Introduction to spectral graph theory

1 Linear Algebra Review

We write $M \in \mathbb{R}^{n\times n}$ to denote that M is an $n \times n$ matrix with real elements, and $v \in \mathbb{R}^n$ to denote that v is a vector of length n. Vectors are usually taken to be column vectors unless otherwise specified. Recall that a real matrix $M \in \mathbb{R}^{n\times n}$ represents a linear operator from $\mathbb{R}^n$ to $\mathbb{R}^n$. In other words, given any vector $v \in \mathbb{R}^n$, we think of M as a function which maps it to another vector $w \in \mathbb{R}^n$, namely the matrix-vector product. Moreover, this function is linear: $M(v_1 + v_2) = Mv_1 + Mv_2$ for any two vectors v_1 and v_2, and $M(\lambda v) = \lambda Mv$ for any real number λ and any vector v. The perspective of matrices as linear operators is important to keep in mind throughout this course.

If there is a $\lambda \in \mathbb{C}$ such that $Mv = \lambda v$, then we say that v is an (right) eigenvector of M corresponding to the eigenvalue λ. Likewise, w is a left eigenvector of M with eigenvalue λ if $w^T M = \lambda w^T$. Here w^T denotes the transpose of w. We will write $M^\dagger$ and $w^\dagger$ to denote the Hermitian conjugates of M and w, namely the complex conjugate of the transpose of M and w respectively. For example,

$$M = \begin{pmatrix} 1 & 2+i \\ 3-i & i \end{pmatrix} \Rightarrow M^\dagger = \begin{pmatrix} 1 & 3+i \\ 2-i & -i \end{pmatrix}.$$

In this course, we will mainly be dealing with real vectors and matrices; for these, the Hermitian conjugate is the same as the transpose.

We now recall some basic facts about eigenvalues and eigenvectors. The eigenvalues of $M \in \mathbb{R}^{n\times n}$ are the zeros of its characteristic polynomial, $\det(M - xI)$. Since this polynomial is of degree n, it has n roots, which may be real or complex, and which may be repeated. Is there an eigenvector corresponding to each eigenvalue? If λ is an eigenvalue, then $\det(M -$

$\lambda I) = 0$ and so the matrix $M - \lambda I$ is singular. Hence, there is at least one non-zero vector v solving $(M - \lambda I)v = 0$, which implies that $Mv = \lambda v$. Thus, v is an eigenvector of M corresponding to the eigenvalue λ. However, if an eigenvalue is repeated, say has multiplicity $k > 1$, then there may not necessarily be k linearly independent eigenvectors corresponding to it. Consider, for example, the matrix

$$M = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix},$$

which has eigenvalues 0 and 0. However, all eigenvectors corresponding to this eigenvalue are of the form $\begin{pmatrix} c \\ 0 \end{pmatrix}$, i.e., there aren't two linearly independent eigenvectors.

Definitions

1. Vectors $v_1, \ldots, v_m \in \mathbb{R}^n$ are called orthonormal if they are mutually orthogonal ($v_i \cdot v_j = 0$ whenever $i \neq j$), and each of them has unit norm ($\|v_i\| = \sqrt{v_i \cdot v_i} = 1$). Here, $v_i \cdot v_j$ denotes the inner product or dot product of v_i and v_j.
2. A square matrix $V \in \mathbb{R}^{n \times n}$ is called orthogonal if $VV^T = I$.

Note that there can be at most n mutually orthogonal vectors in $\mathbb{R}^n$, and that if a set of n vectors is orthogonal, then they constitute a basis of $\mathbb{R}^n$. If a matrix V is orthogonal, then its rows are orthonormal, hence constitute an orthonormal basis (ONB) of $\mathbb{R}^n$. In particular, the rows are linearly independent, so that V has full rank and hence is invertible. Since $VV^T = I$, it follows that $V^T = V^{-1}$ and hence that $V^T V = I$ as well. (In general, V and V^T need not commute, but they do for orthogonal matrices.) Thus, the columns of V also make up an ONB of $\mathbb{R}^n$.

Remark. A matrix $U \in \mathbb{C}^{n \times n}$ is said to be unitary if $UU^\dagger = I$. All the properties stated above for orthogonal matrices also hold for unitary matrices, which are important in their own right but will not play much part in this course.

