Graph partitioning with spectral methods

J. Devevey D. Rochette

26 November 2019

Outline

1. Background
   - Spectral Decomposition of Symmetric Matrices
   - Graph Laplacian
   - Laplacian Systems of Linear Equations
2. Applications to graph partitioning
   - Normalized Laplacian
   - Balanced cuts
   - Computing the second eigenvector
3. Local Graph Partitioning
   - A local variation of SparsestCut

Outline

1. Background
   - Spectral Decomposition of Symmetric Matrices
   - Graph Laplacian
   - Laplacian Systems of Linear Equations

2. Applications to graph partitioning
   - Normalized Laplacian
   - Balanced cuts
   - Computing the second eigenvector

3. Local Graph Partitioning
   - A local variation of SPARSESTCUT

Notations

Let A, a symmetric $n \times n$ matrix of reals, we denote by :

- $\lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_n$ the n eigenvalues of A
- $\mathbf{u}_1, \ldots, \mathbf{u}_n$ corresponding eigenvectors

We can show that all of its eigenvalues are real.

Spectral Decomposition of Symmetric Matrices

Lemma

For a symmetric matrix A, if $\lambda_i \neq \lambda_j$ then $\langle \mathbf{u}_i, \mathbf{u}_j \rangle = 0$.

Note

This implies that one can decompose $\mathbb{R}^n$ into U_i, where U_i is the space spanned by all eigenvectors corresponding to that eigenvalue λ_i.

Spectral Decomposition of Symmetric Matrices

Theorem

For a symmetric matrix A, we can write $A = \sum_{i=1}^{n} \lambda_i \mathbf{u}_i \mathbf{u}_i^T$

This decomposition can be (approximately) computed in time polynomial.

Eigenvalues of real symmetric matrices

Lemma

For an $n \times n$ real symmetric matrix A

$$\lambda_1 = \min_{v \in \mathbb{R}^n \setminus \{0\}} \frac{v^T A v}{v^T v}$$

$$\lambda_n = \max_{v \in \mathbb{R}^n \setminus \{0\}} \frac{v^T A v}{v^T v}$$

Special case: positive (semi)definite matrices

Definition

A symmetric matrix A is called positive semidefinite if $v^T A v \geq 0$ for every nonzero vector v. A is said to be positive definite if $v^T A v$ is positive.

Special case: positive (semi)definite matrices

Definition

A symmetric matrix A is called positive semidefinite if $v^T A v \geq 0$ for every nonzero vector v. A is said to be positive definite if $v^T A v$ is positive.

Lemma

For a symmetric matrix A, if $\lambda_1 \geq 0$ then A is positive semidefinite (equiv. $\lambda_1 > 0 \Rightarrow A$ positive definite).

A-norm

Definition

For a symmetric positive semidefinite matrix A we define the A-norm $\|\cdot\|_A$ such that for a vector v

$$\|v\|_A \stackrel{\text{def}}{=} \sqrt{v^T A v}$$

A-norm

Definition

For a symmetric positive semidefinite matrix A we define the A-norm $\|\cdot\|_A$ such that for a vector v

$$\|v\|_A \stackrel{\text{def}}{=} \sqrt{v^T A v}$$

Lemma

For an $n \times n$ symmetric positive semidefinite matrix A and v, w two vectors the following holds

$$\lambda_1 \|v - w\|^2 \leq \|v - w\|_A^2 \leq \lambda_n \|v - w\|^2$$

Hence the distortion in distances due to the A-norm is at most $\sqrt{\lambda_n / \lambda_1}$.

