

COMPSCI 514: Algorithms for Data Science

Cameron Musco

University of Massachusetts Amherst. Spring 2026.

Lecture 18

- Midterm in class next Thursday 4/23.
- See Piazza/last lecture/midterm study guide for details.
- Will cover material up to this lecture.
- The quiz this week is **optional**. If you take it, it will replace your second lowest regular quiz grade (the lowest is dropped).

Summary

Last Few Classes: Spectral Graph Partitioning

- Focus on separating graphs with small but relatively balanced cuts.
- Connection to second smallest eigenvector of graph Laplacian.
- Provable guarantees for stochastic block model.
- Expectation analysis in class. Quick sketch of full analysis.

$$V^T L V$$

$$V_{n-1}$$

This Class: Computing the SVD/eigendecomposition.

- Efficient algorithms for SVD/eigendecomposition.
- Iterative methods: (power method,) Krylov subspace methods.
- High level: a glimpse into fast methods for linear algebraic computation, which are workhorses behind modern data science/ML.

Efficient Eigendecomposition and SVD

We have talked about the eigendecomposition and SVD as ways to compress data, to embed entities like words and documents, to compress/cluster non-linearly separable data and graphs.

How efficient are these techniques? Can they be run on very large datasets?

Computing the SVD

Basic Algorithm: To compute the SVD of full-rank $X \in \mathbb{R}^{n \times d}$,
 $X = U \Sigma V^T$

- Compute $X^T X - O(nd^2)$ runtime.
- Find eigendecomposition $X^T X = V \Lambda V^T - O(d^3)$ runtime.
- Compute $L = X V - O(nd^2)$ runtime. Note that $L = U \Sigma$.
- Set $\sigma_i = \|L_i\|_2$ and $U_i = L_i / \|L_i\|_2$. - $O(nd)$ runtime.

Total runtime: $O(nd^2 + d^3) = O(nd^2)$ (assume w.l.o.g. $n \geq d$)

$n \geq d$

$$\begin{matrix} n & & d \\ \left[X^T \right] & \left[X \right] & = & \left[\right] \\ d & & d \end{matrix}$$

$$\begin{aligned} X V &= U \Sigma V^T V \\ &= U \Sigma \end{aligned}$$

SVD

Computing the SVD

Basic Algorithm: To compute the SVD of full-rank $\mathbf{X} \in \mathbb{R}^{n \times d}$,
 $\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$:

- Compute $\mathbf{X}^T\mathbf{X} - O(nd^2)$ runtime.
- Find eigendecomposition $\mathbf{X}^T\mathbf{X} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^T - O(d^3)$ runtime.
- Compute $\mathbf{L} = \mathbf{X}\mathbf{V} - O(nd^2)$ runtime. Note that $\mathbf{L} = \mathbf{U}\mathbf{\Sigma}$.
- Set $\sigma_i = \|\mathbf{L}_i\|_2$ and $\mathbf{U}_i = \mathbf{L}_i/\|\mathbf{L}_i\|_2$. - $O(nd)$ runtime.

Total runtime: $O(nd^2 + d^3) = O(nd^2)$ (assume w.l.o.g. $n \geq d$)

- If we have $n = 10$ million images with $200 \times 200 \times 3 = 120,000$ pixel values each, runtime is 1.5×10^{17} operations! = \downarrow

Computing the SVD

Basic Algorithm: To compute the SVD of full-rank $\mathbf{X} \in \mathbb{R}^{n \times d}$,
 $\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$:

- Compute $\mathbf{X}^T\mathbf{X} - O(nd^2)$ runtime.
- Find eigendecomposition $\mathbf{X}^T\mathbf{X} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^T - O(d^3)$ runtime.
- Compute $\mathbf{L} = \mathbf{X}\mathbf{V} - O(nd^2)$ runtime. Note that $\mathbf{L} = \mathbf{U}\mathbf{\Sigma}$.
- Set $\sigma_i = \|\mathbf{L}_i\|_2$ and $\mathbf{U}_i = \mathbf{L}_i/\|\mathbf{L}_i\|_2$. - $O(nd)$ runtime.

Total runtime: $O(nd^2 + d^3) = O(nd^2)$ (assume w.l.o.g. $n \geq d$)

- If we have $n = 10$ million images with $200 \times 200 \times 3 = 120,000$ pixel values each, runtime is 1.5×10^{17} operations!
- The worlds fastest super computers compute at ≈ 100 petaFLOPS = 10^{17} FLOPS (floating point operations per second).
- This is a relatively easy task for them – but no one else.

