

COMPSCI 514: Algorithms for Data Science

Cameron Musco

University of Massachusetts Amherst. Fall 2022.

Lecture 19

- Problem Set 3 is due Monday at 11:59pm.
- No quiz due.

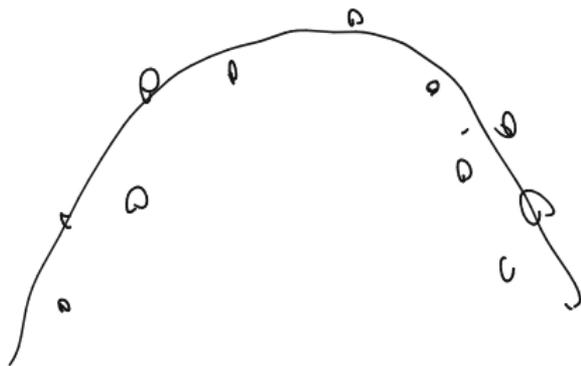
Summary

Last Class: Applications of Low-Rank Approximation

- Matrix completion
- Entity Embeddings.

\Rightarrow high dimensional vectors \rightarrow low approx.

Non-linear dimensionality reduction via low-rank approximation of near-neighbor graphs



Summary

Last Class: Applications of Low-Rank Approximation

- Matrix completion
- Entity Embeddings.
- Non-linear dimensionality reduction via low-rank approximation of near-neighbor graphs

This Class: Spectral Graph Theory and Spectral Clustering

- Start on graph clustering for community detection and non-linear clustering.
- **Spectral clustering**: finding good cuts via Laplacian eigenvectors.
- Start on **stochastic block model**: A simple clustered graph model where we can prove the effectiveness of spectral clustering.

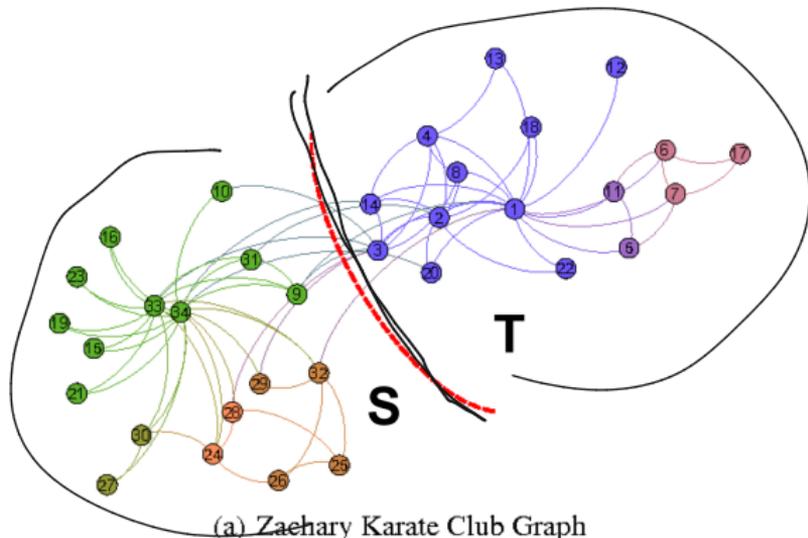
Spectral Clustering

A very common task is to **partition or cluster** vertices in a graph based on similarity/connectivity.

Spectral Clustering

A very common task is to **partition or cluster** vertices in a graph based on similarity/connectivity.

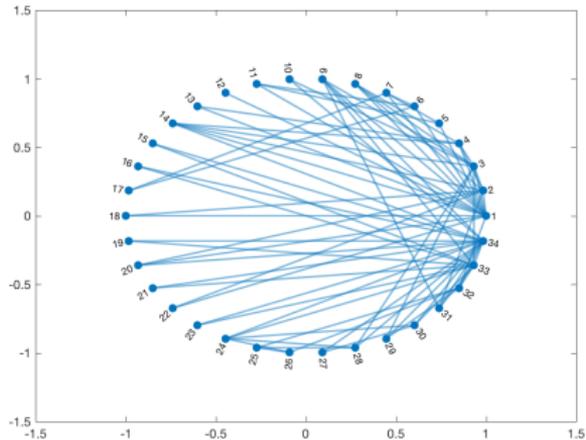
Community detection in naturally occurring networks.



Spectral Clustering

A very common task is to **partition or cluster** vertices in a graph based on similarity/connectivity.

Community detection in naturally occurring networks.



Spectral Clustering

A very common task is to **partition or cluster** vertices in a graph based on similarity/connectivity.

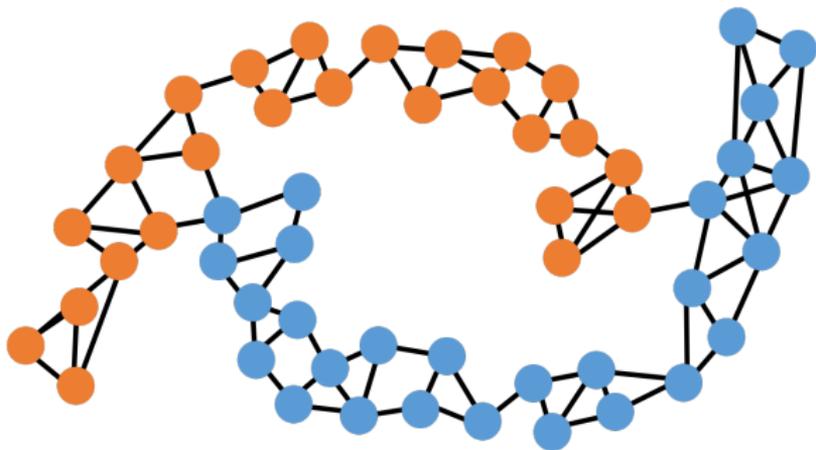
Non-linearly separable data.



