

Advanced Linear Algebra MA500-1: Lecture Notes Semester 1 2016-2017

Dr Rachel Quinlan
School of Mathematics, Statistics and Applied Mathematics, NUI Galway

September 5, 2017

Contents

| | |
|--|-----------|
| Preface | 2 |
| 1 Matrices and Graphs | 2 |
| 1.1 The Adjacency Matrix | 2 |
| 1.2 Some matrix background | 8 |
| 2 Real Symmetric Matrices | 13 |
| 2.1 Special properties of real symmetric matrices | 13 |
| 2.2 Diagonalizability of symmetric matrices | 18 |
| 2.3 Connections to the adjacency spectrum | 21 |
| 3 The Laplacian Matrix of a Graph | 24 |
| 3.1 Introduction to the graph Laplacian | 24 |
| 3.2 Spanning Trees | 28 |
| 3.3 Connectivity and the graph Laplacian | 31 |
| 4 Strongly Regular Graphs | 35 |
| 4.1 Parameters and Properties | 35 |
| 4.2 The Adjacency Spectrum of a strongly regular graph | 37 |
| 4.3 Two classes of strongly regular graphs | 39 |

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 disconnect 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.



2. A directed graph and its adjacency matrix.



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 $(A^k)_{ij}$. \square

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

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

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 $\sum_{v \in V} \text{deg}(v)$.

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 $\text{trace}(A) = \lambda_1 + \lambda_2 + \dots + \lambda_n$. 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 $[-2, -1, -1, 1, 1, 1 + \sqrt{7}, 1 - \sqrt{7}]$.



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.



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 $\text{spec}(A_H) = [0, 0, 0, 2, -2]$.

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 $\text{spec}(A_G) = [0, 0, 0, 2, -2] = \text{spec}(A_H)$.


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 (Erdős-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.



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, \dots, 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_i 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 $\deg u \leq \deg v$. The same argument with the roles of u and v reversed shows that $\deg v \leq \deg u$, 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 :

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

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

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

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

$$\text{spec}(J) = [0, 0, \dots, 0, n] \implies \text{spec}(A^2) = [k - 1, k - 1, \dots, k - 1, n + k - 1].$$

Note that $n + k - 1 = k^2$, so $\text{spec}(A) = [k - 1, k - 1, \dots, k - 1, k^2]$. 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 $\text{trace}(A) = 0$, 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 property has a vertex adjacent to all others, so we have proved the theorem. \square

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 Erdős, 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 (available 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.

1.2 Some matrix background

The goal of this section is to fill in some details about matrices that were used implicitly or explicitly in Section 1.1, in particular the assertions about the trace of a square matrix A and its positive integer powers. For this we require the concept and meaning of similarity. Also in the background are the Rank-Nullity Theorem and the concept of an eigenvector, along with its interpretation for graphs.

First we revisit the process of matrix-vector multiplication. Let $A \in M_{m \times n}(\mathbb{R})$, let $v \in \mathbb{R}^n$ and let u in $(\mathbb{R}^m)^T$. Then

- Av is the column in \mathbb{R}^m that is the linear combination of the columns of A with the entries of v as coefficients.
- uA is the row in \mathbb{R}^n that is the linear combination of the rows of A with the entries of u as coefficients.

Definition 1.2.1. *If A is square, $A \in M_n(\mathbb{R})$, then a non-zero column v is an eigenvector of A if Av is a scalar multiple of v itself (the scalar that turns up here is the eigenvalue of A to which v corresponds).*

In this context, multiplication on the left by A determines a *function* from \mathbb{R}^n to \mathbb{R}^n . An eigenvector of A corresponds to a one-dimensional subspace of \mathbb{R}^n that is mapped into itself by A . In Section 1.1 however, we were considering eigenvectors of the adjacency matrix of a graph. The meaning of this has the following interpretation in graph theory.

Let G be a graph with adjacency matrix A . What is the meaning of an eigenvector of A in terms of the graph G ? For example

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

We saw in Section 1.1 that -2 is an eigenvalue of this $A(G)$. A corresponding eigenvector is $u = (0, 1, -1, 1, -1, 1, -1)$.

In the context of graphs and adjacency matrices, we can think of a column vector u as a function on the vertex set, that assigns a number u_i (the i th entry of u) to each vertex v_i . We can think of u_i as the value of u at vertex v_i . In this context the i th entry of the product $A(G)u_i$ is the sum of those u_j for which v_j is a neighbour of v_i , i.e.

$$(A(G)u)_i = \sum_{j:v_j \sim v_i} u_j.$$

If u is an eigenvector of $A(G)$ corresponding to the eigenvalue λ , it means that for every vertex v of G , the sum of the values of u at the neighbours of v is the value at v itself multiplied by λ . In the example above the function corresponding to u is

and it is easily checked for each vertex in this picture that the sum of the values labelling the neighbouring vertices is -2 multiplied by the value at the vertex itself. Thus the vector $u = (0, 1, -1, 1, -1, 1, -1)$ is an eigenvector of this graph corresponding to the eigenvalue -2 .

Definition 1.2.2. Let \mathbb{R}^n denote the vector space of column vectors of length n over \mathbb{R} . A linear transformation of \mathbb{R}^n is a function $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$ such that

- $T(u + v) = T(u) + T(v) \forall u, v \in \mathbb{R}^n$;
- $T(ku) = kT(u) \forall u \in \mathbb{R}^n, k \in \mathbb{R}$.

Let $\mathcal{C} = \{c_1, \dots, c_n\}$ be a basis of \mathbb{R}^n . Then every element of \mathbb{R}^n has a unique set of \mathcal{C} -coordinates, namely the coefficients of its expression as a linear combination of the elements of \mathcal{C} . Let T be a linear transformation of \mathbb{R}^n .

Definition 1.2.3. The matrix of T with respect to \mathcal{C} is the $n \times n$ matrix $A_{\mathcal{C}}$ whose j th column has the \mathcal{C} -coordinates of $T(c_j)$ as its entries.

If $v \in \mathbb{R}^n$ and $v_{\mathcal{C}}$ is the column vector whose entries are the \mathcal{C} -coordinates of v , then the matrix-vector product $A_{\mathcal{C}}v_{\mathcal{C}}$ is the column vector whose entries are the \mathcal{C} -coordinates of $T(v)$. This follows from the observation that this product is nothing but the linear combination of the columns of A with the entries of $v_{\mathcal{C}}$ as coefficients.

Example 1.2.4. Let T be the linear transformation of \mathbb{R}^2 with $T(e_1) = \begin{pmatrix} 1 \\ 1 \end{pmatrix}$ and $T(e_2) = \begin{pmatrix} -2 \\ 4 \end{pmatrix}$ (where $e_1 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and $e_2 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$ are the elements of the standard basis). The matrix of T with respect to the standard basis is

$$A = \begin{pmatrix} 1 & -2 \\ 1 & 4 \end{pmatrix},$$

and if $v = \begin{pmatrix} a \\ b \end{pmatrix}$ is any element of \mathbb{R}^2 , then

$$T(v) = T(ae_1 + be_2) = aT(e_1) + bT(e_2) = a \begin{pmatrix} 1 \\ 1 \end{pmatrix} + b \begin{pmatrix} -2 \\ 4 \end{pmatrix} = \begin{pmatrix} 1 & -2 \\ 1 & 4 \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix}.$$

Now let \mathcal{C} be the basis of \mathbb{R}^2 consisting of $c_1 = \begin{pmatrix} -2 \\ 1 \end{pmatrix}$ and $c_2 = \begin{pmatrix} 1 \\ -1 \end{pmatrix}$. Then

$$\begin{aligned} T(c_1) &= \begin{pmatrix} 1 & -2 \\ 1 & 4 \end{pmatrix} \begin{pmatrix} -2 \\ 1 \end{pmatrix} = \begin{pmatrix} -4 \\ 2 \end{pmatrix} = 2c_1 \implies [T(c_1)]_{\mathcal{C}} = \begin{pmatrix} 2 \\ 0 \end{pmatrix} \\ T(c_2) &= \begin{pmatrix} 1 & -2 \\ 1 & 4 \end{pmatrix} \begin{pmatrix} 1 \\ -1 \end{pmatrix} = \begin{pmatrix} 3 \\ -3 \end{pmatrix} = 3c_2 \implies [T(c_2)]_{\mathcal{C}} = \begin{pmatrix} 0 \\ 3 \end{pmatrix} \end{aligned}$$

The matrix $A_{\mathcal{C}}$ of T with respect to \mathcal{C} is $\begin{pmatrix} 2 & 0 \\ 0 & 3 \end{pmatrix}$.

Definition 1.2.5. Two matrices in $M_n(\mathbb{R})$ are similar if they represent the same linear transformation with respect to different bases of \mathbb{R}^n .

Two similar matrices can look quite different. When dealing with linear transformations, a general goal is to try to find a basis with respect to which the transformation is easy to describe. The best you can hope for is that there might be a basis consisting entirely of *eigenvectors* for the transformation, i.e. non-zero vectors v for which $T(v) = \lambda v$ for a real number λ (the corresponding eigenvalue of T). The matrix of T with respect to such a basis is diagonal; however such bases might not exist.

Definition 1.2.5 is the essential meaning of similarity, but it has a meaning in terms of matrix algebra also, which is equivalent and often useful. To figure this out, let $\mathcal{B} = \{b_1, \dots, b_n\}$ and $\mathcal{C} = \{c_1, \dots, c_n\}$ be different bases of \mathbb{R}^n , and let the matrix of T with respect to \mathcal{B} be B .

Let P be the matrix whose j th column contains the \mathcal{B} -coordinates of the vector c_j . Then P is invertible, and the \mathcal{B} -coordinates of any vector $v \in \mathbb{R}^n$ are given by the matrix-vector product $Pv_{\mathcal{C}}$, where $v_{\mathcal{C}}$ is the column consisting of the \mathcal{C} -coordinates of v . Thus, for any vector v in \mathbb{R}^n , $v_{\mathcal{B}} = Pv_{\mathcal{C}}$ and equivalently $v_{\mathcal{C}} = P^{-1}v_{\mathcal{B}}$.

Now the \mathcal{C} -coordinates of $T(v)$ may be found by first multiplying the column vector $v_{\mathcal{C}}$ on the left by P (to convert to \mathcal{B} -coordinates), then multiplying on the left by the matrix B (this gives the \mathcal{B} -coordinates of $T(v)$), then multiplying on the left by P^{-1} (to get back to \mathcal{C} -coordinates). This means that the matrix of T with respect to \mathcal{B} is $P^{-1}BP$, prompting the following arithmetic definition of similarity.

Definition 1.2.6. *Two matrices B and C in $M_n(\mathbb{R})$ are similar to each other if*

$$C = P^{-1}BP,$$

for some invertible $P \in M_n(\mathbb{R})$.

Definitions 1.2.5 and 1.2.6 are entirely equivalent to each other and it useful to be able to think of both of them together. The arithmetic version is particularly useful for proving some of the shared properties of similar matrices.

To connect to the content of Section 1.1, we first show that similar matrices have the same trace, which is a consequence of the fact that while the matrix products AB and BA are generally different, they always have the same trace.

Lemma 1.2.7. *Let A and B be matrices in $M_n(\mathbb{R})$. Then $\text{trace}(AB) = \text{trace}(BA)$.*

Proof. The trace of AB is the sum of the diagonal entries of AB . The entry in the (i, i) position of AB is the scalar product of Row i of A with Column i of B . Thus

$$\begin{aligned} \text{trace}(AB) &= \sum_{i=1}^n (AB)_{ii} \\ &= \sum_{i=1}^n \sum_{k=1}^n A_{ik}B_{ki} \\ &= \sum_{k=1}^n \sum_{i=1}^n B_{ki}A_{ik} \\ &= \sum_{k=1}^n (BA)_{kk} \\ &= \text{trace}(BA). \end{aligned}$$

□

Corollary 1.2.8. *Suppose that A and B are similar matrices in $M_n(\mathbb{R})$. Then $\text{trace}(A) = \text{trace}(B)$.*

Proof. Since A and B are similar, $B = P^{-1}AP$ for some invertible $P \in M_n(\mathbb{R})$. Then

$$\text{trace}(B) = \text{trace}(P^{-1}AP) = \text{trace}(APP^{-1}) = \text{trace}(A),$$

by Lemma 1.2.7. □

Our next goal is to show that the trace of a square matrix is the sum of its eigenvalues, something that we used in Section 1.1. We recall a few details about eigenvalues first.

1. The (possibly complex) number λ is an eigenvalue of the square matrix $A \in M_n(\mathbb{R})$ if and only if there exists a non-zero column vector v with $Av = \lambda v$, which means $(\lambda I - A)v = 0$.

2. This means that some non-zero linear combination of the columns of $\lambda I - A$ is the zero vector, which means exactly that these columns are linearly dependent, which occurs exactly if $\det(\lambda I - A) = 0$. So the eigenvalues of A are the roots of the polynomial $\det(\lambda I - A)$ (the characteristic polynomial of A). So similar matrices have the same spectrum.
3. The determinant is a *multiplicative* function on $M_n(\mathbb{R})$, which means that $\det(AB) = \det(A) \det(B)$ for $A, B \in M_n(\mathbb{R})$ (this is the Cauchy-Binet formula, definitely not an obvious thing). If P is an invertible matrix, then it follows from the Cauchy-Binet formula $\det(P) \det(P^{-1}) = \det I_n = 1$ and that A and $P^{-1}AP$ have the same characteristic polynomial for all $A \in M_n(\mathbb{R})$, i.e.

$$\begin{aligned} \det(\lambda I - P^{-1}AP) &= \det(\lambda P^{-1}IP - P^{-1}AP) = \det(P^{-1}(\lambda I - A)P) \\ &= \det(P^{-1}) \det(\lambda I - A) \det(P) = \det(\lambda I - A). \end{aligned}$$

4. Finally suppose that A is an upper triangular matrix (i.e. $A_{ij} = 0$ whenever $i > j$, all entries of A below the main diagonal are zero). Then the determinant of A is just the product of the entries on the main diagonal, and the characteristic polynomial of A is just the product over i of $(\lambda - A_{ii})$ - so the spectrum of A consists of the entries on the main diagonal.