A matrix $M \in \mathbb{R}^{n \times n}$ is called symmetric if $m_{ij} = m_{ji}$ for all i and j, i.e., if $M = M^T$. A matrix $M \in \mathbb{C}^{n \times n}$ is called Hermitian if $m_{ij} = m_{ji}^*$ for all i and j, i.e., if $M = M^\dagger$. Here, z^* denotes the complex conjugate of a complex number z.

Theorem 1 (Spectral theorem for Hermitian matrices) *Let M be an $n \times n$ Hermitian matrix with real or complex elements. Then its eigenvalues $\lambda_1, \ldots, \lambda_n$ are real, and it has an orthonormal basis of eigenvectors $v_1, \ldots, v_n$. Moreover, $M = \sum_{i=1}^{n} \lambda_i v_i v_i^\dagger$.*

Remark. We will present a partial proof of this theorem, showing that all eigenvalues are real, and that eigenvectors corresponding to distinct eigenvalues are orthogonal. If all eigenvalues are distinct, this completes the proof; if some eigenvalues are repeated, then some more work is required. If M is a real matrix, then all eigenvectors can also be chosen to be real, and Hermitian conjugates can be replaced with transposes throughout the following proof. Though we will only be working with real matrices, we present the more general statement and proof as it is no more difficult and has exactly the same structure.

Proof. Suppose λ is an eigenvalue of M and v a corresponding eigenvector, so that $Mv = \lambda v$. Then, $v^\dagger(Mv) = \lambda v^\dagger v$. On the other hand,

$$v^\dagger(Mv) = (v^\dagger M)v = (v^\dagger M^\dagger)v = (Mv)^\dagger v = (\lambda v)^\dagger v = \lambda^* v^\dagger v,$$

where λ^* is the complex conjuage of λ. We've used the fact that $M = M^\dagger$ to obtain the second equality above, and that $(\lambda v)^\dagger = \lambda^* v^\dagger$ to obtain the last equality. We have thus shown that $\lambda v = \lambda^* v$. Hence, $\lambda = \lambda^*$ (as the zero vector is not considered an eigenvector), which implies that λ is real.

Suppose v_1 and v_2 are eigenvectors corresponding to distinct eigenvalues λ_1 and λ_2. Then, $v_1^\dagger(Mv_2) = v_1^\dagger \lambda_2 v_2 = \lambda_2 v_1^\dagger v_2$. On the other hand, using the fact that $M = M^\dagger$, we have

$$(v_1^\dagger M)v_2 = (v_1^\dagger M^\dagger)v_2 = (Mv_1)^\dagger)v_2 = (\lambda_1 v_1)^\dagger v_2 = \lambda_1^* v_1^\dagger v_2.$$

Thus, we obtain that $\lambda_2 v_1^\dagger v_2 = \lambda_1^* v_1^\dagger v_2$, i.e., $(\lambda_2 - \lambda_1^*) v_1^\dagger v_2 = 0$. But $\lambda_1^* = \lambda_1$ since all eigenvalues are real, and is different from λ_2 by assumption. Hence, it follows that $v_1^\dagger v_2 = 0$, i.e., that v_1 and v_2 are orthogonal to each other.

Finally, suppose all n eigenvalues $\lambda_1, \ldots, \lambda_n$ are distinct. Then, the corresponding eigenvectors $v_1, \ldots, v_n$ are mutually orthogonal, which implies that they are linearly independent. (Exercise: prove this.) Hence, they constitute a basis of $\mathbb{R}^n$. Without loss of generality, we may assume that they all have unit norm, so that we obtain an orthonormal basis.

Now, to show that the linear operators M and $\sum_{i=1}^{n} \lambda_i v_i v_i^\dagger$ are equal to each other, we need to show that $Mv = \sum_{i=1}^{n} \lambda_i v_i v_i^\dagger v$ for all $v \in \mathbb{R}^n$. Clearly,

it suffices to show that equality holds for all v in some basis of $\mathbb{R}^n$. Let us consider the basis of eigenvectors above. Then, $Mv_j = \lambda_j v_j$, whereas

$$\sum_{i=1}^{n} \lambda_i v_i v_i^\dagger v_j = \sum_{i=1}^{n} \lambda_i v_i (v_i^\dagger v_j) = \lambda_j v_j (v_j^\dagger v_j) = \lambda_j v_j.$$

