

CS-GY 9223 I: Lecture 10

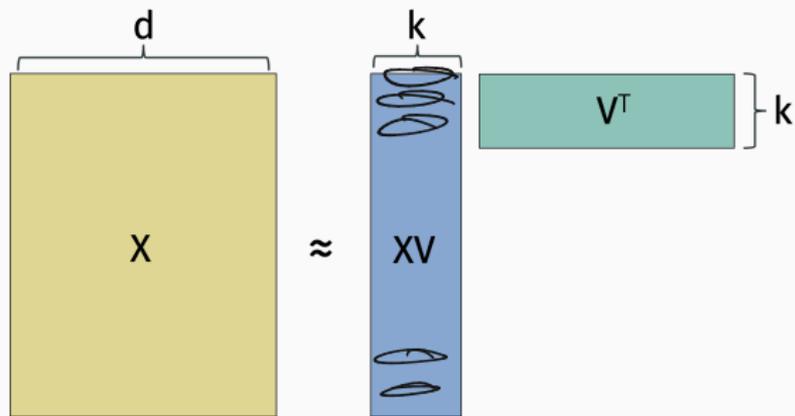
Spectral clustering, spectral graph theory.

NYU Tandon School of Engineering, Prof. Christopher Musco

- Project proposal feedback.
- Problem set.

LOW-RANK APPROXIMATION

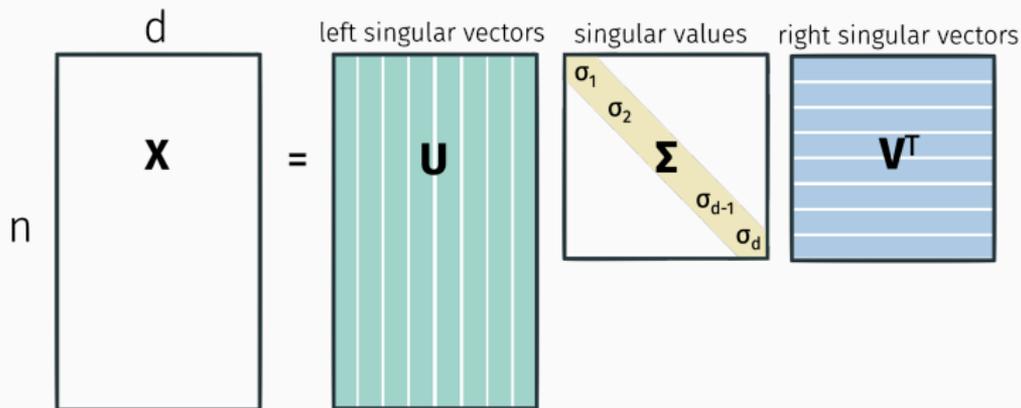
Write \mathbf{X} as a rank k factorization by projecting onto the subspace spanned by an orthonormal matrix $\mathbf{V} \in \mathbb{R}^{d \times k}$



SINGULAR VALUE DECOMPOSITION

One-stop shop for computing optimal low-rank approximations.

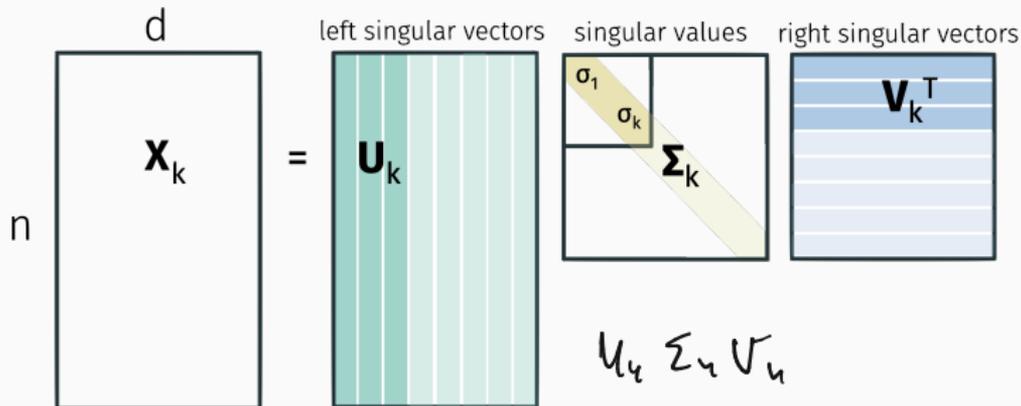
Any matrix X can be written:



Where $U^T U = I$, $V^T V = I$, and $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_d \geq 0$.

SINGULAR VALUE DECOMPOSITION

Can read off optimal low-rank approximations from the SVD:



$$X_k = U_k U_k^T X_k = \underbrace{X_k V_k^T}_{\text{circled}}$$

$$V_k = \arg \min_{\text{orthonormal } V \in \mathbb{R}^{d \times k}} \|X - X V V^T\|_F^2 = \arg \max_{\text{orthonormal } V \in \mathbb{R}^{d \times k}} \|X V V^T\|_F^2$$

POWER METHOD CONVERGENCE

Theorem (Power Method Convergence)

Let $\gamma = \frac{\sigma_1 - \sigma_2}{\sigma_1}$ be parameter capturing the “gap” between the first and second largest singular values. If Power Method is initialized with a random Gaussian vector then, with high probability, after $T = O\left(\frac{\log d/\epsilon}{\gamma}\right)$ steps, we have:

$$\|v_1 - \underline{z}^{(T)}\|_2 \leq \epsilon.$$

Total runtime: $O(T \cdot \text{nnz}(X)) \leq O(T \cdot nd)$

↙
number of non-zeros

BLOCK POWER METHOD FOR LARGER k

Block power method:

$$Z^T Z = I$$

- Choose $G \in \mathbb{R}^{d \times k}$ be a random Gaussian matrix.
 - $Z_0 = \text{orth}(G)$.
 - For $i = 1, \dots, T$
 - $Z^{(i)} = X^T \cdot (XZ^{(i-1)})$
 - $Z^{(i)} = \text{orth}(Z^{(i)})$
- Return $Z^{(T)}$

$\text{Gr}(Z)$

$$\underline{Z^{(T)}} \approx V_k$$

$$\text{rank}(A)$$

$$d \int \approx \int$$

Convergence Guarantee: $T = O\left(\frac{\log d/\epsilon}{\epsilon}\right)$ iterations to obtain a nearly optimal low-rank approximation:

$$\underline{\|A - AZZ^T\|_F^2} \leq (1 + \epsilon) \|A - AV_k V_k^T\|_F^2$$

Runtime: $O(\text{nnz}(X) \cdot k \cdot T) \leq O(\underline{ndk} \cdot T)$.

number of nonzeros in X

Possible to “accelerate” these methods. $1/\sqrt{\epsilon}$

Convergence Guarantee: $T = O\left(\frac{\log d/\epsilon}{\sqrt{\epsilon}}\right)$ iterations to obtain a nearly optimal low-rank approximation:

$$\|\mathbf{A} - \mathbf{AZZ}^T\|_F^2 \leq (1 + \epsilon)\|\mathbf{A} - \mathbf{AV}_k\mathbf{V}_k^T\|_F^2.$$

Runtime: $O(\text{nnz}(\mathbf{X}) \cdot k \cdot T) \leq O(ndk \cdot T)$.

