Graph Theory

Spectral clustering

Instructor: *Nicolò Cesa-Bianchi* version of April 27, 2024

The material in this handout is mostly taken from: Luca Trevisan, Lecture Notes on Graph Partitioning, Expanders and Spectral Methods, 2016.

Intuitively, a graph is clusterable if its vertices can be partitioned (in a non-trivial way) so that the number of edges across the elements of the partition is small. A key notion is therefore that of cut between disjoint subsets of vertices. We study the clusterability of a graph through the algebraic properties of its adiacency matrix.

Given two disjoints subsets S, T of vertices of a graph $G = (V, E)$, let $E(S, T)$ be the set of edges having one endpoint in S and one endpoint in T. Also, let $\neg S = V \setminus S$.

A cut is any partition $(S, \neg S)$ such that $S \not\equiv V$ and $S \not\equiv \emptyset$. The volume $\mathrm{vol}(S)$ of $S \subseteq V$ is the number of edges incident with a node in S. The **conductance** of $S \subseteq V$ is defined by

$$\phi(S) = \frac{\big|E(S, \neg S)\big|}{\min\big\{\mathrm{vol}(S), \mathrm{vol}(\neg S)\big\}}$$

If $\mathrm{vol}(S) \le \mathrm{vol}(\neg S)$, this is the fraction of edges in the cut $(S, \neg S)$ among those incident on S. If a graph is clusterable, then there exists a partition whose each element S has a small conductance. Finally, the conductance of a graph is

$$\phi(G) = \min_{S\,:\,(S,\neg S)\text{ is a cut}} \phi(S) \ .$$

The sparsity of a cut $(S, \neg S)$ is

$$\sigma(S) = \frac{\big|E(S, \neg S)\big|}{|S|\,|\neg S|}$$

This the fraction of edges in the cut among all potential edges between the two subset of vertices. The sparsity of a graph is

$$\sigma(G) = \min_{S\,:\,(S,\neg S)\text{ is a cut}} \sigma(S)$$

In what follows, we focus on d-regular graphs for simplicity, where $\mathrm{vol}(S) = d|S|$ and so

$$\min\big\{\mathrm{vol}(S), \mathrm{vol}(\neg S)\big\} = d\min\{|S|, |\neg S|\}$$

As $\min\{\alpha, 1-\alpha\} \le 2\alpha(1-\alpha)$ for all $\alpha \in [0,1]$, we have

$$\min\big\{|S|, |\neg S|\big\} \le \frac{2}{n}|S|\,|\neg S|$$

Therefore, $\sigma(S) \le 2(d/n)\phi(S)$ for all cuts $(S, \neg S)$. We now study the relationships between conductance and algebraic properties of the adjacency matrix.

Laplacian matrix. The Laplacian matrix of a d-regular graph $G = (V, E)$ is the symmetric matrix $L = I - \frac{1}{d}A$, where A is the adiacency matrix with entries $A_{i,j} = \mathbb{I}\{(i,j) \in E\}$. For any $\boldsymbol{x} \in \mathbb{R}^n$ we have that

$$
\begin{aligned}
\boldsymbol{x}^\top L \boldsymbol{x} &= \sum_{i \in V} x_i^2 - \frac{1}{d} \sum_{i \in V} \sum_{j \in V} A_{i,j} x_i x_j \\
&= \frac{1}{d} \sum_{i \in V} \sum_{j\,:\,(i,j) \in E} x_i^2 - \frac{1}{d} \sum_{i \in V} \sum_{j\,:\,(i,j) \in E} x_i x_j \\
&= \frac{1}{d} \sum_{i \in V} \sum_{j\,:\,(i,j) \in E} \left(x_i^2 - x_i x_j\right) \\
&= \frac{1}{d} \sum_{(i,j) \in E} \left(x_i^2 + x_j^2 - 2 x_i x_j\right) \\
&= \frac{1}{d} \sum_{(i,j) \in E} (x_i - x_j)^2 \geq 0
\end{aligned}
$$