To obtain the second equality above, we have used the fact that $v_i^\dagger v_j = 0$ whenever $i \neq j$, since the vectors v_i, $i = 1, \ldots, n$, are mutually orthogonal. The last equality relies on the fact that $v_j^\dagger v_j = \|v_j\|^2 = 1$. □

Variational principle: Let M be a real, symmetric matrix. Then, its eigenvalues can be expressed in terms of the solution to an optimisation problem. The term 'variational problem' is also used for optimisation problems (especially in infinite dimensions), whence the name for this way of describing eigenvalues. It is also referred to as the Rayleigh-Ritz principle.

Theorem 2 *Let $\lambda_1 \leq \lambda_2 \leq \ldots \leq \lambda_n$ denote the eigenvalues of a symmetric matrix $M \in \mathbb{R}^{n\times n}$. Then,*

$$\lambda_1 = \min_{x\in\mathbb{R}^n\setminus\{\mathbf{0}\}} \frac{x^T M x}{x^T x}, \quad \lambda_n = \max_{x\in\mathbb{R}^n\setminus\{\mathbf{0}\}} \frac{x^T M x}{x^T x},$$

where $\mathbf{0}$ denotes the zero vector in $\mathbb{R}^n$. Moreover, if $v_1, \ldots, v_n$ are an ONB of eigenvectors of M associated with eigenvalues $\lambda_1, \ldots, \lambda_n$, then

$$\lambda_{k+1} = \min_{x\perp v_1,\ldots,v_k} \frac{x^T M x}{x^T x},$$

i.e., the minimum is taken over non-zero vectors which are orthogonal to all the first k eigenvectors (and hence to the k-dimensional subspace spanned by them).

Remark. The ratio $x^T Mx/x^T x$ is referred to as the Rayleigh quotient or the Rayleigh-Ritz ratio. Analogous to the last part of the theorem, λ_k can also be expressed as the maximum of the Rayleigh quotient over the subspace orthogonal to the eigenvectors corresponding to the $n-k$ largest eigenvalues.

Proof. Since the eigenvectors $v_1, \ldots, v_n$ of M constitute an ONB, any $x \in \mathbb{R}^n$ can be expressed as a linear combination of them, $x = \sum_{i=1}^n a_i v_i$. Moreover, $a_i = x^T v_i$. (Why?) Hence, using the spectral decomposition of M from the

previous theorem, we have

$$x^T M x = \sum_{i=1}^{n} x^T \lambda_i v_i v_i^T x = \sum_{i=1}^{n} \lambda_i a_i^2,$$

whereas

$$x^T x = \sum_{i,j=1}^{n} a_i a_j v_i^T v_j = \sum_{i=1}^{n} a_i^2.$$

But $\sum_{i=1}^n \lambda_i a_i^2 \geq \sum_{i=1}^n \lambda_1 a_i^2$, so it follows from the above that $x^T M x \geq \lambda_1 x^T x$ for all $x \in \mathbb{R}^n$. Moreover, equality holds for $x = v_1$. This establishes the first claim of the theorem.

The second claim follows similary, by observing that $\sum_{i=1}^n \lambda_i a_i^2 \leq \sum_{i=1}^n \lambda_n a_i^2$, and hence that $x^T M x \leq \lambda_n x^T x$ for all $x \in \mathbb{R}^n$, with equality for $x = v_n$.

The proof of the last part is very similar, starting from the fact that if $x \in \mathbb{R}^n$ is orthogonal to $v_1, \ldots, v_k$, then the coefficients $a_1, \ldots, a_k$ are all zero. Hence,

$$\sum_{i=1}^{n} \lambda_i a_i^2 = \sum_{i=k+1}^{n} \lambda_i a_i^2 \geq \sum_{i=k+1}^{n} \lambda_{k+1} a_i^2.$$

The rest of the proof is identical and so is omitted. □

2 Matrices associated with graphs

Let $G = (V, E)$ be an undirected graph with vertex set V and edge set E. Its adjacency matrix, A_G is defined as

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

Note that A_G is symmetric. Let $d_i = \sum_{j \in V} A_G(i,j)$. Then d_i is the degree of node i. Recall that node j is called a neighbour of node i if $(i,j) \in E$, and that the degree of a node is defined as the number of neighbours it has. Define

$$D_G = \mathrm{diag}(d_i) = \begin{pmatrix} d_1 & & \\ & \ddots & \\ & & d_n \end{pmatrix}.$$