Outline

1. Background
   - Spectral Decomposition of Symmetric Matrices
   - Graph Laplacian
   - Laplacian Systems of Linear Equations
2. Applications to graph partitioning
   - Normalized Laplacian
   - Balanced cuts
   - Computing the second eigenvector
3. Local Graph Partitioning
   - A local variation of SparsestCut

Graph Laplacian

Let $G = (V, E)$ an undirected graph and its adjacency matrix A.
Consider a function $f : V \to \mathbb{R}$, where $V = \{v_1, \ldots, v_n\}$, then

$$A \begin{pmatrix} f(v_1) \\ \vdots \\ f(v_n) \end{pmatrix} = \begin{pmatrix} \sum\limits_{(1,i)\in E} f(v_i) \\ \vdots \\ \sum\limits_{(n,i)\in E} f(v_i) \end{pmatrix}$$

Graph Laplacian

Definition

Consider $G = (V, E)$ an undirected graph. Let A be its adjacency matrix and D be its degree matrix *(i.e., a diagonal matrix where $D_{i,i}$ is equal to the degree of vertex i)*. The **graph Laplacian** of G is $L \stackrel{\text{def}}{=} D - A$.

Graph Laplacian

Definition

Consider $G = (V, E)$ an undirected graph. Let A be its adjacency matrix and D be its degree matrix *(i.e., a diagonal matrix where $D_{i,i}$ is equal to the degree of vertex i)*. The graph Laplacian of G is $L \stackrel{\text{def}}{=} D - A$.

For a graph function $f : V \to \mathbb{R}$, the graph Laplacian behaves as a discrete analog of the usual Laplacian operator:

$$L \begin{pmatrix} f(v_1) \\ \vdots \\ f(v_n) \end{pmatrix} = \begin{pmatrix} \sum_{(1,i)\in E} f(v_1) - f(v_i) \\ \vdots \\ \sum_{(n,i)\in E} f(v_n) - f(v_i) \end{pmatrix}$$

Graph Laplacian

Lemma

L is a positive semidefinite matrix.

Graph Laplacian

Proof (sketch).

Define the $n \times n$ matrices L_e with $e \in E$ by :

$$L_e(i,j) = \begin{cases} -1 & \text{if } (i,j) = e \\ 1 & \text{if } i = j \text{ and } i \in e \\ 0 & \text{otherwise} \end{cases}$$

One can show that $L = \sum_{e \in E} L_e$. Then for every nonzero vector v

$$\begin{aligned} v^T L v &= \sum_{e \in E} v^T L_e v \\ &= \sum_{(i,j) \in E} (v_i - v_j)^2 \\ &\geq 0 \end{aligned}$$

□

Graph Laplacian: First Eigenvalue

Note

L is not a positive definite matrix:

$$L \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix} = 0 \cdot \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix}$$

hence $\lambda_1 = 0$.

Graph Laplacian: Second Eigenvalue

Theorem

For L the Laplacian of a graph G, $\lambda_2 > 0$ iff G is connected.

Outline

1. Background
   - Spectral Decomposition of Symmetric Matrices
   - Graph Laplacian
   - Laplacian Systems of Linear Equations
2. Applications to graph partitioning
   - Normalized Laplacian
   - Balanced cuts
   - Computing the second eigenvector
3. Local Graph Partitioning
   - A local variation of SparsestCut

Laplacian Systems of Linear Equations

Consider a linear system of equation $Lx = b$ where L is the Laplacian of a graph G.
If G is a connected graph, we know that $\lambda_1 = 0$ and $\lambda_2 > 0$. Since

$$\mathbf{u}_1 = \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix}$$

then $\forall v \in \mathrm{Im}(L), \langle v, \mathbf{u}_1 \rangle = 0$.
We'll see that we can solve such a system if $\langle b, \mathbf{u}_1 \rangle = 0$.

Pseudo-inverse

Definition

Let $A \in \mathbb{R}^{n \times n}$ and u the linear operator associated to A in the canonical basis. Then $u_{|Im(u)}$ is isomorphic and its inverse is denoted by A^+. A^+ is called the pseudo-inverse of A.

Pseudo-inverse

Definition

Let $A \in \mathbb{R}^{n \times n}$ and u the linear operator associated to A in the canonical basis. Then $u_{|Im(u)}$ is isomorphic and its inverse is denoted by A^+. A^+ is called the pseudo-inverse of A.

Remark

Following what precedes, solving $Lx = b$ when $\langle b, \mathbf{u}_1 \rangle = 0$ and G is connected corresponds in fact to computing $L^+ b$.

