

Spectral Theory of Unsigned
and Signed Graphs
Applications to Graph Clustering

Some Slides

Jean Gallier
Department of Computer and Information Science
University of Pennsylvania
Philadelphia, PA 19104, USA
e-mail: jean@cis.upenn.edu

© Jean Gallier

December 18, 2017

Contents

1	Graphs and Graph Laplacians	5
1.1	Directed, Undirected, and Weighted Graphs	5
1.2	Laplacian Matrices of Graphs	22
2	Spectral Graph Drawing	37
2.1	Graph Drawing and Energy Minimization	37
2.2	Examples of Graph Drawings	45
3	Graph Clustering	55
3.1	Graph Clustering Using Normalized Cuts .	55
3.2	Special Case: 2-Way Clustering Using Normalized Cuts	60
3.3	K -Way Clustering Using Normalized Cuts	79
3.4	Discrete Solution Close to a Continuous Approximation	116
4	Signed Graphs	159
4.1	Signed Graphs and Signed Laplacians . . .	159
4.2	Signed Normalized Cuts	169

4.3	Balanced Graphs	176
4.4	K -Way Clustering of Signed Graphs . . .	187
4.5	Signed Graph Drawing	191
5	Graph Clustering Using Ratio Cuts	201
A	Rayleigh Ratios and the Courant-Fischer Theorem	211
	Bibliography	217

Chapter 1


Graphs and Graph Laplacians

1.1 Directed Graphs, Undirected Graphs, Incidence Matrices, Adjacency Matrices, Weighted Graphs

Definition 1.1. A *directed graph* is a pair $G = (V, E)$, where $V = \{v_1, \dots, v_m\}$ is a set of *nodes* or *vertices*, and $E \subseteq V \times V$ is a set of ordered pairs of distinct nodes (that is, pairs $(u, v) \in V \times V$ with $u \neq v$), called *edges*. Given any edge $e = (u, v)$, we let $s(e) = u$ be the *source* of e and $t(e) = v$ be the *target* of e .

Remark: Since an edge is a pair (u, v) with $u \neq v$, self-loops are not allowed.

Also, there is at most one edge from a node u to a node v . Such graphs are sometimes called *simple graphs*.

Figure 1.1: Graph G_1 .

For every node $v \in V$, the *degree* $d(v)$ of v is the number of edges leaving or entering v :

$$d(v) = |\{u \in V \mid (v, u) \in E \text{ or } (u, v) \in E\}|.$$

We abbreviate $d(v_i)$ as d_i . The *degree matrix* $D(G)$, is the diagonal matrix

$$D(G) = \text{diag}(d_1, \dots, d_m).$$

For example, for graph G_1 , we have

$$D(G_1) = \begin{pmatrix} 2 & 0 & 0 & 0 & 0 \\ 0 & 4 & 0 & 0 & 0 \\ 0 & 0 & 3 & 0 & 0 \\ 0 & 0 & 0 & 3 & 0 \\ 0 & 0 & 0 & 0 & 2 \end{pmatrix}.$$

Unless confusion arises, we write D instead of $D(G)$.

Definition 1.2. Given a directed graph $G = (V, E)$, for any two nodes $u, v \in V$, a *path from u to v* is a sequence of nodes (v_0, v_1, \dots, v_k) such that $v_0 = u$, $v_k = v$, and (v_i, v_{i+1}) is an edge in E for all i with $0 \leq i \leq k - 1$. The integer k is the *length* of the path. A path is *closed* if $u = v$. The graph G is *strongly connected* if for any two distinct nodes $u, v \in V$, there is a path from u to v and there is a path from v to u .

Remark: The terminology *walk* is often used instead of *path*, the word path being reserved to the case where the nodes v_i are all distinct, except that $v_0 = v_k$ when the path is closed.

The binary relation on $V \times V$ defined so that u and v are related iff there is a path from u to v and there is a path from v to u is an equivalence relation whose equivalence classes are called the *strongly connected components* of G .