The diffusion operator or random walk operator on G is defined as $W_G = D_G^{-1} A_G$. If $d_i = 0$, we take $W_G(i,i) = 1$ and $W_G(i,j) = 0$ for all $j \neq i$. Why the name? Observe from the definition that all rows of W_G are non-negative and sum to 1, i.e., they are probability distributions on V. Now, consider a discrete time random walk on G where, at each time step, the walker moves to one of the neighbours of its current node, chosen uniformly at random independent of its past. The position of such a random walk is a discrete time Markov chain on the state space V, and the matrix W_G is the transition probability matrix of this Markov chain.

The Laplacian of G is defined as the matrix $L_G = D_G - A_G$. (This is the definition that we will use, though the definition $L_G = I - D_G^{-1/2} A_G D_G^{-1/2}$ is also in vogue. Properties of the Laplacian derived using either definition can be easily related to each other.) The row sums of L_G are all zero. The Laplacian matrix also admits a random walk interpretation. Consider a continuous time random walk where the walker has unit rate of moving along each edge. Thus, if the walker is at a vertex v, it stays there for a random time which is exponentially distributed with parameter d_v (and hence mean $1/d_v$) before moving to one of the neighbours of v chosen uniformly at random, independent of the past and of the time spent in v. This random walk is closely related to the discrete time random walk above, the main difference being that the time spent at each node during a visit is random, and it is smaller on average at high degree nodes than at low degree nodes. The position of this random walk is a continuous time Markov chain on the state space V, with rate matrix or infinitesimal generator $Q = -L_G$.

We shall now look at a few specific examples of graphs and compute the eigenvalues of their Laplacian matrix. Similar calculations can be made for the diffusion matrix. The examples we choose are very simple graphs for which it is possible to calculate the eigenvalue spectrum explicitly in closed form. We will use these specific graphs as recurring examples throughout this course. In general, no such closed form solutions are available for the spectrum of the Laplacian, and eigenvalues have to be computed numerically. The computational complexity of determining the eigenvalues of an $n \times n$ matrix is $O(n^3)$ in general, but can be significantly smaller if the matrix is sparse. While a graph on n nodes can have up to $\binom{n}{2} \approx n^2/2$ edges, many graphs encountered in practice have far fewer edges, of order n or close to n rather than n^2. In this case, their Laplacians are also sparse and the computation of their eigenvalues correspondingly easier.

2.1 Examples

Complete graph

The complete graph on n nodes, denoted K_n, is the graph where the edge between every pair of nodes is present, i.e., $(i, j) \in E$ for all $i, j \in V$ such that $i \neq j$. Thus, each vertex has degree $n-1$. The Laplacian of this graph is thus given by

$$L_{K_n} = \begin{pmatrix} n-1 & -1 & \cdots & -1 & -1 \\ -1 & n-1 & \cdots & -1 & -1 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ -1 & \cdots & -1 & n-1 & \end{pmatrix} = nI - \mathbf{1}\mathbf{1}^T,$$

where I denotes the identity matrix in n dimensions and $\mathbf{1}$ the all-1 column vector of length n.

Claim: The eigenvalues of L_{K_n} are $\lambda_1 = 0$ and $\lambda_2 = \lambda_3 = \ldots = \lambda_n = n$.

Proof. Since all the rows of L_{K_n} sum to zero, the all-1 vector $\mathbf{1}$ is an eigenvector, and the corresponding eigenvalue is 0, i.e., $L_{K_n}\mathbf{1} = \mathbf{0} = 0 \cdot \mathbf{1}$.

Now, let $x \in \mathbb{R}^n$ be any non-zero vector orthogonal to $\mathbf{1}$, i.e, $x^T\mathbf{1} = 0$. Then,

$$L_{K_n}x = nIx - \mathbf{1}\mathbf{1}^T x = nx,$$

since $\mathbf{1}^T x = 0$. But this implies that x is an eigenvector of L_{K_n} with eigenvalue n. Since this holds for all vectors x orthogonal to $\mathbf{1}$, and the subspace of $\mathbb{R}^n$ consisting of such vectors is $(n-1)$-dimensional, it follows that there are $n-1$ mutually orthogonal vectors spanning this subspace. Thus the geometric multiplicity of the eigenvalue n (defined as the number of linearly independent eigenvectors corresponding to this eigenvalue) is $n-1$. It is known that the algebraic multiplicity of any eigenvalue (the number of times it appears as a root of the characteristic equation) is at least as big as the geometric multiplicity.