Spectral Clustering

A very common task is to **partition or cluster** vertices in a graph based on similarity/connectivity.

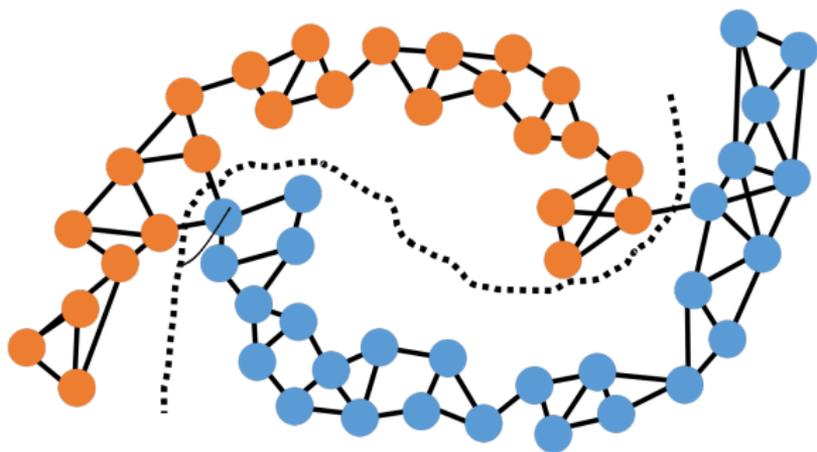
Non-linearly separable data.



Spectral Clustering

A very common task is to **partition or cluster** vertices in a graph based on similarity/connectivity.

Non-linearly separable data.

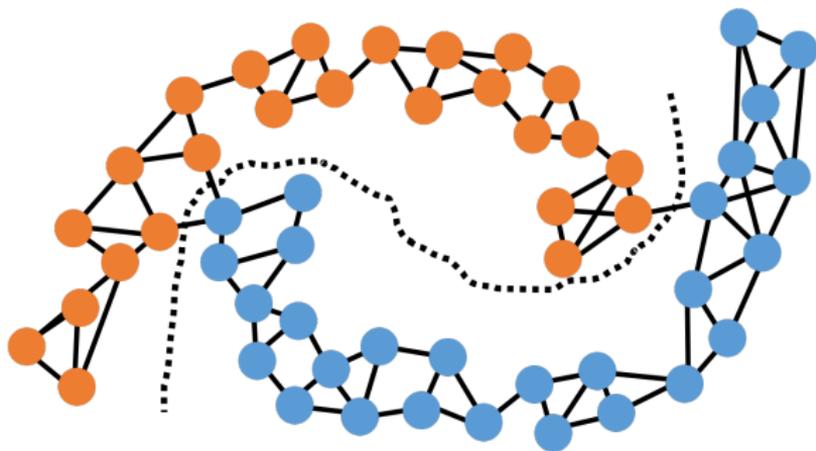


- kernel SVMs,

Spectral Clustering

A very common task is to **partition or cluster** vertices in a graph based on similarity/connectivity.

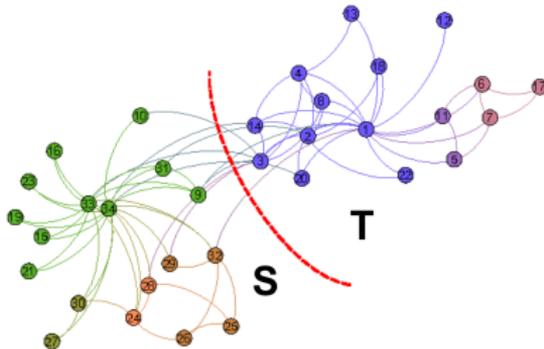
Non-linearly separable data.



Next Few Classes: Find this cut using eigendecomposition. First – motivate why this type of approach makes sense.

Cut Minimization

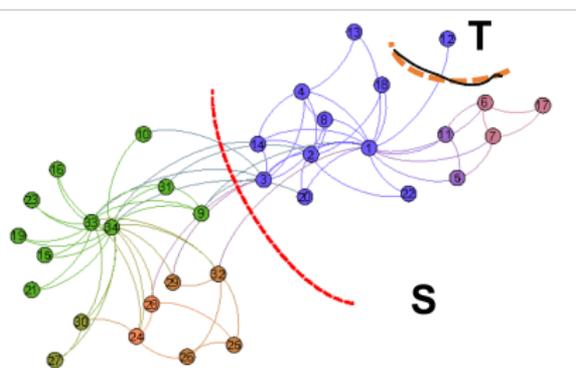
Simple Idea: Partition clusters along minimum cut in graph.