Therefore, the Laplacian matrix is positive semidefinite. Since the rows and columns of L sum to zero (verify that),

$$
\lambda_1 = \min_{\boldsymbol{u} \in \mathbb{R}^n \setminus \{\boldsymbol{0}\}} \frac{\boldsymbol{u}^\top L \boldsymbol{u}}{\boldsymbol{u}^\top \boldsymbol{u}} = 0
$$

where the minimum is attained by $\boldsymbol{u} = \boldsymbol{1}$, where we write $\boldsymbol{1} = (1, \ldots, 1)$. Hence, $\boldsymbol{u}_1 = \frac{1}{\sqrt{n}}\boldsymbol{1}$ is the eigenvector of λ_1, while the remaining eigenvalues of L are all nonnegative because L is positive semidefinite. Note also that any other eigenvector $\boldsymbol{u}_i$ of L with $i > 1$ is such that $\boldsymbol{u}_i^\top \boldsymbol{u}_1 = 0$. This helps us characterize λ_2,

$$
\lambda_2 = \min_{\substack{\boldsymbol{u} \in \mathbb{R}^n \setminus \{\boldsymbol{0}\} \\ \boldsymbol{u}^\top \boldsymbol{1} = 0}} \frac{\boldsymbol{u}^\top L \boldsymbol{u}}{\boldsymbol{u}^\top \boldsymbol{u}} = \min_{\substack{\boldsymbol{u} \in \mathbb{R}^n \setminus \{\boldsymbol{0}\} \\ \boldsymbol{u}^\top \boldsymbol{1} = 0}} \frac{\sum_{(i,j) \in E} (u_i - u_j)^2}{d \sum_{i \in V} u_i^2} .
$$

If $G = (V, E)$ has two connected components $X, Y \subset V$, then we can choose $\boldsymbol{u} \in \mathbb{R}^n$ such that $u_i = 1/|X|$ for all $i \in X$ and $u_j = -1/|Y|$ for all $j \in Y$. This ensures that $\boldsymbol{u}^\top \boldsymbol{1} = 0$. Moreover, $(i,j) \in E$ if and only if $(u_i - u_j)^2 = 0$. So $\boldsymbol{u}^\top L \boldsymbol{u} = 0$ and therefore $\boldsymbol{u} / \|\boldsymbol{u}\|$ is an eigenvector with eigenvalue $\lambda_2 = 0$. More generally, it can be proven that $\lambda_k = 0$ if and only if G has k connected components. We now look at the largest eigenvalue,

$$
\begin{aligned}
\lambda_n &= \max_{\boldsymbol{u} \in \mathbb{R}^n \setminus \{\boldsymbol{0}\}} \frac{\sum_{(i,j) \in E} (u_i - u_j)^2}{d \sum_{i \in V} u_i^2} \\
&= \max_{\boldsymbol{u} \in \mathbb{R}^n \setminus \{\boldsymbol{0}\}} \frac{d \sum_{i \in V} u_i^2 - \sum_{(i,j) \in E} 2 u_i u_j}{d \sum_{i \in V} u_i^2} \\
&= \max_{\boldsymbol{u} \in \mathbb{R}^n \setminus \{\boldsymbol{0}\}} \frac{2d \sum_{i \in V} u_i^2 - d \sum_{i \in V} u_i^2 - \sum_{(i,j) \in E} 2 u_i u_j}{d \sum_{i \in V} u_i^2} \\
&= \max_{\boldsymbol{u} \in \mathbb{R}^n \setminus \{\boldsymbol{0}\}} \frac{2d \sum_{i \in V} u_i^2 - \sum_{(i,j) \in E} (u_i + u_j)^2}{d \sum_{i \in V} u_i^2} \\
&= 2 - \min_{\boldsymbol{u} \in \mathbb{R}^n \setminus \{\boldsymbol{0}\}} \frac{\sum_{(i,j) \in E} (u_i + u_j)^2}{d \sum_{i \in V} u_i^2} .
\end{aligned}
$$

So $\lambda_n \le 2$ and $\lambda_n = 2$ if G has at least a bipartite component (X, Y). Indeed, in this case we can pick $\boldsymbol{u} \in \mathbb{R}^n$ such that $u_i = 1$ for all $i \in X$, $u_j = -1$ for all $j \in Y$, and $u_k = 0$ for all remaining k. Then $(u_i + u_j)^2 = 0$ for all $(i,j) \in E$, and so $\boldsymbol{u}/\|\boldsymbol{u}\|$ is an eigenvector of G with eigenvalue 2.