Hence, the eigenvalue n has (algebraic) multiplicity at least equal to $n-1$, while the eigenvalue 0 has multiplicity at least equal to 1. But the characteristic polynomial is of degree n, and therefore has exactly n zeros. This implies that the eigenvalue n must have multiplicity exactly equal to $n-1$, and the eigenvalue 0 must have multiplicity exactly equal to 1. This completes the proof of the claim. □

Star graph

The star graph on n nodes, denoted S_n, consists of a central hub and $n-1$ leaves. There is an edge between the hub and each leaf, but no edges between leaves. Thus, its adjacency matrix and Laplacian are

$$A_{S_n} = \begin{pmatrix} 0 & \mathbf{1}_{n-1}^T \\ \mathbf{1}_{n-1} &)_{n-1} \end{pmatrix} \text{ and } L_{S_n} = \begin{pmatrix} n-1 & -\mathbf{1}_{n-1}^T \\ -\mathbf{1}_{n-1} & I_{n-1}, \end{pmatrix}$$

where $\mathbf{1}_{n-1}$ denotes the all-1 vector of length $n-1$, while O_{n-1} and I_{n-1} denote the $(n-1)\times(n-1)$ all-zero matrix and identity matrix respectively.

Claim: The eigenvalues of L_{S_n} are $\lambda_1 = 0$, $\lambda_2 = \ldots = \lambda_{n-1} = 1$ and $\lambda_n = n$.

Proof. As before, $\mathbf{1}$ is an eigenvector of the Laplacian with eigenvalue 0, because all row sums of the Laplacian are zero. It can also be easily verified that $v_n = (n-1\ -1 \ldots -1)^T$ is an eigenvector with eigenvalue n.

Let e_i denote the i^{th} unit vector, namely the vector with a 1 in the i^{th} position and zeros elsewhere. Now, for $i \geq 2$, $L_{S_n} e_i = -e_1 + e_i$. So, for $i, j \geq 2$ with $i \neq j$, $L_{S_n}(e_i - e_j) = e_i - e_j$. In other words, all the vectors $e_i - e_j$ with $2 \leq i, j \leq n$ and $i \neq j$ are eigenvectors with eigenvalue 1. There are $(n-1)(n-2)$ such vectors, but clearly they can't all be linearly independent. So, let's first find a maximal linearly indepedent subset of them. It is easy to see that the vectors $e_2 - e_3$, $e_3 - e_4$, …, $e_{n-1} - e_n$ form a linearly independent set, of cardinality $n-2$. Hence, the eigenvalue 1 has multiplicity at least $n-2$. As there are n eigenvalues in total, and we have already found two of them (namely 0 and n), it follows that these are all the eigenvalues, which completes the proof of the claim. □

Cycle graph

The cycle on n nodes, denoted C_n, has n nodes arranged in a ring, with each node being connected to its two immediate neighbours, one on each side. Its adjacency matrix is given by

$$A_{C_n}(i,j) = \begin{cases} 1, & \text{if } |i-j| = 1 \text{ or } n-1, \\ 0, & \text{otherwise} \end{cases}$$

and its Laplacian by $L_{C_n} = 2I - A_{C_n}$ since all nodes have degree 2.

Claim: The eigenvalues of L_{C_n} are $\lambda_k = 2 - 2\cos\frac{2\pi(k-1)}{n}$, $k = 1, \ldots, n$.