For the next theorem we consider matrices over \mathbb{C} , the reason being that the field \mathbb{C} of complex numbers is *algebraically closed*, which means that every polynomial with complex coefficients has a full set of roots in \mathbb{C} . Note that every matrix in $M_n(\mathbb{R})$ is also in $M_n(\mathbb{C})$. All eigenvalues of a real matrix are complex, they are not necessarily all real.

Theorem 1.2.9. *Let $A \in M_n(\mathbb{C})$. Then A is similar in $M_n(\mathbb{C})$ to an upper triangular matrix in $M_n(\mathbb{C})$.*

Proof. By induction on n . The case $n = 1$ is clear, since every 1×1 matrix is upper triangular. Let T be the linear transformation of \mathbb{C}^n determined by $T(v) = Av$, for $v \in \mathbb{C}^n$. Let λ_1 be an eigenvalue of A in \mathbb{C} with corresponding eigenvector v_1 . Expand $\{v_1\}$ to a basis $\{v_1, v_2, \dots, v_n\}$ of \mathbb{C}^n . Then the matrix A_1 of T with respect to this basis has λ_1 in its top left entry and zeros otherwise in its first column, write this matrix as

$$A_1 = \left(\begin{array}{c|ccc} \lambda_1 & * & \dots & * \\ 0 & & & \\ \vdots & & A' & \\ 0 & & & \end{array} \right),$$

where $A' \in M_{n-1}(\mathbb{C})$. By the induction hypothesis, there exists an invertible matrix $Q \in M_{n-1}(\mathbb{C})$ for which $T' = Q^{-1}A'Q$ is upper triangular. Write

$$P = \left(\begin{array}{c|ccc} 1 & 0 & \dots & 0 \\ 0 & & & \\ \vdots & & Q & \\ 0 & & & \end{array} \right).$$

Then

$$P^{-1} = \left(\begin{array}{c|ccc} 1 & 0 & \dots & 0 \\ 0 & & & \\ \vdots & & Q^{-1} & \\ 0 & & & \end{array} \right),$$

and

$$\begin{aligned}
P^{-1}A_1P &= \left(\begin{array}{c|ccc} 1 & 0 & \dots & 0 \\ \hline 0 & & & \\ \vdots & & Q^{-1} & \\ 0 & & & \end{array} \right) \left(\begin{array}{c|ccc} \lambda_1 & * & \dots & * \\ \hline 0 & & & \\ \vdots & & A' & \\ 0 & & & \end{array} \right) \left(\begin{array}{c|ccc} 1 & 0 & \dots & 0 \\ \hline 0 & & & \\ \vdots & & & \\ 0 & & & Q \end{array} \right) \\
&= \left(\begin{array}{c|ccc} \lambda_1 & \square & \dots & \square \\ \hline 0 & & & \\ \vdots & & Q^{-1}A'Q & \\ 0 & & & \end{array} \right) = \left(\begin{array}{c|ccc} \lambda_1 & \square & \dots & \square \\ \hline 0 & & & \\ \vdots & & T' & \\ 0 & & & \end{array} \right).
\end{aligned}$$

Thus A_1 , and hence A , is similar to an upper triangular matrix. \square

Finally we are in a position to prove the following statement.

Theorem 1.2.10. *Let $A \in M_n(\mathbb{R})$, and let $\text{spec}(A) = [\lambda_1, \dots, \lambda_n]$ (a multiset of complex numbers). For every positive integer k ,*

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

Proof. Let T be an upper triangular matrix similar to A in $M_n(\mathbb{C})$. Since A and T have the same spectrum, the diagonal entries of T are $\lambda_1, \dots, \lambda_n$ (in some order). Since A and T have the same trace we have $\text{trace}(A) = \sum \lambda_i$.

Since $A = P^{-1}TP$ for some invertible $P \in M_n(\mathbb{C})$ we have

$$A^k = (PP^{-1}TP)(P^{-1}TP) \dots (P^{-1}TP) = P^{-1}T^kP,$$

so A^k is similar to T^k . Using the mechanism of matrix multiplication and the fact that T is upper triangular, it is straightforward to see that the diagonal entries of T^k are the k th powers of the corresponding diagonal entries of T . Thus

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

\square

Chapter 2

Real Symmetric Matrices

2.1 Special properties of real symmetric matrices

A matrix $A \in M_n(\mathbb{C})$ (or $M_n(\mathbb{R})$) is *diagonalizable* if it is similar to a diagonal matrix. If this happens, it means that there is a basis of \mathbb{C}^n with respect to which the linear transformation of \mathbb{C}^n defined by left multiplication by A has a diagonal matrix. Every element of such a basis is simply multiplied by a scalar when it is multiplied by A , which means exactly that the basis consists of eigenvectors of A .

Lemma 2.1.1. *A matrix $A \in M_n(\mathbb{C})$ is diagonalizable if and only if \mathbb{C}^n possesses a basis consisting of eigenvectors of A .*

Not all matrices are diagonalizable. For example $A = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$ is not. To see this note that 1 (occurring twice) is the only eigenvalue of A , but that all eigenvectors of A are scalar multiples of $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$, so \mathbb{C}^2 (or \mathbb{R}^2) does not contain a basis consisting of eigenvectors of A , and A is not similar to a diagonal matrix.

We note that a matrix can fail to be diagonalizable only if it has repeated eigenvalues, as the following lemma shows.

Lemma 2.1.2. *Let $A \in M_n(\mathbb{R})$ and suppose that A has distinct eigenvalues $\lambda_1, \dots, \lambda_n$ in \mathbb{C} . Then A is diagonalizable.*

Proof. Let v_i be an eigenvector of A corresponding to λ_i . We will show that $S = \{v_1, \dots, v_n\}$ is a linearly independent set. Since v_1 is not the zero vector, we know that $\{v_1\}$ is linearly independent. If S is linearly dependent, let k be the least for which $\{v_1, \dots, v_k\}$ is linearly dependent. This means that $\{v_1, \dots, v_{k-1}\}$ is a linearly independent set and

$$v_k = a_1 v_1 + \dots + a_{k-1} v_{k-1}$$

for some $a_i \in \mathbb{C}$, not all zero. Multiplying this equation on the left separately by A and by λ_k gives

$$\begin{aligned} \lambda_k v_k &= a_1 \lambda_1 v_1 + a_2 \lambda_2 v_2 + a_{k-1} \lambda_{k-1} v_{k-1} \\ a_1 \lambda_k v_1 + a_2 \lambda_k v_2 + a_{k-1} \lambda_k v_{k-1} \\ \implies 0 &= a_1 (\lambda_1 - \lambda_k) v_1 + a_2 (\lambda_2 - \lambda_k) v_2 + \dots + a_{k-1} (\lambda_{k-1} - \lambda_k) v_{k-1}. \end{aligned}$$

Since the complex numbers $\lambda_i - \lambda_k$ are non-zero for $i = 1, \dots, k-1$ and at least one of these a_i is non-zero, the above is an expression for the zero vector as a nontrivial linear combination of v_1, \dots, v_{k-1} , contrary to the choice of k . We conclude that S is linearly independent and hence that it is a basis of \mathbb{C}^n . \square

Definition 2.1.3. A matrix $A \in M_n(\mathbb{R})$ is symmetric if it is equal to its transpose, i.e. if $A_{ij} = A_{ji}$ for all i and j .

Symmetric matrices arise naturally in various contexts, including as adjacency matrices of undirected graphs. Fortunately they have lots of nice properties. To explore some of these we need a slightly more general concept, that of a complex Hermitian matrix.

Definition 2.1.4. Let $A \in M_n(\mathbb{C})$. The Hermitian transpose, or conjugate transpose of A is the matrix A^* obtained by taking the transpose of A and then taking the complex conjugate of each entry. The matrix A is said to be Hermitian if $A = A^*$.

Notes

1. Example: If $A = \begin{pmatrix} 2+i & 4-i \\ 3 & 3-i \end{pmatrix}$, then $A^* = \begin{pmatrix} 2-i & 3 \\ 4+i & 3+i \end{pmatrix}$
2. The Hermitian transpose of A is equal to its (ordinary) transpose if and only if $A \in M_n(\mathbb{R})$. In some contexts the Hermitian transpose is the appropriate analogue in \mathbb{C} of the concept of transpose of a real matrix.
3. If $A \in M_n(\mathbb{C})$, then the trace of the product A^*A is the sum of all the entries of A , each multiplied by its own complex conjugate (check this). This is a non-negative real number and it is zero only if $A = 0$. In particular, if $A \in M_n(\mathbb{R})$, then $\text{trace}(A^T A)$ is the sum of the squares of all the entries of A .
4. Suppose that A and B are two matrices for which the product AB exists. Then $(AB)^T = B^T A^T$ and $(AB)^* = B^* A^*$ (it is routine but worthwhile to prove these statements). In particular, if A is any matrix at all, then

$$(A^T A)^T = A^T (A^T)^T = A^T A, \text{ and } (A^* A)^* = A^* (A^*)^* = A^* A,$$

so $A^T A$ and $A^* A$ are respectively symmetric and Hermitian (so are AA^T and AA^*).

The following theorem is the start of the story of what makes real symmetric matrices so special.

Theorem 2.1.5. The eigenvalues of a real symmetric matrix are all real.

Proof. We will prove the stronger statement that the eigenvalues of a complex Hermitian matrix are all real. Let A be a Hermitian matrix in $M_n(\mathbb{C})$ and let λ be an eigenvalue of A with corresponding eigenvector v . So $\lambda \in \mathbb{C}$ and v is a non-zero vector in \mathbb{C}^n . Look at the product $v^* A v$. This is a complex number.

$$v^* A v = v^* \lambda v = \lambda v^* v.$$

The expression $v^* v$ is a positive real number, since it is the sum of the expressions $\bar{v}_i v_i$ over all entries v_i of v .

We have not yet used the fact that $A^* = A$.

Now look at the Hermitian transpose of the matrix product $v^* A v$.

$$(v^* A v)^* = v^* A^* (v^*)^* = v^* A v.$$

This is saying that $v^* A v$ is a complex number that is equal to its own Hermitian transpose, i.e. equal to its own complex conjugate. This means exactly that $v^* A v \in \mathbb{R}$.

We also know that $v^* A v = \lambda v^* v$, and since $v^* v$ is a non-zero real number, this means that $\lambda \in \mathbb{R}$. □

So the eigenvalues of a real symmetric matrix are real numbers. This means in particular that the eigenvalues of the adjacency matrix of an undirected graph are real numbers, they can be arranged in order and we can ask questions about (for example) the greatest eigenvalue, the least eigenvalue, etc.

Another concept that is often mentioned in connection with real symmetric matrices is that of positive definiteness. We mentioned above that if $A \in M_{m \times n}(\mathbb{R})$, then $A^T A$ is a symmetric matrix. However not every symmetric matrix has the form $A^T A$, since for example the entries on the main diagonal of $A^T A$ do not. It turns out that those symmetric matrices that have the form $A^T A$ (even for a non-square A) can be characterized in another way.

Definition 2.1.6. Let A be a symmetric matrix in $M_n(\mathbb{R})$. Then A is called positive semidefinite (PSD) if $v^T A v \geq 0$ for all $v \in \mathbb{R}^n$. In addition, if $v^T A v$ is strictly positive whenever $v \neq 0$, then A is called positive definite (PD).

Notes

1. The identity matrix I_n is the classical example of a positive definite symmetric matrix, since for any $v \in \mathbb{R}^n$, $v^T I_n v = v^T v = v \cdot v \geq 0$, and $v \cdot v = 0$ only if v is the zero vector.
2. The matrix $\begin{pmatrix} 1 & 2 \\ 2 & 1 \end{pmatrix}$ is an example of a matrix that is *not* positive semidefinite, since

$$\begin{pmatrix} -1 & 1 \end{pmatrix} \begin{pmatrix} 1 & 2 \\ 2 & 1 \end{pmatrix} \begin{pmatrix} -1 \\ 1 \end{pmatrix} = -2.$$

So positive (semi)definite is not the same thing as positive - a symmetric matrix can have all of its entries positive and still fail to be positive (semi)definite.

3. A symmetric matrix can have negative entries and still be positive definite, for example the matrix $A = \begin{pmatrix} 1 & -1 \\ -1 & 2 \end{pmatrix}$ is SPD (symmetric positive definite). To see this observe that for real numbers a and b we have

$$\begin{pmatrix} a & b \end{pmatrix} \begin{pmatrix} 1 & -1 \\ -1 & 2 \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix} = a^2 - ab - ab + 2b^2 = (a - b)^2 + b^2.$$

Since $(a - b)^2 + b^2$ cannot be negative and is zero only if both a and b are equal to zero, the matrix A is positive definite.

The importance of the concept of positive definiteness is not really obvious at first glance, it takes a little bit of discussion. We will defer this discussion for now, and mention two observations related to positive (semi)definiteness that have a connection to spectral graph theory.

Lemma 2.1.7. The eigenvalues of a real symmetric positive semidefinite matrix are non-negative (positive if positive definite).