Faster Algorithms

To speed up SVD computation we will take advantage of the fact that we typically only care about computing the **top (or bottom) k singular vectors** of a matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$ for $k \ll d$.

- Suffices to compute $\mathbf{V}_k \in \mathbb{R}^{d \times k}$ and then compute $\mathbf{U}_k \mathbf{\Sigma}_k = \mathbf{XV}_k$.
- Use an **iterative approximation algorithm** to compute an approximation to the top k singular vectors \mathbf{V}_k (the top k eigenvectors of $\mathbf{X}^T \mathbf{X}$.)
- Runtime will be roughly $O(\underline{ndk})$ instead of $O(\underline{nd^2})$.

$$k \ll d$$

Faster Algorithms

To speed up SVD computation we will take advantage of the fact that we typically only care about computing the **top (or bottom) k singular vectors** of a matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$ for $k \ll d$.

- Suffices to compute $\mathbf{V}_k \in \mathbb{R}^{d \times k}$ and then compute $\mathbf{U}_k \boldsymbol{\Sigma}_k = \mathbf{XV}_k$.
- Use an **iterative approximation algorithm** to compute an approximation to the top k singular vectors \mathbf{V}_k (the top k eigenvectors of $\mathbf{X}^T \mathbf{X}$.)
- Runtime will be roughly $O(ndk)$ instead of $O(nd^2)$.

Sparse (iterative) vs. Direct Method.

eig

elgs

svd

svds.

slow (direct method)

fast, iterative

Power Method

Power Method: The most fundamental iterative method for approximate SVD/eigendecomposition. Applies to computing $k = 1$ eigenvectors, but can be generalized to larger k .

Goal: Given symmetric $\mathbf{A} \in \mathbb{R}^{d \times d}$, with eigendecomposition $\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^T$, find $\vec{z} \approx \vec{v}_1$. I.e., the top eigenvector of \mathbf{A} .

$$\mathbf{A}\mathbf{v}_1 = \lambda_1 \mathbf{v}_1$$

Power Method

Power Method: The most fundamental iterative method for approximate SVD/eigendecomposition. Applies to computing $k = 1$ eigenvectors, but can be generalized to larger k .

Goal: Given symmetric $\mathbf{A} \in \mathbb{R}^{d \times d}$, with eigendecomposition $\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^T$, find $\vec{z} \approx \vec{v}_1$. I.e., the top eigenvector of \mathbf{A} .

Apply to $\mathbf{A} = \mathbf{X}^T\mathbf{X}$ to approximate the top right singular vector of \mathbf{X} .

Power Method

Power Method: The most fundamental iterative method for approximate SVD/eigendecomposition. Applies to computing $k = 1$ eigenvectors, but can be generalized to larger k .

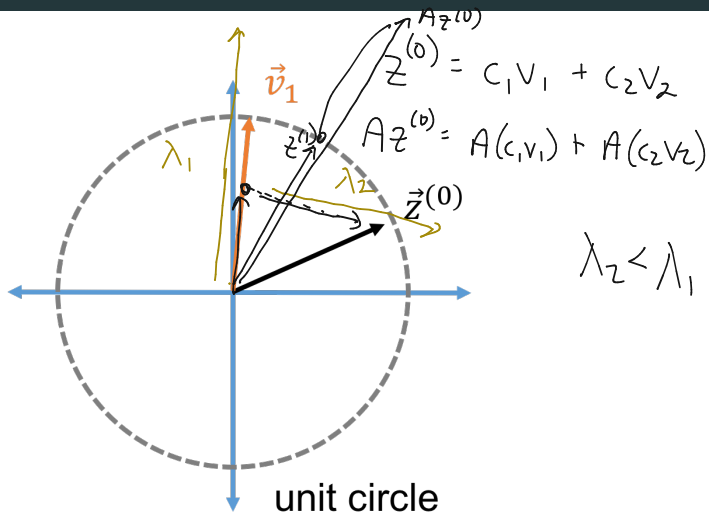


Goal: Given symmetric $\mathbf{A} \in \mathbb{R}^{d \times d}$, with eigendecomposition $\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^T$, find $\vec{z} \approx \vec{v}_1$. I.e., the top eigenvector of \mathbf{A} .

Apply to $\mathbf{A} = \mathbf{X}^T\mathbf{X}$ to approximate the top right singular vector of \mathbf{X} .