Cheeger's inequalities. While winimizing conductance over all exponentially many cuts is NP-hard, the proof of Cheeger's inequalities provides an efficient approximation of $\phi(G)$. These inequalities connect the second eigenvalue with the conductance, revealing the key role of λ_2 in clustering,

$$\frac{\lambda_2}{2} \le \phi(G) \le \sqrt{2\lambda_2}$$

The second inequality is proven via an efficient algorithm that finds a cut $(S_F, \neg S_F)$ such that $\phi(S_F) \le \sqrt{2\lambda_2}$. Together with the first inequality, this implies that $\phi(S_F) \le \sqrt{2\phi(G)}$, which shows how we can efficiently approximate conductance. We begin by proving the first inequality. From now on we write $\sum_{i=1}^n$ instead of $\sum_{i \in V}$.

Lemma 1 *For any connected and d-regular graph G, $\lambda_2 \le 2\phi(G)$.*

Proof. We start noticing that, for any $\boldsymbol{u} \in \mathbb{R}^d$ such that $\boldsymbol{u}^\top \boldsymbol{1} = 0$,

$$\sum_{i=1}^n \sum_{j=1}^n (u_i - u_j)^2 = 2n \sum_{i=1}^n u_i^2 - 2 \sum_{i=1}^n \sum_{j=1}^n u_i u_j = 2n \sum_{i=1}^n u_i^2 - 2\left(\sum_{i=1}^n u_i\right)^2 = 2n \sum_{i=1}^n u_i^2 \tag{1}$$

Therefore, we have that

$$\begin{aligned}
\lambda_2 &= \min_{\substack{\boldsymbol{u} \in \mathbb{R}^n \setminus \{\boldsymbol{0}\} \\ \boldsymbol{u}^\top \boldsymbol{1} = 0}} \frac{\sum_{(i,j) \in E} (u_i - u_j)^2}{d \sum_{i=1}^n u_i^2} \\
&= \min_{\substack{\boldsymbol{u} \in \mathbb{R}^n \setminus \{\boldsymbol{0}\} \\ \boldsymbol{u}^\top \boldsymbol{1} = 0}} \frac{\sum_{(i,j) \in E} (u_i - u_j)^2}{\frac{d}{2n} \sum_{i=1}^n \sum_{j=1}^n (u_i - u_j)^2} && (2) \\
&= \min_{\boldsymbol{u} \in \mathbb{R}^n \setminus \{\boldsymbol{0}, \boldsymbol{1}\}} \frac{\sum_{(i,j) \in E} (u_i - u_j)^2}{\frac{d}{2n} \sum_{i=1}^n \sum_{j=1}^n (u_i - u_j)^2} && (3)
\end{aligned}$$

To understand the last equality: if $\boldsymbol{u} \in \mathbb{R}^n \setminus \{\boldsymbol{0}\}$ and $\boldsymbol{u}^\top \boldsymbol{1} = 0$, then $\boldsymbol{u} \ne \boldsymbol{1}$ and (3) is not larger than (2). Vice versa, if $\boldsymbol{u} \notin \{\boldsymbol{0}, \boldsymbol{1}\}$, then $\boldsymbol{u}'$ defined by $u_i' = u_i - \frac{1}{n}\sum_j u_j$ satisfies $\boldsymbol{u}' \ne \boldsymbol{0}$ and $(\boldsymbol{u}')^\top \boldsymbol{1} = 0$. Hence, the value of (2) is not larger than (3) because the shift by $\frac{1}{n}\sum_j u_j$ cancels out in the numerator and the denominator of the obective function.