Proof. Let λ be an eigenvalue of the real symmetric positive semidefinite matrix A , and let $v \in \mathbb{R}^n$ be a corresponding eigenvector. Then

$$0 \leq v^T A v = v^T \lambda v = \lambda v^T v.$$

Thus λ is nonnegative since $v^T v$ is a positive real number. □

Lemma 2.1.8. Let $B \in M_{n \times m}(\mathbb{R})$ for some positive integers m and n . Then the symmetric matrix $A = B B^T$ in $M_n(\mathbb{R})$ is positive semidefinite.

Proof. Let $u \in \mathbb{R}^n$. Then


$$u^T A u = u^T B B^T u = (u^T B)(B^T u) = (B^T u)^T (B^T u) = (B^T u) \cdot (B^T u) \geq 0,$$

so A is positive semidefinite. □

The next section will contain a more detailed discussion of positive (semi)definiteness, including the converses of the two statements above. First we digress to look at an application of what we know so far to spectral graph theory.

Definition 2.1.9. Let G be a graph. The line graph of G , denoted by $L(G)$, has a vertex for every edge of G , and two vertices of $L(G)$ are adjacent if and only if their corresponding edges in G share an incident vertex.

Example 2.1.10. K_4 (left) and its line graph (right).



Choose an edge of K_4 . Since each of its incident vertices has degree 3, there are four other edges with which it shares a vertex. So the vertex that represents it in $L(K_4)$ has degree 4. In general, if a graph G is regular of degree k , then $L(G)$ will be regular of degree $2k - 2$. For any graph G , a vertex of degree d in G corresponds to a copy of the complete graph K_d within $L(G)$.

Not every graph can be a line graph. For example, a vertex of degree 3 in a line graph $L(G)$ must have the property that at least two of its neighbours are adjacent to each other, because it corresponds to an edge e of the graph G that shares a vertex with three other edges. At least two of these three must be incident with the same vertex of e . Thus (for example) $L(G)$ can be a tree or forest only if $L(G)$ has no vertex of degree exceeding 2, which means that G is a collection of disjoint paths (in this case $L(G)$ is also a collection of disjoint paths, of lengths one less than the paths of G itself). The cycle C_n is its own line graph. The line graph of the path P_n is P_{n-1} . The line graph of the star on n vertices (which has one vertex of degree $n - 1$ and $n - 1$ of degree 1) is the complete graph K_n .

Line graphs all share the following spectral property, which is remarkable easy to prove using what we already know about positive semidefinite matrices.


Theorem 2.1.11. Let $L(G)$ be the line graph of a graph G , and let $A(L(G))$ be the adjacency matrix of $L(G)$. Then every eigenvalue of $L(G)$ is at least equal to -2 .

To prove Theorem 2.1.14 we need one more device that links matrices to graphs.

Definition 2.1.12. Let G be a graph with n vertices and m edges. The incidence matrix of G , denoted $B(G)$, is the $m \times n$ $(0,1)$ -matrix with rows indexed by the edges of G and columns by the vertices of G , that has a 1 in the (i, j) -position if and only if the edge labelling Row i is incident with the vertex labelling Column j .

The incidence matrix depends on a choice of ordering of both the vertices and the edges obviously.

Example 2.1.13. A graph and its incidence matrix.



Each row of an incidence matrix has two 1s, and the number of 1s in a column is the degree of the corresponding vertex.

Now suppose that B is the incidence matrix of a graph G , and consider the positive semidefinite matrices BB^T and B^TB .

The rows and columns of BB^T are indexed by the edges of G . The entry in the (i, j) position of BB^T is the scalar product of Rows i and j of B , each of which has exactly two entries equal to 1. If $i = j$ then the entry in the (i, i) position of BB^T is 2. If $i \neq j$, then Rows i and j of B are different since they represent different edges e_i and e_j respectively of G . In this case $(BB^T)_{ij}$ is equal to 1 if the edges e_i and e_j have a vertex in common, and 0 otherwise. Thus an off-diagonal entries of BB^T is 1 or 0 according as the edges of G labelling its row and column share an incident vertex or not. The diagonal entries are all 2. Thus $BB^T - 2I$ is exactly the adjacency matrix of the line graph of G , or

$$BB^T = 2I + A(L(G)).$$

Theorem 2.1.14. *Let $L(G)$ be the line graph of a graph G , and let λ be the least eigenvalue of the adjacency matrix of $L(G)$. The $\lambda \geq -2$.*

Proof. From the above description of BB^T we know that $2I + A(L(G))$ is a positive semidefinite matrix and so its eigenvalues are all non-negative. Moreover the spectrum of $2I + A(L(G))$ are obtained by adding 2 to each item in the spectrum of $A(L(G))$, so $\lambda + 2 \geq 0 \implies \lambda \geq -2$. \square

It is not true unfortunately that a graph must be a line graph if all eigenvalues of its adjacency matrix are at least -2 .

Now we turn to the matrix B^TB . The rows and columns of this matrix are labelled by the vertices of G and the entry in the (i, j) position is the scalar product of Column i and Column j of B . If $i = j$, this is the degree of Vertex i . If $i \neq j$, then Column i and Column j have a 1 in the same position if and only if Vertex i and Vertex j belong to the same edge. This happens in exactly one position if the vertices i and j are adjacent in G , and in no position if they are non-adjacent. Thus the entry in the off-diagonal position (i, j) of B^TB is 1 if Vertices i and j are adjacent in G and 0 otherwise. This means that away from the main diagonal, B^TB coincides with the adjacency matrix of G . Putting all of this together gives

$$B^TB = \Delta + A(G),$$

where $A(G)$ is the adjacency matrix of G and Δ is the diagonal matrix whose entry in the (i, i) -position is the degree of vertex v_i . We have shown the the matrix $\Delta + A(G)$ is positive semidefinite for every graph G . In the special case where G is regular of degree k , this shows that every eigenvalue of G is at least $-k$.

2.2 Diagonalizability of symmetric matrices

The main theorem of this section is that every real symmetric matrix is not only diagonalizable but *orthogonally* diagonalizable. Two vectors u and v in \mathbb{R}^n are *orthogonal* to each other if $u \cdot v = 0$ or equivalently if $u^T v = 0$. This is sometimes written as $u \perp v$. A matrix A in $M_n(\mathbb{R})$ is called *orthogonal* if

- $u \cdot v = 0$ if u and v are distinct columns of A (the columns of A are pairwise orthogonal to each other), and
- $u \cdot u = 1$ for each column u of A (each column of A is a vector of length 1 in \mathbb{R}^n).

Another way to say this is that the columns of A form an *orthonormal basis* of \mathbb{R}^n , which means a basis consisting of mutually orthogonal unit vectors. Note that for any matrix $B \in M_{m \times n}(\mathbb{R})$, $B^T B$ is the $n \times n$ matrix whose entry in the (i, j) position is the scalar product of Columns i and j of B . Putting this together with the above definition of an orthogonal matrix, it is saying that the square matrix $A \in M_n(\mathbb{R})$ is orthogonal if and only if

$$(A^T A)_{ij} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases} ,$$

i.e. $A \in M_n(\mathbb{R})$ is orthogonal if and only if $A^T A = I_n$.

Definition 2.2.1. A matrix in $M_n(\mathbb{R})$ is orthogonal if and only if its inverse is equal to its transpose.

We note that the set of orthogonal matrices in $M_n(\mathbb{R})$ forms a group under multiplication, called the orthogonal group and written $O_n(\mathbb{R})$. The use of the term “orthogonal” for square matrices differs from its use for vectors - a vector can’t just be orthogonal, it can be orthogonal to another vector, but a matrix can be orthogonal by itself. An example of an orthogonal matrix in $M_2(\mathbb{R})$ is $\begin{pmatrix} 1/2 & -\sqrt{3}/2 \\ \sqrt{3}/2 & 1/2 \end{pmatrix}$.

The following is our main theorem of this section.

Theorem 2.2.2. Let A be a symmetric matrix in $M_n(\mathbb{R})$. Then there exists an orthogonal matrix P for which $P^T A P$ is diagonal.

Note that this is saying that \mathbb{R}^n has a basis consisting of eigenvectors of A that are all orthogonal to each other, something that is true only for symmetric matrices. If we have a basis consisting of orthogonal eigenvectors, we can normalize its elements so that our basis consists of unit vectors as required. After we prove Theorem 2.2.2 we will deduce some consequences about positive (semi)definiteness and then look at some applications to graph spectra in the next section.

The following theorem is one of the two keys to the proof of Theorem 2.2.2, and it takes care of the case where the eigenvalues of A are distinct.

Theorem 2.2.3. Let A be a real symmetric matrix. Let λ and μ be distinct eigenvalues of A , with respective eigenvectors u and v in \mathbb{R}^n . Then $u^T v = 0$.

Note that $u^T v$ is just the ordinary scalar product of u and v (u^T is just u written as a row). So this theorem is saying that eigenvectors of a real symmetric matrix that correspond to different eigenvalues are orthogonal to each other under the usual scalar product.

Proof. The matrix product $u^T A v$ is a real number (a 1×1 matrix). We can write

$$u^T A v = u^T \mu v = \mu u^T v.$$

On the other hand, being a 1×1 matrix, $u^T A v$ is equal to its own transpose, so

$$u^T A v = (u^T A v)^T = v^T A^T (u^T)^T = v^T A u = v^T \lambda u = \lambda v^T u.$$

Now $v^T u = u^T v$ since both are equal to the scalar product $u \cdot v$ (or because they are 1×1 matrices that are transposes of each other). So what we are saying is

$$\mu u^T v = \lambda u^T v.$$

Since $\mu \neq \lambda$, it follows that $u^T v = 0$. □

From Theorem 2.2.3 and Lemma 2.1.2, it follows that if the symmetric matrix $A \in M_n(\mathbb{R})$ has distinct eigenvalues, then $A = P^{-1}AP$ (or P^TAP) for some orthogonal matrix P . It remains to consider symmetric matrices with repeated eigenvalues. We need a few observations relating to the ordinary scalar product on \mathbb{R}^n .

Definition 2.2.4. Let U be a subspace of \mathbb{R}^n . Then the orthogonal complement of U , denoted U^\perp , is defined by

$$U^\perp = \{v \in \mathbb{R}^n : v \cdot u = 0 \forall u \in U\}.$$

Notes

1. For example, if $U = \langle e_1, e_2 \rangle$ in \mathbb{R}^n , then $U^\perp = \langle e_3, \dots, e_n \rangle$.
2. It is easily checked that U^\perp is a subspace of \mathbb{R}^n , not just a subset.
3. For any subspace U of \mathbb{R}^n , $U \cap U^\perp = \{0\}$, since element of $U \cap U^\perp$ must be orthogonal to itself under the usual scalar product. However the scalar product of any non-zero vector in \mathbb{R}^n with itself is the sum of the squares of its entries, which is a positive real number.
4. Suppose that U has dimension k and let $\{u_1, \dots, u_k\}$ be a basis of u . Let A_U be the $k \times n$ matrix that has u_1^T, \dots, u_k^T as its k rows. Then A_U has rank k since its rows are linearly independent, and by definition U^\perp is just the right nullspace of A_U . It follows from the rank-nullity theorem that the dimension of U^\perp is $n - k$.
5. Suppose that $\{u_1, \dots, u_k\}$ is a linearly independent set of vectors in \mathbb{R}^n whose elements are mutually orthogonal, so that $u_i \cdot u_j = 0$ whenever $i \neq j$. Let $U = \langle u_1, \dots, u_k \rangle$. If $k < n$, let $v_{k+1} \in U^\perp$ and note that $\{u_1, \dots, u_k, v_{k+1}\}$ is a linearly independent set, since $v_{k+1} \notin U$. If the span of these $k+1$ elements is still not all of \mathbb{R}^n , we can add an element of $\langle u_1, \dots, u_k, v_{k+1} \rangle^\perp$ to obtain a larger linearly independent set of mutually orthogonal vectors in \mathbb{R}^n . Continuing in this way we can extend $\{u_1, \dots, u_k\}$ to a basis of \mathbb{R}^n consisting of mutually orthogonal elements (we can normalize these if we wish to obtain an orthonormal basis). We have the following useful fact: *any linearly independent set of mutually orthogonal unit vectors in \mathbb{R}^n can be extended to an orthonormal basis of \mathbb{R}^n .*

The following lemma is the last ingredient needed for the proof of Theorem 2.2.2. This lemma would not be true without the hypothesis that A is symmetric. When you are studying the proof, make sure that you are attentive to how the symmetry of A is used. Note the statement that U is A -invariant means that $Au \in U$ whenever $u \in U$.

Lemma 2.2.5. Let $A \in M_n(\mathbb{R})$ be symmetric and suppose that U is an A -invariant subspace of \mathbb{R}^n . Then U^\perp is also A -invariant.

Proof. Suppose that $v \in U^\perp$. We need to show that $Av \in U^\perp$ also, i.e. that $u^T Av = 0$ for all $u \in U$. So let $u \in U$ and observe that

$$(u^T Av)^T = v^T A^T u = v^T Au.$$

Since $Au \in U$ and $v \in U^\perp$, we know that $v^T Au = 0$. Thus $u^T Av = 0$ also, for all $u \in U$. This means exactly that $Av \in U^\perp$, as required. □

We are now ready to complete the proof of Theorem 2.2.2.

Proof. The proof proceeds by induction on n , but Lemma 2.2.5 is the key ingredient. The case $n = 1$ is trivial, since all 1×1 matrices are diagonal.

