

COMPSCI 514: Algorithms for Data Science

Cameron Musco

University of Massachusetts Amherst. Fall 2024.

Lecture 18

- Problem Set 3 solutions have been posted.
- Problem Set 4 will be released soon.

Last Class: SVD and Applications of Low-Rank Approximation

$$X = UV^T$$

SVD and connections to eigendecomposition and optimal low-rank approximation.

- Matrix completion
- Entity Embeddings.

words \rightarrow vectors

Last Class: SVD and Applications of Low-Rank Approximation

- SVD and connections to eigendecomposition and optimal low-rank approximation.
- Matrix completion
- Entity Embeddings.

This Class: Linear Algebraic Techniques for Graph Analysis

- Start on graph clustering for community detection and non-linear clustering.
- Spectral clustering: finding good cuts via Laplacian eigenvectors.

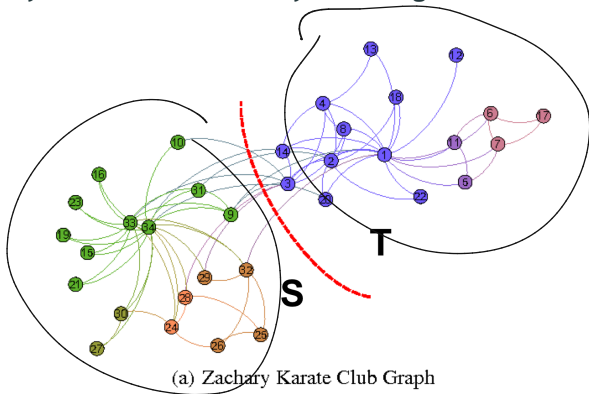
Spectral Clustering

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.



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.



(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.

$$\vec{v} = \begin{bmatrix} -1 \\ -1 \\ -1 \\ 1 \\ -1 \\ 1 \end{bmatrix}$$



(a) Zachary Karate Club Graph

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 -1 s. I.e.,


$$\sum_{i=1}^n v(i) = \vec{v}^T \vec{1} \approx 0.$$

The Laplacian View

$$\vec{v}^T L \vec{v} = \sum_{(i,j) \in E} (\vec{v}(i) - \vec{v}(j))^2 = \vec{v}^T L \vec{v} = 4$$

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

$n \times n$



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

$$\vec{v} \in \mathbb{R}^4$$

$$\vec{v} = \begin{bmatrix} 1 \\ 1 \\ 2 \\ 2 \end{bmatrix}$$

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

$$\vec{v}^T L \vec{v} = \underbrace{0^2}_{1} + \underbrace{(-1)^2}_{15} + \underbrace{(-1)^2}_{100} + \underbrace{0^2}_{5} = 2$$


$$\vec{v}^T L \vec{v} = \sum_{i=1}^n \vec{v}(i)^2 - \sum_{(i,j) \in E} \vec{v}(i)\vec{v}(j)$$

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

$$\begin{aligned} \vec{v}^T L \vec{v} &= 2^2 + 2^2 \\ &= 8 \end{aligned}$$



$$V = \begin{pmatrix} 1 \\ -1 \\ -1 \\ 1 \end{pmatrix}$$

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

want both
small

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

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

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:

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

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



$$v = \begin{bmatrix} 1 \\ \vdots \\ 1 \end{bmatrix}$$

$$v^T L v = 0$$

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 $L = D - A$.

Second Smallest Laplacian Eigenvector

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

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

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 $\mathbf{L} = \mathbf{D} - \mathbf{A}$. S, T : vertex sets on different sides of cut.

Second Smallest Laplacian Eigenvector

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

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

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

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

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

$$V_{n-1}^T V_n = 0$$

$$V_{n-1}^T \vec{1} = 0 \quad \left[\begin{array}{c} 1 \\ 1 \end{array} \right]$$

$$\sum_{i=1}^n V_{n-1}(i) = 0$$

$$V_n = \begin{bmatrix} 1 \\ \vdots \\ 1 \end{bmatrix} \cdot \frac{1}{\sqrt{6}}$$

$$V_{n-1} = \begin{bmatrix} 1 \\ \vdots \\ -1 \end{bmatrix} \cdot \frac{1}{\sqrt{6}}$$



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

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 $\mathbf{L} = \mathbf{D} - \mathbf{A}$. S, T : vertex sets on different sides of cut.

Cutting With the Second Laplacian Eigenvector

Find a good partition of the graph by computing

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

cut size
balance

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


$$\begin{bmatrix} -0.58 \\ 1 \\ 0.2 \\ -0.03 \\ \vdots \end{bmatrix}$$

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 $\vec{v}_2(i) \geq 0$.




Cutting With the Second Laplacian Eigenvector

Find a good partition of the graph by computing

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

generalize to
other thresholds

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



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 $\bar{\mathbf{L}} = \mathbf{D}^{-1/2} \mathbf{L} \mathbf{D}^{-1/2}$.


$$\mathbf{V}^T \mathbf{L} \mathbf{V}$$

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

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 $\bar{L} = D^{-1/2}LD^{-1/2}$.

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



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

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

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

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:

$$k = 2$$

\vec{v}_{n-1}	\vec{v}_{n-2}
-1	.2
-.7	.03
0	⋮
.3	⋮
1	⋮

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

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:

$$\sqrt{\mathbf{L}}\mathbf{V}$$

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

Laplacian Embedding

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


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

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

agglomerative clustering
kernel method
pick two kernel

Original Data: (not linearly separable)

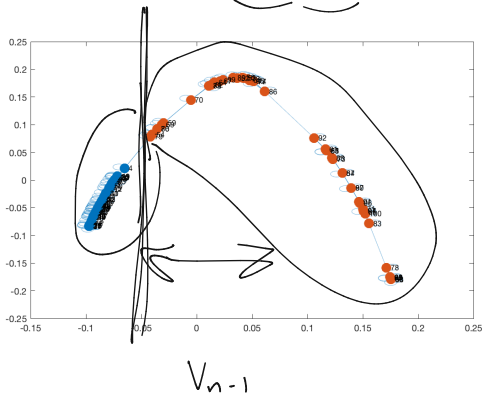


k -Nearest Neighbors Graph:



Laplacian Embedding

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



Generative Models

So Far: Have argued that spectral clustering partitions a graph effectively, along a small cut that separates the graph into large pieces. But it is difficult to give any formal guarantee on the 'quality' of the partitioning in general graphs.

Generative Models

So Far: Have argued that spectral clustering partitions a graph effectively, along a small cut that separates the graph into large pieces. But it is difficult to give any formal guarantee on the 'quality' of the partitioning in general graphs.

Common Approach: Give a natural generative model for random inputs and analyze how the algorithm performs on inputs drawn from this model.

- Very common in algorithm design for data analysis/machine learning (can be used to justify least squares regression, k -means clustering, PCA, etc.)
- We'll do this next time, introducing the Stochastic Block Model.

λ_{n-1}
(Fiedler value)
small = well partitioned
big = not well-partitioned