Laplacian Systems of Linear Equations

Theorem

There is an algorithm LSOLVE *which given a linear system of equation* $Lx = b$ *and* $\epsilon > 0$*, returns* x *such that*

$$\|x - L^+ b\|_L \leq \epsilon \|L^+ b\|_L$$

where L^+ *is the pseudo-inverse of* L*.*
It runs in $O(\tau \log 1/\epsilon)$*, with* τ *the number of nonzero entries in* L*.*

Outline

1 Background
- Spectral Decomposition of Symmetric Matrices
- Graph Laplacian
- Laplacian Systems of Linear Equations

2 Applications to graph partitioning
- Normalized Laplacian
- Balanced cuts
- Computing the second eigenvector

3 Local Graph Partitioning
- A local variation of SparsestCut

Graph Conductance

Definition: Cut

A cut of $G = (V, E)$ is a partition $(S, \bar{S})$ of V.

Graph Conductance

Definition: Cut

A cut of $G = (V, E)$ is a partition $(S, \bar{S})$ of V.

Definition: Cutting edges

We denote by $|E(S, \bar{S})|$ the number of edges that have one end point in S and the other in $\bar{S}$.

Graph Conductance

Definition: Cut

A cut of $G = (V, E)$ is a partition $(S, \bar{S})$ of V.

Definition: Cutting edges

We denote by $|E(S, \bar{S})|$ the number of edges that have one end point in S and the other in $\bar{S}$.

Definition: Volume and conductance of a cut

Let $(S, \bar{S})$ a cut of $G = (V, E)$, we define the volume of the cut $Vol(S) = \sum_{i \in S} d_i$, and the conductance

$$\Phi(S) = \frac{|E(S, \bar{S})|}{\min(Vol(S), Vol(\bar{S}))}$$

Definition: The SPARSEST-CUT problem

Given $G = (V, E)$ compute $\Phi(G) = \min_{\emptyset \neq S \subset V} \Phi(S)$.

A little help?

Problem 1

SPARSEST-CUT is NP-hard! (Can be reduced from MAX-CUT)

A little help?

Problem 1

SPARSEST-CUT is NP-hard! (Can be reduced from MAX-CUT)

Definition: *h*-value

Let $(S, \bar{S})$ be a cut. The *h*-value of the cut is defined as:

$$h(S) = \frac{|E(S, \bar{S})|}{Vol(S) \cdot Vol(\bar{S})} \cdot Vol(G)$$

Lemma

For all possible cut $(S, \bar{S})$, we have: $\Phi(S) \leq h(S) \leq 2\Phi(S)$.

A little help?

Problem 1

SPARSEST-CUT is NP-hard! (Can be reduced from MAX-CUT)

Definition: h-value

Let $(S, \bar{S})$ be a cut. The h-value of the cut is defined as:

$$h(S) = \frac{|E(S, \bar{S})|}{Vol(S) \cdot Vol(\bar{S})} \cdot Vol(G)$$

Lemma

For all possible cut $(S, \bar{S})$, we have: $\Phi(S) \leq h(S) \leq 2\Phi(S)$.

Problem 2

Computing h is no easier than Φ!

Two probability measures

Definition

We define a probability measure on the edges and one on the vertices the following way:

$$\nu(e) = \frac{1}{|E|}, \forall e \in E \qquad \mu(i) = \frac{d_i}{Vol(G)} = \frac{d_i}{2|E|}, \forall i \in V$$

The two of them can be naturally extended to subsets of E (resp. V).

Two probability measures

Definition

We define a probability measure on the edges and one on the vertices the following way:

$$\nu(e) = \frac{1}{|E|}, \forall e \in E \qquad \mu(i) = \frac{d_i}{Vol(G)} = \frac{d_i}{2|E|}, \forall i \in V$$

The two of them can be naturally extended to subsets of E (resp. V).