Let $\lambda_1, \dots, \lambda_k$ be the *distinct* eigenvalues of A , and let u_i be an eigenvector (of length 1) corresponding to λ_i . Note that $k \geq 1$ since A has at least one eigenvalue. If $k = n$, then by Theorem 2.2.3 and Lemma 2.1.2, there is nothing to do. So we assume that $k < n$ and write $U = \langle u_1, \dots, u_k \rangle \subseteq \mathbb{R}^n$. Then U is A -invariant, since Au_i is a scalar multiple of u_i for each i . Moreover, the u_i are mutually orthogonal by Theorem 2.2.3, and $\dim U = k$ by Lemma 2.1.2.

Now as in item 5. in the notes above, we can extend $\{u_1, \dots, u_k\}$ to an orthonormal basis $\{u_1, \dots, u_k, v_{k+1}, \dots, v_n\}$, where $U^\perp = \{v_{k+1}, \dots, v_n\}$. Let Q be the orthogonal matrix whose columns are $u_1, \dots, u_k, v_{k+1}, \dots, v_n$. Then $Q^{-1}AQ$ is symmetric, since $Q^{-1} = Q^T$. Moreover, because u_1, \dots, u_k are eigenvectors of A and because U^\perp is A -invariant, the matrix $Q^T A Q$ has $\lambda_1, \dots, \lambda_k$ in the first k diagonal positions, has a symmetric $(n-k) \times (n-k)$ block A_1 in the lower right, and is otherwise full of zeros.

By the induction hypothesis, there exists an orthogonal matrix $Q_1 \in M_{n-k}(\mathbb{R})$ for which $Q_1^{-1}A_1Q_1$ is diagonal. Let $P \in M_n(\mathbb{R})$ be the orthogonal matrix that has I_k in the upper left $k \times k$ block, Q_1 in the lower right $(n-k) \times (n-k)$ block, and zeros elsewhere. Then

$$P^{-1}Q^{-1}AQP = (QP)^{-1}A(QP)$$

is diagonal. Moreover QP is orthogonal since

$$(QP)^{-1} = P^{-1}Q^{-1} = P^T Q^T = (QP)^T.$$

So A is orthogonally diagonalizable as required. \square

Two consequences of Theorem 2.2.2 are the following two characterizations of symmetric positive semidefinite matrices.

Theorem 2.2.6. *Let A be a symmetric matrix in $M_n(\mathbb{R})$. Then the following conditions are equivalent.*

1. A is positive semidefinite.
2. All eigenvalues of A are non-negative.
3. $A = BB^T$ for some $B \in M_n(\mathbb{R})$.

We have seen some of the implications of this theorem already in Section 2.1, where we proved that 1. \implies 2 and 3. \implies 1. We complete the proof by using Theorem 2.2.2 to show that 2. \implies 3.

Proof. First assume 2., that the eigenvalues $\lambda_1, \dots, \lambda_n$ of A are all non-negative. Then, by Theorem 2.2.2, the matrix $D = \text{diag}(\lambda_1, \dots, \lambda_n)$ satisfies

$$D = P^T A P,$$

for some orthogonal matrix $A \in M_n(\mathbb{R})$. Then $A = P D P^T$. Let D_1 be the diagonal matrix in $M_n(\mathbb{R})$ whose diagonal entries are the non-negative square roots in \mathbb{R} of $\lambda_1, \dots, \lambda_n$. Then D_1 is symmetric and $D_1^2 = D$. We use this to deduce 3. as follows:

$$A = P D P^T = P (D_1)^2 P^T = (P D_1) (D_1 P^T) = (P D_1) (D_1^T P^T) = (P D_1) (P D_1)^T.$$

Thus A satisfies 3., and we now have the implications 1. \implies 2., 2. \implies 3. and 3. \implies 1, which means that any of the three conditions of Theorem 2.2.6 follows from any of the others. \square

We will look at some consequences for graphs in the next section.

2.3 Connections to the adjacency spectrum

In this section we consider some of the consequences for spectral graph theory of the properties of symmetric and positive semidefinite real matrices that were established in Section 2.2. First we consider the connection between the adjacency spectrum of a regular graph and that of its complement.


Recall that a graph is *regular* if all of its vertices have the same degree. Also, the complement of a graph G is the graph \bar{G} that has the same vertex set as G and whose edges are precisely the non-edges of G . The adjacency matrix of \bar{G} has 1s precisely in the off-diagonal positions where the adjacency matrix of G has zeros. So the adjacency matrices of any graph and its complement are related by the equation

$$A(G) + A(\bar{G}) + I_n = J,$$

where as usual J denotes the matrix in which every entry is 1.

If G is regular of degree k , then \bar{G} is regular of degree $n - 1 - k$, where n is the order (number of vertices) of G .

Example



Note that if G is a k -regular graph, then k is an eigenvalue of $A(G)$, corresponding the eigenvector whose entries are all 1. This follows from the fact that the sum of the entries in each row of $A(G)$ is k .

Theorem 2.3.1. *Let G be a k -regular graph of order n , and let the spectrum of $A(G)$ be $[k, \theta_2, \dots, \theta_n]$. Then the spectrum of $A(\bar{G})$ is $[n - k - 1, -1 - \theta_2, \dots, -1 - \theta_n]$. Furthermore $A(G)$ and $A(\bar{G})$ have the same eigenvectors.*

Outline of proof: By Theorem 2.2.2 we may choose a basis $\{v_1, \dots, v_n\}$ consisting of \mathbb{R}^n of mutually orthogonal eigenvectors of $A(G)$, where v_1 is the all-1 vector (corresponding to the eigenvalue k) and v_i corresponds to θ_i for $i \geq 2$. Note that $v_i \perp v_1$, in particular this means that $Jv_i = 0$ for $i \geq 2$. Now consider the product


$$A(\bar{G})v_i = (J - A(G) - I_n)v_i.$$

Note: this is Problem 9 in Problem Sheet 1.

Now we consider how the adjacency spectrum of a graph G relates to the spectra of some of some of its subgraphs. We recall some definitions.

Definition 2.3.2. *Let G be a graph with vertex set V and edge set E . A subgraph of G is a graph whose vertex set is a subset of V and whose edge set is a subset of E . An induced subgraph of G is a subgraph H whose edge set consists of all edges of G that involve two vertices of H .*

Example A graph G , a (non-induced) subgraph H_1 and an induced subgraph H_2 .



If H is an *induced subgraph* of a graph G , then the adjacency matrix of H consists of the entries of the rows and columns of $A(G)$ that label those vertices that belong to H . These form a *principal submatrix* of $A(G)$. In general a principal submatrix of a square matrix M is a square submatrix whose main diagonal coincides with that of M . In the case of adjacency matrices, principal submatrices correspond to induced subgraphs. If H' is *any* subgraph of G , then the adjacency matrix of H' is obtained from the relevant principal submatrix of $A(G)$ by possibly replacing some symmetrically opposite pairs of 1s with zeros. The adjacency matrices of the graphs G , H_1 and H_2 (corresponding to a vertex labelling that starts in the top right and proceeds anticlockwise) are given below.

$$A(G) = \begin{pmatrix} 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 \\ 1 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 & 1 \\ 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 \end{pmatrix}, \quad A(H_1) = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \end{pmatrix}, \quad A(H_2) = \begin{pmatrix} 0 & 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & 0 & 1 & 0 \end{pmatrix}.$$

The following lemma notes a useful property of positive semidefinite matrices.

Lemma 2.3.3. *Suppose that A is a symmetric PSD matrix. Then every principal submatrix of A is PSD. (If A is positive definite then every principal submatrix of A is positive definite).*

Proof. Let A_1 be the principal submatrix of A consisting determined by Rows and Columns i_1, i_2, \dots, i_k . Let V be the subspace of \mathbb{R}^n consisting of all those column vectors that have zeros outside of positions i_1, \dots, i_k . Then for every $v \in V$, let v_1 denote the vector in \mathbb{R}^k whose entries are the entries from positions i_1, \dots, i_k of v . Then for $v \in V$ we have

$$v^T A v = v_1^T A_1 v_1.$$

Since $v^T A v \geq 0$ for all $v \in V$, it follows that $v_1^T A_1 v_1 \geq 0$ for all $v_1 \in \mathbb{R}^k$, hence that A_1 is positive semidefinite as required. \square

Note that a particular consequence of Lemma 2.3.3 is that the diagonal entries of a symmetric PSD matrix must be non-negative (positive if the matrix is PD).

For any simple undirected graph G , the eigenvalues of $A(G)$ are real numbers. We write $\lambda_{\max}(G)$ for the maximum eigenvalue of G , and $\lambda_{\min}(G)$ for the minimum eigenvalue of G . The following theorem is related to Lemma 2.3.3.

Theorem 2.3.4. *Let H be an induced subgraph of order k of a graph G of order n . Then*

$$\lambda_{\min}(G) \leq \lambda_{\min}(H) \leq \lambda_{\max}(H) \leq \lambda_{\max}(G).$$

Proof. Write A for the adjacency matrix of G and θ and μ respectively for the maximum and minimum eigenvalues of A . Suppose that σ is an eigenvalue of the matrix $\theta I_n - A$. Then

$$(\theta I - A)v = \sigma v \implies Av = (\theta - \sigma)v$$

for some non-zero $v \in \mathbb{R}^n$. So the eigenvalues of $\theta I_n - A$ are obtained by subtracting the eigenvalues of A from θ , hence they are all non-negative and $\theta I_n - A$ is positive semidefinite. It follows that $\theta I_k - A(H)$ is positive semidefinite, since it is a principal submatrix of $\theta I_n - A$. This means that $\theta - \rho \geq 0$ for every eigenvalue ρ of $A(H)$, so $\lambda_{\max}(H) \leq \theta$.

On the other hand every eigenvalue of $A - \mu I$ is of the form $\sigma - \mu$ for some eigenvalue σ of A , and is therefore non-negative. Thus $A - \mu I_n$ is a positive semidefinite matrix and so is its principal submatrix $A(H) - \mu I_k$, which means that $\rho - \mu \geq 0$ for every eigenvalue ρ of $A(H)$, and in particular $\lambda_{\min}(A(H)) \geq \mu$. \square

In fact it is true for any subgraph H of G that

$$\lambda_{\min}(G) \leq \lambda_{\min}(H) \leq \lambda_{\max}(H) \leq \lambda_{\max}(G),$$

but to prove this for non-induced subgraphs requires the Perron-Frobenius Theorem. More on this later.

Chapter 3

The Laplacian Matrix of a Graph

3.1 Introduction to the graph Laplacian

Definition 3.1.1. Let G be a graph. The Laplacian matrix of G , denoted $L(G)$, is defined by $L(G) = \Delta(G) - A(G)$, where $A(G)$ is the adjacency matrix of G and $\Delta(G)$ is the diagonal matrix whose (i, i) entry is equal to the degree of the i th vertex of G .

The Laplacian matrix of a graph carries the same information as the adjacency matrix obviously, but has different useful and important properties, many relating to its spectrum. We start with a few examples.

Examples

1. Complete graphs If $G = K_4$ then $L(G) = \begin{pmatrix} 3 & -1 & -1 & -1 \\ -1 & 3 & -1 & -1 \\ -1 & -1 & 3 & -1 \\ -1 & -1 & -1 & 3 \end{pmatrix}$. We can observe that

$v_1 = (1 \ 1 \ 1 \ 1)^T$ is an eigenvector of $L(G)$ corresponding to the eigenvalue 0, since the row sums in $L(G)$ are all equal to zero. This is true of the Laplacian matrix of any graph, and it follows from the fact that in each row we have the degree of the corresponding vertex on the diagonal, along with a -1 for each of its incident edges. At this point we don't know the multiplicity of the zero eigenvalue, but we know from Theorem 2.2.2 that any eigenvector corresponding to a non-zero eigenvalue must be orthogonal to v_1 , which means that the sum of its entries must be zero. So suppose that $a + b + c + d = 0$ (with a, b, c, d not all zero) and consider the equation

$$\begin{pmatrix} 3 & -1 & -1 & -1 \\ -1 & 3 & -1 & -1 \\ -1 & -1 & 3 & -1 \\ -1 & -1 & -1 & 3 \end{pmatrix} \begin{pmatrix} a \\ b \\ c \\ d \end{pmatrix} = \lambda \begin{pmatrix} a \\ b \\ c \\ d \end{pmatrix},$$

with $\lambda \neq 0$. This says

$$\begin{aligned} (3 - \lambda)a &= b + c + d &\implies (3 - \lambda)a &= -a \\ (3 - \lambda)b &= a + c + d &\implies (3 - \lambda)b &= -b \\ (3 - \lambda)c &= a + b + d &\implies (3 - \lambda)c &= -c \\ (3 - \lambda)d &= a + b + c &\implies (3 - \lambda)d &= -d \end{aligned}$$

Any choice of a, b, c, d with $a + b + c + d = 0$ satisfies these equations, with $3 - \lambda = -1$, so $\lambda = 4$. The 3-dimensional subspace $\langle v_1 \rangle^\perp$ of \mathbb{R}^4 consists entirely of eigenvectors of $L(G)$ corresponding to the eigenvalue 4, so this eigenvalue occurs with multiplicity 3 and

$\text{spec}L(G) = [0, 4, 4, 4]$. Note that the sum of the eigenvalues is 3×4 which is also the trace as expected.

In general, $\text{spec}L(K_n) = [0, \underbrace{n, \dots, n}_{n-1}]$.

For the complete graphs, all the non-zero eigenvalues coincide. The greatest is n which is also the graph order.

2. Cycles

Let C_4 be the cycle of length 4. Then $L(G) = \begin{pmatrix} 2 & -1 & 0 & -1 \\ -1 & 2 & -1 & 0 \\ 0 & -1 & 2 & -1 \\ -1 & 0 & -1 & 2 \end{pmatrix}$.

