number of walks of length $k-1$ from v_i to v_r. If v_r is not adjacent to v_j, it is zero. In either case it is $(A^{k-1})_{ir}A_{rj}$, since A_{rj} is 1 or 0 according as v_r is adjacent to v_j or not. Thus the total number of walks of length k from v_i to v_j is the sum of the expressions $(A^{k-1})_{ir}A_{rj}$ over all vertices v_r of G, which is exactly $(A^k)_{ij}$. □

An immediate consequence of Theorem 1.1.3 is that the trace of the matrix A^2 (i.e. the sum of the diagonal entries) is the sum over all vertices v_i of the number of walks of length 2 from v_i to v_i. The number of walks of length 2 from a vertex to itself is just the number of edges at that vertex, or the vertex *degree*. So

$$\text{trace}(A^2) = \sum_{v \in V} \deg(v) = 2|E|.$$

It is a well known and very useful fact that in a graph without loops, the sum of the vertex degrees is twice the number of edges - essentially this is the number of "ends of edges" - every edge contributes twice to $\sum_{v \in V} \deg(v)$.

In A^3, the entry in the (i,i)-position is the number of walks of length 3 from v_i to itself. This is twice the number of 3-cycles in G that include the vertex v_i (why twice?). Thus, in calculating the trace of A^3, every 3-cycle (or triangle) in the graph, contributes six times - twice for each of its three vertices. Thus

$$\text{trace}(A^3) = 6 \times (\text{number of triangles in } G).$$

Exercise: this interpretation of the trace of A^k as counting certain types of walks in G does not work so well from $k = 4$ onwards - why is that?

A reason for focussing on the trace of powers of the adjacency matrix at this stage is that it opens a door to the subject of spectral graph theory. Recall the following facts about the trace of a $n \times n$ matrix A (these will be justified in the next section).

1. Let the eigenvalues of A (i.e. the roots of the polynomial $\det(\lambda I_n - A)$) be $\lambda_1, \dots, \lambda_n$ (not necessarily distinct). Then $\text{trace}(A) = \lambda_1 + \lambda_2 + \cdots + \lambda_n$. So the sum of the eigenvalues is equal to the sum of the diagonal entries. The eigenvalues are generally not equal to the diagonal entries, but they are for example if A is upper or lower triangular.

2. The eigenvalues of A^2 are $\lambda_1^2, \lambda_2^2, \dots, \lambda_n^2$, and the trace of A^2 is the *sum of the squares of the eigenvalues of* A.

3. In general, for a positive integer k,

$$\text{trace}(A^k) = \sum_{i=1}^{n} (\lambda_i)^k.$$

The central question of spectral graph theory asks what the spectrum (i.e. the list of eigenvalues) of the adjacency matrix A of a graph G tells us about the graph G itself. The observations above tell us that the answer is not nothing. We know that if $\text{spec}(A) = [\lambda_1, \dots, \lambda_n]$, then

- $\sum_{i=1}^{n} \lambda_i^2$ is twice the number of edges in G.
- $\sum_{i=1}^{n} \lambda_i^3$ is six times the number of triangles in G.

This means that the adjacency spectrum of a graph G "knows" the number of edges in G and the number of triangles in G (and obviously the number of vertices in G). To put that another way, if two graphs of order n have *the same spectrum*, they must have the same number of edges and the same number of triangles.

Definition 1.1.4. *Two graphs* G *and* H *are called* cospectral *if their adjacency matrices have the same spectrum.*

Below is a pair of cospectral graphs that do not have the same number of cycles of length 4; G has 5 and H has 6. Each has 7 vertices, 12 edges and 6 triangles. Each has spectrum $[-2, -1, -1, 1, 1, 1 + \sqrt{7}, 1 - \sqrt{7}]$.

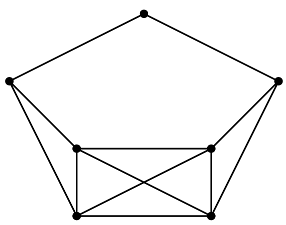

G

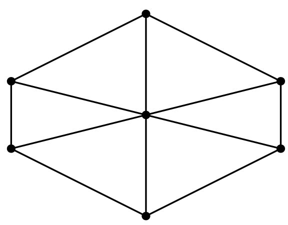

H

Exercise: Why does it not follow from the reasoning for edges and triangles above that cospectral graphs must have the same number of 4-cycles?