APPLICATION OF PARTIAL SVD: ENTITY EMBEDDINGS

Corpus of Documents

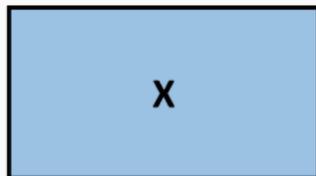


Term Document Matrix X

	car	loan	house	...	dog	cat			
doc_1	0	0	1	0	0	1	1	0	0
doc_2	0	0	0	1	0	1	0	0	0
...	1	1	0	1	0	0	0	1	0
doc_n	1	0	0	0	0	0	0	1	1



Low-Rank Approximation via SVD



\approx



EXAMPLE: LATENT SEMANTIC ANALYSIS

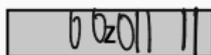
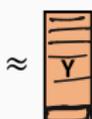
Corpus of Documents



Term Document Matrix X

	dog	tree	house	...	dog	cat
doc_1	0	0	1	0	0	1
doc_2	0	0	1	0	1	0
...	1	1	0	1	0	0
...	0	0	0	0	0	1
doc_n	1	0	0	0	0	1

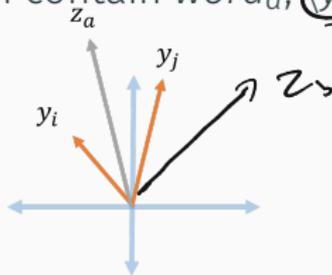
Low-Rank Approximation via SVD



≈

y_i row of Y

- $\langle \vec{y}_i, \vec{z}_a \rangle \approx 1$ when doc_i contains word a .
- If doc_i and doc_j both contain word a , $\langle \vec{y}_i, \vec{z}_a \rangle \approx \langle \vec{y}_j, \vec{z}_a \rangle = 1$.



EXAMPLE: LATENT SEMANTIC ANALYSIS

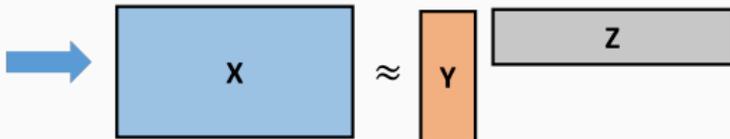
XX^T

Term Document Matrix X

	car	loan	house	...	dog	cat
doc_1	0	0	1	0	0	1
doc_2	0	0	0	1	0	1
⋮	1	1	0	1	0	0
⋮	0	0	0	0	0	1
doc_n	1	0	0	0	0	1

X^T  X 

Low-Rank Approximation via SVD



- The columns $\vec{z}_1, \vec{z}_2, \dots$ give representations of words, with \vec{z}_i and \vec{z}_j tending to have high dot product if $word_i$ and $word_j$ appear in many of the same documents.
- Z corresponds to the top k right singular vectors: the eigenvectors of XX^T . Intuitively, what is XX^T ?

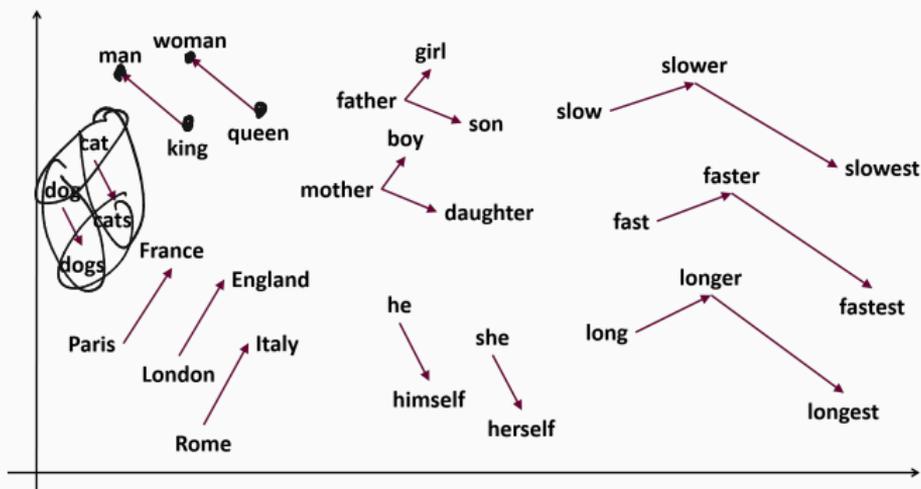
• $(XX^T)_{i,j} =$  $= \langle X_i, X_j \rangle$
for i, j column of X

EXAMPLE: WORD EMBEDDINGS

Not obvious how to convert a word into a feature vector that captures the meaning of that word. Approach suggested by LSA: build a $d \times d$ symmetric “similarity matrix” \mathbf{M} between words, and factorize: $\mathbf{M} \approx \mathbf{F}\mathbf{F}^T$ for rank k \mathbf{F} .

- **Similarity measures:** How often do $word_i, word_j$ appear in the same sentence, in the same window of w words, in similar positions of documents in different languages?
- Replacing $\mathbf{X}\mathbf{X}^T$ with these different metrics (sometimes appropriately transformed) leads to popular word embedding algorithms: **word2vec**, **GloVe**, etc.

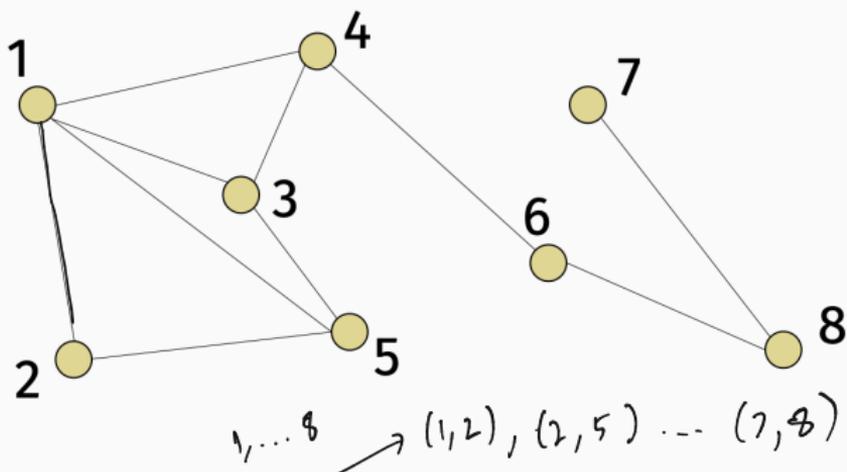
EXAMPLE: WORD EMBEDDINGS



word2vec was originally described as a neural-network method, but Levy and Goldberg show that it is simply low-rank approximation of a specific similarity matrix. (*Neural word embedding as implicit matrix factorization.*)

SPECTRAL GRAPH THEORY

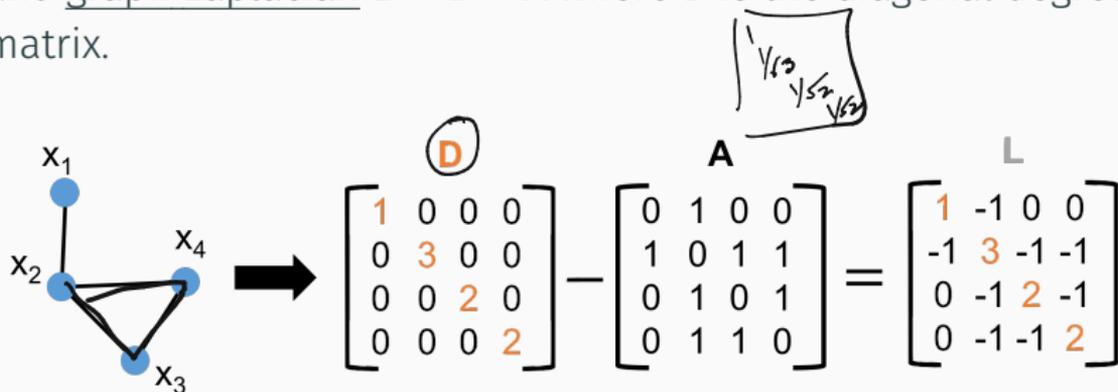
Main idea: Understand graph data by constructing natural matrix representations, and studying that matrix's spectrum (eigenvalues/eigenvectors).



