

COMPSCI 514: ALGORITHMS FOR DATA SCIENCE

Cameron Musco

University of Massachusetts Amherst. Spring 2020.

Lecture 18

- Problem Set 3 is due Monday 4/13 at 8pm.
- Office hours zoom link has changed.
Get from class homepage.

Last Class: Applications of Low-Rank Approximation

- Low-rank matrix completion (predicting missing measurements using low-rank structure).
- Entity embeddings (e.g., LSA, word embeddings). View as low-rank approximation of a similarity matrix.
- Start on spectral graph theory.

Last Class: Applications of Low-Rank Approximation

- Low-rank matrix completion (predicting missing measurements using low-rank structure).
- Entity embeddings (e.g., LSA, word embeddings). View as low-rank approximation of a similarity matrix.
- Start on spectral graph theory.


This Class: More Spectral Graph Theory & Spectral Clustering.

- Using eigendecomposition to partition graphs into clusters.
- Clustering non-linearly separable data.
- Application to the *stochastic block model* and community detection.

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

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


Community detection in naturally occurring networks.



(a) Zachary Karate Club Graph


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

Community detection in naturally occurring networks.



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


Linearly separable data.



SVM
logistic regress.
k-means

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


Non-linearly separable data k -nearest neighbor graph.



SPECTRAL CLUSTERING


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

Non-linearly separable data k -nearest neighbor graph.




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

Non-linearly separable data k -nearest neighbor graph.




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

Non-linearly separable data k -nearest neighbor graph.



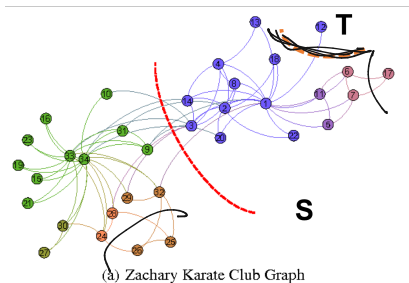
This Class: Find this cut using eigendecomposition. First – motivate why this type of approach makes sense.

Simple Idea: Partition clusters along minimum cut in graph.




(a) Zachary Karate Club Graph

Simple Idea: Partition clusters along minimum cut in graph.



Small cuts are often not informative.

Simple Idea: Partition clusters along minimum cut in graph.



balanced min-cut.


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.

2 partitions
clusters have
roughly equal
size



$$\sum_{i=1}^n v(i) = |S| - |T|$$


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 \mathbf{A} and degree matrix \mathbf{D} , $\mathbf{L} = \mathbf{D} - \mathbf{A}$ is the **graph Laplacian**.



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} ,

$$\vec{v}^T L \vec{v} = \vec{v}^T D \vec{v} - \vec{v}^T A \vec{v} = \sum_{i=1}^n d(i) \vec{v}(i)^2 - \sum_{i=1}^n \sum_{j=1}^n A(i, j) \cdot \vec{v}(i) \cdot \vec{v}(j)$$


0 when i, j aren't connected

note: i appears in $d(i)$ different edges

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

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 $i \in T$:

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



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

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

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

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

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

Next Step: See how this dual minimization problem is naturally solved by eigendecomposition.

SMALLEST LAPLACIAN EIGENVECTOR

$$v_1 \dots \dots \underline{v_n} \quad \lambda_1 > \lambda_2 > \dots > \lambda_n$$

The smallest eigenvector of the Laplacian is:

$$\vec{v}_n = \frac{1}{\sqrt{n}} \cdot \vec{1} = \underset{v \in \mathbb{R}^n \text{ with } \|\vec{v}\|=1}{\operatorname{arg\,min}} \quad \vec{v}^T L \vec{v}$$

with eigenvalue $\vec{v}_n^T L \vec{v}_n = 0$.

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

Laplacian is positive semidefinite
all nonnegative eigenvalues

$$\begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix} \frac{1}{\sqrt{n}}$$

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 $L = A - D$.

SMALLEST LAPLACIAN EIGENVECTOR

The smallest eigenvector of the Laplacian is:

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

with eigenvalue $\vec{v}_n^T L \vec{v}_n = 0$. Why?

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

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 $L = A - D$.


SMALLEST LAPLACIAN EIGENVECTOR

The smallest eigenvector of the Laplacian is:

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

with eigenvalue $\vec{v}_n^T L \vec{v}_n = 0$. Why?

λ_n



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 $L = A - D$.

SECOND SMALLEST LAPLACIAN EIGENVECTOR

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

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

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:

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

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

$\vec{v}_{n-1}^T L \vec{v}_{n-1} = \text{cut}(S, T)$ as small as possible given that
 $\vec{v}_{n-1}^T \vec{1} = |T| - |S| = 0.$

$$V_{n-1}^T V_n = 0 \rightarrow V_{n-1}^T \frac{1}{\sqrt{n}} \vec{1} = 0 \rightarrow V_{n-1}^T \vec{1} = 0$$

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:

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

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

- $\vec{v}_{n-1}^T L \vec{v}_{n-1} = \operatorname{cut}(S, T)$ as small as possible given that $\vec{v}_{n-1}^T \vec{1} = |T| - |S| = 0$.
- 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:

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

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

- $\vec{v}_{n-1}^T L \vec{v}_{n-1} = \text{cut}(S, T)$ as small as possible given that $\vec{v}_{n-1}^T \vec{1} = |T| - |S| = 0$.
- 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, $A \in \mathbb{R}^{n \times n}$: 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.

CUTTING WITH THE SECOND LAPLACIAN EIGENVECTOR

Find a good partition of the graph by computing

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

→ second smallest
of Laplacian

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


CUTTING WITH THE SECOND LAPLACIAN EIGENVECTOR

Find a good partition of the graph by computing

$$\vec{v}_2 = \underset{v \in \mathbb{R}^d \text{ with } \|\vec{v}\|=1, \vec{v}_2^T \vec{1}=0}{\text{arg min}} \quad \vec{v}^T L \vec{v}$$

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

V_{n-1}



CUTTING WITH THE SECOND LAPLACIAN EIGENVECTOR

Find a good partition of the graph by computing


$$\vec{v}_2 = \arg \min_{v \in \mathbb{R}^d \text{ with } \|\vec{v}\|=1, \vec{v}^T \vec{1}=0} \vec{v}^T L \vec{v}$$

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

$$V_{n-1}^T \vec{1} = 0 \quad \text{balance condition}$$

$$V_{n-1}^T V_n = 0 \quad \text{ortho. of eigs.}$$

$$V^T L V$$




orthogonality \rightarrow balance cut

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 $\mathbf{L} = \mathbf{A} - \mathbf{D}$.

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 $\mathbf{L} = \mathbf{A} - \mathbf{D}$.

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 $\mathbf{L} = \mathbf{A} - \mathbf{D}$.

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 $\bar{\mathbf{L}}$.

smallest

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 $\mathbf{L} = \mathbf{A} - \mathbf{D}$.

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:

$$n \begin{bmatrix} \vec{v}_{n-1} & \dots & \vec{v}_{n-k} \end{bmatrix} \begin{matrix} k \\ \end{matrix}$$

- 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 rows 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 $\mathbf{L} = \mathbf{A} - \mathbf{D}$.

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 rows 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

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



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

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

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)

For Partitioning graphs:
1. small cut
2. balanced cut
Introduced Laplacian
- showed that second smallest eig of L gives a small balanced cut
↓
Spectral clustering

Original Data: (not linearly separable)



k -Nearest Neighbors Graph:

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