Proof. As before, $\mathbf{1}$ is an eigenvector with eigenvalue 0. Next, for $k =$

$2, \ldots, n$, define $z_k \in \mathbb{C}^n$ to be the vector with components

$$z_k(\ell) = \exp(i\frac{2\pi(k-1)\ell}{n}), \quad \ell = 1, \ldots, n,$$

where $i = \sqrt{-1}$. We have

$$\begin{aligned}(L_{C_n} z_k)(\ell) &= 2e^{i2\pi(k-1)\ell/n} - e^{i2\pi(k-1)(\ell-1)/n} - e^{i2\pi(k-1)(\ell+1)/n} \\ &= e^{i2\pi(k-1)\ell/n}\Big(2 - e^{\frac{i2\pi(k-1)}{n}} - e^{-\frac{i2\pi(k-1)}{n}}\Big) \\ &= z_k(\ell)\Big(2 - 2\cos\frac{2\pi(k-1)}{n}\Big).\end{aligned}$$

Hence, z_k is an eigenvector of L_{C_n} with eigenvalue $2 - 2\cos\frac{2\pi(k-1)}{n}$ for each $k = 2, \ldots, n$. We have thus identified n distinct eigenvalues. This completes the proof of the claim. □

2.2 Properties of the Laplacian

In all the examples above, we saw that $\mathbf{1}$ was an eigenvector of the Laplacian with eigenvalue 0. Is this true of the Laplacian of any graph? The answer is yes, and follows immediately from the fact that rows of a Laplacian always sum to zero. This, in turn, follows from the definition $L_G = D_G - A_G$, as $D_G(i,i) = d_i = \sum_{j \in V} A_G(i,j)$.

In the examples above, we also saw that all the other eigenvalues were strictly positive. Is this also a general property valid for all graphs? Not quite, as we shall in the next subsection, but all eigenvalues are non-negative. To see why, observe that for all $x \in \mathbb{R}^n$,

$$\begin{aligned}\sum_{(i,j)\in E} (x_i - x_j)^2 &= \frac{1}{2}\sum_{i\in V}\sum_{j\in V:(i,j)\in E} (x_i - x_j)^2 \\ &= \frac{1}{2}\sum_{i\in V}\sum_{j\in V:(i,j)\in E} (x_i^2 - 2x_i x_j + x_j^2) \\ &= \frac{1}{2}\Big(\sum_{i\in V} x_i^2 d_i + \sum_{j\in V} x_j^2 d_j\Big) - \sum_{i\in V}\sum_{j\in V} x_i a_{ij} x_j \\ &= \sum_{i\in V} d_i x_i^2 - \sum_{i,j\in V} a_{ij} x_i x_j \;=\; x^T L_G x.\end{aligned}$$

The first equality holds because the sum on the right counts each edge twice, once as (i,j) and once as (j,i). The third equality uses the definition of a_{ij}

as 1 if $(i,j) \in E$ and 0 otherwise, and the fact that $\sum_{j:(i,j)\in E} 1 = d_i$. The last equality follows from the definition of L_G as $D_G - A_G$.

Thus, we have shown that $x^T L_G x$ is a weighted sum of squares. Hence, it is non-negative for all $x \in \mathbb{R}^n$. Therefore, by the variational characterisation of eigenvalues,

$$\lambda_2 = \min_{x \perp \mathbf{1}} \frac{x^T L_G x}{x^T x} \geq 0.$$

As the eigenvalues are arranged in increasing order, it follows that all other eigenvalues are also non-negative.

Definition: A real symmetric matrix M is said to be positive semi-definite (p.s.d.) if $x^T M x \geq 0$ for all $x \in \mathbb{R}^n$, and positive definite (p.d.) if $x^T M x > 0$ for all $x \neq \mathbf{0}$, where $\mathbf{0}$ denotes the zero vector.

A real symmetric matrix M is p.s.d. (respectively p.d.) if and only if all its eigenvalues are non-negative (respectively positive). We have shown above that the Laplacian matrix L_G is always positive semi-definite.

2.3 Eigenvalues and connectivity

In this section, we show that Laplacian of a graph has as many eigenvalues equal to zero as there are connected components in the graph. In particular, if the graph is connected, i.e., it consists of a single connected component, then the eigenvalue zero is simple, and all other eigenvalues are strictly positive.

Let $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ be two graphs. Their union is defined as the graph $G = (V, E)$ with $V = V_1 \cup V_2$ and $E = E_1 \cup E_2$. We say that it is a disjoint union if V_1 and V_2 are disjoint sets.

