University of Edinburgh

INFR11156: Algorithmic Foundations of Data Science (2019)

Lecture 14: Spectral Sparsification of Graphs (1)

Graph sparsification is the procedure of approximating a graph G by a sparse graph H such that certain quantities between G and H are approximately preserved, see Figure 1 for example. Over the past three decades, several notions of graph sparsification have been proposed and have led numerous applications in designing approximation algorithms. Among these, a spectral sparsifier is a sparse subgraph H of an original graph G that maintains spectral properties between the Laplacian matrices of G and H. Over the past 15 years, spectral sparsification has become one of the most central components in designing efficient algorithms for a number of important optimisation problems. The formal definition of spectral sparsifiers is as follows:

Definition 1. *For any undirected graph G with n vertices and m edges, we call a subgraph H of G, with proper reweighting of the edges, a $(1+\varepsilon)$-spectral sparsifier if*

$$(1-\varepsilon)x^{\intercal}L_G x \le x^{\intercal}L_H x \le (1+\varepsilon)x^{\intercal}L_G x \tag{1}$$

holds for any $x \in \mathbb{R}^n$, where L_G and L_H are the respective Laplacian matrices of G and H.

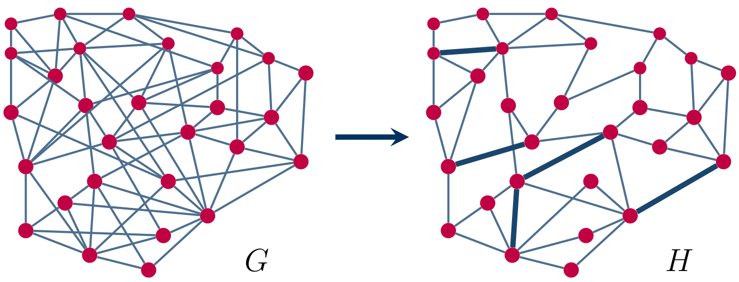

Figure 1: The graph sparsification is a reweighted subgraph H of an original graph G such that certain properties are preserved. These subgraphs are sparse, and are more space-efficient to be stored than the original graphs. The picture above uses the thickness of edges in H to represent their weights.

We will first prove that some important properties of an input graph $G(V,E,w)$ are approximately preserved in its spectral sparsifier H.

Lemma 2. *The following statements hold: (1) For $S \subseteq V$, it holds that*

$$(1-\varepsilon)\cdot w_G(S, V\setminus S) \le w_H(S,V\setminus S) \le (1+\varepsilon)\cdot w_G(S,V\setminus S),$$

where $w_G(S,V\setminus S)$ and $w_H(S,V\setminus S)$ are the cut values between S and $V\setminus S$ in G and H respectively. (2) It holds for any $1 \le i \le n$ that

$$(1-\varepsilon)\cdot\lambda_i(L_G) \le \lambda_i(L_H) \le (1+\varepsilon)\cdot\lambda_i(L_G).$$

Proof. For the first statement, let $S \subseteq V$ be an arbitrary set of vertices and let χ_S be the indicator vector of S, i.e., $\chi_S(u) = 1$ if $u \in S$, and $\chi_S(u) = 0$ otherwise. Then, it holds that

$$\chi_S^\intercal L_G \chi_S = \sum_{u \sim v} w(u,v) \left(\chi_S(u) - \chi_S(v)\right)^2 = \sum_{\substack{u \in S, v \in V \setminus S \\ u \sim v}} w(u,v) = w_G(S, V \setminus S).$$

By Definition 1 we have

$$(1-\varepsilon)\chi_S^\intercal L_G \chi_S \leq \chi_S^\intercal L_H \chi_S \leq (1+\varepsilon)\chi_S^\intercal L_G \chi_S,$$

the statement holds.