Definition 1.3. Given a directed graph $G = (V, E)$, with $V = \{v_1, \dots, v_m\}$, if $E = \{e_1, \dots, e_n\}$, then the *incidence matrix* $B(G)$ of G is the $m \times n$ matrix whose entries b_{ij} are given by


$$b_{ij} = \begin{cases} +1 & \text{if } s(e_j) = v_i \\ -1 & \text{if } t(e_j) = v_i \\ 0 & \text{otherwise.} \end{cases}$$

Here is the incidence matrix of the graph G_1 :

$$B = \begin{pmatrix} 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ -1 & 0 & -1 & -1 & 1 & 0 & 0 \\ 0 & -1 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & -1 & -1 \\ 0 & 0 & 0 & 0 & -1 & 1 & 0 \end{pmatrix}.$$

Again, unless confusion arises, we write B instead of $B(G)$.

Remark: Some authors adopt the opposite convention of sign in defining the incidence matrix, which means that their incidence matrix is $-B$.

Figure 1.2: The undirected graph G_2 .

Undirected graphs are obtained from directed graphs by forgetting the orientation of the edges.

Definition 1.4. A *graph* (or *undirected graph*) is a pair $G = (V, E)$, where $V = \{v_1, \dots, v_m\}$ is a set of *nodes* or *vertices*, and E is a set of two-element subsets of V (that is, subsets $\{u, v\}$, with $u, v \in V$ and $u \neq v$), called *edges*.

Remark: Since an edge is a set $\{u, v\}$, we have $u \neq v$, so self-loops are not allowed. Also, for every set of nodes $\{u, v\}$, there is at most one edge between u and v .

As in the case of directed graphs, such graphs are sometimes called *simple graphs*.

For every node $v \in V$, the *degree* $d(v)$ of v is the number of edges incident to v :

$$d(v) = |\{u \in V \mid \{u, v\} \in E\}|.$$

The degree matrix D is defined as before.

Definition 1.5. Given a (undirected) graph $G = (V, E)$, for any two nodes $u, v \in V$, a *path from u to v* is a sequence of nodes (v_0, v_1, \dots, v_k) such that $v_0 = u$, $v_k = v$, and $\{v_i, v_{i+1}\}$ is an edge in E for all i with $0 \leq i \leq k-1$. The integer k is the *length* of the path. A path is *closed* if $u = v$. The graph G is *connected* if for any two distinct nodes $u, v \in V$, there is a path from u to v .

Remark: The terminology *walk* or *chain* is often used instead of *path*, the word *path* being reserved to the case where the nodes v_i are all distinct, except that $v_0 = v_k$ when the path is closed.

The binary relation on $V \times V$ defined so that u and v are related iff there is a path from u to v is an equivalence relation whose equivalence classes are called the *connected components* of G .

The notion of incidence matrix for an undirected graph is not as useful as in the case of directed graphs

Definition 1.6. Given a graph $G = (V, E)$, with $V = \{v_1, \dots, v_m\}$, if $E = \{e_1, \dots, e_n\}$, then the *incidence matrix* $B(G)$ of G is the $m \times n$ matrix whose entries b_{ij} are given by

$$b_{ij} = \begin{cases} +1 & \text{if } e_j = \{v_i, v_k\} \text{ for some } k \\ 0 & \text{otherwise.} \end{cases}$$

Unlike the case of directed graphs, the entries in the incidence matrix of a graph (undirected) are nonnegative. We usually write B instead of $B(G)$.

The notion of adjacency matrix is basically the same for directed or undirected graphs.

Definition 1.7. Given a directed or undirected graph $G = (V, E)$, with $V = \{v_1, \dots, v_m\}$, the *adjacency matrix* $A(G)$ of G is the symmetric $m \times m$ matrix (a_{ij}) such that

(1) If G is directed, then

$$a_{ij} = \begin{cases} 1 & \text{if there is some edge } (v_i, v_j) \in E \\ & \text{or some edge } (v_j, v_i) \in E \\ 0 & \text{otherwise.} \end{cases}$$

(2) Else if G is undirected, then

$$a_{ij} = \begin{cases} 1 & \text{if there is some edge } \{v_i, v_j\} \in E \\ 0 & \text{otherwise.} \end{cases}$$

As usual, unless confusion arises, we write A instead of $A(G)$.