For any $S \subseteq V$, let $\boldsymbol{u} \in \{0,1\}^n$ be the incidence vector of the set S, that is $u_i = \mathbb{I}\{i \in S\}$ for $i = 1, \dots, n$. Then $\left|E(S, \neg S)\right| = \sum_{(i,j) \in E} (u_i - u_j)^2$. Also, using $u_i = u_i^2$ for all i,

$$|S|\,|\neg S| = \left(\sum_{i=1}^n u_i^2\right)\left(n - \sum_{j=1}^n u_j^2\right) = n \sum_{i=1}^n u_i^2 - \sum_{i=1}^n \sum_{j=1}^n u_i u_j = \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n (u_i - u_j)^2$$

Therefore,

$$\sigma(G) = \min_{S \subsetneq V \,:\, S \ne \emptyset} \frac{\left|E(S, \neg S)\right|}{|S|\,|\neg S|} = \min_{\boldsymbol{u} \in \{0,1\}^n \setminus \{\boldsymbol{0}, \boldsymbol{1}\}} \frac{\sum_{(i,j) \in E} (u_i - u_j)^2}{\frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n (u_i - u_j)^2}$$

which implies $\frac{d}{n}\lambda_2 = \sigma(G)$. Since $\sigma(G) \leq \frac{2d}{n}\phi(G)$, the proof is concluded. $\square$

The proof of the second inequality of Cheeger is based on the analysis of Fiedler's algorithm, the simplest algorithm for spectral clustering. The algorithm finds a cut of small conductance by looking at the $n-1$ cuts induced by the ranked components of the input vector $\boldsymbol{x}$. As we see in the analysis, the algorithm works well when $\boldsymbol{x}$ is the eigenvector of λ_2.

Algorithm 1 (Fiedler)

Input: Graph $G = (V, E)$, vector $\boldsymbol{x} \in \mathbb{R}^n \setminus \{\mathbf{0}\}$.

1: Sort V according to the components of $\boldsymbol{x}$ and let $v_1 \leq \cdots \leq v_n$ be the vertices of V after sorting
2: Find $k \in \{1, \ldots, n-1\}$ minimizing the conductance $\phi\big(\{v_1, \ldots, v_k\}\big)$

Output: $\{v_1, \ldots, v_k\}$

Note that Fiedler's algorithm can be implemented in time $\mathcal{O}(|E| + |V| \ln |V|)$, because it takes time $\mathcal{O}(|V| \ln |V|)$ to sort the vertices, and the cut of minimal expansion that respects the sorted order can be found in time $\mathcal{O}(|E|)$.

We move on to the analysis of the algorithm, which gives us the second inequality of Cheeger as an immediate consequence. Let

$$R_L(\boldsymbol{x}) = \frac{\sum_{(i,j)\in E}(x_i - x_j)^2}{d\sum_{i=1}^n x_i^2}$$

be the Rayleigh quotient for L evaluated at $\boldsymbol{x} \in \mathbb{R}^n$, and recall that

$$\lambda_2 = \min_{\substack{\boldsymbol{x}\in\mathbb{R}^n\setminus\{\mathbf{0}\} \\ \boldsymbol{x}^\top\mathbf{1}=0}} R_L(\boldsymbol{x})$$

We now prove the following result, which implies $\phi(G) \leq \sqrt{2\lambda_2}$.

Theorem 2 *Let $\boldsymbol{x} \in \mathbb{R}^n \setminus \{\mathbf{0}\}$ be such that $\boldsymbol{x}^\top\mathbf{1} = 0$, and let $S_F \subset V$ be the cut found by Fiedler's algorithm with input $\boldsymbol{x}$. Then $\phi(S_F) \leq \sqrt{2R_L(\boldsymbol{x})}$.*

Indeed, when the input $\boldsymbol{x}$ is the eigenvector of λ_2 we get that

$$\phi(G) \leq \phi(S_F) \leq \sqrt{2\lambda_2}$$

In order to prove Theorem 2, we need to prove two auxiliary lemmas first.