For now assume $\underline{G} = (\underline{V}, \underline{E})$ is an undirected, unweighted graph with n nodes.

MATRIX REPRESENTATIONS OF GRAPHS

Two most common representations: $n \times n$ adjacency matrix A and graph Laplacian $L = D - A$ where D is the diagonal degree matrix.

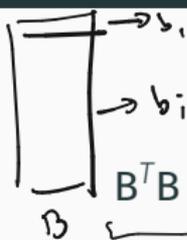


Also common to look at normalized versions of both of these: $\bar{A} = D^{-1/2}AD^{-1/2}$ and $\bar{L} = I - D^{-1/2}AD^{-1/2}$.



- If \mathbf{L} have k eigenvalues equal to 0, then G has k connected components.
- Sum of cubes of \mathbf{A} 's eigenvalues is equal to number of triangles in the graph.
- Sum of eigenvalues to the power q is the number of q cycles.

THE LAPLACIAN VIEW



PSD

$$x^T B^T B x = x^T L x = \|Bx\|_2^2$$

$$B^T B = b_1 b_1^T + b_2 b_2^T + \dots + b_m b_m^T$$

where b_i is the i^{th} row of B (each row corresponds to a single edge).

$$\begin{bmatrix} 1 \\ -1 \end{bmatrix} \begin{bmatrix} 1 & -1 \end{bmatrix} = \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}$$

$$\begin{bmatrix} 1 \\ -1 \end{bmatrix} \begin{bmatrix} 1 & -1 \end{bmatrix} = \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}$$

$$\begin{bmatrix} 1 \\ -1 \end{bmatrix} \begin{bmatrix} -1 & 1 \end{bmatrix} = -1 \cdot -1 + 1 \cdot 1 = 2$$

$$= b_i b_i^T$$

Conclusions from $L = B^T B$

- L is positive semidefinite: $x^T L x \geq 0$ for all x .
- $L = \underline{V \Sigma^2 V^T}$ where $U \Sigma V^T$ is B 's SVD. Columns of V are eigenvectors of L .

- For any vector $x \in \mathbb{R}^n$,

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

"smoothness of x "

$= \|Bx\|_2^2$

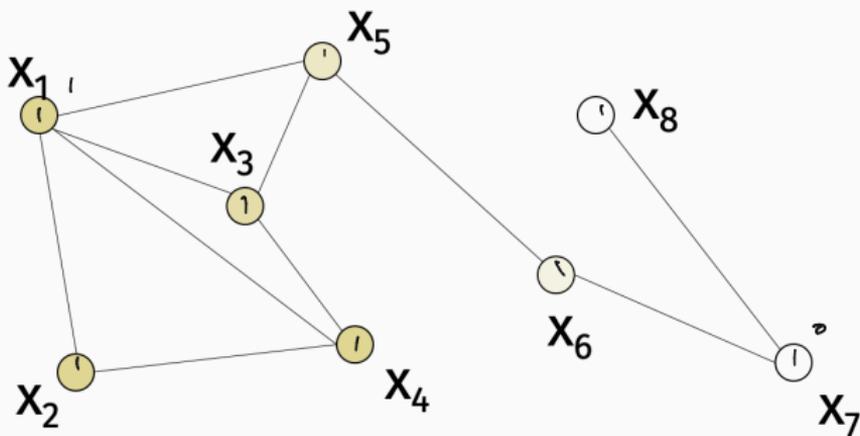
$x \in \mathbb{R}^n$

$x^T L x$

THE LAPLACIAN VIEW

$x^T L x$ is small if x is a “smooth” function with respect to the graph.

$$x^T L x = 0 \quad \text{if} \quad x = \underline{\underline{\vec{1}}}$$



Eigenvectors of the Laplacian with small eigenvalues correspond to smooth functions over the graph.

Courant–Fischer min-max principle

Let $\mathbf{V} = [\mathbf{v}_1, \dots, \mathbf{v}_n]$ be the eigenvectors of \mathbf{L} .

$$\mathbf{v}_n = \arg \min_{\|\mathbf{v}\|=1} \mathbf{v}^T \mathbf{L} \mathbf{v}$$

$$\mathbf{v}_{n-1} = \arg \min_{\|\mathbf{v}\|=1, \mathbf{v} \perp \mathbf{v}_n} \mathbf{v}^T \mathbf{L} \mathbf{v}$$

$$\mathbf{v}_{n-2} = \arg \min_{\|\mathbf{v}\|=1, \mathbf{v} \perp \mathbf{v}_n, \mathbf{v}_{n-1}} \mathbf{v}^T \mathbf{L} \mathbf{v}$$

$$\vdots$$

$$\mathbf{v}_1 = \arg \min_{\|\mathbf{v}\|=1, \mathbf{v} \perp \mathbf{v}_n, \dots, \mathbf{v}_2} \mathbf{v}^T \mathbf{L} \mathbf{v}$$

Courant–Fischer min-max principle

Let $\mathbf{V} = [\mathbf{v}_1, \dots, \mathbf{v}_n]$ be the eigenvectors of \mathbf{L} .

$$\underline{\mathbf{v}}_1 = \arg \max_{\|\mathbf{v}\|=1} \underline{\mathbf{v}}^T \underline{\mathbf{L}} \underline{\mathbf{v}}$$

$$\mathbf{v}_2 = \arg \max_{\|\mathbf{v}\|=1, \mathbf{v} \perp \mathbf{v}_1} \mathbf{v}^T \mathbf{L} \mathbf{v}$$

$$\mathbf{v}_3 = \arg \max_{\|\mathbf{v}\|=1, \mathbf{v} \perp \mathbf{v}_1, \mathbf{v}_2} \mathbf{v}^T \mathbf{L} \mathbf{v}$$

⋮

$$\mathbf{v}_n = \arg \max_{\|\mathbf{v}\|=1, \mathbf{v} \perp \mathbf{v}_1, \dots, \mathbf{v}_{n-1}} \mathbf{v}^T \mathbf{L} \mathbf{v}$$

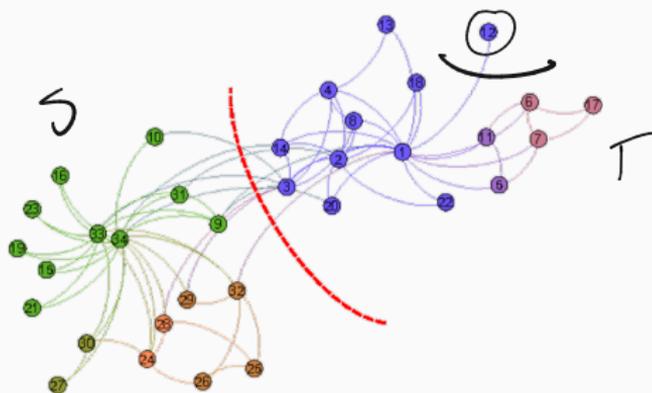
EXAMPLE APPLICATION OF SPECTRAL GRAPH THEORY

- Study graph partitioning problem important in 1) understanding social networks 2) nonlinear clustering in unsupervised machine learning (spectral clustering).
- See how this problem can be solved approximately using Laplacian eigenvectors.
- Give a full analysis of the method for a common random graph model.
- Use two tools: matrix concentration and eigenvector perturbation bounds.

BALANCED CUT

Common goal: Given a graph $G = (V, E)$, partition nodes along a cut that:

- Has few crossing edges: $|\{(u, v) \in E : u \in S, v \in T\}|$ is small.
- Separates large partitions: $|S|, |T|$ are not too small.