Lemma

Let $(S, \bar{S})$ be a cut of $G = (V, E)$ and $\mathbb{1}_S : i \mapsto 1$ if $i \in S$, 0 otherwise.
Then:

$$h(S) = \frac{\nu(E(S, \bar{S}))}{2\mu(S)\mu(\bar{S})} = \frac{\mathbb{E}_{(i,j)\leftarrow\nu}\left((\mathbb{1}_S(i) - \mathbb{1}_S(j))^2\right)}{\mathbb{E}_{i\leftarrow\mu}\mathbb{E}_{j\leftarrow\mu}\left((\mathbb{1}_S(i) - \mathbb{1}_S(j))^2\right)}$$

$$h(G) = \min_{\mathbf{x}\in\{0,1\}^n} \frac{\mathbb{E}_{(i,j)\leftarrow\nu}\left((x_i - x_j)^2\right)}{\mathbb{E}_{i\leftarrow\mu}\mathbb{E}_{j\leftarrow\mu}\left((x_i - x_j)^2\right)}$$

Real conductance

Definition: Real Conductance

Let $G = (V, E)$ be a graph. We define the real conductance of G as:

$$h_{\mathbb{R}}(G) = \min_{\mathbf{x} \in \mathbb{R}^n} \frac{\mathbb{E}_{ij \leftarrow \nu}\left((x_i - x_j)^2\right)}{\mathbb{E}_{i,j \leftarrow \mu \times \mu}\left((x_i - x_j)^2\right)}$$

Lemma

We immediately have the following result:

$$h_{\mathbb{R}}(G) \leq h(G) \leq 2\Phi(G)$$

Computing the real conductance

Definition: Normalized Laplacian

The normalized Laplacian $\mathcal{L}$ of a graph G is defined as:

$$\mathcal{L} = D^{-1/2} L D^{-1/2}$$

Where D is the diagonal matrix with $D_{i,i} = d_i, \forall i \in V$.

Theorem

$$\lambda_2(\mathcal{L}) = h_{\mathbb{R}}(G)$$

Computing the real conductance

Definition: Normalized Laplacian

The normalized Laplacian $\mathcal{L}$ of a graph G is defined as:

$$\mathcal{L} = D^{-1/2} L D^{-1/2}$$

Where D is the diagonal matrix with $D_{i,i} = d_i, \forall i \in V$.

Theorem

$$\lambda_2(\mathcal{L}) = h_{\mathbb{R}}(G)$$

Conclusion

We know how to compute this eigenvalue! We are now able to compute a 2-approximation to the Sparsest-Cut problem.

Conclusion: final bounds

Theorem (Cheeger's inequality) (admitted)

$$\frac{\lambda_2(\mathcal{L})}{2} \leq \Phi(G) \leq 2\sqrt{\lambda_2(\mathcal{L})}$$

Proof (sketch of the right hand side)

- Show that $\Phi(G) = \min_{\mathbf{y}\in\mathbb{R}^n, \mu_{1/2}(\mathbf{y})=0} \frac{\mathbb{E}_{ij\leftarrow\nu}\left(|y_i - y_j|\right)}{\mathbb{E}_{i\leftarrow\mu}(|y_i|)}$ where $\mu_{1/2}(\mathbf{y})$ is the median over the measure μ of $\mathbf{y}$.
- Show that if there is an $\mathbf{x} \in \mathbb{R}^n$ such that $\frac{\mathbb{E}_{ij\leftarrow\nu}\left((x_i - x_j)^2\right)}{\mathbb{E}_{i,j\leftarrow\mu\times\mu}\left((x_i - x_j)^2\right)} = \epsilon$ then there exists a $\mathbf{y} \in \mathbb{R}^n$ with $\mu_{1/2}(\mathbf{y}) = 0$ such that $\frac{\mathbb{E}_{ij\leftarrow\nu}\left(|y_i - y_j|\right)}{\mathbb{E}_{i\leftarrow\mu}(|y_i|)} \leq 2\sqrt{\epsilon}$.
- Conclude.