Lemma 3 *Let $\boldsymbol{x} \in \mathbb{R}^n \setminus \{\mathbf{0}\}$ be such that $\boldsymbol{x}^\top\mathbf{1} = 0$. Then there exists a nonnegative vector $\boldsymbol{y}$ such that $R_L(\boldsymbol{y}) \leq R_L(\boldsymbol{x})$. Furthermore, for every $0 < t \leq \max_{v\in V} y_v$, the cut*

$$\Big(\{v \in V \,:\, y_v \geq t\}\,, \{v \in V \,:\, y_v < t\}\Big)$$

is one of the cuts considered in line 2 of Fiedler's algorithm on input $\boldsymbol{x}$.

PROOF. Let m be the median value of the entries of $\boldsymbol{x}$. Let $\boldsymbol{x}^+, \boldsymbol{x}^-$ have components $x_v^+ = \big[x_v - m\big]_+$ and $x_v^- = \big[m - x_v\big]_+$, where $[z]_+ = z\,\mathbb{I}\{z > 0\}$. Note that $\boldsymbol{x}^+, \boldsymbol{x}^-$ are both nonnegative. Now, for every $t > 0$,

$$\big\{v \in V \,:\, x_v^+ \geq t\big\} = \Big\{v \in V \,:\, \big[x_v - m\big]_+ \geq t\Big\} = \{v \in V \,:\, x_v \geq m + t\}$$

is one of the cuts considered by Fiedler's algorithm on input $\boldsymbol{x}$. Similarly, for every $t > 0$,

$$\big\{v \in V \,:\, x_v^- \geq t\big\} = \Big\{v \in V \,:\, \big[m - x_v\big]_+ \geq t\Big\} = \{v \in V \,:\, x_v \leq m - t\}$$

is also one of the cuts considered by Fiedler's algorithm on input $\boldsymbol{x}$. It remains to show that $R_L(\boldsymbol{y}) \leq R_L(\boldsymbol{x})$ for some nonnegative $\boldsymbol{y} \in \mathbb{R}^n$. We set

$$\boldsymbol{y} = \underset{\boldsymbol{z} \in \{\boldsymbol{x}^+, \boldsymbol{x}^-\}}{\operatorname{argmin}} \; R_L(\boldsymbol{z})$$

Let $\boldsymbol{x}' = \boldsymbol{x} - m\,\mathbf{1} = \boldsymbol{x}^+ - \boldsymbol{x}^-$ and observe that, for every constant c, $R_L(\boldsymbol{x} + c\,\mathbf{1}) \leq R_L(\boldsymbol{x})$. Indeed, the numerator of $R_L(\boldsymbol{x}+c\,\mathbf{1})$ and the numerator of $R_L(\boldsymbol{x})$ are the same. Moreover, the denominator of $R_L(\boldsymbol{x} + c\,\mathbf{1})$ is $\|\boldsymbol{x} + c\,\mathbf{1}\|^2 = \|\boldsymbol{x}\|^2 + \|c\,\mathbf{1}\|^2 \geq \|\boldsymbol{x}\|^2$. Therefore $R_L(\boldsymbol{x}') \leq R_L(\boldsymbol{x})$ and we are left to show that $R_L(\boldsymbol{y}) \leq R_L(\boldsymbol{x}')$. To this end we write