(a) Zachary Karate Club Graph

Cut Minimization

Simple Idea: Partition clusters along minimum cut in graph.



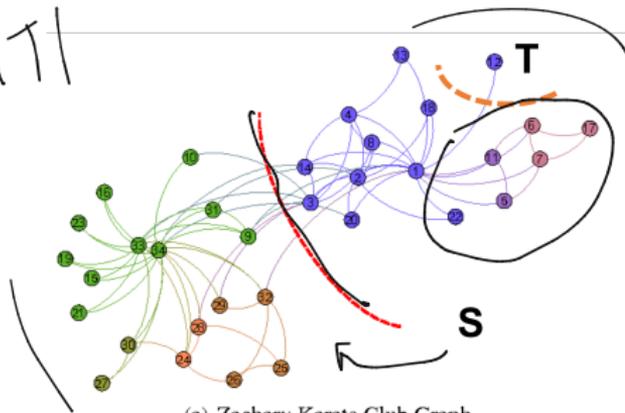
(a) Zachary Karate Club Graph

Small cuts are often not informative.

Cut Minimization

Simple Idea: Partition clusters along minimum cut in graph.

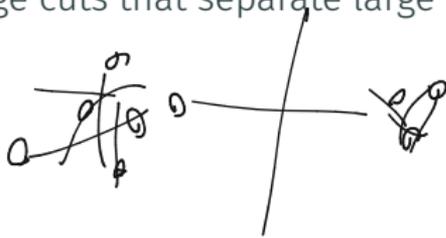
$$|S| = |T|$$



(a) Zachary Karate Club Graph

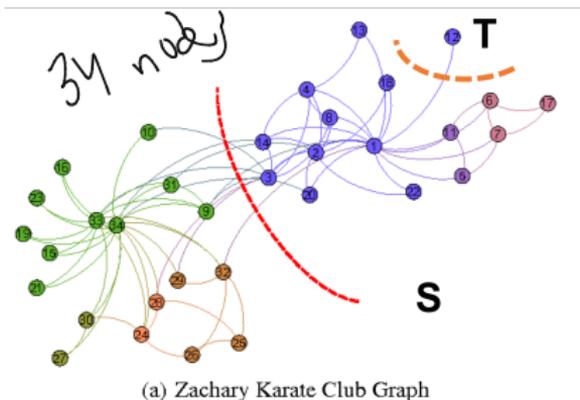
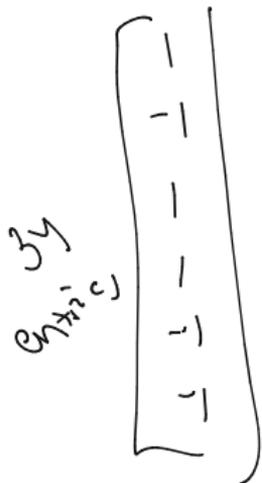
Small cuts are often not informative.

Solution: Encourage cuts that separate large sections of the graph.



Cut Minimization

Simple Idea: Partition clusters along minimum cut in graph.



$$\begin{aligned} \sqrt{1} &= \langle \mathbf{v}, \mathbf{1} \rangle \\ &= \sum_{i=1}^n 1 \cdot v(i) = \sum_{i=1}^n v(i) \end{aligned}$$

Small cuts are often not informative.

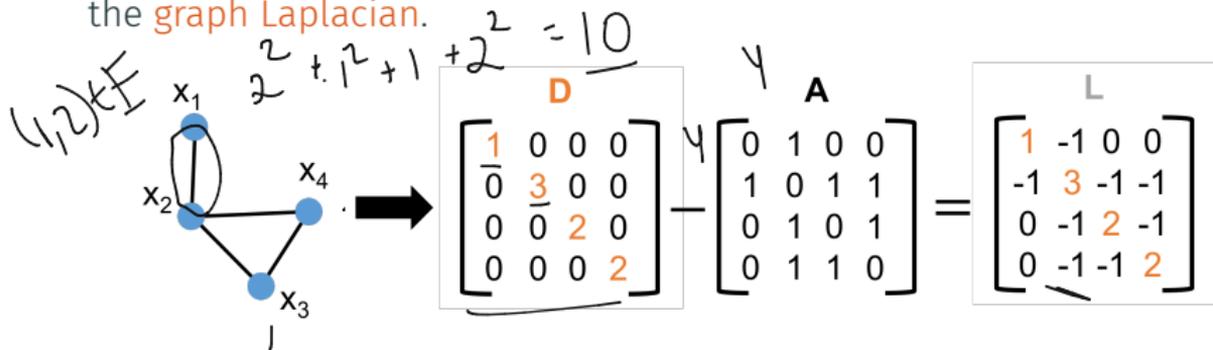
Solution: Encourage cuts that separate large sections of the graph.

- Let $\vec{v} \in \mathbb{R}^n$ be a **cut indicator**: $\vec{v}(i) = 1$ if $i \in S$. $\vec{v}(i) = -1$ if $i \in T$.
Want \vec{v} to have roughly equal numbers of 1s and -1s. I.e.,

$$\vec{v}^T \vec{1} \approx 0.$$

The Laplacian View

For a graph with adjacency matrix A and degree matrix D , $L = D - A$ is the **graph Laplacian**.