As above, an eigenvector of $L(G)$ corresponding to a non-zero eigenvalue λ is a non-zero vector whose entries sum to zero and satisfy

$$\begin{pmatrix} 2 & -1 & 0 & -1 \\ -1 & 2 & -1 & 0 \\ 0 & -1 & 2 & -1 \\ -1 & 0 & -1 & 2 \end{pmatrix} \begin{pmatrix} a \\ b \\ c \\ d \end{pmatrix} = \lambda \begin{pmatrix} a \\ b \\ c \\ d \end{pmatrix},$$

This means

$$\begin{aligned} (2 - \lambda)a &= b + d \implies (2 - \lambda)a = b + d \\ (2 - \lambda)b &= a + c \implies (2 - \lambda)b = a + c \\ (2 - \lambda)c &= b + d \implies (2 - \lambda)c = b + d \\ (2 - \lambda)d &= a + c \implies (2 - \lambda)d = a + c \end{aligned}$$

Adding the first two equations gives $(2 - \lambda)(a + b) = 0$, which means that $\lambda = 2$ or $a + b = 0$. The possibility that $\lambda = 2$ gives two independent (and orthogonal) eigenvectors

$$v_2 = \begin{pmatrix} 1 \\ 0 \\ -1 \\ 0 \end{pmatrix}, \quad v_3 = \begin{pmatrix} 0 \\ 1 \\ 0 \\ -1 \end{pmatrix}.$$

The remaining possibility $a + b = 0$ gives $c + d = 0$ also. Putting $a = 1$ we find the eigenvector

$$v_4 = \begin{pmatrix} 1 \\ -1 \\ 1 \\ -1 \end{pmatrix}$$

corresponding to the eigenvalue 4. So the Laplacian spectrum of C_4 is $[0, 2, 2, 4]$. Again the greatest eigenvalue is 4 (equal to the order) and the least positive eigenvalue is 2 this time.

In general the Laplacian spectrum of C_n is $[2 - 2\cos(\frac{2\pi k}{n}), k = 0 \dots n - 1]$. All eigenvalues are in the range 0 to 4, and the least positive eigenvalue approaches 0 as n increases, and occurs with multiplicity 2. The greatest eigenvalue is 4 exactly if n is even (when π is an integer multiple of $\frac{2\pi k}{n}$).

3. Stars

Let $G = S_n$, the star on n vertices. This graph has one vertex that is adjacent to all others, which have degree 1. We take $n = 4$ as an example. Then

$$L(G) = \begin{pmatrix} 3 & -1 & -1 & -1 \\ -1 & 1 & 0 & 0 \\ -1 & 0 & 1 & 0 \\ -1 & 0 & 0 & 1 \end{pmatrix}.$$

As above we consider

$$\begin{pmatrix} 3 & -1 & -1 & -1 \\ -1 & 1 & 0 & 0 \\ -1 & 0 & 1 & 0 \\ -1 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} a \\ b \\ c \\ d \end{pmatrix} = \lambda \begin{pmatrix} a \\ b \\ c \\ d \end{pmatrix},$$

with $a + b + c + d = 0$ and $\lambda \neq 0$. Then

$$\begin{aligned} (3 - \lambda)a &= b + c + d \implies (3 - \lambda)a = -a \\ (1 - \lambda)b &= a \\ (1 - \lambda)c &= a \\ (1 - \lambda)d &= a \end{aligned}$$

If $a \neq 0$, then $3 - \lambda = -1$ and $\lambda = 4$. We then find $-b = -c = -d = a$, so we obtain the eigenvector

$$\begin{pmatrix} 3 \\ -1 \\ -1 \\ -1 \end{pmatrix}$$

corresponding to the eigenvalue 4.

Alternatively if $a = 0$ then remaining eigenvectors are in the 2-dimensional space of vectors satisfying $b + c + d = 0$. We find that all elements of this space are eigenvectors of $L(S_4)$ corresponding to $\lambda = 1$. So the spectrum of $L(S_4)$ is $[0, 1, 1, 4]$. The greatest eigenvalue is 4 again, and the least positive eigenvalue is 1, which occurs twice.

In general the Laplacian spectrum of S_n is $[0, \underbrace{1, \dots, 1}_{n-2}, n]$ (this is not too hard to check). The minimum positive eigenvalue is 1 this time, and it occurs with multiplicity $n - 2$.

Theorem 3.1.2. *The Laplacian matrix of a graph G is a positive semidefinite matrix.*

Proof. Let B be the incidence matrix of G , in which rows are labelled by the edges of G , columns by the vertices of G , and the entry in the (i, j) position is 1 or 0 according to whether vertex j is incident with edge i or not. Thus each row of B has exactly two 1s, and the number of 1s in Column j of B is the degree of vertex j . Now adjust B by changing the first 1 in each row to -1 and leaving all other entries alone. (Now B_1 is what is called an “oriented incidence matrix” for G , writing the two non-zero entries in Row i as 1 and -1 can be interpreted as assigning a direction to edge i).

The square matrix $B_1^T B_1$ has rows and columns labelled by the vertices v_1, \dots, v_n of G . Its entry in the (i, j) position is the scalar product of Columns i and j of B_1 . This is $\deg(v_i)$ if $i = j$, and if $i \neq j$ it is 0 *unless* there is a row in which both Column i and Column j have nonzero entries. There can be at most one such row and it occurs when $v_i v_j$ is an edge of G , in which case the scalar product of Columns i and j of B_1 is $(1)(-1) = -1$. Thus

$$(B_1^T B_1)_{ij} = \begin{cases} \deg(v_i) & \text{if } i = j \\ 0 & \text{if } i \neq j \\ -1 & \text{if } i \sim j \end{cases}$$

Thus $B_1^T B_1 = L(G)$ and $L(G)$ is positive semidefinite by Lemma 2.1.8. □

Thus all eigenvalues of the Laplacian matrix of a graph are non-negative, and the zero eigenvalue occurs with multiplicity at least 1, since the row sums are all zero. Our next main result is that the multiplicity of the zero eigenvalue tells us the number of connected components.

Theorem 3.1.3. *Let G be a graph. Then the dimension of the nullspace of $L(G)$ is the number of connected components of G .*

Since the matrix $L(G)$ is symmetric and therefore diagonalizable, the multiplicity of zero as a root of its characteristic polynomial is the same as the dimension of the right nullspace of $L(G)$ which is the geometric multiplicity of zero as an eigenvalue of $L(G)$. So in order to prove Theorem 3.1.3 it is enough to consider the right nullspace of $L(G)$.

Proof. Let B be an oriented incidence matrix of G and write $L(G) = B^T B$, where $L(G)$ and B are written with respect to the ordering v_1, \dots, v_n of the vertices of G . Suppose that $x \in \mathbb{R}^n$ is an eigenvector of $L(G)$ corresponding to 0, i.e. that $L(G)x = 0$. Then

$$B^T Bx = 0 \implies x^T B^T Bx = 0 \implies (Bx)^T Bx = 0.$$

Thus Bx is a self-orthogonal vector in \mathbb{R}^n which means $Bx = 0$, and it is enough to consider the right nullspace of B . If the column vector x is orthogonal to every row of B , it means that the components in the $x_i = x_j$ whenever the vertices v_i and v_j are adjacent in G . Thus x_i and x_j must be equal whenever there is a path from v_i to v_j in G , i.e. whenever v_i and v_j belong to the same component of G .

Let C_1, \dots, C_k be the connected components of G , and for $i \in \{1, \dots, k\}$ let u_i be the vector that has 1s in the positions corresponding to the vertices of C_i and zeros elsewhere. Then u_i is easily confirmed to be in the right nullspace of $L(G)$, and by the above argument every element of this nullspace is a linear combination of u_1, \dots, u_k . Since these vectors are linearly independent, they form a basis of the zero eigenspace of $L(G)$, and the dimension of this space is k , the number of components. \square

An alternative version of Theorem 3.1.3 expresses the rank of $L(G)$ as $n - k$ (the number of components). Thus the graph Laplacian provides a feasible means for determining the number of components in a graph. There is no direct way of reading this number from the adjacency matrix.

Just as the multiplicity of the zero eigenvalue of $L(G)$ carries information about the number of connected components of G , the appearance and/or multiplicity of the eigenvalue n tells us about components of the complement \overline{G} .

Theorem 3.1.4. *Suppose that G is a graph of order n and that n occurs c times as an eigenvalue of $L(G)$, where $c \geq 0$. Then the number of connected components of \overline{G} is $c + 1$.*

Example We have seen that n occurs $n - 1$ times as an eigenvalue of $L(K_n)$. The complement of K_n has n isolated vertices and so has n connected components. The star S_n has n appearing once as an eigenvalue, and its complement has two components - an isolated vertex and a copy of K_{n-1} .

Theorem 3.1.4 is a consequence of the following lemma which explains a complementarity between the Laplacian spectra of a graph G and its complement.

Lemma 3.1.5. *Let G be a graph and let $0, \lambda_2, \dots, \lambda_n$ be the eigenvalues of $L(G)$, listed in increasing order. Then the eigenvalues of $L(\overline{G})$ are $0, n - \lambda_n, n - \lambda_{n-1}, \dots, n - \lambda_2$.*

Proof. That 0 is an eigenvalue of $L(\overline{G})$ is clear. Note that $L(G) + L(\overline{G}) = nI - J$, which is the Laplacian matrix of K_n . Suppose that $i \geq 2$ and let v be an eigenvector of G corresponding to λ_i . We may assume that $v \perp \mathbf{1}$ (where $\mathbf{1}$ denotes the all-1 vector). Thus the sum of the entries of v is zero. Then

$$L(\overline{G})v = (nI - J - L(G))v = nv - \lambda_i v = (n - \lambda_i)v.$$

Thus $n - \lambda_i$ is an eigenvalue of $L(\overline{G})$ whose eigenspace is the same as the $L(G)$ -eigenspace of λ_i . \square

From Lemma 3.1.5 it is immediate that n is an eigenvalue of $L(G)$ if and only if 0 occurs at least twice as an eigenvalue of $L(\overline{G})$, i.e. of and only if \overline{G} is disconnected. The multiplicity of n as an eigenvalue of $L(G)$ is one less than the multiplicity of 0 as an eigenvalue of $L(\overline{G})$, i.e. one less than the number of connected components of \overline{G} .

Exercise: If n is an eigenvalue of $L(G)$ for some graph G , prove that 0 occurs only once as an eigenvalue of $L(G)$.

3.2 Spanning Trees

A *tree* is a connected graph with no cycle. Some basic properties of trees are noted below.

- A tree of order n has exactly $n - 1$ edges.
- Every tree has at least two *leaves*, i.e. vertices of degree 1.
- A graph is a tree if and only if it contains a *unique* path between any pair of its vertices.
- Trees can be regarded as minimally connected in the sense that the deletion of any edge would result in a disconnected graph.
- If the maximum degree of a vertex in a tree T is Δ , then T has at least Δ leaves.

Definition 3.2.1. Let G be a connected graph. A spanning tree of G is a subgraph T which is a tree and whose vertex set is the full vertex set of G .

Every connected graph has at least one spanning tree, since one may be obtained by repeating the step of deleting an edge that belongs to a cycle until none remain. Such a step will never disconnect a graph.

We also introduce the following piece of matrix notation: if u is a vertex of a graph G with Laplacian matrix $L(G)$, we denote by $L(G)[u]$ the matrix obtained from $L(G)$ by deleting the row and column corresponding to u . If G has order n , then $L(G)[u]$ is a principal $(n - 1) \times (n - 1)$ submatrix of $L(G)$. Similarly, if u and v are both vertices of G , we denote by $L(G)[u, v]$ the principal $(n - 2) \times (n - 2)$ submatrix of $L(G)$ obtained by deleting the rows and columns labelled by u and v . The main theorem of this section is the statement that, for any vertex u , $\det(L(G)[u])$ counts the spanning trees in G .

We have noted this already in the case of complete graphs and stars, and we also remark that if G is disconnected then L has rank at most $n - 2$, so $L[u]$ is singular for all u , which is consistent with the theorem since the number of spanning trees in a disconnected graph is zero.

Theorem 3.2.2. Let G be a graph with Laplacian matrix L . Let u be any vertex of G . Then $\det(L[u])$ is the number of spanning trees in G .

The mechanism of the proof is induction on the number of edges. The induction step relies on the following two distinct methods of moving from G to a graph with one fewer edge.

- Let e be an edge of G . Then $G \setminus e$ is the graph obtained from G by deleting e (but not the vertices with which e is incident). Note that $G \setminus e$ need not be connected even if G is.
- Let e be an edge of G . Then G/e is obtained from G by *contracting* the edge e , which means identifying the two vertices of e together. If $e = uv$ then u and v are identified as a single vertex (still called either u or v) of G/e , and all neighbours of u or v in G are neighbours of this merged vertex in G/e . Note that if u and v have common neighbours in G , then G/e has multiple edges. If G is connected, then G/e is also connected.
- The number of edges in both $G \setminus e$ and G/e is one less than the number in G .

For any graph G , we denote the number of spanning trees of G by $\tau(G)$. The proof of Theorem 3.2.2 is presented as a series of lemmas.

Lemma 3.2.3. Let G be a (connected) graph and let e be an edge of G . Then

$$\tau(G) = \tau(G \setminus e) + \tau(G/e).$$

Proof. Every spanning tree of G that contains e contracts to a spanning tree of G/e when the edge e is contracted, and every spanning tree of G/e may be expanded to a spanning tree of G by reintroducing the edge e . Thus the number of spanning trees of G that contain e is $\tau(G/e)$.

On the other hand any spanning tree of G that does not contain e is a spanning tree of $G \setminus e$.