(a) Zachary Karate Club Graph

Important in understanding community structure in social networks.

Wayne W. Zachary (1977). An Information Flow Model for Conflict and Fission in Small Groups.

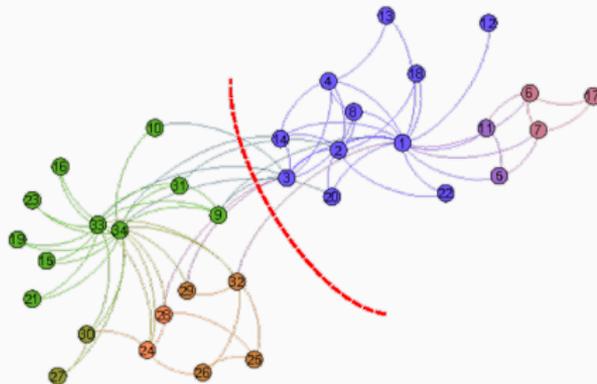
“The network captures 34 members of a karate club, documenting links between pairs of members who interacted outside the club. During the study a conflict arose between the administrator “John A” and instructor “Mr. Hi” (pseudonyms), which led to the split of the club into two. Half of the members formed a new club around Mr. Hi; members from the other part found a new instructor or gave up karate. Based on collected data Zachary correctly assigned all but one member of the club to the groups they actually joined after the split.” – Wikipedia

Beautiful paper – definitely worth checking out!

BALANCED CUT

Common goal: Given a graph $G = (V, E)$, partition nodes along a cut that:

- Has few crossing edges: $|\{(u, v) \in E : u \in S, v \in T\}|$ is small.
- Separates large partitions: $|S|, |T|$ are not too small.

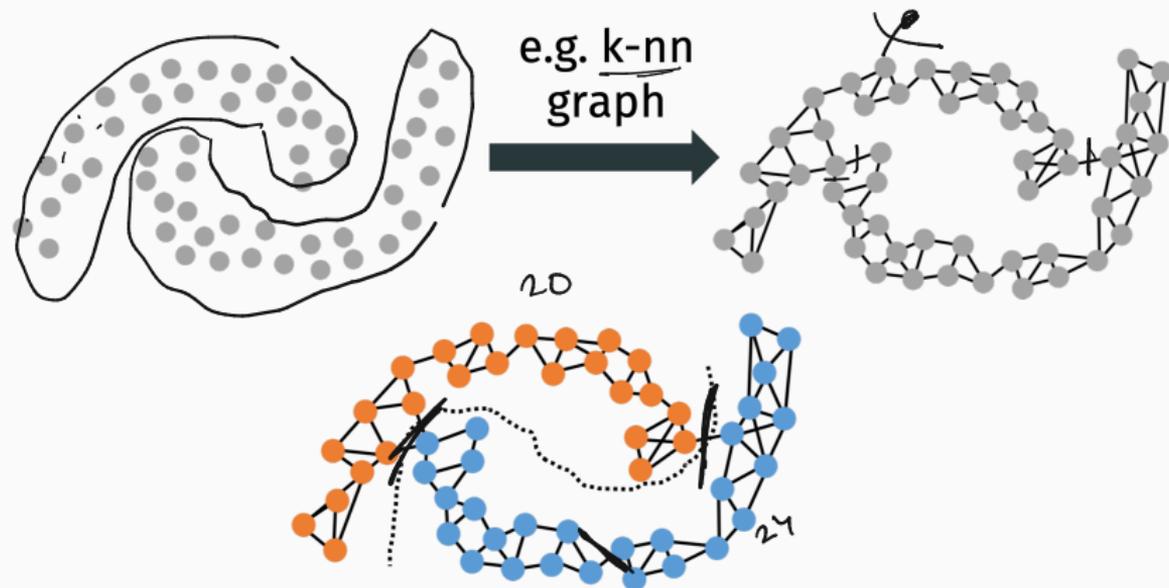


(a) Zachary Karate Club Graph

Important in understanding community structure in social networks.

SPECTRAL CLUSTERING

Idea: Construct synthetic graph for data that is hard to cluster.



Spectral Clustering, Laplacian Eigenmaps, Locally linear embedding, Isomap, etc.

There are many ways to formalize Zachary's problem:

Sparsest Cut:

$$\min_S \frac{[\text{cut}(S, V \setminus S)]}{\min(|S|, |V \setminus S|)} \rightarrow 2D$$



β -Balanced Cut:

$$|S|$$

$$\beta \leq 1/2$$

$$\min_S \text{cut}(S, V \setminus S) \quad \text{such that} \quad \min(|S|, |V \setminus S|) \geq \beta n$$

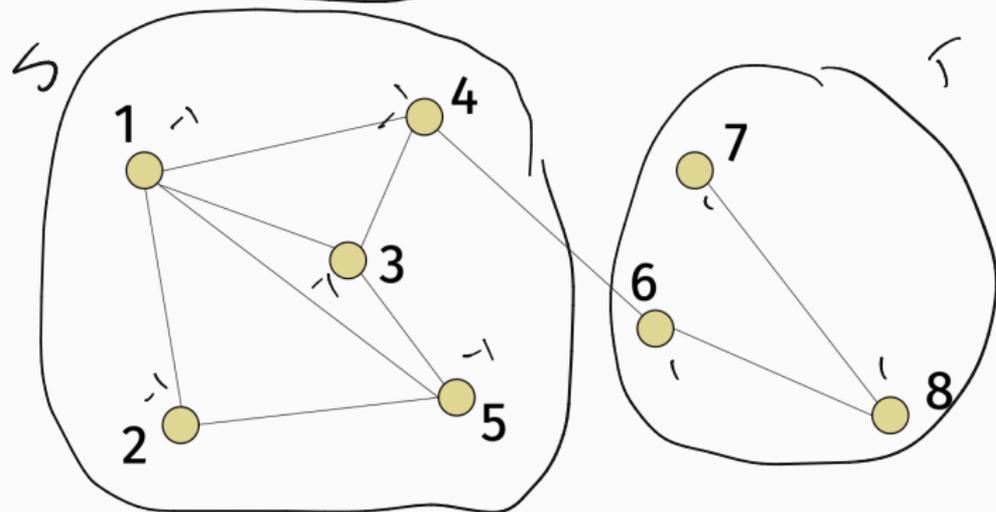
Most formalizations lead to computationally hard problems. Lots of interest in designing polynomial time approximation algorithms, but tend to be slow. In practice, much simpler methods based on the graph spectrum are used.

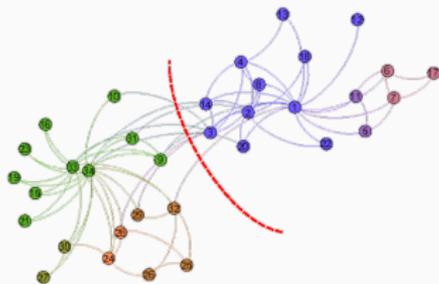
THE LAPLACIAN VIEW

Another conclusion from $L = B^T B$:

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

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





(a) Zachary Karate Club Graph

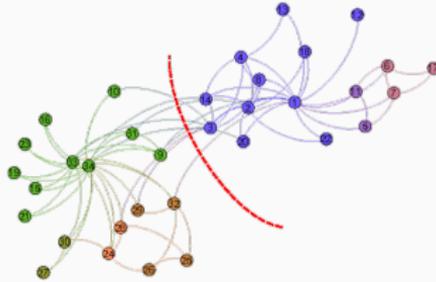
For a cut indicator vector $\mathbf{c} \in \{-1, 1\}^n$ with $\mathbf{c}(i) = -1$ for $i \in S$ and $\mathbf{c}(i) = 1$ for $i \in T$:

- $\mathbf{c}^T \mathbf{L} \mathbf{c} = 4 \cdot \text{cut}(S, T)$.