- **Initialize:** Choose $\vec{z}^{(0)}$ randomly. E.g. $\vec{z}^{(0)}(i) \sim \mathcal{N}(0, 1)$.
- For $i = 1, \dots, t$
 - $\vec{z}^{(i)} := \mathbf{A} \cdot \vec{z}^{(i-1)}$
 - $\vec{z}_\theta^{(i)} := \frac{\vec{z}^{(i)}}{\|\vec{z}^{(i)}\|_2}$ (unit vector)
- Return $\vec{z}_\theta^{(t)}$

Power Method


$$2 \begin{bmatrix} & 2 \\ & A \end{bmatrix}$$



Power Method



Power Method



Power Method Analysis

Power method:

- **Initialize:** Choose $\vec{z}^{(0)}$ randomly. E.g. $\vec{z}^{(0)}(i) \sim \mathcal{N}(0, 1)$.
- For $i = 1, \dots, t$
 - $\vec{z}^{(i)} := \mathbf{A} \cdot \vec{z}^{(i-1)}$
 - $\vec{z}_i := \frac{\vec{z}^{(i)}}{\|\vec{z}^{(i)}\|_2}$
- Return \vec{z}_t .

Power Method Analysis

Power method:

- **Initialize:** Choose $\vec{z}^{(0)}$ randomly. E.g. $\vec{z}^{(0)}(i) \sim \mathcal{N}(0, 1)$.
- For $i = 1, \dots, t$
 - $\vec{z}^{(i)} := \mathbf{A} \cdot \vec{z}^{(i-1)}$
 - $\vec{z}_i := \frac{\vec{z}^{(i)}}{\|\vec{z}^{(i)}\|_2}$
- Return \vec{z}_t .

Theoretically equivalent to:

- For $i = 1, \dots, t$
 - $\vec{z}^{(i)} := \mathbf{A} \cdot \vec{z}^{(i-1)}$
- $\vec{z}_i := \frac{\vec{z}^{(i)}}{\|\vec{z}^{(i)}\|_2}$.
- Return \vec{z}_t .

Power Method Analysis

Write $\vec{z}^{(0)}$ in \mathbf{A} 's eigenvector basis:

$$\vec{z}^{(0)} = c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_d\vec{v}_d = \mathbf{V}\mathbf{c}.$$

Hand-drawn diagram illustrating the equation $\vec{z}^{(0)} = \mathbf{V}\mathbf{c}$. A square matrix \mathbf{V} is shown with a downward arrow on its left and a superscript d above it. To its right is a column vector \mathbf{c} with elements c_1 , c_2 , and dots, and a downward arrow on its right.

$\mathbf{A} \in \mathbb{R}^{d \times d}$: input matrix with eigendecomposition $\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^T$. \vec{v}_1 : top eigenvector, being computed, $\vec{z}^{(i)}$: iterate at step i , converging to \vec{v}_1 .

Power Method Analysis

Write $\vec{z}^{(0)}$ in \mathbf{A} 's eigenvector basis:

$$\vec{z}^{(0)} = c_1 \vec{v}_1 + c_2 \vec{v}_2 + \dots + c_d \vec{v}_d = \underline{\mathbf{V}} \mathbf{c}.$$

$$\mathbf{A} \vec{z} \\ (\mathbf{V}(\mathbf{V}^T \vec{z}))$$

Update step: $\vec{z}^{(i)} = \mathbf{A} \cdot \vec{z}^{(i-1)} = \underline{\mathbf{V} \mathbf{\Lambda} \mathbf{V}^T} \cdot \vec{z}^{(i-1)}$ (then normalize)

1. $\mathbf{V}^T \vec{z}^{(0)} = \mathbf{V}^T \mathbf{V} \mathbf{c} = \mathbf{c} = \begin{bmatrix} c_1 \\ \vdots \\ c_d \end{bmatrix}$

2. $\mathbf{\Lambda} \mathbf{V}^T \vec{z}^{(0)} = \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_d \end{bmatrix} \begin{bmatrix} c_1 \\ \vdots \\ c_d \end{bmatrix} = \begin{bmatrix} \lambda_1 c_1 \\ \lambda_2 c_2 \\ \vdots \\ \lambda_d c_d \end{bmatrix}$

3. $\vec{z}^{(1)} = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^T \cdot \vec{z}^{(0)} = \begin{bmatrix} \mathbf{V} \\ \mathbf{c} \end{bmatrix} \begin{bmatrix} \lambda_1 c_1 \\ \vdots \\ \lambda_d c_d \end{bmatrix} = v_1 \cdot \lambda_1 c_1 + v_2 \cdot \lambda_2 c_2 + \dots + v_d \cdot \lambda_d c_d$

$\mathbf{A} \in \mathbb{R}^{d \times d}$: input matrix with eigendecomposition $\mathbf{A} = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^T$. \vec{v}_1 : top eigenvector, being computed, $\vec{z}^{(i)}$: iterate at step i , converging to \vec{v}_1 .