For any vector \vec{v} , its 'smoothness' over the graph is given by:

$V = \begin{bmatrix} 2 \\ 0 \\ -1 \\ 1 \end{bmatrix}$

$\sum_{(i,j) \in E} (\vec{v}(i) - \vec{v}(j))^2 = \vec{v}^T L \vec{v}$

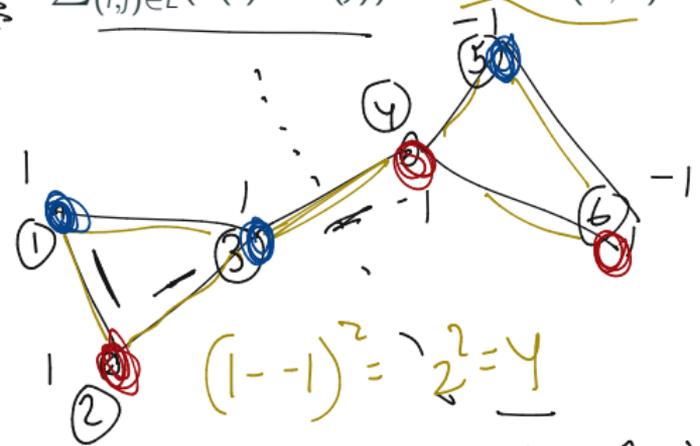
$= \sum_{i,j} \underbrace{v(i)^2 + v(j)^2}_{\text{degree}} - \underbrace{2v_i v_j}_{\text{adjacency}}$

$[V^T] [L] [V]$

The Laplacian View

For a cut indicator vector $\vec{v} \in \{-1, 1\}^n$ with $\vec{v}(i) = -1$ for $i \in S$ and $\vec{v}(i) = 1$ for $i \in T$:

1. $\vec{v}^T L \vec{v} = \sum_{(i,j) \in E} (\vec{v}(i) - \vec{v}(j))^2 = 4 \cdot \text{cut}(S, T)$.



$$\vec{v}^T L \vec{v}$$



$$\vec{v}^T L \vec{v} = 2^2 + 2^2 + 2^2 + 2^2 + 2^2 = 20$$

The Laplacian View

E edges

For a cut indicator vector $\vec{v} \in \{-1, 1\}^n$ with $\vec{v}(i) = -1$ for $i \in S$ and $\vec{v}(i) = 1$ for $i \in T$:

$$1. \vec{v}^T \mathbf{L} \vec{v} = \sum_{(i,j) \in E} (\vec{v}(i) - \vec{v}(j))^2 = 4 \cdot \text{cut}(S, T).$$

$$2. \underbrace{|\vec{v}^T \vec{1}| = ||T|| - |S|}$$

The Laplacian View

For a cut indicator vector $\vec{v} \in \{-1, 1\}^n$ with $\vec{v}(i) = -1$ for $i \in S$ and $\vec{v}(i) = 1$ for $i \in T$:

1. $\vec{v}^T \mathbf{L} \vec{v} = \sum_{(i,j) \in E} (\vec{v}(i) - \vec{v}(j))^2 = 4 \cdot \text{cut}(S, T)$.
2. $\vec{v}^T \vec{1} = |V| - |S|$.

Want to minimize both $\vec{v}^T \mathbf{L} \vec{v}$ (cut size) and $\vec{v}^T \vec{1}$ (imbalance).

The Laplacian View

For a cut indicator vector $\vec{v} \in \{-1, 1\}^n$ with $\vec{v}(i) = -1$ for $i \in S$ and $\vec{v}(i) = 1$ for $i \in T$:

1. $\vec{v}^T \mathbf{L} \vec{v} = \sum_{(i,j) \in E} (\vec{v}(i) - \vec{v}(j))^2 = 4 \cdot \text{cut}(S, T).$

2. $\vec{v}^T \vec{1} = |V| - |S|.$

Want to minimize both $\vec{v}^T \mathbf{L} \vec{v}$ (cut size) and $|\vec{v}^T \vec{1}|$ (imbalance).

Next Step: See how this dual minimization problem is naturally solved (sort of) by eigendecomposition.

Smallest Laplacian Eigenvector

The smallest eigenvector of the Laplacian is:

$$\min_{v: \|v\|_2=1} v^T L v$$

$$L v = \lambda v \quad \|v\|_2=1$$

$$v_n = \frac{1}{\sqrt{n}} \cdot \vec{1} = \arg \min_{v \in \mathbb{R}^n \text{ with } \|v\|=1} v^T L v$$

$$0 = v_n^T L v_n = \lambda \cdot v_n^T v_n = \lambda$$

with eigenvalue $\lambda_n(L) = \underline{\underline{v_n^T L v_n}} = 0$. Why?

$$v_n = \begin{bmatrix} \frac{1}{\sqrt{n}} \\ \frac{1}{\sqrt{n}} \\ \vdots \\ \frac{1}{\sqrt{n}} \end{bmatrix}$$

$$\|v_n\|_2 = 1$$

$$\begin{bmatrix} 3 & -1 & -1 & -1 \\ -1 & 2 & -1 & \\ -1 & -1 & & \\ -1 & & & \end{bmatrix} \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

$$v_n^T L v_n = \sum_{(i,j) \in E} [v_n(i) - v_n(j)]^2 = 0$$

n : number of nodes in graph, $A \in \mathbb{R}^{n \times n}$: adjacency matrix, $D \in \mathbb{R}^{n \times n}$: diagonal degree matrix, $L \in \mathbb{R}^{n \times n}$: Laplacian matrix $L = A - D$.