- $|\mathbf{c}^T \mathbf{1}| = \left| |T| - |S| \right|$

$$\mathbf{c}^T \mathbf{1} = \sum_{i=1}^n c_i$$

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



(a) Zachary Karate Club Graph

Equivalent formulation if we divide everything by \sqrt{n} so that \mathbf{c} has norm 1. Then $\mathbf{c} \in \left\{ -\frac{1}{\sqrt{n}} \frac{1}{\sqrt{n}} \right\}^n$ and:

$$\bullet \mathbf{c}^T \mathbf{L} \mathbf{c} = \frac{4}{n} \cdot \text{cut}(S, T).$$

$$\|\mathbf{c}\|_2 = 1$$

$$\bullet \mathbf{c}^T \mathbf{1} = \frac{1}{\sqrt{n}} (|T| - |S|).$$

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

SMALLEST LAPLACIAN EIGENVECTOR

The smallest eigenvector/singular vector \mathbf{v}_n satisfies:

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

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

SECOND SMALLEST LAPLACIAN EIGENVECTOR

By Courant-Fischer, \mathbf{v}_{n-1} is given by:

$$\underline{\mathbf{v}_{n-1}} = \underset{\|\mathbf{v}\|=1, \underline{\mathbf{v}_n^T \mathbf{v}}=0}{\operatorname{arg\,min}} \mathbf{v}^T \mathbf{L} \mathbf{v}$$

If \mathbf{v}_{n-1} were binary $\{-\frac{1}{\sqrt{n}}, \frac{1}{\sqrt{n}}\}^n$ it would have:

- $\mathbf{v}_{n-1}^T \mathbf{L} \mathbf{v}_{n-1} = \frac{4}{n} \operatorname{cut}(S, T)$ as small as possible **given that**
 $\mathbf{v}_{n-1}^T \mathbf{1} = |T| - |S| = \mathbf{0}$.
- \mathbf{v}_{n-1} would indicate the smallest perfectly balanced cut.

$\mathbf{v}_{n-1} \in \mathbb{R}^n$ is not generally binary, but a natural approach is to 'round' the vector to obtain a cut.

CUTTING WITH THE SECOND LAPLACIAN EIGENVECTOR

Find a good partition of the graph by computing

$$V \Lambda V^T$$

$$v_{n-1} = \operatorname{argmin}_{v \in \mathbb{R}^n \text{ with } \|v\|=1, v^T \mathbf{1}=0} \underbrace{v^T L v}_{v_n}$$

$$v_{n-1}^T v_n = 0$$

Set S to be all nodes with $v_{n-1}(i) < 0$, and T to be all with

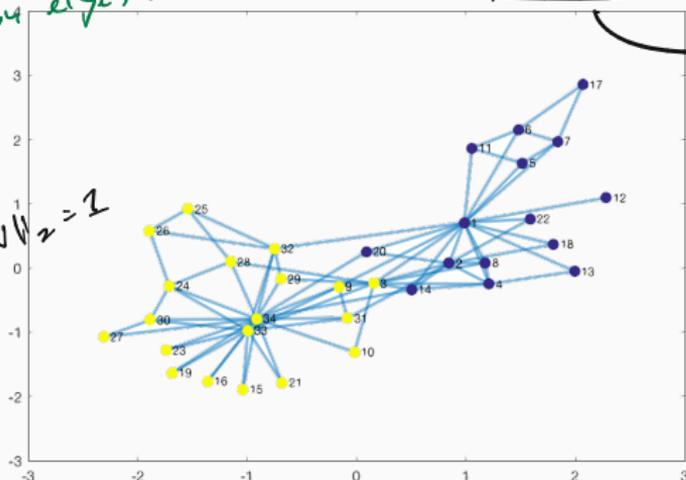
$$v_{n-1}(i) \geq 0.$$

$$v^T L v = 0$$

Assume v is an eigenvector:

$$L v = \lambda v$$

$$v^T L v = \lambda v^T v \rightarrow \|v\|_2^2 = 1$$



$$v^T B^T B v = \|B v\|_2^2 \geq 0$$

L is PSD

$$x = \mathbf{1}$$

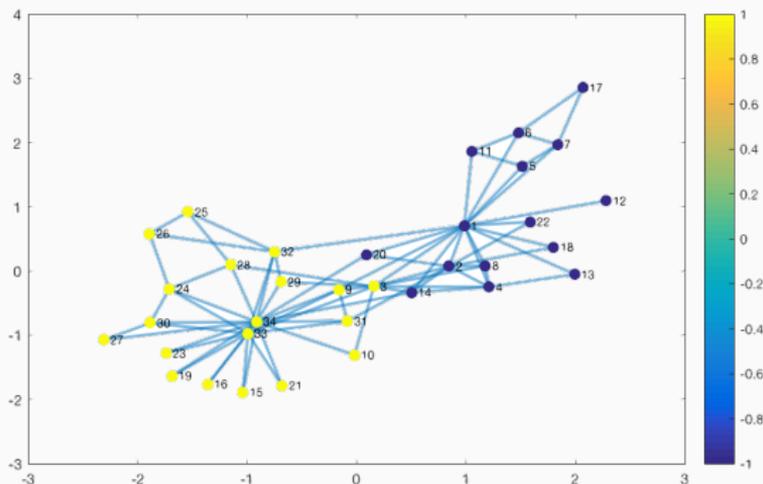
$$x^T L x = x^T B^T B x = \sum_{(i,j) \in E} (x(i) - x(j))^2$$

CUTTING WITH THE SECOND LAPLACIAN EIGENVECTOR

Find a good partition of the graph by computing

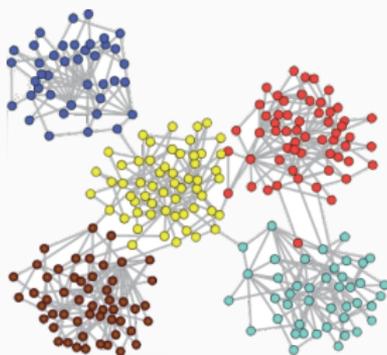
$$\mathbf{v}_{n-1} = \underset{\mathbf{v} \in \mathbb{R}^n \text{ with } \|\mathbf{v}\|=1, \mathbf{v}^T \mathbf{1}=0}{\operatorname{argmin}} \quad \mathbf{v}^T \mathbf{L} \mathbf{v}$$

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

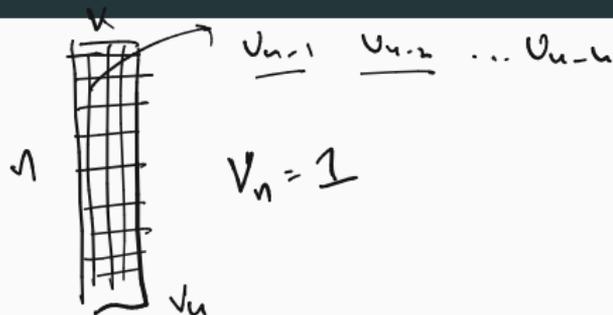


The Shi-Malik normalized cuts algorithm is one of the most commonly used variants of this approach, using the normalized Laplacian $\bar{\mathbf{L}} = \mathbf{D}^{-1/2} \mathbf{L} \mathbf{D}^{-1/2}$.