Spectral cut Algorithm

Algorithm

Algorithm 1: Spectral Cut Algorithm

Data: $G = (V, E)$ an undirected graph and $V = \{1, \dots, n\}$

Result: A subset $S \subset V$ with $\Phi(S) \leq 2\sqrt{\Phi(G)}$

$D \leftarrow$ degree matrix of G

$L \leftarrow$ Laplacian of G

$\mathcal{L} \leftarrow D^{-1/2} L D^{-1/2}$

$\mathbf{y}^* \leftarrow \arg\min_{\mathbf{y} \in \mathbb{R}^n, \langle \mathbf{y}, D^{1/2}\mathbf{1} \rangle = 0} \frac{\mathbf{y}^T \mathcal{L} \mathbf{y}}{\mathbf{y}^T \mathbf{y}}$

$\mathbf{x}^* \leftarrow D^{-1/2} \mathbf{y}^*$

Re-index V so that $x_1^* \leq x_2^* \leq \ldots \leq x_n^*$

$S_i \leftarrow \{1, \dots, i\}$ for $i = 1, 2, \dots, n-1$

Return S_i that has minimum conductance from among $1, \dots, n-1$

Outline

1 Background
- Spectral Decomposition of Symmetric Matrices
- Graph Laplacian
- Laplacian Systems of Linear Equations

2 Applications to graph partitioning
- Normalized Laplacian
- Balanced cuts
- Computing the second eigenvector

3 Local Graph Partitioning
- A local variation of SparsestCut

A new problem

Definition: Balanced Edge-Separator

Given $G = (V, E)$ and $b \in (0, \frac{1}{2}]$, the Balanced Edge-Separator problem is to find a cut $(S, \bar{S})$ such that

$$\Phi(S) = \min_{(S', \bar{S}') \text{ s.t. } \min(\mu(S'), \mu(\bar{S}')) \geq b} \Phi(S')$$

Remark

- The case $b = 0$ is in fact the Sparsest-Cut problem!
- This problem is NP-Hard.

Tackling Balanced Edge-Separator

Idea

We can call SpectralCut recursively to try and find a cut which is $\frac{b}{2}$-balanced. As long as the cut isn't $\frac{b}{2}$ balanced, we can take the smallest part of the cut out of the graph, and call SpectralCut on what is left of the graph. Moreover we can at each step check that the second eigenvalue of the normalized laplacian of the leftover graph isn't too big, and by doing so we will be able to give a bound on the conductance of the cut we get (or certify that all cuts have at least a certain conductance).

BALANCEDCUT Algorithm

Algorithm 2: BALANCEDCUT

Data: a graph $G = (\{1, 2, \dots, n\}, E)$, $b \in (0, \frac{1}{2}]$ and target conductance $\gamma \in [0, 1]$
Result: a cut $(S, \bar{S})$ or certify that all b-balanced cut have conductance at least γ

```
μ_G ← degree measure of G
H ← G
S' ← ∅
repeat
    ℒ_H ← normalized laplacian of H
    if λ_2(ℒ_H) > 4γ then
        return NO
    end
    S ← SPECTRALCUT(H)
    S' ← arg min(μ_G(S), μ_G(V(H) \ S)) ∪ S'
    Swap S and V(H) \ S if the argmin was V(H) \ S
    if min(μ_G(S'), μ_G(V(G) \ S')) ≥ b/2 then
        return (S', S̄') and (S, S̄)
    else
        V' ← V(H) \ S
        E' ← {ij s.t. i, j ∈ V'} ∪ {ii s.t. ij ∈ E(V', S)}
        H ← (V', E')
    end
until;
```

Analysis of BalancedCut

No case

If No is returned then every b-balanced cut in G has conductance more than γ.

Other case

In the other case, BalancedCut returns a $\frac{b}{2}$-balanced cut in G of conductance at most $O(\sqrt{\gamma})$.

Outline

1 Background
- Spectral Decomposition of Symmetric Matrices
- Graph Laplacian
- Laplacian Systems of Linear Equations

2 Applications to graph partitioning
- Normalized Laplacian
- Balanced cuts
- Computing the second eigenvector

3 Local Graph Partitioning
- A local variation of SparsestCut

Computing the largest eigenvalue: Power Method

Lemma

Let $A \in \mathbb{R}^{n \times n}$, $\epsilon > 0$, $k > \frac{1}{2}\epsilon \log\left(\frac{9n}{4}\right)$ such that $k \in \mathbb{N}$. Then the following holds with probability at least $\frac{1}{2}$ over $\mathbf{v}$ chosen uniformly at random from the n-dimensional unit sphere:

$$|\lambda_n(A)| \geq \frac{\|A^{k+1}\mathbf{v}\|}{\|A^k\mathbf{v}\|} \geq (1-\epsilon)|\lambda_n(A)|$$