Power Method Analysis

Claim 1: Writing $\vec{z}^{(0)} = c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_d\vec{v}_d$,

$$\vec{z}^{(1)} = c_1 \cdot \lambda_1 \vec{v}_1 + c_2 \cdot \lambda_2 \vec{v}_2 + \dots + c_d \cdot \lambda_d \vec{v}_d.$$

$\mathbf{A} \in \mathbb{R}^{d \times d}$: input matrix with eigendecomposition $\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^T$. \vec{v}_1 : top eigenvector, being computed, $\vec{z}^{(i)}$: iterate at step i , converging to \vec{v}_1 .

Power Method Analysis

Claim 1: Writing $\vec{z}^{(0)} = c_1 \vec{v}_1 + c_2 \vec{v}_2 + \dots + c_d \vec{v}_d$,

$$\vec{z}^{(1)} = c_1 \cdot \lambda_1 \vec{v}_1 + c_2 \cdot \lambda_2 \vec{v}_2 + \dots + c_d \cdot \lambda_d \vec{v}_d.$$

$$\vec{z}^{(2)} = \mathbf{A} \vec{z}^{(1)} = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^T \vec{z}^{(1)} = c_1 \lambda_1^2 \vec{v}_1 + \dots + c_d \lambda_d^2 \vec{v}_d$$

$\mathbf{A} \in \mathbb{R}^{d \times d}$: input matrix with eigendecomposition $\mathbf{A} = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^T$. \vec{v}_1 : top eigenvector, being computed, $\vec{z}^{(i)}$: iterate at step i , converging to \vec{v}_1 .

Power Method Analysis

Claim 1: Writing $\vec{z}^{(0)} = c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_d\vec{v}_d$,

$$\vec{z}^{(1)} = c_1 \cdot \lambda_1 \vec{v}_1 + c_2 \cdot \lambda_2 \vec{v}_2 + \dots + c_d \cdot \lambda_d \vec{v}_d.$$

$$A^+ \sqrt{(\cdot)}$$
$$\left(\lambda_i^+ \right)$$

$$\vec{z}^{(2)} = A\vec{z}^{(1)} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^T\vec{z}^{(1)} =$$

Claim 2:

$$\vec{z}^{(t)} = c_1 \cdot \lambda_1^t \vec{v}_1 + c_2 \cdot \lambda_2^t \vec{v}_2 + \dots + c_d \cdot \lambda_d^t \vec{v}_d.$$

huge

$A \in \mathbb{R}^{d \times d}$: input matrix with eigendecomposition $A = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^T$. \vec{v}_1 : top eigenvector, being computed, $\vec{z}^{(i)}$: iterate at step i , converging to \vec{v}_1 .

Power Method Convergence


After t iterations, we have 'powered' up the eigenvalues, making the component in the direction of v_1 much larger, relative to the other components.

$$\vec{z}^{(0)} = c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_d\vec{v}_d \implies \vec{z}^{(t)} = c_1\lambda_1^t\vec{v}_1 + c_2\lambda_2^t\vec{v}_2 + \dots + c_d\lambda_d^t\vec{v}_d$$

Power Method Convergence

After t iterations, we have ‘powered’ up the eigenvalues, making the component in the direction of v_1 much larger, relative to the other components.


$$\vec{z}^{(0)} = c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_d\vec{v}_d \implies \vec{z}^{(t)} = c_1\lambda_1^t\vec{v}_1 + c_2\lambda_2^t\vec{v}_2 + \dots + c_d\lambda_d^t\vec{v}_d$$



Power Method Convergence

After t iterations, we have ‘powered’ up the eigenvalues, making the component in the direction of v_1 much larger, relative to the other components.


$$\vec{z}^{(0)} = c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_d\vec{v}_d \implies \vec{z}^{(t)} = c_1\lambda_1^t\vec{v}_1 + c_2\lambda_2^t\vec{v}_2 + \dots + c_d\lambda_d^t\vec{v}_d$$