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



SPECTRAL PARTITIONING IN PRACTICE



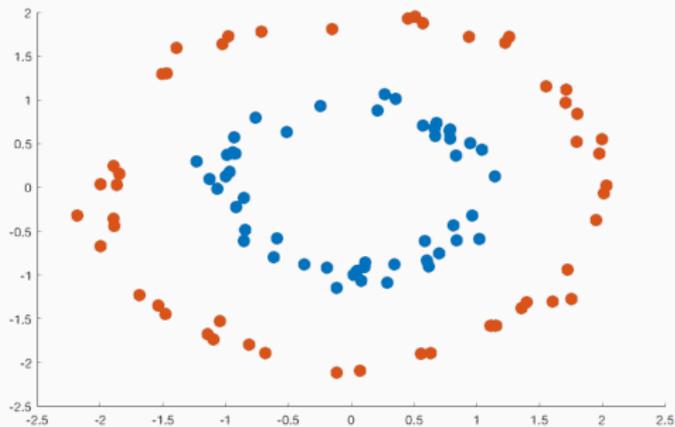
Spectral Clustering:

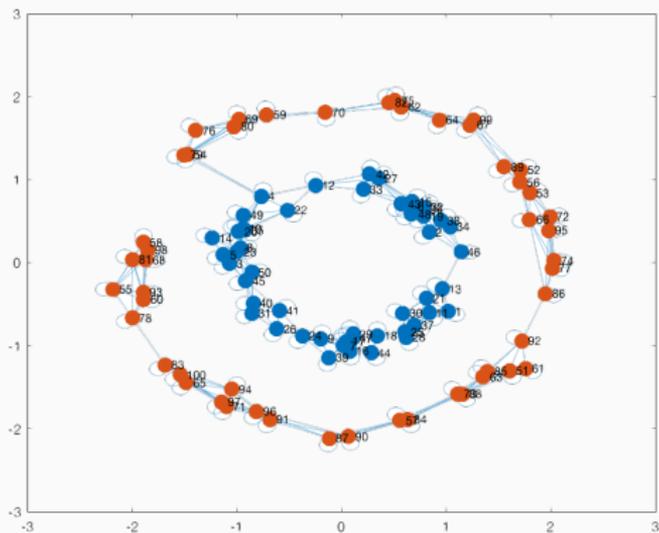
- Compute smallest k eigenvectors $\mathbf{v}_{n-1}, \dots, \mathbf{v}_{n-k}$ of \mathbf{L} .
- Represent each node by its corresponding row in $\mathbf{V} \in \mathbb{R}^{n \times k}$ whose rows are $\mathbf{v}_{n-1}, \dots, \mathbf{v}_{n-k}$.
- Cluster these rows using k -means clustering (or really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

$$\min_{\mathbf{V}} \mathbf{V}^T \mathbf{L} \mathbf{V} \quad \hat{\mathbf{1}}^T \mathbf{L} \hat{\mathbf{1}} = \sum_{(i,j) \in E} (1_i - 1_j)^2 = 0$$

$$\mathbf{V}^T \mathbf{L} \mathbf{V} = \mathbf{0}$$

Original Data: (not linearly separable)

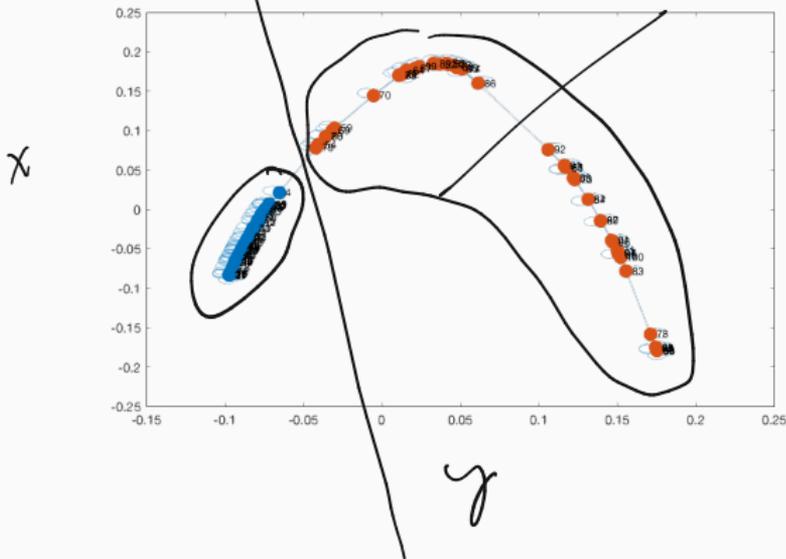


k -Nearest Neighbors Graph:

LAPLACIAN EMBEDDING

v_{n-1}

Embedding with eigenvectors v_{n-1} , v_{n-2} : (linearly separable)

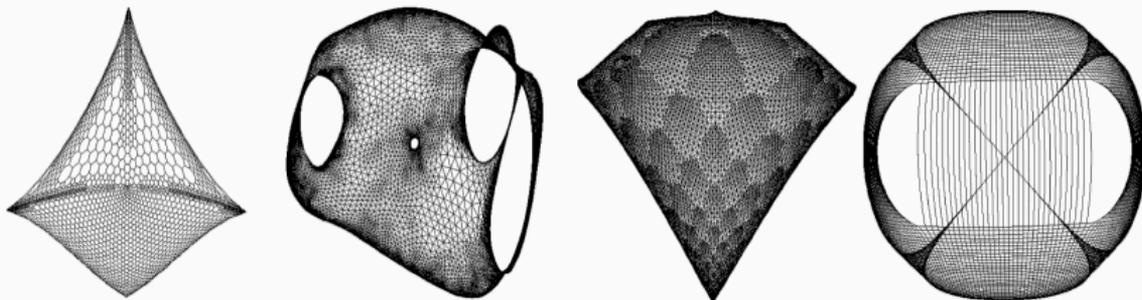


WHY DOES THIS WORK?

Intuitively, since $\mathbf{v} \in \mathbf{v}_1, \dots, \mathbf{v}_\beta$ are smooth over the graph,

$$\sum_{i,j \in E} (\mathbf{v}[i] - \mathbf{v}[j])^2$$

is small for each coordinate. I.e. this embedding explicitly encourages nodes connected by an edge to be placed in nearby locations in the embedding.



Also useful e.g., in graph drawing.

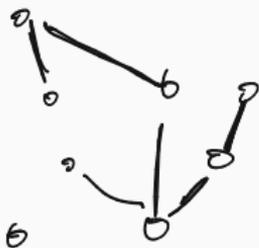
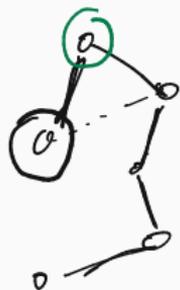
So far: Showed that spectral clustering partitions a graph along a small cut between large pieces.

- No formal guarantee on the ‘quality’ of the partitioning.
- Would be difficult to analyze for general input graphs.

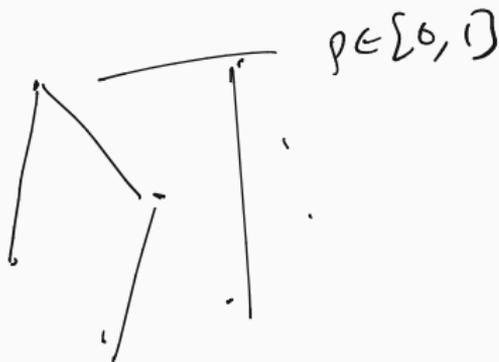
Common approach: Design a natural **generative model** that produces random but realistic inputs and analyze how the algorithm performs on inputs drawn from this model.

- Very common in algorithm design and analysis. Great way to start approaching a problem.
- This is also the whole idea behind Bayesian Machine Learning (can be used to justify ℓ_2 linear regression, k -means clustering, PCA, etc.)