1st attempt

Idea

Notice how $I_n - \mathcal{L} = D^{-1/2} A D^{-1/2}$ with A the adjacency matrix and D the degree matrix. Also $\forall i \in \{1, \ldots, n\}, \lambda_i(I_n - \mathcal{L}) = 1 - \lambda_{n-i+1}(\mathcal{L})$. We deduce that the largest eigenvalue of $I_n - \mathcal{L}$ is 1 and an eigenvector is $D^{1/2}\mathbf{1}$.

1st attempt

Idea

Notice how $I_n - \mathcal{L} = D^{-1/2} A D^{-1/2}$ with A the adjacency matrix and D the degree matrix. Also $\forall i \in \{1, \ldots, n\}, \lambda_i(I_n - \mathcal{L}) = 1 - \lambda_{n-i+1}(\mathcal{L})$. We deduce that the largest eigenvalue of $I_n - \mathcal{L}$ is 1 and an eigenvector is $D^{1/2}\mathbf{1}$.

Theorem

By applying the Power Method orthogonally to $D^{1/2}\mathbf{1}$ and with $\epsilon = \lambda_2(\mathcal{L})/2$ and $k = \Theta\left(\frac{\log(n)}{\lambda_2(\mathcal{L})}\right)$ we can compute $\lambda_2(\mathcal{L})$ up by a factor 2 in time $O\left(\frac{m \log(n)}{\lambda_2(\mathcal{L})}\right)$.

2nd attempt

Idea

Remark that $1/\lambda_2(\mathcal{L})$ is the largest eigenvalue of $\mathcal{L}^+$, the pseudo-inverse of $\mathcal{L}$. Then applying the power method with $\epsilon = 1/2$ and $k = \Theta(\log(n))$ on $\mathcal{L}^+$ is sufficient to compute $\lambda_2(\mathcal{L})$ up to a factor 2.

2nd attempt

Idea

Remark that $1/\lambda_2(\mathcal{L})$ is the largest eigenvalue of $\mathcal{L}^+$, the pseudo-inverse of $\mathcal{L}$. Then applying the power method with $\epsilon = 1/2$ and $k = \Theta(\log(n))$ on $\mathcal{L}^+$ is sufficient to compute $\lambda_2(\mathcal{L})$ up to a factor 2.

Inverse computation?

Computing $\mathcal{L}^+$ explicitly is hard.

Solution

We use the linear system solver we discussed before, LSolver.

Computation of second eigenvector

Theorem

Given an undirected, unweighted graph $G = (V, E)$ with $n = |V|$ and $m = |E|$, there exist an algorithm that finds a vector $\mathbf{x}$ in time $O(m + n)$ such that:

$$\frac{\mathbb{E}_{ij \leftarrow \nu}\left((x_i - x_j)^2\right)}{\mathbb{E}_{(i,j) \leftarrow \mu \times \mu}\left((x_i - x_j)^2\right)} \leq 2\lambda_2(\mathcal{L})$$

Note

Recall that the condition in the BALANCEDCUT algorithm was $\lambda_2(\mathcal{L}_H) > 4\gamma$ and can now be rewritten as $\frac{\mathbb{E}_{ij \leftarrow \nu}\left((x_i - x_j)^2\right)}{\mathbb{E}_{(i,j) \leftarrow \mu \times \mu}\left((x_i - x_j)^2\right)} > 2\gamma$.

Outline

1 Background
- Spectral Decomposition of Symmetric Matrices
- Graph Laplacian
- Laplacian Systems of Linear Equations

2 Applications to graph partitioning
- Normalized Laplacian
- Balanced cuts
- Computing the second eigenvector

3 Local Graph Partitioning
- A local variation of SparsestCut

Global vs Local

Remark

Eigenvalues are a global property of the graph. But a cut is by definition a local part of the graph, even more when we look for small cuts.
This is one of the limitation of the previous part.