Here is the adjacency matrix of both graphs G_1 and G_2 :

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

If $G = (V, E)$ is a directed or an undirected graph, given a node $u \in V$, any node $v \in V$ such that there is an edge (u, v) in the directed case or $\{u, v\}$ in the undirected case is called *adjacent to v* , and we often use the notation

$$u \sim v.$$

Observe that the binary relation \sim is symmetric when G is an undirected graph, but in general it is not symmetric when G is a directed graph.

If $G = (V, E)$ is an undirected graph, the adjacency matrix A of G can be viewed as a linear map from \mathbb{R}^V to \mathbb{R}^V , such that for all $x \in \mathbb{R}^m$, we have

$$(Ax)_i = \sum_{j \sim i} x_j;$$

that is, the value of Ax at v_i is the sum of the values of x at the nodes v_j adjacent to v_i .

The adjacency matrix can be viewed as a *diffusion operator*.

This observation yields a geometric interpretation of what it means for a vector $x \in \mathbb{R}^m$ to be an eigenvector of A associated with some eigenvalue λ ; we must have

$$\lambda x_i = \sum_{j \sim i} x_j, \quad i = 1, \dots, m,$$

which means that the the sum of the values of x assigned to the nodes v_j adjacent to v_i is equal to λ times the value of x at v_i .

Definition 1.8. Given any undirected graph $G = (V, E)$, an *orientation* of G is a function $\sigma: E \rightarrow V \times V$ assigning a source and a target to every edge in E , which means that for every edge $\{u, v\} \in E$, either $\sigma(\{u, v\}) = (u, v)$ or $\sigma(\{u, v\}) = (v, u)$. The *oriented graph* G^σ obtained from G by applying the orientation σ is the directed graph $G^\sigma = (V, E^\sigma)$, with $E^\sigma = \sigma(E)$.

Proposition 1.1. *Let $G = (V, E)$ be any undirected graph with m vertices, n edges, and c connected components. For any orientation σ of G , if B is the incidence matrix of the oriented graph G^σ , then $c = \dim(\text{Ker}(B^\top))$, and B has rank $m - c$. Furthermore, the nullspace of B^\top has a basis consisting of indicator vectors of the connected components of G ; that is, vectors (z_1, \dots, z_m) such that $z_j = 1$ iff v_j is in the i th component K_i of G , and $z_j = 0$ otherwise.*

Following common practice, we denote by $\mathbf{1}$ the (column) vector whose components are all equal to 1. Observe that

$$B^\top \mathbf{1} = 0.$$

According to Proposition 1.1, the graph G is connected iff B has rank $m - 1$ iff the nullspace of B^\top is the one-dimensional space spanned by $\mathbf{1}$.

In many applications, the notion of graph needs to be generalized to capture the intuitive idea that two nodes u and v are linked with a degree of certainty (or strength).

Thus, we assign a nonnegative weight w_{ij} to an edge $\{v_i, v_j\}$; the smaller w_{ij} is, the weaker is the link (or similarity) between v_i and v_j , and the greater w_{ij} is, the stronger is the link (or similarity) between v_i and v_j .

Definition 1.9. A *weighted graph* is a pair $G = (V, W)$, where $V = \{v_1, \dots, v_m\}$ is a set of *nodes* or *vertices*, and W is a symmetric matrix called the *weight matrix*, such that $w_{ij} \geq 0$ for all $i, j \in \{1, \dots, m\}$, and $w_{ii} = 0$ for $i = 1, \dots, m$. We say that a set $\{v_i, v_j\}$ is an edge iff $w_{ij} > 0$. The corresponding (undirected) graph (V, E) with $E = \{\{v_i, v_j\} \mid w_{ij} > 0\}$, is called the *underlying graph* of G .

Remark: Since $w_{ii} = 0$, these graphs have no self-loops. We can think of the matrix W as a generalized adjacency matrix. The case where $w_{ij} \in \{0, 1\}$ is equivalent to the notion of a graph as in Definition 1.4.

We can think of the weight w_{ij} of an edge $\{v_i, v_j\}$ as a degree of similarity (or affinity) in an image, or a cost in a network.