$$\begin{aligned}
R_L(\boldsymbol{y}) &= \min\big\{R_L(\boldsymbol{x}^+), R_L(\boldsymbol{x}^-)\big\} \\
&\leq \frac{\|\boldsymbol{x}^+\|^2\, R_L(\boldsymbol{x}^+) + \|\boldsymbol{x}^-\|^2\, R_L(\boldsymbol{x}^-)}{\|\boldsymbol{x}^+\|^2 + \|\boldsymbol{x}^-\|^2} && \text{(using } \min\{a,b\} \leq \alpha a + (1-\alpha) b\text{)} \\
&= \frac{\sum_{(i,j)\in E}(x_i^+ - x_j^+)^2 + \sum_{(i,j)\in E}(x_i^- - x_j^-)^2}{\|\boldsymbol{x}^+\|^2 + \|\boldsymbol{x}^-\|^2} \\
&\leq \frac{\sum_{(i,j)\in E}\Big((x_i^+ - x_j^+) - (x_i^- - x_j^-)\Big)^2}{\|\boldsymbol{x}^+\|^2 + \|\boldsymbol{x}^-\|^2} && \text{(this is shown below)} \\
&= \frac{\sum_{(i,j)\in E}(x_i' - x_j')^2}{\|\boldsymbol{x}'\|^2} && \text{(using } \boldsymbol{x}' = \boldsymbol{x}^+ + \boldsymbol{x}^- \text{ and } (\boldsymbol{x}^+)^\top \boldsymbol{x}^- = 0\text{)} \\
&= R_L(\boldsymbol{x}')
\end{aligned}$$

To finish the proof, we need to verify that for each $(i,j) \in E$,

$$(x_i^+ - x_j^+)^2 + (x_i^- - x_j^-)^2 \leq \Big((x_i^+ - x_j^+) - (x_i^- - x_j^-)\Big)^2 \tag{4}$$

By computing the square on the right-hand side, the two squares on the left-hand side cancel out with the corresponding squares on the right-hand side. Hence proving (4) is equivalent to proving

$$(x_i^+ - x_j^+)(x_i^- - x_j^-) \leq 0 \qquad \Longleftrightarrow \qquad x_i^+ x_i^- - x_i^+ x_j^- - x_j^+ x_i^- + x_j^+ x_j^- \leq 0$$

The proof is concluded by observing that $x_i^+ x_i^- = x_j^+ x_j^- = 0$ by definition, whereas $x_i^+ x_j^- \geq 0$ and $x_j^+ x_i^- \geq 0$ holds because all the fours factors are nonnegative by definition. □

The following observation is used in the proof of the next lemma.

Fact 4 *For all random variables X, Y such that $Y > 0$ and $\mathbb{E}[X], \mathbb{E}[Y] < \infty$,*

$$\mathbb{P}\left(\frac{X}{Y} \leq \frac{\mathbb{E}[X]}{\mathbb{E}[Y]}\right) > 0$$

PROOF. Let $r = \mathbb{E}[X]/\mathbb{E}[Y]$. Because of linearity of expectation, $\mathbb{E}\big[X - rY\big] = 0$. Since the expected value is zero, the random variable $X - rY$ must be nonpositive with probability bigger than zero, $\mathbb{P}\big(X - rY \leq 0\big) > 0$. Dividing both sides of $X - rY \leq 0$ by $Y > 0$, we get the desired result. □

We are now ready to prove the second auxiliary lemma. Define the expansion of a set $S \subset V$ by

$$\operatorname{xpn}(S) = \frac{\big|E(S, \neg S)\big|}{\operatorname{vol}(S)} = \frac{\big|E(S_t, \neg S_t)\big|}{d|S_t|} \qquad \text{(for regular graphs)}$$

Note that, for regular graphs, $\operatorname{xpn}(S) = \phi(S)$ when $|S| \leq |\neg S|$.

Lemma 5 *For all nonnegative vectors $\boldsymbol{y} \in \mathbb{R}^n$ there exists $0 < t \leq \max_v y_v$ such that*

$$\operatorname{xpn}(S_t) \leq \sqrt{2R_L(\boldsymbol{y})}$$

where $S_t = \{v \in V \,:\, y_v \geq t\}$.