Something else that the adjacency spectrum does not "know" about a graph is its number of components. The following is an example of a pair of cospectral graphs of order 5, one is connected and one is not.

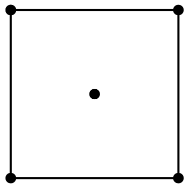

G

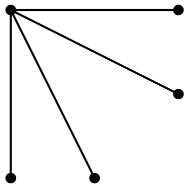

H

The adjacency matrices are

$$A_G = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & 1 & 0 \end{pmatrix}, \quad A_H = \begin{pmatrix} 0 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \end{pmatrix}$$

It is easily observed that both A_G and A_H have rank 2, so each has zero occurring at least three times as an eigenvalue. By considering

$$A_H v = \begin{pmatrix} 0 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} a \\ b \\ c \\ d \\ e \end{pmatrix} = \begin{pmatrix} b + c + d + e \\ a \\ a \\ a \\ a \end{pmatrix},$$

we find that $A_H v = \lambda v$ only if $a = \lambda b = \lambda c = \lambda d = \lambda e$ and $b + c + d + e = \lambda a$. If $\lambda \neq 0$ this means $b = c = d = e$ and $\lambda a = \lambda^2 b = 4b$. So $\lambda^2 = 4$ and the possible values of λ are 2 and -2. Thus $\mathrm{spec}(A_H) = [0, 0, 0, 2, -2]$.

On the other hand A_G also has rank 2 and so has zero occurring (at least) three times as an eigenvalue. Because the first row of A_G is a zero row, and the other row sums in A_G are all equal to 2, it follows that 2 is an eigenvalue of A_G, with corresponding eigenvector having 0 in the first position and 1 in the other four. Since the sum of the eigenvalues is the trace of A_G which is zero, the fifth eigenvalue must be -2. So $\mathrm{spec}(A_G) = [0, 0, 0, 2, -2] = \mathrm{spec}(A_H)$.

We have shown that A_H and A_G are cospectral, but G has two connected components and H has one. So the number of connected components in a graph is not determined by the adjacency spectrum.

We finish off this section with a famous example of a theorem in graph theory that can be proved using analysis of the spectrum of an adjacency matrix.

Theorem 1.1.5 (Erdös-Rényi-Sós, the Friendship Theorem (1966))**.** *Let G be a finite graph on at least three vertices, in which every pair of vertices has exactly one common neighbour (the "friendship property"). Then there is a vertex in G that is adjacent to all the others.*

Remarks

1. The theorem is called the Friendship Theorem because it can be expressed by the statement that in a group of people in which every pair has exactly one mutual friend, there is a person who is friends with everyone (the "politician").

2. After we prove the theorem it is relatively easy to describe the finite graphs which have the property - they are the "windmills", also called "friendship graphs". The windmill W_r has $2r + 1$ vertices and consists of r triangles, all sharing one vertex but otherwise disjoint.

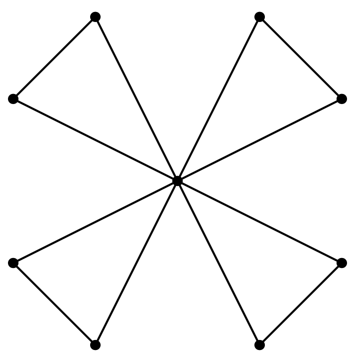

W_4

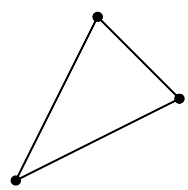

W_1

3. A graph is *regular* if all of its vertices have the same degree.

Proof. Our proof has two steps - the first is to show that a counterexample to the theorem would have to be a regular graph, and the second is to consider what the hypotheses would say about the square of the its adjacency matrix.

Let G be a graph satisfying the hypothesis of the theorem, and suppose that no vertex of G is adjacent to all others. Let u and v be two non-adjacent vertices of G. Write $k = \deg(u)$ and let $x_1, \ldots, x_k$ be the neighbours of u, where x_1 is the unique common neighbour of u and v. For each i in the range 1 to k, let y_1 be the unique common neighbour of v and x_i. The y_i are all distinct, since if two of them coincided then this vertex would have more than one common neighbour with u. Thus v has degree at least k and $\deg u \leqslant \deg v$. The same argument with the roles of u and v reversed shows that $\deg v \leqslant \deg u$, so we conclude that $\deg v = k$, and that $\deg u' = \deg v'$ whenever u' and v' are non-adjacent vertices of G.