An example of a weighted graph is shown in Figure 1.3. The thickness of an edge corresponds to the magnitude of its weight.




Figure 1.3: A weighted graph.

For every node $v_i \in V$, the *degree* $d(v_i)$ of v_i is the sum of the weights of the edges adjacent to v_i :

$$d(v_i) = \sum_{j=1}^m w_{ij}.$$

Note that in the above sum, only nodes v_j such that there is an edge $\{v_i, v_j\}$ have a nonzero contribution. Such nodes are said to be *adjacent* to v_i , and we write $v_i \sim v_j$.

The degree matrix D is defined as before, namely by $D = \text{diag}(d(v_1), \dots, d(v_m))$.

The weight matrix W can be viewed as a linear map from \mathbb{R}^V to itself. For all $x \in \mathbb{R}^m$, we have

$$(Wx)_i = \sum_{j \sim i} w_{ij}x_j;$$

that is, the value of Wx at v_i is the weighted sum of the values of x at the nodes v_j adjacent to v_i .

Observe that $W\mathbf{1}$ is the (column) vector $(d(v_1), \dots, d(v_m))$ consisting of the degrees of the nodes of the graph.

Given any subset of nodes $A \subseteq V$, we define the *volume* $\text{vol}(A)$ of A as the sum of the weights of all edges adjacent to nodes in A :

$$\text{vol}(A) = \sum_{v_i \in A} d(v_i) = \sum_{v_i \in A} \sum_{j=1}^m w_{ij}.$$




Figure 1.4: Degree and volume.

Observe that $\text{vol}(A) = 0$ if A consists of isolated vertices, that is, if $w_{ij} = 0$ for all $v_i \in A$. Thus, it is best to assume that G does not have isolated vertices.

Given any two subset $A, B \subseteq V$ (not necessarily distinct), we define $\text{links}(A, B)$ by

$$\text{links}(A, B) = \sum_{v_i \in A, v_j \in B} w_{ij}.$$

Since the matrix W is symmetric, we have

$$\text{links}(A, B) = \text{links}(B, A),$$

and observe that $\text{vol}(A) = \text{links}(A, V)$.

The quantity $\text{links}(A, \bar{A}) = \text{links}(\bar{A}, A)$, where $\bar{A} = V - A$ denotes the complement of A in V , measures *how many links escape from A (and \bar{A})*, and the quantity $\text{links}(A, A)$ measures *how many links stay within A itself*.

The quantity

$$\text{cut}(A) = \text{links}(A, \bar{A})$$

is often called the *cut* of A , and the quantity

$$\text{assoc}(A) = \text{links}(A, A)$$

is often called the *association* of A . Clearly,

$$\text{cut}(A) + \text{assoc}(A) = \text{vol}(A).$$




Figure 1.5: A Cut involving the set of nodes in the center and the nodes on the perimeter.

We now define the most important concept of these notes: The Laplacian matrix of a graph. Actually, as we will see, it comes in several flavors.

1.2 Laplacian Matrices of Graphs

Let us begin with directed graphs, although as we will see, graph Laplacians are fundamentally associated with undirected graph.

The key proposition below shows how given an undirected graph G , for any orientation σ of G , $B^\sigma(B^\sigma)^\top$ relates to the adjacency matrix A (where B^σ is the incidence matrix of the directed graph G^σ).

We reproduce the proof in Gallier [5] (see also Godsil and Royle [7]).

Proposition 1.2. *Given any undirected graph G , for any orientation σ of G , if B^σ is the incidence matrix of the directed graph G^σ , A is the adjacency matrix of G^σ , and D is the degree matrix such that $D_{ii} = d(v_i)$, then*

$$B^\sigma(B^\sigma)^\top = D - A.$$

Consequently, $L = B^\sigma(B^\sigma)^\top$ is independent of the orientation σ of G , and $D - A$ is symmetric and positive semidefinite; that is, the eigenvalues of $D - A$ are real and nonnegative.

The matrix $L = B^\sigma(B^\sigma)^\top = D - A$ is called the *(un-normalized) graph Laplacian* of the graph G^σ .