For the second statement, we apply (1) and the min-max theorem, and obtain that

$$(1-\varepsilon) \cdot x_G^\intercal L_G x_G \leq (1-\varepsilon) \cdot x_H^\intercal L_G x_H \leq x_H^\intercal L_H x_H \leq x_G^\intercal L_H x_G \leq (1+\varepsilon) \cdot x_G^\intercal L_G x_G,$$

where x_G and x_H are the eigenvectors of L_G and L_H corresponding to the i-th smallest eigenvalue. □

1 Electrical flows

Given an undirected graph G, we fix an arbitrary orientation of the edges (The specific choice of the orientation does not affect anything in our discussion). We introduce the incidence matrix $B \in \mathbb{R}^{m \times n}$ of G: the incidence matrix of G is the matrix $B \in \mathbb{R}^{m \times n}$, where the rows and columns of B are indexed by the edges and vertices of graph G, and for any edge $e \in E[G]$ and vertex v

$$B_{e,v} = \begin{cases} 1 & \text{if } v \text{ is } e\text{'s head} \\ -1 & \text{if } v \text{ is } e\text{'s tail} \\ 0 & \text{otherwise.} \end{cases}$$

We also define the diagonal matrix $W \in \mathbb{R}^{m \times m}$, where $W_{e,e}$ equals to the weight of edge e. Sometimes, we write the weight $w(u,v)$ of edge $e = \{u,v\}$ as $w(e)$ to simply the notation.

Lemma 3. *It holds that $L = B^\intercal W B$.*

Proof. For any indices u, v, we have

$$\left(B^\intercal W B\right)_{u,v} = \sum_e B^\intercal_{u,e} \left(WB\right)_{e,v} = \sum_e B_{e,u} W_{e,e} B_{e,v}. \tag{2}$$

Now we first assume that $u = v$. Then, it holds that

$$\sum_e B_{e,u} W_{e,e} B_{e,u} = \sum_e \left(B_{e,u}\right)^2 W_{e,e} = \deg(u) = L_{u,u},$$

where the last line holds by the fact that $B_{e,u} \neq 0$ iff e is adjacent to vertex u. Secondly, when u, v are connected by an edge, it holds that

$$\sum_e B_{e,u} W_{e,e} B_{e,v} = -w(u,v) = L_{u,v},$$

since there is exactly one edge e connecting u and v and $B_{e,u} B_{e,u} = -1$. Finally, when u and v are not connected, we know that $B_{e,u} B_{e,v} = 0$ for any edge e, and $L_{u,v} = 0$. Combining the three cases together proves the lemma. □

Now we treat the graph G as a resistor network by replacing each edge e with a resistor of resistance $1/w(e)$. In other words, let's think of $w(e)$ as the conductance of edge e. We can then study how electricity flows in the network. To this end we write two underlying properties of electrical flows. Suppose we send one unit of flow from s to t, and let $b \in \mathbb{R}^n$ be the vector which indicates how much current is going in at each vertex. The *Kirchoff's law* states that the difference between the outgoing and incoming current on the edges adjacent to each vertex equals to the external current input at that vertex. To formalise this mathematically, for every edge e let $i(e)$ be the flow along edge e, that is $i(e)$ is non-negative if electricity is going from u to v and it is non-positive otherwise. Then the Kirchoff's law can be written as

$$B^\intercal i = b. \tag{3}$$

The second property of electrical flows is the *Ohm's law*, which states that the current in an edge equals the potential difference across its ends times its conductance, i.e.,

$$i = WBv, \tag{4}$$

where $v \in \mathbb{R}^n$ expresses the potentials of the vertices. Based on (3) and (4), we have

$$b = B^\intercal WBv = Lv,$$

and

$$v = L^{-1}b.$$

This formation is a bit problematic at the first thought, as the matrix L does not have an inverse. However, notice that we are only interested in the solutions for b such that $\langle b, \mathbf{1} \rangle = 0$, i.e., the total amount of current injected is equal to the total amount extracted, and the solution exists as long as the graph G is connected. Hence, we introduce the *pseudo-inverse* $L^\dagger$ of L as follows:

Definition 4. *The pseudo-inverse $L^\dagger$ of L is the matrix that has the same span as L and that satisfies $LL^\dagger = \Pi$, where Π is the symmetric projection onto the span of L.*

Based on the definition of $L^\dagger$, we have

$$v = L^\dagger b.$$

2 Effective resistance

The effective resistance between vertices u and v is the potential difference induced between them when a unit current is injected at one vertex and extracted at the other. Let us derive an algebraic expression for the effective resistance in terms of $L^\dagger$. To inject and extract a unit current across the endpoints of an edge $u \sim v$, we set $b_{u,v} = (\delta_u - \delta_v)$, where $\delta_u \in \{0,1\}^n$ is the indicator vector of vertex u. Then the potentials induced by $b_{u,v}$ at the vertices are given by $L^\dagger b_{u,v}$. To measure the potential difference across $u \sim v$, we simply multiply by $(\delta_u - \delta_v)^\intercal$ on the left, hence the effective resistance $\mathrm{Reff}(u,v)$ of edge $u \sim v$ can be written as