Now let w be any vertex of G, other than x_1. Since u and v have only one common neighbour, w is not adjacent to both u and v, so there is a vertex of degree k to which it is not adjacent. Thus $\deg w = k$ by the above argument. Now all vertices of G have degree k except possibly x_1. If there is a vertex of G to which x_1 is non-adjacent, then this vertex has degree k and hence so does x_1. The alternative is that x_1 is adjacent to all other vertices of G which means that the conclusion of the theorem is satisfied. We have shown that any counterexample to the statement of the theorem would have to be a regular graph.

Now we assume that G is such a counterexample and that G is regular of degree k. Let n be the order (number of vertices) of G. Let u be a vertex of G. Each of the other $n-1$ vertices of G is reachable from u by a unique path of length 2. The number of such paths emanating from u is $k(k-1)$, since there are k choices for the first edge and then $k-1$ for the second. It follows that we can express n in terms of k:

$$k(k-1) = n - 1 \Longrightarrow n = k^2 - k + 1.$$

Now let A be the adjacency matrix of G and consider the matrix A^2. Each entry on the diagonal of A^2 is k, the number of walks of length 2 from a vertex to itself. Each entry away from the main diagonal is 1 - the number of walks of length 2 between two distinct vertices. Thus

$$A^2 = (k-1)I + J,$$

where I is the identity matrix and J is the matrix whose entries are all equal to 1 (this is fairly standard notation in combinatorics). We consider the eigenvalues of A^2. These are the roots of the characteristic polynomial

$$\det(\lambda I - (k-1)I - J) = \det\left((\lambda - k + 1)I - J\right).$$

Thus the number λ_1 is an eigenvalue of A^2 if and only if $\lambda_1 - k + 1$ is an eigenvalue of J, and these respective eigenvalues of A^2 and J occur with the same multiplicities. We can obtain the spectrum of A^2 by adding $k-1$ to every element in the spectrum of J. The spectrum of J is easy to determine directly - it has rank 1 and so has 0 occurring as an eigenvalue $n-1$ times. Its row sums are all equal to n and so it has n occurring (once) as an eigenvalue. Thus

$$\operatorname{spec}(J) = [0, 0, \ldots, 0, n] \Longrightarrow \operatorname{spec}(A^2) = [k-1, k-1, \ldots, k-1, n+k-1].$$

Note that $n + k - 1 = k^2$, so $\operatorname{spec}(A) = [k-1, k-1, \ldots, k-1, k^2]$. Now the eigenvalues of A are square roots of the eigenvalues of A^2. We know that k is an eigenvalue of A since every row sum in A is equal to k; this occurs once. Every other eigenvalue of A is either $\sqrt{k-1}$ or $-\sqrt{k-1}$. Say that $\sqrt{k-1}$ occurs r times and $-\sqrt{k-1}$ occurs s times, where $r + s = n - 1$. Finally we make use of the fact that $\operatorname{trace}(A) = 0$, which means

$$k + (r - s)\sqrt{(}k-1) = 0.$$

Rearranging this equation gives $k^2 = (s-r)^2(k-1)^2$, which means that $k-1$ divides k^2. Since $k-1$ also divides $k^2 - 1$, it follows that $k - 1 = 1$ which means that $k = 2$ and $n = k^2 - k + 1 = 3$.

In this case G is the graph K_3 (or W_1) consisting of a single triangle. This is the only regular graph satisfying the hypothesis of the theorem, and it also satisfies the conclusion (and it is a windmill). By the first half of the proof, every non-regular graph that possesses the friendship propery has a vertex adjacent to all others, so we have proved the theorem. □

The Friendship Theorem is a famous example of the use of matrix and specifically spectral techniques to solve a purely combinatorial problem. The proof here is essentially the original one of Erdös, Rényi and Sós. There are several proofs in the literature, most of which involve consideration of matrix spectra in some way. For many years there was interest in finding a "purely combinatorial" proof. Some do exist now in the literature, see for example "The Friendship Theorem" by Craig Huneke, in the February 2002 volume of the American Mathematical Monthly (avaiable on JSTOR). Another interesting feature of this theorem is that it is no longer true if the condition that G is finite is dropped - there exist examples of infinite "friendship graphs" with no politician.