For example, the graph Laplacian of graph G_1 is

$$L = \begin{pmatrix} 2 & -1 & -1 & 0 & 0 \\ -1 & 4 & -1 & -1 & -1 \\ -1 & -1 & 3 & -1 & 0 \\ 0 & -1 & -1 & 3 & -1 \\ 0 & -1 & 0 & -1 & 2 \end{pmatrix}.$$

The *(unnormalized) graph Laplacian* of an undirected graph $G = (V, E)$ is defined by

$$L = D - A.$$

Observe that each row of L sums to zero (because $(B^\sigma)^\top \mathbf{1} = 0$). Consequently, the vector $\mathbf{1}$ is in the nullspace of L .

Remark: With the unoriented version of the incidence matrix (see Definition 1.6), it can be shown that

$$BB^{\top} = D + A.$$

The natural generalization of the notion of graph Laplacian to weighted graphs is this:

Definition 1.10. Given any weighted graph $G = (V, W)$ with $V = \{v_1, \dots, v_m\}$, the (*unnormalized*) *graph Laplacian* $L(G)$ of G is defined by

$$L(G) = D(G) - W,$$

where $D(G) = \text{diag}(d_1, \dots, d_m)$ is the degree matrix of G (a diagonal matrix), with

$$d_i = \sum_{j=1}^m w_{ij}.$$

As usual, unless confusion arises, we write L instead of $L(G)$.

The graph Laplacian can be interpreted as a linear map from \mathbb{R}^V to itself. For all $x \in \mathbb{R}^V$, we have

$$(Lx)_i = \sum_{j \sim i} w_{ij}(x_i - x_j).$$

It is clear that each row of L sums to 0, so the vector $\mathbf{1}$ is in the nullspace of L , but it is less obvious that L is positive semidefinite. One way to prove it is to generalize slightly the notion of incidence matrix.

Definition 1.11. Given a weighted graph $G = (V, W)$, with $V = \{v_1, \dots, v_m\}$, if $\{e_1, \dots, e_n\}$ are the edges of the underlying graph of G (recall that $\{v_i, v_j\}$ is an edge of this graph iff $w_{ij} > 0$), for any oriented graph G^σ obtained by giving an orientation to the underlying graph of G , the *incidence matrix* B^σ of G^σ is the $m \times n$ matrix whose entries b_{ij} are given by

$$b_{ij} = \begin{cases} +\sqrt{w_{ij}} & \text{if } s(e_j) = v_i \\ -\sqrt{w_{ij}} & \text{if } t(e_j) = v_i \\ 0 & \text{otherwise.} \end{cases}$$

For example, given the weight matrix

$$W = \begin{pmatrix} 0 & 3 & 6 & 3 \\ 3 & 0 & 0 & 3 \\ 6 & 0 & 0 & 3 \\ 3 & 3 & 3 & 0 \end{pmatrix},$$

the incidence matrix B corresponding to the orientation of the underlying graph of W where an edge (i, j) is oriented positively iff $i < j$ is

$$B = \begin{pmatrix} 1.7321 & 2.4495 & 1.7321 & 0 & 0 \\ -1.7321 & 0 & 0 & 1.7321 & 0 \\ 0 & -2.4495 & 0 & 0 & 1.7321 \\ 0 & 0 & -1.7321 & -1.7321 & -1.7321 \end{pmatrix}.$$

The reader should verify that $BB^{\top} = D - W$. This is true in general, see Proposition 1.3.

It is easy to see that Proposition 1.1 applies to the underlying graph of G .

For any oriented graph G^σ obtained from the underlying graph of G , the rank of the incidence matrix B^σ is equal to $m - c$, where c is the number of connected components of the underlying graph of G , and we have $(B^\sigma)^\top \mathbf{1} = 0$.

We also have the following version of Proposition 1.2 whose proof is immediately adapted.

Proposition 1.3. *Given any weighted graph $G = (V, W)$ with $V = \{v_1, \dots, v_m\}$, if B^σ is the incidence matrix of any oriented graph G^σ obtained from the underlying graph of G and D is the degree matrix of W , then*

$$B^\sigma (B^\sigma)^\top = D - W = L.$$

Consequently, $B^\sigma (B^\sigma)^\top$ is independent of the orientation of the underlying graph of G and $L = D - W$ is symmetric, positive, semidefinite; that is, the eigenvalues of $L = D - W$ are real and nonnegative.