$$\mathrm{Reff}(u,v) = b_{u,v}^\intercal L^\dagger b_{u,v}.$$

The laws of series and parallel resistance are applicable for the graph setting. Namely, if a path consists of edges of resistance $r_{1,2}, \ldots, r_{n-1,n}$, then the effective resistance between the two extreme vertices is $r_{1,2} + \cdots + r_{n-1,n}$. If instead we have k parallel edges between two vertices s and t of resistances $r_1, \ldots, r_k$, then the effective resistance is

$$\mathrm{Reff}(s,t) = \frac{1}{1/r_1 + \cdots + 1/r_k}.$$

Effective resistance as a metric. Next we'll prove that the effective resistance is a metric. Recall that a metric on a set X is a function $d : X \times X \to [0, \infty)$ such that the following holds for any $x, y, z \in X$: (1) $d(x, x) = 0$; (2) $d(x, y) \geq 0$; (3) $d(x, y) = d(y, x)$, and (4) $d(x, z) \leq d(x, y) + d(y, z)$. It's straightforward to see that the first three conditions hold for the effective resistance, and we only need to prove the triangle inequality for the effective resistance.

Lemma 5 (Metric property)**.** *For any triple of vertices s, t, u,*

$$\mathrm{Reff}(s,t) + \mathrm{Reff}(t,u) \geq \mathrm{Reff}(s,u).$$

Proof. By definition we have

$$\begin{aligned}
\mathrm{Reff}(s,u) &= b_{s,u}^{\intercal} L^{\dagger} b_{s,u} \\
&= (b_{s,t} + b_{t,u})^{\intercal} L^{\dagger} (b_{s,t} + b_{t,u}) \\
&= b_{s,t}^{\intercal} L^{\dagger} b_{s,t} + b_{t,u}^{\intercal} L^{\dagger} b_{t,u} + 2 b_{s,t}^{\intercal} L^{\dagger} b_{t,u} \\
&= \mathrm{Reff}(s,t) + \mathrm{Reff}(t,u) + 2 b_{s,t}^{\intercal} L^{\dagger} b_{t,u}.
\end{aligned}$$

To prove the statement, it suffices to show $2 b_{s,t}^{\intercal} L^{\dagger} b_{t,u} \leq 0$. Notice that $2 b_{s,t}^{\intercal} L^{\dagger} b_{t,u}$ is equal to $v(t) - v(u)$ when we send one unit of flow from s to t. But this means that t has the lowest potential in the network, hence the statement holds. □

Effective resistance as an embedding. We can also view effective resistances as an embedding $F : V \to \mathbb{R}^n$. For any vertex u, let $F(u) \triangleq (L^{\dagger})^{1/2} \mathbf{1}_u$. Then, for any pair of vertices u, v,

$$\begin{aligned}
\|F(u) - F(v)\|^2 &= \left\| (L^{\dagger})^{1/2} (\mathbf{1}_u - \mathbf{1}_v) \right\|^2 \\
&= b_{u,v}^{\intercal} L^{\dagger} b_{u,v} \\
&= \mathrm{Reff}(u,v).
\end{aligned}$$

Hence, by Lemma 5 the square of the distances in this mapping is a metric. Exactly computing the effective resistances of all the edges requires the computation of the inverse of the Laplacian that can take $O(n^3)$ time. It is known that, based on the so-called Laplacian solvers and the Johnson-Lindenstrauss lemma, the effective resistances of all the edges can be approximately computed in $O(m \log^c n)$ time for some constant $c \geq 0$.

From effective resistance to leverage score. The leverage score of an edge e is defined by $\ell_e \triangleq w_e \mathrm{Reff}(e)$, and serves an a measure of how important the edge is. For example, if removing an edge disconnects the graph, then $\mathrm{Reff}(e) = 1/w_e$, as all current flowing between its endpoints must use the edge itself, and $\ell_e = 1$.

Finally, we briefly discuss some applications of leverage scores. Consider the problem of sampling a random spanning tree with probability proportional to the product of the weights of its edges. The theorem below shows that the probability that edge e appears in the tree is exactly its leverage score.

Theorem 6. *If we choose a spanning tree T with probability proportional to the product of its edge weights. Then it holds for any edge e that*

$$\mathbb{P}[e \in T] = \ell(e).$$