Since every spanning tree of G either contains e or does not contain e , it follows that

$$\tau(G) = \tau(G/e) + \tau(G \setminus e),$$

as required. \square

The next lemma deals with the relationship between the Laplacian matrices of G , G/e and $G \setminus e$. For a graph G and a designated ordering of its vertices, we denote by E_{vv} the matrix that has a 1 in the row and column labelled by the vertex v , and zeros elsewhere.

Lemma 3.2.4. *Let G be a graph and let $e = uv$. Then*

$$L(G)[u] = L(G \setminus e)[u] + E_{vv}.$$

Proof. The Laplacian matrix of $G \setminus e$ differs from that of G only in the entries in positions (u, u) , (v, v) , (u, v) and (v, u) . Three of these entries, since they belong to the row and column labelled by u , are absent from $L(G)[u]$ and from $L(G \setminus e)[u]$. The fourth represents the degree of v , which is greater by one in G than in $G \setminus e$. Hence the $(n-1) \times (n-1)$ matrices $L(G)[u]$ and $L(G \setminus e)[u]$ are related by the equation

$$L(G)[u] = L(G \setminus e)[u] + E_{vv}.$$

\square

The next lemma relates $(n-2) \times (n-2)$ principal submatrices of $L(G)$ to those of $L(G/e)$. In the statement of the lemma, we interpret that when the edge $e = uv$ is contracted to form G/e , it is the vertex u that is “absorbed” into v and the vertex v that survives.

Lemma 3.2.5. *Let G be a graph and let $e = uv$ be an edge of G . Then*

$$L(G)[u, v] = L(G/e)[v].$$

Proof. Let w and z be vertices of G (other than u and v). If $w \neq z$, the entry in the (w, z) position of $L(G)[u, v]$ is -1 or 0 according as w is adjacent to z or not. Since w and z are adjacent in G if and only if they are adjacent in G/e , the (w, z) -entry of $L(G/e)[v]$ is the same as that of $L(G)[u, v]$. The (w, w) -entry of $L(G)[u, v]$ is the degree in G of w . This is the same as the degree of w in G/e (bearing in mind that a double edge from w to v in G/e contributes 2 to the degree of w in this graph). So the diagonal entries of $L(G)[u, v]$ also coincide with those of $L(G/e)[v]$. \square

We are now in a position to complete the proof of Theorem 3.2.2, by induction on the number of edges.

Base: If G is a graph with a single edge, then the number of spanning trees of G is 1 if G has order 2 and zero if the order of G exceeds 2. If the order of G is 2, then $L(G) = \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}$ and every 1×1 principal submatrix of $L(G)$ has determinant 1. If the order of G is 3 or greater, then $L(G)$ has one 2×2 principal submatrix equal to $\begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}$ and is otherwise full of zeros. In this case $L(G)$ has rank 1 and all of its $(n-1) \times (n-1)$ principal submatrices have determinant zero, which is the number of spanning trees in G . So the theorem is true for graphs with one edge.

Induction hypothesis: We consider a particular graph G and assume that the theorem holds for all graphs with fewer edges than G (in particular for $G \setminus e$ and G/e for any edge e of G).

Induction step: Choose a vertex u of G . We need to show that $\det L(G)[u] = \tau(G)$. If u is isolated in G , then $\tau(G)$ and $\det L(G)[u]$ are both equal to zero and the theorem holds. If not, then there is an edge $e = uv$ in G . By Lemma 3.2.3,

$$\tau(G) = \tau(G \setminus e) + \tau(G/e).$$

By the induction hypothesis,

- $\tau(G \setminus e) = \det(L(G \setminus e)[u])$, and
- $\tau(G/e) = \det(L(G/e)[v])$.

From Lemma 3.2.5 we have $L(G/e)[v] = L(G)[u, v]$. From Lemma 3.2.4 we have

$$L(G)[u] = L(G \setminus e)[u] + E_{vv}.$$

Consider applying a determinant calculation using cofactor expansion by the Row of v to the above description of $L(G)[u]$. Where C_{vw} denotes the cofactor of the entry in the (v, w) -position of $L(G)[u]$, we have

$$\begin{aligned} \det(L(G)[u]) &= \sum_{w \in V(G) \setminus \{u\}} (L(G)[u])_{vw} C_{vw} \\ &= \sum_{w \neq u, v} L(G \setminus e)[u]_{vw} C_{vw} + (L(G)[u])_{vv} C_{vv} \\ &= \sum_{w \neq u} L(G \setminus e)[u]_{vw} C_{vw} + C_{vv}, \end{aligned}$$

where the last line is an application of Lemma 3.2.4. Finally we have

$$\begin{aligned} \sum_{w \neq u} L(G \setminus e)[u]_{vw} C_{vw} + C_{vv} &= \det(L(G \setminus e)[u]) + \det(L(G)[u, v]) \\ &= \tau(G \setminus e) + \tau(G/e), \end{aligned}$$

by the induction hypothesis. Application of Lemma 3.2.3 now completes the proof.

Corollary 3.2.6. *The number of spanning trees of the complete graph K_n is n^{n-2} .*

Proof. Since $\text{spec}(L(K_n)) = [0, \underbrace{n, \dots, n}_{n-1}]$, the sum of all of the $(n-1) \times (n-1)$ principal minors of $L(K_n)$ is n^{n-1} . Since it follows from Theorem 3.2.2 that all of these n principal minors are equal, each of them is equal to n^{n-2} . \square

In fact, a slightly stronger statement than Theorem 3.2.2 can be proved. Not only are all of the principal $(n-1) \times (n-1)$ principal minors of $L(G)$ equal to $\tau(G)$, but the cofactor of *every entry* of $L(G)$ is equal to $\tau(G)$.

For a square matrix A the *cofactor* C_{ij} of the entry A_{ij} is the $(-1)^{i+j} \times \det(A[Ri, Cj])$, where $A[Ri, Cj]$ is the matrix obtained by deleting Row i and Column j from A . The *adjugate* of A , denoted $\text{adj}A$, is the transpose of the matrix of cofactors of A , i.e. it is the matrix whose (i, j) -entry is A_{ji} . The relationship between A and its adjugate is

$$A \text{adj}(A) = \det(A) I_n = \text{adj}(A) A.$$

If A is invertible, this means that $A^{-1} = \frac{1}{\det A} \text{adj}(A)$.

Lemma 3.2.7. *Let G be a graph with Laplacian matrix $L(G)$. Then for all (i, j) , the cofactor C_{ij} of $L(G)$ is equal to $\tau(G)$.*

Proof. First suppose that G is disconnected. Then $\tau(G) = 0$, and the rank of $L(G)$ is less than $n-1$, so all $(n-1) \times (n-1)$ minors of $L(G)$ are zero. Thus all cofactors of $L(G)$ are equal to zero which is $\tau(G)$.

If G is connected, then since $L(G)$ is singular we know that $L(G) \text{adj}(L(G)) = 0_{n \times n}$. Furthermore, $L(G)$ has rank $n-1$ which means that the right nullspace of $L(G)$ is 1-dimensional, spanned by the all-1 vector $\mathbf{1}$. Thus every column of $\text{adj}(L(G))$ is a scalar multiple of $\mathbf{1}$. Since $L(G)$ is symmetric, so also is $\text{adj}(L(G))$, so every row of $\text{adj}(L(G))$ is a scalar multiple of $\mathbf{1}^T$. Since the diagonal entries of $\text{adj}(L(G))$ are all equal to $\tau(G)$, it follows that every entry of $\text{adj}(L(G))$ is equal to $\tau(G)$, as required. \square

3.3 Connectivity and the graph Laplacian

Let G be a connected graph of order n . If S is a set of vertices of G , then $G \setminus S$ is the graph obtained from G by deleting the vertices of G and their incident edges.

Definition 3.3.1. A vertex-cutset of G is set S of vertices of G for which the graph $G \setminus S$ is disconnected. The vertex connectivity of G , denoted $\kappa(G)$, is defined to be the minimum number of vertices in a vertex-cutset.

Notes

1. $\kappa(G) = 0$ if and only if G is disconnected.
2. If G is a tree on more than 2 vertices, then $\kappa(G) = 1$.
3. If $\kappa(G) = 1$, it means that G has vertices x and y with the property that every path from x to y goes via a particular vertex v (a *cut vertex*).
4. If $\kappa(G) = k$ and S is a vertex-cutset with k elements, it means that there is a pair of vertices x and y in G (not in S) with the property that every path from x to y in G is via a vertex of S (but there is no set of fewer than k vertices for which there exists such a pair).
5. G is said to be t -vertex-connected if it cannot be disconnected by the deletion of fewer than t vertices. So if $\kappa(G) = t$, it means that G is t -vertex-connected but not $(t + 1)$ -vertex-connected.
6. The complete graph K_n cannot be disconnected by the deletion of vertices, so its vertex connectivity is not defined.
7. If G is a non-complete graph of order n , then its vertex connectivity is at most $n - 2$, since it can be disconnected by the removal of all vertices except for some non-adjacent pair.

Definition 3.3.2. An edge cutset of a connected graph G is a set of edges of G whose deletion would disconnect G . The edge-connectivity of G , denoted $\epsilon(G)$, is the minimum number of edges in an edge-cutset.

Notes

1. A single edge whose deletion would disconnect G is called a *bridge* or a *cut edge*. If $e = uv$ is a bridge in G , then e is the unique path between u and v in G .
2. If $\epsilon(G) = 2$ and $\{e_1, e_2\}$ is an edge-cutset in G , it means that every cycle in G that contains e_1 also contains e_2 , or that every cycle containing e_2 also contains e_1 .
3. In any connected graph G , let $\delta(G)$ be the minimum of the vertex degrees in G , and let v be a vertex with $\deg(v) = \delta$. Then G can be disconnected by the deletion of the δ edges incident with v , and so $\epsilon(G) \leq \delta$.
4. It follows that the edge connectivity of a non-complete graph of order n is at most $n - 2$.

Lemma 3.3.3. Let G be a (non-complete) connected graph on n vertices. Then $\kappa(G) \leq \epsilon(G)$.

Proof. Let $m = \epsilon(G)$ (note $m \leq n - 2$) and suppose that $S = \{e_1, \dots, e_m\}$ is a set of edges of G whose deletion disconnects G . Since the removal of a single edge can break a connected graph into at most two components, and since S is a minimal set whose deletion disconnects G , it follows that $G \setminus S$ has exactly two connected components. Furthermore, since the restoration of any of the edges e_i would reconnect $G \setminus S$, it follows that for each i , the vertices x_i and y_i of e_i belong to different components of $G \setminus S$. We may label these vertices so that x_1, \dots, x_m all belong to the component C_1 of $G \setminus S$, and y_1, \dots, y_m all belong to the other component C_2 (note that the x_i are not necessarily distinct, same for the y_i).

If the x_i are not all of the vertices of C_1 , then deleting the vertices x_i and their incident edges disconnects G . Since the number of x_i is at most m , this shows that $\kappa(G) \leq m$.

On the other hand, if $\{x_1, \dots, x_m\}$ is the full vertex set of C_1 , then deletion of this set may not disconnect G . In this case the only possible neighbours of x_1 in G are those x_j for which $x_j \neq x_1$ and those y_l for which $x_l = x_1$. The total number of such vertices is at most m , and their deletion (along with their incident edges) leaves a graph G' in which the vertex x_1 is isolated. Since $m \leq n - 2$, the graph G' possesses at least two more vertices, hence it is disconnected and $\kappa(G) \leq m$. \square

Now we return to the Laplacian spectrum of a graph. Let G be a graph of order n with Laplacian matrix $L(G)$, and let $0, \lambda_2, \dots, \lambda_n$ be the eigenvalues of $L(G)$, in increasing order. We know that $\lambda_2 > 0$ if and only if G is connected. We show now that λ_2 is bounded above by $\kappa(G)$. To see this we need the following lemma. In the statement of this lemma, the column vector x is considered as a real-valued function on the vertex set of G , its value on the vertex u is the component x_u .

Lemma 3.3.4. *Let G be a graph with Laplacian matrix L . Let $x \in \mathbb{R}^n$. Then*

$$x^T L x = \sum_{uv \in E(G)} (x_u - x_v)^2.$$

Proof. Let B be an oriented incidence matrix for G , and recall that $L = B^T B$. Then

$$x^T L x = x^T B^T B x = (Bx)^T (Bx) = (Bx) \cdot (Bx).$$

The rows of B (and hence the entries of Bx) are labelled by the edges of G . If $e = uv$ is an edge of G , then the entry of Bx in the position corresponding to e is $\pm(x_u - x_v)$. Thus

$$x^T L x = \sum_{uv \in E(G)} (x_u - x_v)^2.$$

\square

We now show that $\lambda_2(G)$ is the minimum over all unit vectors x satisfying $x \perp \mathbf{1}$ of the expression $\sum_{uv \in E(G)} (x_u - x_v)^2$.