Another way to prove that L is positive semidefinite is to evaluate the quadratic form $x^\top Lx$.

Proposition 1.4. *For any $m \times m$ symmetric matrix $W = (w_{ij})$, if we let $L = D - W$ where D is the degree matrix associated with W , then we have*

$$x^\top Lx = \frac{1}{2} \sum_{i,j=1}^m w_{ij}(x_i - x_j)^2 \quad \text{for all } x \in \mathbb{R}^m.$$

Consequently, $x^\top Lx$ does not depend on the diagonal entries in W , and if $w_{ij} \geq 0$ for all $i, j \in \{1, \dots, m\}$, then L is positive semidefinite.

Proposition 1.4 immediately implies the following facts:
For any weighted graph $G = (V, W)$,

1. The eigenvalues $0 = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_m$ of L are real and nonnegative, and there is an orthonormal basis of eigenvectors of L .
2. The smallest eigenvalue λ_1 of L is equal to 0, and $\mathbf{1}$ is a corresponding eigenvector.

It turns out that the dimension of the nullspace of L (the eigenspace of 0) is equal to the number of connected components of the underlying graph of G . This is an immediate consequence of Proposition Proposition 1.1 and the fact that $L = BB^\top$.

Proposition 1.5. *Let $G = (V, W)$ be a weighted graph. The number c of connected components K_1, \dots, K_c of the underlying graph of G is equal to the dimension of the nullspace of L , which is equal to the multiplicity of the eigenvalue 0. Furthermore, the nullspace of L has a basis consisting of indicator vectors of the connected components of G , that is, vectors (f_1, \dots, f_m) such that $f_j = 1$ iff $v_j \in K_i$ and $f_j = 0$ otherwise.*

Proposition 1.5 implies that if the underlying graph of G is connected, then the second eigenvalue λ_2 of L is strictly positive.

Remarkably, the eigenvalue λ_2 contains a lot of information about the graph G (assuming that $G = (V, E)$ is an undirected graph).

This was first discovered by Fiedler in 1973, and for this reason, λ_2 is often referred to as the *Fiedler number*.

For more on the properties of the Fiedler number, see Godsil and Royle [7] (Chapter 13) and Chung [3].

More generally, the spectrum $(0, \lambda_2, \dots, \lambda_m)$ of L contains a lot of information about the combinatorial structure of the graph G . Leverage of this information is the object of *spectral graph theory*.

It turns out that normalized variants of the graph Laplacian are needed, especially in applications to graph clustering.

These variants make sense only if G has no isolated vertices, which means that every row of W contains some strictly positive entry.

In this case, the degree matrix D contains positive entries, so it is invertible and $D^{-1/2}$ makes sense; namely

$$D^{-1/2} = \text{diag}(d_1^{-1/2}, \dots, d_m^{-1/2}),$$

and similarly for any real exponent α .

Definition 1.12. Given any weighted directed graph $G = (V, W)$ with no isolated vertex and with $V = \{v_1, \dots, v_m\}$, the *(normalized) graph Laplacians L_{sym} and L_{rw} of G* are defined by

$$\begin{aligned} L_{\text{sym}} &= D^{-1/2} L D^{-1/2} = I - D^{-1/2} W D^{-1/2} \\ L_{\text{rw}} &= D^{-1} L = I - D^{-1} W. \end{aligned}$$

Observe that the Laplacian $L_{\text{sym}} = D^{-1/2} L D^{-1/2}$ is a symmetric matrix (because L and $D^{-1/2}$ are symmetric) and that

$$L_{\text{rw}} = D^{-1/2} L_{\text{sym}} D^{1/2}.$$

The reason for the notation L_{rw} is that this matrix is closely related to a random walk on the graph G .