Ideas for a generative model for **social network graphs** that would allow us to understand partitioning?



Erdős - Rényi

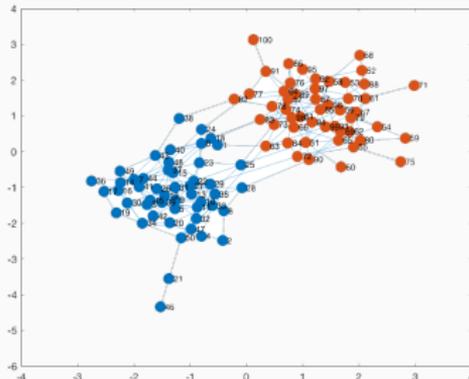


STOCHASTIC BLOCK MODEL

Stochastic Block Model (Planted Partition Model):

Let $G_n(p, q)$ be a distribution over graphs on n nodes, split equally into two groups B and C , each with $n/2$ nodes.

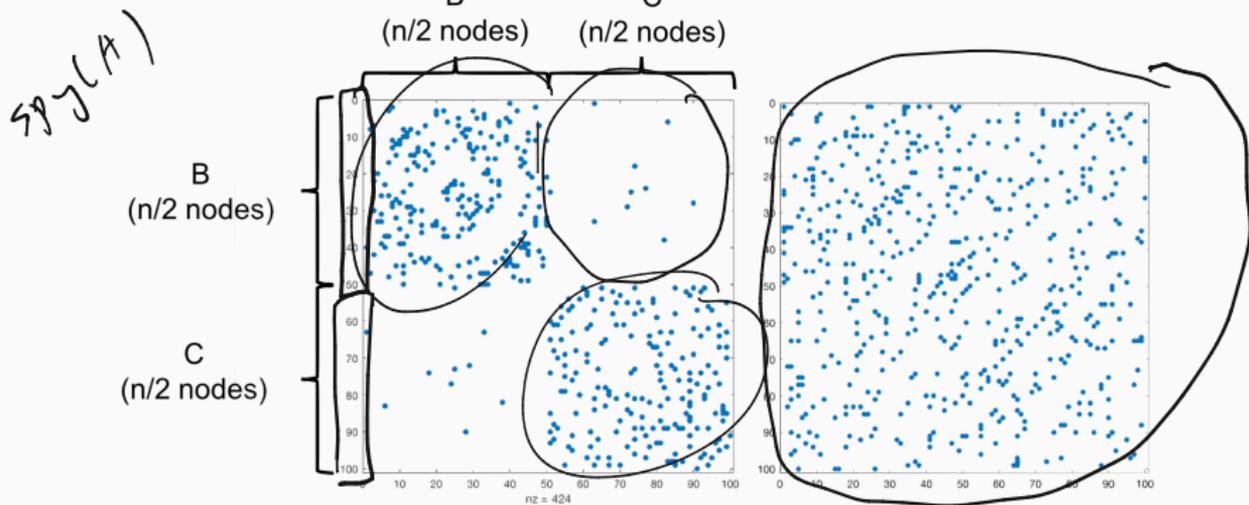
- Any two nodes in the **same group** are connected with probability p (including self-loops).
- Any two nodes in **different groups** are connected with prob $q < p$.



LINEAR ALGEBRAIC VIEW

Let G be a stochastic block model graph drawn from $G_n(p, q)$.

- Let $A \in \mathbb{R}^{n \times n}$ be the adjacency matrix of G . What is $\mathbb{E}[A]$?

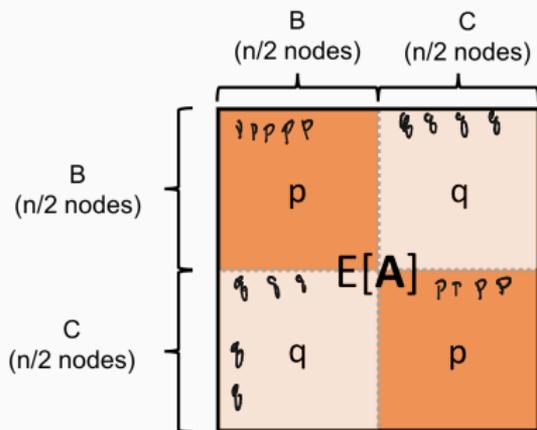


Note that we are arbitrarily ordering the nodes in A by group. In reality A would look "scrambled" as on the right.

EXPECTED ADJACENCY SPECTRUM

$$L = D - A \quad A = D - L \quad \mathbb{E}[A] = \mathbb{E}[D] - \mathbb{E}[L]$$

Letting G be a stochastic block model graph drawn from $G_n(p, q)$ and $A \in \mathbb{R}^{n \times n}$ be its adjacency matrix. $(\mathbb{E}[A])_{i,j} = p$ for i, j in same group, $(\mathbb{E}[A])_{i,j} = q$ otherwise.



What are the eigenvectors and eigenvalues of $\mathbb{E}[A]$?

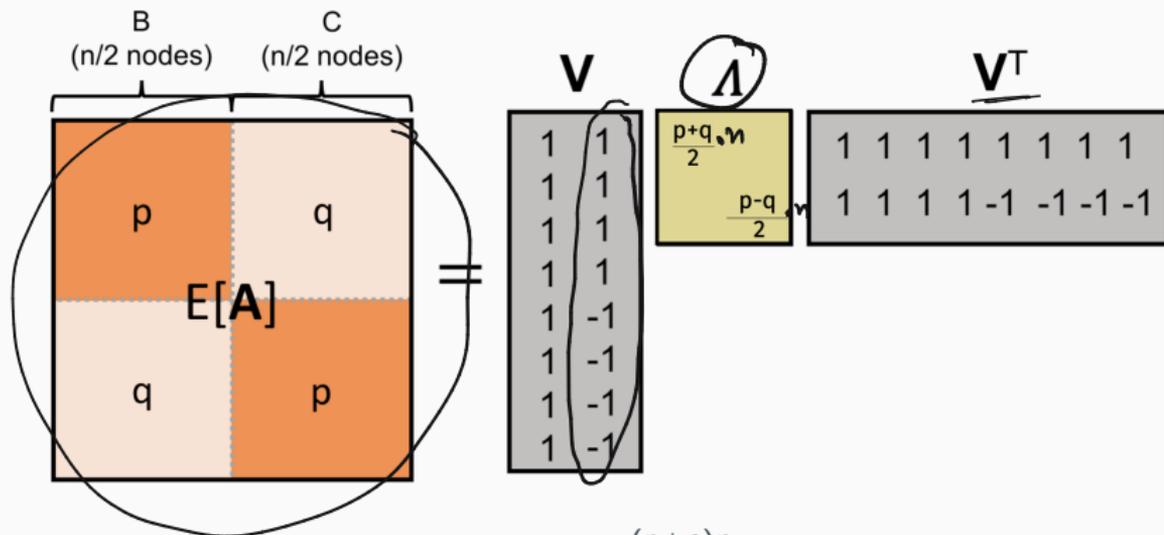
$$\left[\begin{array}{c} (\frac{p+q}{2}) \mathbf{1} \\ \vdots \end{array} \right]$$

What is the expected Laplacian $G_n(p, q)$?

A and **L** have the same eigenvectors and eigenvalue are equal up to a shift.

Letting G be a stochastic block model graph drawn from $G_n(p, q)$ and $\mathbf{A} \in \mathbb{R}^{n \times n}$ be its adjacency matrix, what are the eigenvectors and eigenvalues of $\mathbb{E}[\mathbf{A}]$?