Theorem 3.3.5. *Let G be a graph of order n with Laplacian spectrum $0 \leq \lambda_2 \leq \dots \leq \lambda_n$. Then λ_2 is the minimum over all unit vectors x that belong to $\mathbf{1}^\perp$ of the expression $x^T L(G)x$, and this minimum is attained if and only if x is an eigenvector of $L(G)$ corresponding to λ_2 .*

Proof. Let $\{v_1, v_2, \dots, v_n\}$ be an orthonormal basis of \mathbb{R}^n consisting of eigenvectors of $L(G)$, where $v_1 = \frac{1}{\sqrt{n}}\mathbf{1}$ and otherwise v_i corresponds to the eigenvalue λ_i . Let x be a unit vector in \mathbb{R}^n that is orthogonal to $\mathbf{1}$. Then

$$x = a_2 v_2 + \dots + a_n v_n$$

for some real numbers a_i with $\sum a_i^2 = 1$.

$$\begin{aligned} x^T L(G)x &= (a_2 v_2 + \dots + a_n v_n)^T (a_2 \lambda_2 v_2 + \dots + a_n \lambda_n v_n) \\ &= \lambda_2 a_2^2 + \dots + \lambda_n a_n^2 \\ &= \lambda_2 (a_2^2 + \dots + a_n^2) + \sum_{i=3}^n (\lambda_i - \lambda_2) a_i^2 \\ &= \lambda_2 + \sum_{i=3}^n (\lambda_i - \lambda_2) a_i^2. \end{aligned}$$

Since $\lambda_i - \lambda_2$ is non-negative whenever $i \geq 3$, the last line above says that $x^T L(G)x \geq \lambda_2$ and that equality occurs here if and only if $a_i = 0$ whenever $\lambda_i > \lambda_2$, which means that x is an eigenvector of $L(G)$ corresponding to λ_2 . \square

Lemma 3.3.4 and Theorem 3.3.5 together give us the necessary ingredients to prove that $\lambda_2(G)$ is bounded above by the vertex connectivity of G .

Theorem 3.3.6. *Let G be a (non-complete) connected graph. Then $\lambda_2(G) \leq \kappa(G)$.*

Proof. Write k for $\kappa(G)$ and let S be a set of vertices of G whose deletion disconnects G . Since the graph $G \setminus S$ is disconnected, $\lambda_2(G \setminus S) = 0$, and there exists a unit vector x' in \mathbb{R}^{n-k} that is a zero eigenvector of $L(G \setminus S)$ and is orthogonal to $\mathbf{1}_{n-k}$. Let x be the unit vector in \mathbb{R}^n that coincides with x' on the vertices of $G \setminus S$ and has zeros in the positions corresponding to vertices of S . By Theorem 3.3.5

$$\lambda_2(G) \leq x^T L(G)x = \sum_{uv \in E(G)} (x_u - x_v)^2.$$

We now break this sum into separate components, one involving the edges of $G \setminus S$ and one involving edges for which one vertex is in S and the other is not. There is no need to consider edges whose vertices are both in S since $x_w = 0$ whenever $w \in S$. Thus

$$\sum_{uv \in E(G)} (x_u - x_v)^2 \leq \sum_{uv \in E(G \setminus S)} (x_u - x_v)^2 + \sum_{u \in S} \sum_{v \notin S} x_v^2.$$

Since x' is a zero eigenvector of $L(G \setminus S)$ (which means x is constant on the components of $G \setminus S$), the term $\sum_{uv \in E(G \setminus S)} (x_u - x_v)^2$ is zero. The other term is bounded above by

$$\sum_{u \in S} \sum_{v \in V(G)} x_v^2 = \sum_{u \in S} \|x\|^2 = |S| = k.$$

Thus $\lambda_2(G) \leq k$ as required. \square

Definition 3.3.7. *For a connected graph G , $\lambda_2(G)$ is called the algebraic connectivity of G .*

The following example shows that $\lambda_2(G)$ may be equal to the vertex connectivity of G , or quite far away from it.

Example For the cycle C_n of length n , $\lambda_2(C_n) = 2 - 2 \cos(\frac{2\pi}{n})$. So $\lambda_2(C_4) = 2$ which is equal to the vertex connectivity of C_4 . For all $n \geq 4$, $\kappa(C_n) = 2$. However, as $n \rightarrow \infty$, $\lambda_2(C_n) \rightarrow 0$. This reflects the fact that as n increases, C_n becomes more “flimsily” vertex-connected - to disconnect it, an ever decreasing proportion of its vertices needs to be removed. This observation is a heuristic not a theorem - there are several classes of graphs for which the algebraic connectivity behaves in this manner. Graphs with low values of λ_2 tend to have high diameter compared to their order.

Despite this remark another conclusion from looking at the graphs C_n is that the “gap” between the algebraic connectivity and the vertex connectivity of a graph is not always easy to interpret. Another way of looking at the meaning of λ_2 is to consider how it changes when adjustments are made to a graph. The step of adding an edge between two non-adjacent vertices in a graph cannot decrease the vertex connectivity, and can leave it unchanged or increase it by 1. Our next theorem shows that something like this is true for the algebraic connectivity.

Theorem 3.3.8. *Let G be a (non-complete) graph of order n , and let H be obtained from G by adding an edge between two non-adjacent vertices. Then*

$$\lambda_2(G) \leq \lambda_2(H) \leq \lambda_2(G) + 2.$$

Proof. Let r and s be the two vertices that are adjacent in H and not in G . Let x be a unit vector in \mathbb{R}^n , orthogonal to $\mathbf{1}$, for which

$$\lambda_2(H) = \sum_{uv \in E(H)} (x_u - x_v)^2 = \sum_{uv \in E(G)} (x_u - x_v)^2 + (x_r - x_s)^2.$$

Since $\lambda_2(G)$ is the minimum over all unit vectors y (orthogonal to $\mathbf{1}$) of the sum of $(y_u - y_v)^2$ over all edges uv of G , it follows that $\lambda_2(G)$ is at most equal to $\sum_{uv \in E(H)} (x_u - x_v)^2$ and in particular $\lambda_2(G) \leq \lambda_2(H)$.

On the other hand, let z be a unit eigenvector of $L(G)$ corresponding to $\lambda_2(G)$. Then

$$\lambda_2(H) \leq \sum_{uv \in E(G)} (z_u - z_v)^2 + (z_x - z_y)^2.$$

Note that

$$(z_x + z_y)^2 = z_x^2 + z_y^2 + 2z_x z_y \geq 0 \implies -2z_x z_y \leq z_x^2 + z_y^2.$$

Furthermore, since z_x and z_y are components of a unit vector, we know that $z_x^2 + z_y^2 \leq 1$. Thus

$$0 \leq (z_x - z_y)^2 = -2z_x z_y + (z_x^2 + z_y^2) \leq 1 + 1 = 2.$$

Thus

$$\lambda_2(H) \leq \sum_{uv \in E(G)} (z_u - z_v)^2 + (z_x - z_y)^2 \implies \lambda_2(H) \leq \lambda_2(G) + 2.$$

□

Chapter 4

Strongly Regular Graphs

4.1 Parameters and Properties

Recall that a (simple, undirected) graph is *regular* if all of its vertices have the same degree. This is a strong property for a graph to have, and it can be recognized easily from the adjacency matrix, since it means that all row sums are equal, and that $\mathbf{1}$ is an eigenvector.

If a graph G of order n is regular of degree k , it means that kn must be even, since this is twice the number of edges in G . If $k \leq n - 1$ and kn is even, then there does exist a graph of order n that is regular of degree k (showing that this is true is an exercise worth thinking about).

Regularity is a strong property for a graph to have, and it implies a kind of symmetry, but there are examples of regular graphs that are not particularly “symmetric”, such as the disjoint union of two cycles of different lengths, or the connected example below.

Various properties of graphs that are stronger than regularity can be considered, one of the most interesting of which is strong regularity.

Definition 4.1.1. *A graph G of order n is called strongly regular with parameters (n, k, λ, μ) if*

- every vertex of G has degree k ;
- if u and v are adjacent vertices of G , then the number of common neighbours of u and v is λ (every edge belongs to λ triangles);
- if u and v are non-adjacent vertices of G , then the number of common neighbours of u and v is μ ;
- $1 \leq k < n - 1$ (so the complete graph and the null graph of n vertices are not considered to be strongly regular).

So a srg (strongly regular graph) is a regular graph in which the number of common neighbours of a pair of vertices depends only on whether that pair forms an edge or not).

Examples

1. C_4 is strongly regular with parameters $(4, 2, 0, 2)$.
2. C_5 is strongly regular with parameters $(5, 2, 0, 1)$.
3. Apart from those two examples, C_n is not strongly regular: C_1, C_2 and C_3 are ruled out because they are complete graphs, and for $n \geq 6$, a pair of non-adjacent vertices may have either 1 common neighbour or none.

Strongly regular graphs are elusive and somewhat mysterious objects that have connections to various combinatorial constructions and to algebra over finite fields. A couple of slightly more complicated general families are described below.

Example 4.1.2. Recall that the *line graph* of K_n has vertices given by the $\binom{n}{2}$ edges of K_n , and a pair of vertices is adjacent if and only if the corresponding edges in K_n share a vertex.

- Let x be a vertex of $L(K_n)$, corresponding to the edge uv of K_n . The degree of x in $L(K_n)$ is the number of edges in K_n (other than uv) that are incident with either u or v . This is $2 \times (n - 2) = 2n - 4$.
- Suppose that x and y are adjacent vertices in $L(K_n)$, corresponding to the edges uv and uw of K_n . The number of common neighbours of x and y in $L(K_n)$ is the number of edges of K_n (other than uv and uw) that share a vertex with both uv and uw . There are $n - 2$ of these: vw and the $n - 3$ remaining edges involving u .
- Suppose that x and y are non-adjacent vertices in $L(K_n)$, corresponding to edges uv and wt of K_n . Then the number of common neighbours of x and y in $L(K_n)$ is the number of edges of K_n that are incident with one vertex in $\{u, v\}$ and one in $\{w, t\}$. There are four of these: uw , ut , vw and vt . So x and y have four common neighbours in $L(K_n)$.

The conclusion is that $L(K_n)$ is a strongly regular graph with parameters $(\binom{n}{2}, 2n - 4, n - 2, 4)$.

Example 4.1.3. Let $K_{n,n}$ denote the complete bipartite graph with n vertices in each part. The line graph $L(K_{n,n})$ has n^2 vertices, all of degree $2n - 2$. If two of these vertices are adjacent, they have $n - 2$ common neighbours. If two are non-adjacent, they have 2 common neighbours. So $L(K_{n,n})$ is a strongly regular graph with parameters $(n, 2n - 2, n - 2, 2)$.

Note that the complement of $L(K_n)$ is the *Kneser graph* $Kn(n, 2)$. This is the graph whose vertices are the 2-element subsets of a set with n elements, and in which two vertices are adjacent if and only if the subsets that they represent are disjoint. The Kneser graph $Kn(4, 2)$ consists of three isolated edges, and $Kn(5, 2)$ is the famous *Petersen graph*. In general $Kn(n, 2)$ is a strongly regular graph with parameters $(\binom{n}{2}, \binom{n-2}{2}, \binom{n-4}{2}, \binom{n-3}{2})$.

It is true in general that the complement of a strongly regular graph is strongly regular and the relationship between their parameters can be figured out without too much trouble.

Theorem 4.1.4. Let G be a strongly regular graph with parameters (n, k, λ, μ) . Then \overline{G} is a strongly regular graph with parameters $(n, n - k - 1, \lambda, \mu)$.

Proof. It is straightforward to observe that \overline{G} has n vertices and is regular of degree $n - k - 1$.

Let uv be an edge of \overline{G} . The number of triangles to which uv belongs in \overline{G} is the number of vertices in G that are adjacent to neither u nor v . In G , uv is not an edge, each of u and v has k neighbours, and μ vertices are common neighbours of u and v . So the number of vertices that are adjacent to at least one of u and v is $2k - \mu$. Thus the number of vertices (other than u and v) that is adjacent to neither u nor v is $n - 2 - 2k + \mu$. This is the number of triangles to which the edge uv belongs in \overline{G} .

Now suppose that u and v are non-adjacent edges in \overline{G} . Then uv is an edge of G . Each of u and v has $k - 1$ additional neighbours in G , and they have λ common neighbours, so $2k - 2 - \lambda$ is the number of vertices (other than u and v themselves) that are adjacent in G to at least one of u and v . That leaves $n - 2 - (2k - 2 - \lambda)$ or $n - 2k + \lambda$ vertices in G that are adjacent to neither u nor v . This is the number of common neighbours of u and v in \overline{G} .

We conclude that \overline{G} is a strongly regular graph with parameters $(n, n - k - 1, \lambda, \mu)$. \square

Our final theorem in this section presents a compatibility condition on the parameters in a strongly regular graph. Not surprisingly, not every set of four non-negative integers is a candidate for being the set of parameters. After looking at the adjacency spectrum of a srg in the next section, we will obtain some more constraints of this nature.

Theorem 4.1.5. Let G be a strongly regular graph with parameters (n, k, λ, μ) . Then $k(k - \lambda - 1) = (n - k - 1)\mu$.

Proof. We count, in two ways, the number of ordered triples of the form (u, v, w) in G with the property that v is adjacent to both u and w and that u and w are not adjacent to each other.

Suppose we choose v first - we have n choices here. Regardless of how this choice is made, the number of choices available for a neighbour u of v is k . Having chosen u , the final step is to choose a vertex w that is adjacent to v but is not a common neighbour of u and v . There are $k - 1$ neighbours of v from which w may be chosen, but λ of these are also neighbours of u . So the number of choices for w is $k - 1 - \lambda$. Hence the number of choices for the triple (u, v, w) is $nk(k - \lambda - 1)$.

On the other hand suppose we choose u first. We have n choices for u , and then we may choose w from among the $n - k - 1$ non-neighbours of u . Having done this we have μ choices for v among the common neighbours of u and w . So the number of choices for the triple (u, v, w) is $n(n - k - 1)\mu$.

Putting these two counts together we find

$$nk(k - \lambda - 1) = n(n - k - 1)\mu \implies k(k - \lambda - 1) = (n - k - 1)\mu.$$

□

4.2 The Adjacency Spectrum of a strongly regular graph

In this section we use the defining properties of a strongly regular graph to show that the adjacency matrix of such a graph satisfies a particular quadratic equation, from which we deduce that the adjacency spectrum can have at most three distinct elements. It is clear that the adjacency matrix of a k -regular graph has k as an eigenvalue with corresponding eigenvector $\mathbf{1}$, since the row sums are all equal to k .