Lemma 1 *Suppose G is the disjoint union of the graphs $G_1, G_2, \ldots, G_k$ and that each of the G_i is non-empty (contains at least one vertex). Then, $sp(L_G) = sp(L_{G_1} \cup \ldots \cup sp(L_{G_k})$, where $sp(M)$ denotes the set of eigenvalues of M. The union above is taken with repeated eigenvalues being retained as distinct. (Thus, for example, $\{0,1\} \cup \{1,2\} = \{0,1,1,2\}$.)*

Proof. By the assumtion that the graphs $G_1, \ldots, G_k$ are disjoint, the Lapla-

cian L_G is block-diagonal,

$$L_G = \begin{pmatrix} L_{G_1} & 0 & \dots & 0 \\ 0 & L_{G_2} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & \dots & 0 & L_{G_k} \end{pmatrix}.$$

Let v_{1i} and v_{2j} be eigenvectors of L_{G_1} and L_{G_2} respectively, corresponding to the eigenvalues λ_{1i} and λ_{2j} (which could be the same). Then, it is clear that

$$\begin{pmatrix} v_{1i} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{pmatrix} \text{ and } \begin{pmatrix} \mathbf{0} \\ v_{2j} \\ \vdots \\ \mathbf{0} \end{pmatrix}$$

are eigenvectors of the matrix L_G corresponding to the eigenvalues λ_{1i} and λ_{2j}, and that they are orthogonal to each other. (Here, the same notation $\mathbf{0}$ is used to denote all-zero vectors of various sizes corresponding to the different subgraphs $G_1, \dots, , G_k$.)

By the same argument, any eigenvector of any of the L_{G_i} can be modified (by padding zeros appropriately) to yield an eigenvector of L_G with the same eigenvalue. Moreover, if $|V_i| = n_i$, then L_{G_i} has an ONB of n_i eigenvectors by the spectral theorem. Thus, the above procedure yields a set of $n_1 + n_2 + \dots + n_k = n$ orthonormal eigenvectors of L_G and n associated eigenvalues. Since L_G has exactly n eigenvalues, we have identified all of them, and they are given by the union of the eigenvalues of the different L_{G_i} (respecting multiplicity), as claimed. □

Our next result says that, if a graph is connected, all the eigenvalues of its Laplacian except the smallest one (which is always zero) are strictly positive.

Lemma 2 *Suppose G is connected, and let $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ denote the eigenvalues of its Laplacian. Then $\lambda_2 > 0$.*

Proof. We saw earlier that $x^T L_G x = \sum_{(i,j)\in E}(x_i - x_j)^2$ and also that $\mathbf{1}$ is an eigenvector corresponding to the smallest eigenvalue, $\lambda_1 = 0$. Hence, by the variational characterisation of eigenvalues,

$$\lambda_2 = \min_{x \perp \mathbf{1}} \frac{x^T L_G x}{x^T x} = \min_{x \perp \mathbf{1}, \|x\|=1} x^T L_G x.$$

To obtain the second equality, note that the Rayleigh quotient is unchanged if x is multiplied by any scalar $c \neq 0$, as both numerator and denominator get scaled by the same factor c^2. Hence, we may restrict the minimum to the subset satisfying $\|x\| = 1$ without loss of generality. But, for such x, $x^T x = \|x\|^2 = 1$, and so the ratio becomes $x^T L_G x$.

Now, if x is non-zero and orthogonal to $\mathbf{1}$, then x can't be a constant vector. Hence, there must be some $(i, j) \in E$ for which $x_i \neq x_j$. (This is where we use the fact that the graph is connected.) But this implies that $\sum_{(i,j)\in E}(x_i - x_j)^2 > 0$, and this is true for any vector x which isn't a constant. It follows that $\min_{x\perp\mathbf{1},\|x\|=1} x^T L_G x$ is strictly positive, i.e., that $\lambda_2 > 0$. (Formally, we are using the fact that we are seeking the minimum of a continous function over a compact set, and hence that the minimum is attained at some x in this set. If you haven't taken a course in Analysis that covers this result, please take it for granted. Hopefully, it is intuitive enough.) □

We now put together the above two lemmas to describe the spectrum of the Laplacian of an arbitrary graph. The proof is omitted as it is immediate from the above two lemmas.

Lemma 3 *Let $\lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_n$ denote the eigenvalues of the Laplacian of a graph G. If G consists of k connected components, then $\lambda_1 = \lambda_2 = \ldots = \lambda_k = 0$, and all other eigenvalues are strictly positive.*