$$v^T L v \geq 0$$

Second Smallest Laplacian Eigenvector

By Courant-Fischer, the second smallest eigenvector is given by:

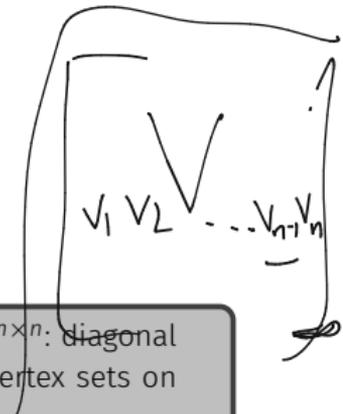
$$\vec{v}_{n-1} = \arg \min_{v \in \mathbb{R}^n \text{ with } \|\vec{v}\|=1, \vec{v}_n^T \vec{v}=0} \vec{v}^T L \vec{v}.$$

$$v^T v = [1]$$

$$V_n^T V = 0$$

$$V^T V = I$$

$$\frac{1}{\sqrt{n}} \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix} \leftrightarrow \mathbb{1}^T v_{n-1} = 0$$



n : number of nodes in graph, $A \in \mathbb{R}^{n \times n}$: adjacency matrix, $D \in \mathbb{R}^{n \times n}$: diagonal degree matrix, $L \in \mathbb{R}^{n \times n}$: Laplacian matrix $L = A - D$. S, T : vertex sets on different sides of cut.

Second Smallest Laplacian Eigenvector

By Courant-Fischer, the second smallest eigenvector is given by:

$$\vec{v}_{n-1} = \arg \min_{v \in \mathbb{R}^n \text{ with } \|\vec{v}\|=1, \vec{v}_n^T \vec{v}=0} \vec{v}^T L \vec{v}.$$

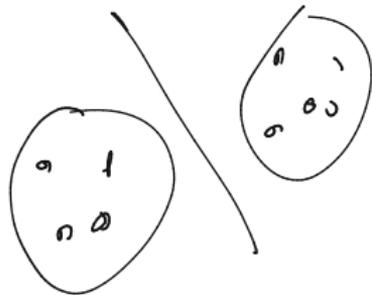
If \vec{v}_{n-1} were in $\left\{ -\frac{1}{\sqrt{n}}, \frac{1}{\sqrt{n}} \right\}^n$ it would have:

• $\vec{v}_{n-1}^T L \vec{v}_{n-1} = \frac{4}{n} \cdot \text{cut}(S, T)$ as small as possible given that

$$\vec{v}_{n-1}^T \vec{v}_n = \frac{1}{\sqrt{n}} \vec{v}_{n-1}^T \vec{1} = \frac{|T| - |S|}{n} = 0.$$

$$\vec{v}^T L \vec{v} = \sum_{(i,j) \in E} \left(\frac{v(i) - v(j)}{\sqrt{n}} \right)^2$$

$$\sum_{i=1}^n v_{n-1}(i)$$



$$\begin{bmatrix} -1/\sqrt{n} \\ \dots \\ 1/\sqrt{n} \end{bmatrix}$$

n : number of nodes in graph, $A \in \mathbb{R}^{n \times n}$: adjacency matrix, $D \in \mathbb{R}^{n \times n}$: diagonal degree matrix, $L \in \mathbb{R}^{n \times n}$: Laplacian matrix $L = A - D$. S, T : vertex sets on different sides of cut.

Second Smallest Laplacian Eigenvector

By Courant-Fischer, the second smallest eigenvector is given by:

$$\vec{v}_{n-1} = \underset{v \in \mathbb{R}^n \text{ with } \|\vec{v}\|=1, \vec{v}_n^T \vec{v}=0}{\arg \min} \vec{v}^T \mathbf{L} \vec{v}.$$

If \vec{v}_{n-1} were in $\left\{-\frac{1}{\sqrt{n}}, \frac{1}{\sqrt{n}}\right\}^n$ it would have:

- $\vec{v}_{n-1}^T \mathbf{L} \vec{v}_{n-1} = \frac{4}{\sqrt{n}} \cdot \text{cut}(S, T)$ as small as possible given that $\vec{v}_{n-1}^T \vec{v}_n = \frac{1}{\sqrt{n}} \vec{v}_{n-1}^T \vec{1} = \frac{|T|-|S|}{n} = 0$.
- I.e., \vec{v}_{n-1} would indicate the smallest perfectly balanced cut.

n : number of nodes in graph, $\mathbf{A} \in \mathbb{R}^{n \times n}$: adjacency matrix, $\mathbf{D} \in \mathbb{R}^{n \times n}$: diagonal degree matrix, $\mathbf{L} \in \mathbb{R}^{n \times n}$: Laplacian matrix $\mathbf{L} = \mathbf{A} - \mathbf{D}$. S, T : vertex sets on different sides of cut.