EXPECTED ADJACENCY SPECTRUM



- $\mathbf{v}_1 \sim \mathbf{1}$ with eigenvalue $\lambda_1 = \frac{(p+q)n}{2}$.
- $\mathbf{v}_2 \sim \chi_{B,C}$ with eigenvalue $\lambda_2 = \frac{(p-q)n}{2}$.
- $\chi_{B,C}(i) = 1$ if $i \in B$ and $\chi_{B,C}(i) = -1$ for $i \in C$.

If we compute \mathbf{v}_2 then we recover the communities B and C!

Upshot: The second smallest eigenvector of $\mathbb{E}[\mathbf{L}]$, equivalently the second largest of $\mathbb{E}[\mathbf{A}]$, is $\chi_{B,C}$ – the indicator vector for the cut between the communities.

- If the random graph G (equivalently \mathbf{A} and \mathbf{L}) were exactly equal to its expectation, partitioning using this eigenvector would exactly recover communities B and C .

How do we show that a matrix (e.g., \mathbf{A}) is close to its expectation? **Matrix concentration inequalities.**

- Analogous to scalar concentration inequalities like Markovs, Chebyshevs, Bernsteins.

Matrix Concentration Inequality: If $p \geq O\left(\frac{\log^4 n}{n}\right)$, then with high probability

$$\|A - \mathbb{E}[A]\|_2 \leq O(\sqrt{pn}).$$

where $\|\cdot\|_2$ is the matrix **spectral** norm (operator norm).

For $X \in \mathbb{R}^{n \times d}$, $\|X\|_2 = \max_{z \in \mathbb{R}^d: \|z\|_2=1} \|Xz\|_2$.

Exercise: Show that $\|X\|_2$ is equal to the largest singular value of X . For symmetric X (like $A - \mathbb{E}[A]$) show that it is equal to the magnitude of the largest magnitude eigenvalue.

For the stochastic block model application, we want to show that the second eigenvectors of A and $\mathbb{E}[A]$ are close. How does this relate to their difference in spectral norm?

EIGENVECTOR PERTURBATION

Davis-Kahan Eigenvector Perturbation Theorem: Suppose $\mathbf{A}, \bar{\mathbf{A}} \in \mathbb{R}^{d \times d}$ are symmetric with $\|\mathbf{A} - \bar{\mathbf{A}}\|_2 \leq \epsilon$ and eigenvectors v_1, v_2, \dots, v_d and $\bar{v}_1, \bar{v}_2, \dots, \bar{v}_d$. Letting $\theta(v_i, \bar{v}_i)$ denote the angle between v_i and \bar{v}_i , for all i :

$$\sin[\theta(v_i, \bar{v}_i)] \leq \frac{\epsilon}{\min_{j \neq i} |\lambda_i - \lambda_j|}$$

eigenvalues of A.

where $\lambda_1, \dots, \lambda_d$ are the eigenvalues of $\bar{\mathbf{A}}$.

The error gets larger if there are eigenvalues with similar magnitudes.

EIGENVECTOR PERTURBATION

$$\begin{array}{ccc} \mathbf{A} & \mathbf{\bar{A}} & \mathbf{A} - \mathbf{\bar{A}} \\ \begin{bmatrix} 1+\varepsilon & 0 \\ 0 & 1 \end{bmatrix} & \begin{bmatrix} 1 & 0 \\ 0 & 1+\varepsilon \end{bmatrix} & = \begin{bmatrix} \varepsilon & 0 \\ 0 & \varepsilon \end{bmatrix} \end{array}$$

$\|\mathbf{A} - \mathbf{\bar{A}}\| = \varepsilon$

APPLICATION TO STOCHASTIC BLOCK MODEL

Claim 1 (Matrix Concentration): For $p \geq O\left(\frac{\log^4 n}{n}\right)$,

$$\|\mathbf{A} - \mathbb{E}[\mathbf{A}]\|_2 \leq O(\sqrt{pn}).$$

Claim 2 (Davis-Kahan): For $p \geq O\left(\frac{\log^4 n}{n}\right)$,

$$\sin \theta(v_2, \bar{v}_2) \leq \frac{O(\sqrt{pn})}{\min_{j \neq i} |\lambda_i - \lambda_j|} \leq \frac{O(\sqrt{pn})}{(p-q)n/2} = O\left(\frac{\sqrt{p}}{(p-q)\sqrt{n}}\right)$$

Recall: $\mathbb{E}[\mathbf{A}]$, has eigenvalues $\lambda_1 = \frac{(p+q)n}{2}$, $\lambda_2 = \frac{(p-q)n}{2}$, $\lambda_i = 0$ for $i \geq 3$.

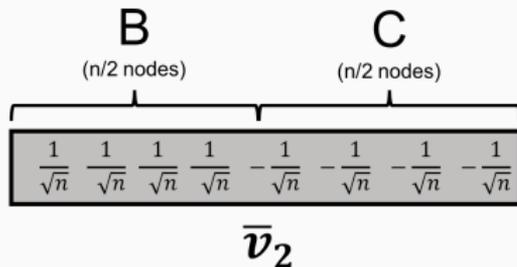
$$\min_{j \neq i} |\lambda_i - \lambda_j| = \min\left(qn, \frac{(p-q)n}{2}\right).$$

Assume $\left|\frac{(p-q)n}{2} - 0\right|$ will be the minimum of the two gaps. I.e.
smaller than $\left|\frac{(p+q)n}{2} - \frac{(p-q)n}{2}\right| = qn$.

APPLICATION TO STOCHASTIC BLOCK MODEL

So Far: $\sin \theta(v_2, \bar{v}_2) \leq O\left(\frac{\sqrt{p}}{(\rho-q)\sqrt{n}}\right)$. What does this give us?

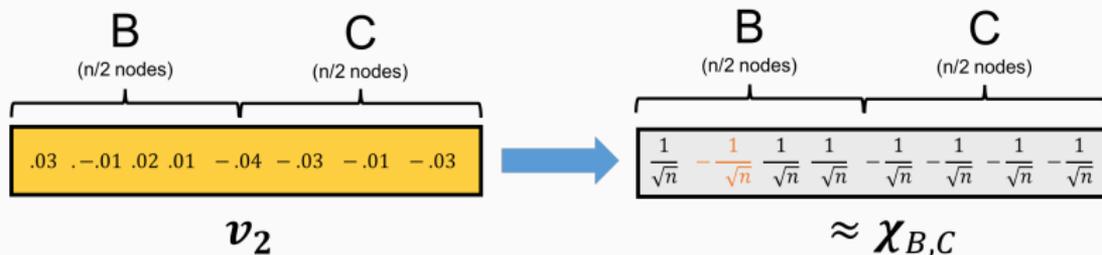
- Can show that this implies $\|v_2 - \bar{v}_2\|_2^2 \leq O\left(\frac{p}{(\rho-q)^2 n}\right)$ (exercise).
- \bar{v}_2 is $\frac{1}{\sqrt{n}}\chi_{B,C}$: the community indicator vector.



- Every i where $v_2(i), \bar{v}_2(i)$ differ in sign contributes $\geq \frac{1}{n}$ to $\|v_2 - \bar{v}_2\|_2^2$.
- So they differ in sign in at most $O\left(\frac{p}{(\rho-q)^2}\right)$ positions.

APPLICATION TO STOCHASTIC BLOCK MODEL

Upshot: If G is a stochastic block model graph with adjacency matrix A , if we compute its second large eigenvector v_2 and assign nodes to communities according to the sign pattern of this vector, we will correctly assign all but $O\left(\frac{p}{(p-q)^2}\right)$ nodes.



- Why does the error increase as q gets close to p ?
- Even when $p - q = O(1/\sqrt{n})$, assign all but an $O(n)$ fraction of nodes correctly. E.g., assign 99% of nodes correctly.