New motivation

New situation: assume this time we know where we want to look for a cut. How can we take that into account?

A local variation of SPARSESTCUT

Definition

Let $G = (V, E)$, $u \in V$ and $k \in \mathbb{N}$. The LOCALSPARSESTCUT problem is to compute:

$$\Phi(u, k) = \min_{T \subset V : u \in T, Vol(T) \leq k} \Phi(T)$$

Note

We won't be able to tackle this problem using global tools. How can we adapt what we did before to solve this problem?

The LocalSpectral Problem

Definition

On the left hand side is the optimization problem to compute $\lambda_2(L)$. On the right hand side is the new local problem $\textsc{LocalSpectral}(\mathbf{s}, \kappa)$, where $\mathbf{s}$ is called the seed vector and $0 \leq \kappa < 1$ is the correlation parameter:

$$\begin{aligned} &\min \mathbf{x}^T L \mathbf{x} \\ &\text{s.t. } \mathbf{x}^T D \mathbf{x} = 1 \\ &(\mathbf{x}^T D \mathbf{1})^2 = 0 \\ &\mathbf{x} \in \mathbb{R}^n \end{aligned} \qquad\qquad \begin{aligned} &\min \mathbf{x}^T L \mathbf{x} \\ &\text{s.t. } \mathbf{x}^T D \mathbf{x} = 1 \\ &(\mathbf{x}^T D \mathbf{1})^2 = 0 \\ &(\mathbf{x}^T D \mathbf{s})^2 \geq \kappa \\ &\mathbf{x} \in \mathbb{R}^n \end{aligned}$$

Solving LocalSpectral

Following theorems are hard, so we admit them

Theorem: Solution Characterization (admitted)

Let $\mathbf{s} \in \mathbb{R}^n$ such that $\mathbf{s}^T D\mathbf{1} = 0$, $\mathbf{s}^T D\mathbf{s} = 1$ and $\mathbf{s}D\mathbf{v}_2 \neq 0$ where $\mathbf{v}_2$ is the normalized eigenvector of $\lambda_2(\mathcal{L})$. Also let $0 \leq \kappa < 1$. Then the optimal solution $\mathbf{x}^*$ to $\textsc{LocalSpectral}(\mathbf{s}, \kappa)$ is of the form $\mathbf{x}^* = c(L - \gamma D)^+ D\mathbf{s}$, where c is chosen such that norm of x^* is 1 and γ is tuned such that $\mathbf{s}^T D\mathbf{x}^* = \kappa$.

Theorem: Solution Computation (admitted)

For any $\epsilon > 0$, a solution of value at most $(1+\epsilon)\lambda(G, \mathbf{s}, \kappa)$ can be computed in time $\tilde{O}(m/\sqrt{\lambda_2(L)}\log(1/\epsilon))$ using the Conjugate Gradient Method.

Approximating LocalSparsestCut

Lemma

For $u \in V$, $\textsc{LocalSpectral}(G, v_{\{u\}}, 1/k)$ is a relaxation of $\textsc{LocalSparsestCut}$ i.e. $\lambda(G, v_{\{u\}}, 1/k) \leq \Phi(u, k)$

Theorem

Given an unweighted graph $G = (V, E)$ a vertex $u \in V$ and a positive integer k, we can find a cut in G of conductance at most $O(\sqrt{\Phi(u, k)})$ in nearly linear time in the size of the graph.

Improving the approximation

Theorem: Cut Improvement (admitted)

Let $G = (V, E)$ and $\mathbf{s} \in \mathbb{R}^n$ such that $\mathbf{s}^T D\mathbf{1} = 0$. Let also $\kappa \geq 0$ be a correlation parameter. Then for all sets $T \subset V$, let $\kappa' = (\mathbf{s}^T D\mathbf{v}_T)^2$ and we have:

$$\Phi(T) \geq \lambda(G, \mathbf{s}, \kappa)$$

if $\kappa \leq \kappa'$, else

$$\Phi(T) \geq \frac{\kappa'}{\kappa}\lambda(G, \mathbf{s}, \kappa)$$