Second Smallest Laplacian Eigenvector

By Courant-Fischer, the second smallest eigenvector is given by:

$$\vec{v}_{n-1} = \arg \min_{v \in \mathbb{R}^n \text{ with } \|\vec{v}\|=1, \vec{v}_n^T \vec{v}=0} \vec{v}^T L \vec{v}.$$

Handwritten diagram showing a vector v in the plane spanned by $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$. A line representing the constraint $v_n^T v = 0$ is drawn, and the vector v is shown to be orthogonal to this line.

If \vec{v}_{n-1} were in $\left\{ -\frac{1}{\sqrt{n}}, \frac{1}{\sqrt{n}} \right\}^n$ it would have:

- $\vec{v}_{n-1}^T L \vec{v}_{n-1} = \frac{4}{\sqrt{n}} \cdot \text{cut}(S, T)$ as small as possible given that

$$\vec{v}_{n-1}^T \vec{v}_n = \frac{1}{\sqrt{n}} \vec{v}_{n-1}^T \vec{1} = \frac{|T| - |S|}{n} = 0.$$

- I.e., \vec{v}_{n-1} would indicate the smallest perfectly balanced cut.

- The eigenvector $\vec{v}_{n-1} \in \mathbb{R}^n$ is not generally binary, but still satisfies a 'relaxed' version of this property.

n : number of nodes in graph, $\mathbf{A} \in \mathbb{R}^{n \times n}$: adjacency matrix, $\mathbf{D} \in \mathbb{R}^{n \times n}$: diagonal degree matrix, $\mathbf{L} \in \mathbb{R}^{n \times n}$: Laplacian matrix $\mathbf{L} = \mathbf{A} - \mathbf{D}$. S, T : vertex sets on different sides of cut.

Cutting With the Second Laplacian Eigenvector

Find a good partition of the graph by computing

$$\vec{v}_{n-1} = \underset{v \in \mathbb{R}^d \text{ with } \|\vec{v}\|=1, \vec{v}^T \vec{1}=0}{\text{arg min}} \quad \vec{v}^T \mathbf{L} \vec{v}.$$

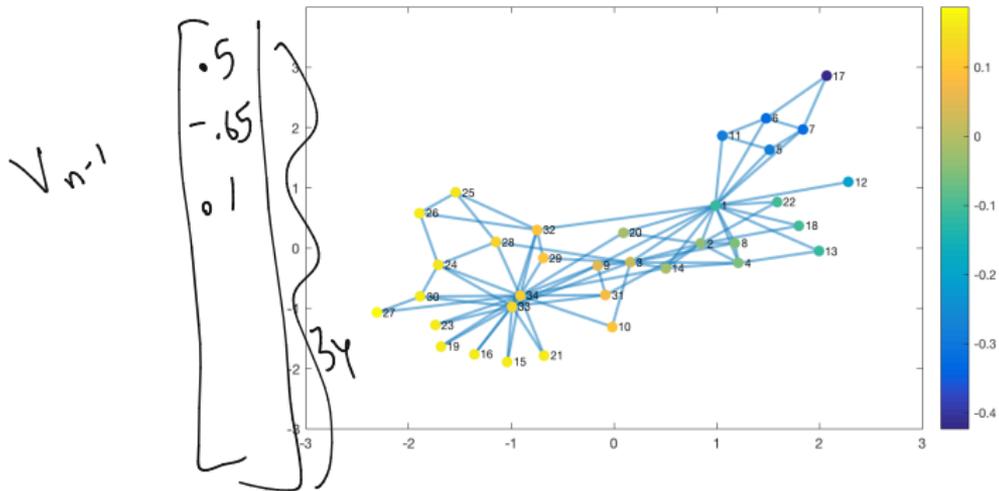
Set S to be all nodes with $\vec{v}_{n-1}(i) < 0$, T to be all with $\vec{v}_2(i) \geq 0$.

Cutting With the Second Laplacian Eigenvector

Find a good partition of the graph by computing

$$\vec{v}_{n-1} = \arg \min_{v \in \mathbb{R}^d \text{ with } \|\vec{v}\|=1, \vec{v}^T \vec{1}=0} \vec{v}^T L \vec{v}.$$

Set S to be all nodes with $\vec{v}_{n-1}(i) < 0$, T to be all with $\vec{v}_{n-1}(i) \geq 0$.



Cutting With the Second Laplacian Eigenvector

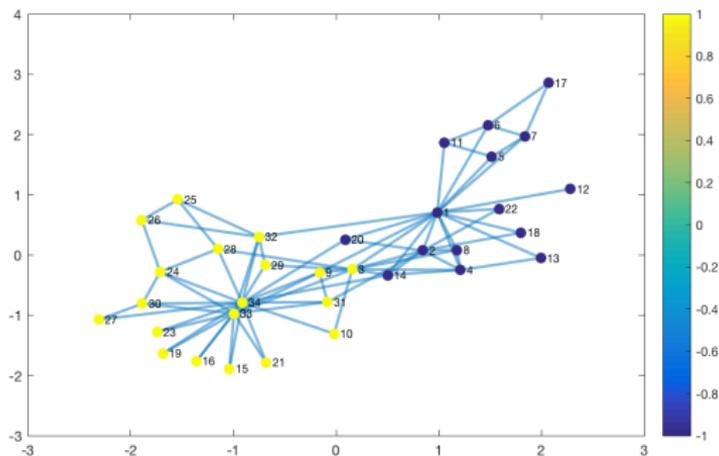
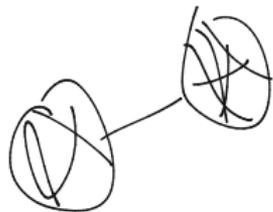
Find a good partition of the graph by computing

$$\vec{v}_{n-1} = \arg \min_{v \in \mathbb{R}^d \text{ with } \|\vec{v}\|=1, \vec{v}^T \vec{1}=0} \vec{v}^T L \vec{v}.$$