Since the unnormalized Laplacian L can be written as $L = BB^\top$, where B is the incidence matrix of any oriented graph obtained from the underlying graph of $G = (V, W)$, if we let

$$B_{\text{sym}} = D^{-1/2}B,$$

we get

$$L_{\text{sym}} = B_{\text{sym}}B_{\text{sym}}^\top.$$

In particular, for any singular decomposition $B_{\text{sym}} = U\Sigma V^\top$ of B_{sym} (with U an $m \times m$ orthogonal matrix, Σ a “diagonal” $m \times n$ matrix of singular values, and V an $n \times n$ orthogonal matrix), the eigenvalues of L_{sym} are the squares of the top m singular values of B_{sym} , and the vectors in U are orthonormal eigenvectors of L_{sym} with respect to these eigenvalues (the squares of the top m diagonal entries of Σ).

Computing the SVD of B_{sym} generally yields more accurate results than diagonalizing L_{sym} , especially when L_{sym} has eigenvalues with high multiplicity.

Proposition 1.6. *Let $G = (V, W)$ be a weighted graph without isolated vertices. The graph Laplacians, L , L_{sym} , and L_{rw} satisfy the following properties:*

(1) *The matrix L_{sym} is symmetric, positive, semidefinite. In fact,*

$$x^\top L_{\text{sym}} x = \frac{1}{2} \sum_{i,j=1}^m w_{ij} \left(\frac{x_i}{\sqrt{d_i}} - \frac{x_j}{\sqrt{d_j}} \right)^2 \quad x \in \mathbb{R}^m.$$

(2) *The normalized graph Laplacians L_{sym} and L_{rw} have the same spectrum*

($0 = \nu_1 \leq \nu_2 \leq \dots \leq \nu_m$), and a vector $u \neq 0$ is an eigenvector of L_{rw} for λ iff $D^{1/2}u$ is an eigenvector of L_{sym} for λ .

(3) *The graph Laplacians, L and L_{sym} are symmetric, positive, semidefinite.*

(4) *A vector $u \neq 0$ is a solution of the generalized eigenvalue problem $Lu = \lambda Du$ iff $D^{1/2}u$ is an eigenvector of L_{sym} for the eigenvalue λ iff u is an eigenvector of L_{rw} for the eigenvalue λ .*

- (5) *The graph Laplacians, L and L_{rw} have the same nullspace. For any vector u , we have $u \in \text{Ker}(L)$ iff $D^{1/2}u \in \text{Ker}(L_{\text{sym}})$.*
- (6) *The vector $\mathbf{1}$ is in the nullspace of L_{rw} , and $D^{1/2}\mathbf{1}$ is in the nullspace of L_{sym} .*
- (7) *For every eigenvalue ν_i of the normalized graph Laplacian L_{sym} , we have $0 \leq \nu_i \leq 2$. Furthermore, $\nu_m = 2$ iff the underlying graph of G contains a nontrivial connected bipartite component.*
- (8) *If $m \geq 2$ and if the underlying graph of G is not a complete graph, then $\nu_2 \leq 1$. Furthermore the underlying graph of G is a complete graph iff $\nu_2 = \frac{m}{m-1}$.*
- (9) *If $m \geq 2$ and if the underlying graph of G is connected then $\nu_2 > 0$.*
- (10) *If $m \geq 2$ and if the underlying graph of G has no isolated vertices, then $\nu_m \geq \frac{m}{m-1}$.*

A version of Proposition 1.5 also holds for the graph Laplacians L_{sym} and L_{rw} .

This follows easily from the fact that Proposition 1.1 applies to the underlying graph of a weighted graph. The proof is left as an exercise.

Proposition 1.7. *Let $G = (V, W)$ be a weighted graph. The number c of connected components K_1, \dots, K_c of the underlying graph of G is equal to the dimension of the nullspace of both L_{sym} and L_{rw} , which is equal to the multiplicity of the eigenvalue 0. Furthermore, the nullspace of L_{rw} has a basis consisting of indicator vectors of the connected components of G , that is, vectors (f_1, \dots, f_m) such that $f_j = 1$ iff $v_j \in K_i$ and $f_j = 0$ otherwise. For L_{sym} , a basis of the nullspace is obtained by multiplying the above basis of the nullspace of L_{rw} by $D^{1/2}$.*