PROOF. Since rescaling does not affect the Rayleigh quotient, we may assume $\max_v y_v = 1$. The proof uses the probabilistic method. Let T be a random variable such that $\mathbb{P}\big(T \leq \sqrt{a}\big) = a$, which means that T^2 is uniformly distributed in $[0, 1]$ and $\mathbb{P}\big(T \leq 0\big) = 0$, which implies $T > 0$ with probability 1. Because S_t is nonempty for all $t \in (0, 1]$, we can write

$$\operatorname{xpn}(S_T) = \frac{\big|E(S_T, \neg S_T)\big|}{d|S_T|} \leq \frac{\mathbb{E}\Big[\big|E(S_T, \neg S_T)\big|\Big]}{d\,\mathbb{E}\big[|S_T|\big]} \qquad \text{(with probability } > 0\text{, by Fact 4)}$$

This implies that there exists some $t \in (0, 1]$ such that the above holds. To conclude the proof, we show that

$$\frac{\mathbb{E}\Big[\big|E(S_T, \neg S_T)\big|\Big]}{d\,\mathbb{E}\big[|S_T|\big]} \leq \sqrt{2R_L(\boldsymbol{y})}$$

We start to bound the denominator. Using that T is uniformly distributed in $[0, 1]$,

$$E\big[|S_T|\big] = \sum_{i=1}^{n} \mathbb{P}(i \in S_T) = \sum_{i=1}^{n} \mathbb{P}(T \leq y_i) = \sum_{i=1}^{n} y_i^2 \tag{5}$$

Now pick any $(i, j) \in E$ and assume $y_j \leq y_i$. Then

$$\begin{aligned}
\mathbb{P}\big(i \in S_T,\, j \in \neg S_T\big) &= \mathbb{P}\big(y_j < T \leq y_i\big) \\
&= \Big(\mathbb{P}(T \leq y_i) - \mathbb{P}(T \leq y_j)\Big) \\
&= y_i^2 - y_j^2
\end{aligned}$$

Therefore,

$$\begin{aligned}
\mathbb{E}\Big[\big|E(S_T,\neg S_T)\big|\Big] &= \sum_{(i,j)\in E} \Big(\big(y_i^2-y_j^2\big)\mathbb{I}\{y_j \le y_i\} + \big(y_j^2-y_i^2\big)\mathbb{I}\{y_i \le y_j\}\Big) \\
&= \sum_{(i,j)\in E} \big|y_i^2-y_j^2\big| \\
&= \sum_{(i,j)\in E} \big|y_i-y_j\big|(y_i+y_j) \\
&\le \sqrt{\sum_{(i,j)\in E}(y_i-y_j)^2}\sqrt{\sum_{(i,j)\in E}(y_i+y_j)^2}
\end{aligned}$$

where we applied the Cauchy-Schwartz inequality $\boldsymbol{u}^\top \boldsymbol{v} \le \|\boldsymbol{u}\|\,\|\boldsymbol{v}\|$ in the last step. Using now the elementary inequality $(a+b)^2 \le 2(a^2+b^2)$ we may write

$$\sum_{(i,j)\in E}(y_i+y_j)^2 \le 2\sum_{(i,j)\in E}\big(y_i^2+y_j^2\big) = 2d\sum_{i=1}^n y_i^2$$

Combining the above with (5) we obtain

$$\frac{\mathbb{E}\Big[\big|E(S_T,\neg S_T)\big|\Big]}{d\,\mathbb{E}\big[|S_T|\big]} \le \frac{\sqrt{\Big(\sum_{(i,j)\in E}(y_i-y_j)^2\Big)\left(2d\sum_{i=1}^n y_i^2\right)}}{d\sum_{i=1}^n y_i^2} = \sqrt{\frac{2\sum_{(i,j)\in E}(y_i-y_j)^2}{d\sum_{i=1}^n y_i^2}}$$

concluding the proof. □

We can now prove Theorem 2.