Fiedler vector

v_{n-1}

Set S to be all nodes with $\vec{v}_{n-1}(i) < 0$, T to be all with $\vec{v}_{n-1}(i) \geq 0$.



$$v_{n-1}^T \vec{1} = 0$$

Spectral Partitioning in Practice

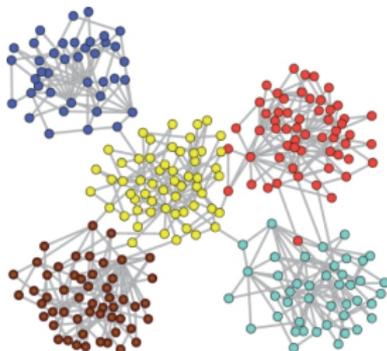
The Shi-Malik normalized cuts algorithm is one of the most commonly used variants of this approach, using the normalized Laplacian $\bar{\mathbf{L}} = \mathbf{D}^{-1/2}\mathbf{L}\mathbf{D}^{-1/2}$.

n : number of nodes in graph, $\mathbf{A} \in \mathbb{R}^{n \times n}$: adjacency matrix, $\mathbf{D} \in \mathbb{R}^{n \times n}$: diagonal degree matrix, $\mathbf{L} \in \mathbb{R}^{n \times n}$: Laplacian matrix $\mathbf{L} = \mathbf{A} - \mathbf{D}$.

Spectral Partitioning in Practice

The Shi-Malik normalized cuts algorithm is one of the most commonly used variants of this approach, using the normalized Laplacian $\bar{\mathbf{L}} = \mathbf{D}^{-1/2}\mathbf{L}\mathbf{D}^{-1/2}$.

Important Consideration: What to do when we want to split the graph into more than two parts?



n : number of nodes in graph, $\mathbf{A} \in \mathbb{R}^{n \times n}$: adjacency matrix, $\mathbf{D} \in \mathbb{R}^{n \times n}$: diagonal degree matrix, $\mathbf{L} \in \mathbb{R}^{n \times n}$: Laplacian matrix $\mathbf{L} = \mathbf{A} - \mathbf{D}$.

Spectral Partitioning in Practice

The Shi-Malik normalized cuts algorithm is one of the most commonly used variants of this approach, using the normalized Laplacian $\bar{\mathbf{L}} = \mathbf{D}^{-1/2}\mathbf{L}\mathbf{D}^{-1/2}$.

Important Consideration: What to do when we want to split the graph into more than two parts?

Spectral Clustering:

n : number of nodes in graph, $\mathbf{A} \in \mathbb{R}^{n \times n}$: adjacency matrix, $\mathbf{D} \in \mathbb{R}^{n \times n}$: diagonal degree matrix, $\mathbf{L} \in \mathbb{R}^{n \times n}$: Laplacian matrix $\mathbf{L} = \mathbf{A} - \mathbf{D}$.

Spectral Partitioning in Practice

The Shi-Malik normalized cuts algorithm is one of the most commonly used variants of this approach, using the normalized Laplacian $\bar{\mathbf{L}} = \mathbf{D}^{-1/2}\mathbf{L}\mathbf{D}^{-1/2}$.

Important Consideration: What to do when we want to split the graph into more than two parts?

Spectral Clustering:

- Compute smallest k nonzero eigenvectors $\vec{v}_{n-1}, \dots, \vec{v}_{n-k}$ of $\bar{\mathbf{L}}$.

n : number of nodes in graph, $\mathbf{A} \in \mathbb{R}^{n \times n}$: adjacency matrix, $\mathbf{D} \in \mathbb{R}^{n \times n}$: diagonal degree matrix, $\mathbf{L} \in \mathbb{R}^{n \times n}$: Laplacian matrix $\mathbf{L} = \mathbf{A} - \mathbf{D}$.

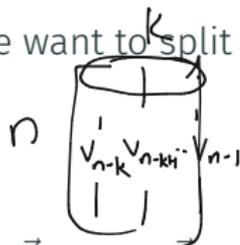
Spectral Partitioning in Practice

The Shi-Malik normalized cuts algorithm is one of the most commonly used variants of this approach, using the normalized Laplacian $\bar{\mathbf{L}} = \mathbf{D}^{-1/2}\mathbf{L}\mathbf{D}^{-1/2}$.

Important Consideration: What to do when we want to split the graph into more than two parts?