Lemma 4.2.1. *Let G be a k -regular graph on n vertices, with adjacency matrix A . Then the multiplicity of k as an eigenvalue of A is the number of connected components of G .*

Proof. The spectrum of A is the list of roots of the polynomial $\det(A - xI)$. This may be rewritten as

$$\det(A - kI_kI - xI) = \det((k - x)I - (kI - A)) = \det((k - x)I - L),$$

where L is the Laplacian matrix of A . Thus x is an eigenvalue of A if and only if $k - x$ is an eigenvalue of L with the same multiplicity. In particular the multiplicity of k as an eigenvalue of A is the multiplicity of 0 as an eigenvalue of L , which is the number of connected components of G , by Theorem 3.1.3. □

Now let G be a connected strongly regular graph with parameters (n, k, λ, μ) , and with adjacency matrix A . Then A_{uv}^2 is the number of walks of length 2 from u to v in G , so

$$A_{uv}^2 = \begin{cases} k & \text{if } u = v \\ \lambda & \text{if } uv \text{ is an edge of } G \\ k & \text{if } u = v \\ k & \text{if } u \text{ and } v \text{ are not adjacent in } G \end{cases}$$

Thus

$$A^2 = kI + \lambda A + \mu(J - I - A) \implies A^2 - (\lambda - \mu)A - (k - \mu)I = \mu J.$$

Now let v be an eigenvector of A corresponding to an eigenvalue θ with $\theta \neq k$. Then v is orthogonal to $\mathbf{1}$ and so $Jv = 0$. Multiplying both sides of the above equation on the right by v gives

$$\theta^2 v - (\lambda - \mu)\theta v - (k - \mu)v = 0.$$

Since v is not the zero vector and since θ, λ and μ are all real numbers, this means that

$$\theta^2 - (\lambda - \mu)\theta - (k - \mu) = 0.$$

The roots of this quadratic equation are

$$\frac{(\lambda - \mu) \pm \sqrt{(\lambda - \mu)^2 + 4(k - \mu)}}{2}.$$

For convenience we write Δ for the expression $(\lambda - \mu)^2 + 4(k - \mu)$. The eigenvalues of A then are

$$k, \theta_1 = \frac{(\lambda - \mu) + \sqrt{\Delta}}{2}, \theta_2 = \frac{(\lambda - \mu) - \sqrt{\Delta}}{2}.$$

Note that Δ is positive. It is clear that $k - \mu$ cannot be negative, since the number of neighbours of a vertex v of G cannot be fewer than its number of common neighbours with another vertex u . It is possible that $k - \mu = 0$, this happens for example in the case of complete bipartite graphs. However if $k = \mu$ then λ cannot also be equal to μ , since an edge in a k -regular graph can belong to at most $k - 1$ triangles.

Since G is connected, k has multiplicity 1. Let m_1 and m_2 be the respective multiplicities of θ_1 and θ_2 as eigenvalues of A . Since the trace of A is 0, we have the following equations

$$m_1 + m_2 = n - 1, \quad m_1\theta_1 + m_2\theta_2 + k = 0.$$

Solving these equations gives

$$m_1 = -\frac{(n - 1)\theta_2 + k}{\theta_1 - \theta_2}, \quad m_2 = \frac{(n - 1)\theta_1 + k}{\theta_1 - \theta_2}.$$

Note that $\theta_1 - \theta_2 = \sqrt{\Delta}$. Entering the expressions for θ_1 and θ_2 in terms of the parameters of G to the equations above, we find

$$\begin{aligned} m_1 &= \frac{1}{2} \left[(n - 1) - \frac{2k + (n - 1)(\lambda - \mu)}{\sqrt{\Delta}} \right] \\ m_2 &= \frac{1}{2} \left[(n - 1) + \frac{2k + (n - 1)(\lambda - \mu)}{\sqrt{\Delta}} \right] \end{aligned}$$

Since $m_1 + m_2 = n - 1$ it is clear that if one of m_1, m_2 is an integer then so is the other. That m_1 is an integer requires either that Δ is a square or that $2k + (n - 1)(\lambda - \mu) = 0$. In the latter case n must be odd.

Example 4.2.2. Let G be the Petersen graph, with parameters $(10, 3, 0, 1)$. Then

$$k = 3, \Delta = 1 + 4(2) = 9, \theta_1 = 1, \theta_2 = -2, m_1 = 4, m_2 = 5.$$

We have shown that a connected strongly regular graph has exactly three distinct eigenvalues, k with multiplicity 1 and θ_1 and θ_2 with multiplicities adding to $n - 1$. The product of θ_1 and θ_2 is non-positive (it can be zero in the case $k = \mu$), and θ_1 and θ_2 are distinct. We now show that a connected regular graph with exactly three distinct eigenvalues must be strongly regular.

Theorem 4.2.3. Let G be a connected graph of order n that is regular of degree k , and let A be its adjacency matrix. Suppose that A has just three distinct eigenvalues: k, α and β . Then A is strongly regular.

Proof. Since G is connected, A has k just once as an eigenvalue, with corresponding eigenvector $\mathbf{1}$. Every eigenvector of A corresponding to α or β is in the right nullspace of the matrix $A' = (A - \alpha I)(A - \beta I)$ (to see this note that $(A - \alpha I)$ and $(A - \beta I)$ commute with each other). Since the eigenspaces of A and B together account for the subspace of \mathbb{R}^n of dimension $n - 1$ consisting of all vectors orthogonal to $\mathbf{1}$. It follows that A' has rank 1 and that every row of A' has all entries equal. Finally

$$A'\mathbf{1} = (A - \alpha I)(A - \beta I)\mathbf{1} = (A - \alpha I)(k - \beta)\mathbf{1} = (k - \alpha)(k - \beta)\mathbf{1},$$

which means that

$$A' = (A - \alpha I)(A - \beta I) = \frac{1}{n(k - \alpha)(k - \beta)} J.$$

In particular then A^2 is a linear combination of A , I and J , hence of A , I and $J - I - A$. This means that A^2 has the same entry in every position on the diagonal, the same entry in all positions corresponding to edges of G , and the same entry in all positions corresponding to non-edges of G . Since A is a $(0, 1)$ -matrix, these entries are all non-negative integers and it follows that G is strongly regular. \square

4.3 Two classes of strongly regular graphs

Let G is a strongly regular graph with parameters (n, k, λ, μ) , and assume that $k \leq \frac{n-1}{2}$; there is no real loss of generality in this assumption since either G or its complement has this property. We have seen that the eigenvalues of G occur with multiplicities

$$1, m_1 = \frac{1}{2} \left[(n-1) - \frac{2k + (n-1)(\lambda - \mu)}{\sqrt{\Delta}} \right], m_2 = \frac{1}{2} \left[(n-1) + \frac{2k + (n-1)(\lambda - \mu)}{\sqrt{\Delta}} \right].$$

The condition that m_1 and m_2 are integers means that one of the following two cases occurs:

1. $2k + (n-1)(\lambda - \mu) \neq 0$ and Δ is an integer square; $m_1 \neq m_2$ in this case.
2. $2k + (n-1)(\lambda - \mu) \neq 0$, and $m_1 = m_2 = \frac{1}{2}(n-1)$ (this is referred to as the "half case" for this reason). In this case n must be odd obviously. Furthermore, since $2k \leq n-1$, the condition that

$$2k = (n-1)(\mu - \lambda)$$

can be satisfied only if $2k = n-1$ and $\mu - \lambda = 1$, so $\lambda = \mu - 1$. Moreover we know from Theorem 4.1.5 that $k(k - \lambda - 1) = (n-1-k)\mu$. Since $n-1-k = k$ and $\lambda + 1 = \mu$, this means that $k - \mu = \mu$ or $k = 2\mu$. Finally $n = 2k + 1 = 4\mu + 1$ and G has parameters

$$(4\mu + 1, 2\mu, \mu - 1, \mu)$$

for some positive integer μ . A strongly regular graph of this type is called a *conference graph*.

We look briefly at some examples of both types. The Kneser graph $\text{Kn}(n, 2)$ (the complement of the line graph of K_n) is an example of the first type. In the case $n = 5$, this is the Petersen graph which has parameters $(10, 3, 0, 1)$, with

$$\Delta = 1^2 + 4(3-1) = 9, \theta_1 = \frac{-1+3}{2} = 1, \theta_2 = \frac{-1-3}{2} = -2.$$

$$m_1 = \frac{1}{2} \left[9 - \frac{6+9(-1)}{3} \right] = 5, m_2 = \frac{1}{2} \left[9 + \frac{6+9(-1)}{3} \right] = 4.$$

In general the Kneser graph $\text{Kn}(n, 2)$ has parameters

$$\left(\binom{n}{2}, \binom{n-2}{2}, \binom{n-4}{2}, \binom{n-3}{2} \right).$$

Recall that for any integer m , $\binom{m+1}{2} - \binom{m}{2} = m$ (easily verified by a calculation or by a counting exercise). Thus $\lambda - \mu = 4 - n$ for the $\text{Kn}(n, 2)$, and $k - \mu = n - 3$. Then

$$\Delta = (4-n)^2 + 4(n-3) = n^2 - 8n + 16 + 4n - 12 = n^2 - 4n + 4 = (n-2)^2,$$

so Δ is a square. The eigenvalues are

$$\theta_1 = \frac{(4-n) + (n-2)}{2} = 1, \theta_2 = \frac{(4-n) - (n-2)}{2} = 3-n.$$

The multiplicities are given by

$$\begin{aligned}
m_1 + m_2 &= \binom{n}{2} - 1 \\
m_1(1) + m_2(3 - n) &= -\binom{n-2}{2} \\
\implies m_2(n-2) &= \binom{n}{2} + \binom{n-2}{2} - 1 \\
&= \frac{n(n-1) + (n-2)(n-3) - 2}{2} \\
&= \frac{2n^2 - 6n + 4}{2} \\
&= n^2 - 3n + 2 \\
&= (n-2)(n-1) \implies m_2 = n-1.
\end{aligned}$$

Then $m_1 = \binom{n}{2} - n$.

Families of examples of the second type are a bit harder to construct, although one example is the cycle C_5 , which has parameters $(5, 2, 0, 1)$. In this graph $\Delta = 1 + 4(2 - 1) = 5$ is not a square, and

$$2k + (n-1)(\lambda - \mu) = 4 + 4(-1) = 0.$$

The eigenvalues are $\frac{-1 \pm \sqrt{5}}{2}$, which are irrational, both appearing with multiplicity 2.

The graph C_5 does belong to an infinite family of strongly regular graphs known as the *Paley* graphs, which are constructed from finite fields. A Paley graph on p vertices exists for every p with the property that p is a power of some prime and $p \equiv 1 \pmod{4}$. We will only consider the case where p is prime, examples of primes of the form $4t + 1$ are 5, 13, 17 etc.

For such a prime p , let \mathbb{F}_p denote the finite field $\mathbb{Z}/p\mathbb{Z}$ of integers modulo p . The elements of \mathbb{F}_p are $0, 1, \dots, p-1$, with addition and multiplication modulo p . The non-zero elements form a group under multiplication, and this group is cyclic of order $p-1 = 4\mu$, because all finite subgroups of multiplicative groups of fields are cyclic. This means that there is an element x of \mathbb{F}_p , with the property that $x^{p-1} = 1$ and the powers x, x^2, \dots, x^{p-1} are the distinct non-zero elements of \mathbb{F}_p in some order. Note that

$$(x^{\frac{p-1}{2}})^2 = 1,$$

which means that $x^{\frac{p-1}{2}}$ is a square root of 1 in \mathbb{F}_p that is different from 1, so it is -1 . Then (since $p-1$ is a multiple of 4), we have that $x^{\frac{p-1}{4}}$ is an element of \mathbb{F}_p whose square is -1 . Thus -1 is a square in \mathbb{F}_p if (and only if) $p \equiv 1 \pmod{4}$. The squares in \mathbb{F}_p are the even powers of x (and 0), they account for $\frac{p+1}{2}$ of the p elements. Note also that the set of squares in \mathbb{F}_p is closed under multiplication, since the product of two squares is a square.

Example 4.3.1. If $p = 13$, the squares in \mathbb{F}_{13} are $-4, -3, -1, 0, 1, 3, 4$:

$$0^2, 1^2 = 1, 2^2 = 4, 3^2 = -4, 4^2 = 3, 5^2 = -1, 6^2 = -3, 7^2 = -3, 8^2 = -1, 9^2 = 3, 10^2 = -4, 11^2 = 4, 12^2 = 1$$

Definition 4.3.2. For $p = 4\mu + 1$, the *Paley graph* $P(p)$ is defined to be the graph whose vertices are labelled by the elements of \mathbb{F}_p and in which two vertices are adjacent if the difference of the corresponding elements of \mathbb{F}_p is a square.

The fact that -1 is a square means that an element is a square if and only if its negative is, so this adjacency condition does not depend on which element is subtracted from the other to form the difference. The degree of each vertex of $P(p)$ is $\frac{p-1}{2}$, and $P(p)$ is a strongly regular graph with

parameters $(4\mu + 1, 2\mu, \mu - 1, \mu)$. For example $P(13)$ has parameters $(13, 6, 6, 3)$. The adjacency matrix of $P(13)$ (with rows and columns labelled 0 through 12) is

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

This is a *circulant* matrix (every row is obtained from the previous one by shifting the entries one step to the right and then wrapping the last entry to the front).

Note: it is not true that conference graphs exist of all orders n with $n \equiv 1 \pmod{4}$. For example there is no conference graph on 21 vertices.