Proof of Theorem 2. Let $\boldsymbol{x} \in \mathbb{R}^n$ be such that $\boldsymbol{x}^\top \mathbf{1} = 0$ and let $\big(S_F, \neg S_F\big)$ be the cut found by Fiedler's algorithm on input $\boldsymbol{x}$. Lemma 3 states that:

1. there exists a nonnegative vector $\boldsymbol{y}$ such that $R_L(\boldsymbol{y}) \le R_L(\boldsymbol{x})$;
2. for this $\boldsymbol{y}$ and for any $0 < t \le \max_{v\in V} y_v$, the set $S_t = \{v \in V : y_v \ge t\}$ has at most $\frac{n}{2}$ vertices (because $\boldsymbol{y}$ has at most $\frac{n}{2}$ nonzero components, as we defined it using the median) implying $\phi(S_t) = \mathrm{xpn}(S_t)$;
3. the cut $(S_t, \neg S_t)$ is one of the cuts considered by Fiedler's algorithm on input $\boldsymbol{x}$, which implies $\phi(S_F) \le \phi(S_t)$ for all t.

Then, Lemma 5 ensures there exists a threshold $0 < t \le \max_{v\in V} y_v$ such that $\mathrm{xpn}(S_t) \le \sqrt{2R_L(\boldsymbol{y})}$. We can thus write

$$\phi(S_F) \le \phi(S_t) = \mathrm{xpn}(S_t) \le \sqrt{2R_L(\boldsymbol{y})} \le \sqrt{2R_L(\boldsymbol{x})}$$

concluding the proof. □

Nonregular graphs. What is the correct generalization of the Laplacian matrix $I - \frac{1}{d}A$ when G is not d-regular? As we want to preserve the spectral properties, we look at the Rayleigh quotient for the d-regular case:

$$\frac{\sum_{(i,j)\in E}(x_i - x_j)^2}{d\sum_{i=1}^n x_i^2}$$

The natural generalization to nonregular graphs is then

$$\frac{\sum_{(i,j)\in E}(x_i - x_j)^2}{\sum_{i=1}^n d(i)x_i^2} = \frac{\sum_{(i,j)\in E}(x_i - x_j)^2}{\sum_{i=1}^n \left(\sqrt{d(i)}x_i\right)\left(\sqrt{d(i)}x_i\right)} = \frac{\boldsymbol{x}^\top (D - A)\boldsymbol{x}}{\left(D^{1/2}\boldsymbol{x}\right)^\top \left(D^{1/2}\boldsymbol{x}\right)}$$

where where $D^{1/2} = \operatorname{diag}\left(\sqrt{d(1)}, \ldots, \sqrt{d(n)}\right)$ and $d(i)$ is the degree of i. If we now set $\boldsymbol{u} = D^{1/2}\boldsymbol{x}$, the above becomes

$$\frac{\left(D^{-1/2}\boldsymbol{u}\right)^\top (D - A)\left(D^{-1/2}\boldsymbol{u}\right)}{\boldsymbol{u}^\top \boldsymbol{u}} = \frac{\boldsymbol{u}^\top D^{-1/2}(D - A)D^{-1/2}\boldsymbol{u}}{\boldsymbol{u}^\top \boldsymbol{u}} = \frac{\boldsymbol{u}^\top \left(I - D^{-1/2}AD^{-1/2}\right)\boldsymbol{u}}{\boldsymbol{u}^\top \boldsymbol{u}}$$

where we assumed $d(v) > 0$ for all v (there are no isolated vertices) and used $D^{-1/2}DD^{-1/2} = I$. The matrix $L_{\text{norm}} = I - D^{-1/2}AD^{-1/2}$ whose components are

$$L_{\text{norm}}(i,j) = \begin{cases} 1 & \text{if } i = j \\ -A(i,j)/\sqrt{d(i)d(j)} & \text{otherwise} \end{cases}$$

is known as the normalized Laplacian. All the spectral properties which we proved for d-regular graphs, including Cheeger's inequalities, continue to hold for the normalized Laplacian of arbitrary graphs.

Acknowledgements. Thanks to Alberto Boggio, Alessia Galdeman, and Francesco Agrimonti for flagging mistakes and typos in earlier versions of this handout. Thanks to Dario Moschetti for simplifying the proof of Lemma 3.