Spectral Clustering:

- Compute smallest k nonzero eigenvectors $\vec{v}_{n-1}, \dots, \vec{v}_{n-k}$ of $\bar{\mathbf{L}}$.
- Represent each node by its corresponding row in $\mathbf{V} \in \mathbb{R}^{n \times k}$ whose columns are $\vec{v}_{n-1}, \dots, \vec{v}_{n-k}$.



n : number of nodes in graph, $\mathbf{A} \in \mathbb{R}^{n \times n}$: adjacency matrix, $\mathbf{D} \in \mathbb{R}^{n \times n}$: diagonal degree matrix, $\mathbf{L} \in \mathbb{R}^{n \times n}$: Laplacian matrix $\mathbf{L} = \mathbf{A} - \mathbf{D}$.

Spectral Partitioning in Practice

The Shi-Malik normalized cuts algorithm is one of the most commonly used variants of this approach, using the normalized Laplacian $\bar{\mathbf{L}} = \mathbf{D}^{-1/2}\mathbf{L}\mathbf{D}^{-1/2}$.

Important Consideration: What to do when we want to split the graph into more than two parts?

Spectral Clustering:

$$\begin{pmatrix} -1 \\ .5 \\ .6 \end{pmatrix}$$

- Compute smallest k nonzero eigenvectors $\vec{v}_{n-1}, \dots, \vec{v}_{n-k}$ of $\bar{\mathbf{L}}$.
- Represent each node by its corresponding row in $\mathbf{V} \in \mathbb{R}^{n \times k}$ whose columns are $\vec{v}_{n-1}, \dots, \vec{v}_{n-k}$.
- Cluster these rows using k -means clustering (or really any clustering method).

n : number of nodes in graph, $\mathbf{A} \in \mathbb{R}^{n \times n}$: adjacency matrix, $\mathbf{D} \in \mathbb{R}^{n \times n}$: diagonal degree matrix, $\mathbf{L} \in \mathbb{R}^{n \times n}$: Laplacian matrix $\mathbf{L} = \mathbf{A} - \mathbf{D}$.

Laplacian Embedding

The smallest eigenvectors of $\mathbf{L} = \mathbf{D} - \mathbf{A}$ give the orthogonal 'functions' that are smoothest over the graph. I.e., minimize

$$\vec{v}^T \mathbf{L} \vec{v} = \sum_{(i,j) \in E} [\vec{v}(i) - \vec{v}(j)]^2.$$

Laplacian Embedding

The smallest eigenvectors of $\mathbf{L} = \mathbf{D} - \mathbf{A}$ give the orthogonal 'functions' that are smoothest over the graph. I.e., minimize

$$\vec{v}^T \mathbf{L} \vec{v} = \sum_{(i,j) \in E} [\vec{v}(i) - \vec{v}(j)]^2.$$

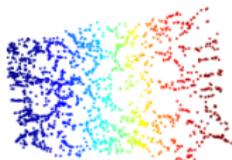
Embedding points with coordinates given by $[\vec{v}_{n-1}(j), \vec{v}_{n-2}(j), \dots, \vec{v}_{n-k}(j)]$ ensures that coordinates connected by edges have minimum total squared Euclidean distance.

Laplacian Embedding

The smallest eigenvectors of $\mathbf{L} = \mathbf{D} - \mathbf{A}$ give the orthogonal 'functions' that are smoothest over the graph. I.e., minimize

$$\vec{v}^T \mathbf{L} \vec{v} = \sum_{(i,j) \in E} [\vec{v}(i) - \vec{v}(j)]^2.$$

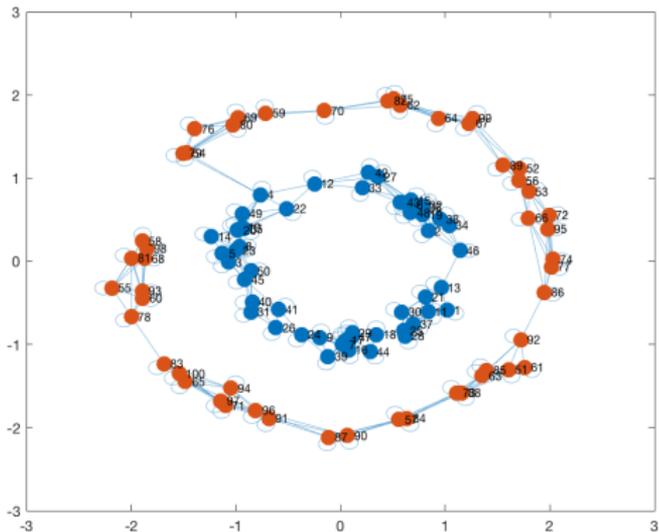
Embedding points with coordinates given by $[\vec{v}_{n-1}(j), \vec{v}_{n-2}(j), \dots, \vec{v}_{n-k}(j)]$ ensures that coordinates connected by edges have minimum total squared Euclidean distance.



- Spectral Clustering
- Laplacian Eigenmaps
- Locally linear embedding
- Isomap
- Node2Vec, DeepWalk, etc.
(variants on Laplacian)

Laplacian Embedding

k -Nearest Neighbors Graph:



Laplacian Embedding

Embedding with eigenvectors $\vec{v}_{n-1}, \vec{v}_{n-2}$: (linearly separable)