Power Method Convergence

After t iterations, we have ‘powered’ up the eigenvalues, making the component in the direction of v_1 much larger, relative to the other components.


$$\vec{z}^{(0)} = c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_d\vec{v}_d \implies \vec{z}^{(t)} = c_1\lambda_1^t\vec{v}_1 + c_2\lambda_2^t\vec{v}_2 + \dots + c_d\lambda_d^t\vec{v}_d$$



Power Method Convergence

After t iterations, we have ‘powered’ up the eigenvalues, making the component in the direction of v_1 much larger, relative to the other components.


$$\vec{z}^{(0)} = c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_d\vec{v}_d \implies \vec{z}^{(t)} = c_1\lambda_1^t\vec{v}_1 + c_2\lambda_2^t\vec{v}_2 + \dots + c_d\lambda_d^t\vec{v}_d$$



Power Method Convergence

After t iterations, we have ‘powered’ up the eigenvalues, making the component in the direction of v_1 much larger, relative to the other components.


$$\vec{z}^{(0)} = c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_d\vec{v}_d \implies \vec{z}^{(t)} = c_1\lambda_1^t\vec{v}_1 + c_2\lambda_2^t\vec{v}_2 + \dots + c_d\lambda_d^t\vec{v}_d$$



Power Method Convergence

After t iterations, we have ‘powered’ up the eigenvalues, making the component in the direction of v_1 much larger, relative to the other components.


$$\vec{z}^{(0)} = c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_d\vec{v}_d \implies \vec{z}^{(t)} = c_1\lambda_1^t\vec{v}_1 + c_2\lambda_2^t\vec{v}_2 + \dots + c_d\lambda_d^t\vec{v}_d$$



Power Method Convergence

After t iterations, we have 'powered' up the eigenvalues, making the component in the direction of v_1 much larger, relative to the other components.


$$\vec{z}^{(0)} = c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_d\vec{v}_d \implies \vec{z}^{(t)} = c_1\lambda_1^t\vec{v}_1 + c_2\lambda_2^t\vec{v}_2 + \dots + c_d\lambda_d^t\vec{v}_d$$



Power Method Convergence

After t iterations, we have ‘powered’ up the eigenvalues, making the component in the direction of v_1 much larger, relative to the other components.


$$\vec{z}^{(0)} = c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_d\vec{v}_d \implies \vec{z}^{(t)} = c_1\lambda_1^t\vec{v}_1 + c_2\lambda_2^t\vec{v}_2 + \dots + c_d\lambda_d^t\vec{v}_d$$



Power Method Convergence

After t iterations, we have ‘powered’ up the eigenvalues, making the component in the direction of \vec{v}_1 much larger, relative to the other components.


$$\vec{z}^{(0)} = c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_d\vec{v}_d \implies \vec{z}^{(t)} = c_1\lambda_1^t\vec{v}_1 + c_2\lambda_2^t\vec{v}_2 + \dots + c_d\lambda_d^t\vec{v}_d$$



Power Method Convergence

After t iterations, we have ‘powered’ up the eigenvalues, making the component in the direction of \vec{v}_1 much larger, relative to the other components.


$$\vec{z}^{(0)} = c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_d\vec{v}_d \implies \vec{z}^{(t)} = c_1\lambda_1^t\vec{v}_1 + c_2\lambda_2^t\vec{v}_2 + \dots + c_d\lambda_d^t\vec{v}_d$$



Power Method Convergence

After t iterations, we have 'powered' up the eigenvalues, making the component in the direction of \vec{v}_1 much larger, relative to the other components.

$$\vec{z}^{(0)} = c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_d\vec{v}_d \implies \vec{z}^{(t)} = c_1\lambda_1^t\vec{v}_1 + c_2\lambda_2^t\vec{v}_2 + \dots + c_d\lambda_d^t\vec{v}_d$$



When will convergence be slow?

$$\lambda_1 \approx \lambda_2$$

Power Method Slow Convergence


Slow Case: A has eigenvalues: $\lambda_1 = 1, \lambda_2 = .99, \lambda_3 = .9, \lambda_4 = .8, \dots$

$$\vec{z}^{(0)} = c_1 \vec{v}_1 + c_2 \vec{v}_2 + \dots + c_d \vec{v}_d \implies \vec{z}^{(t)} = c_1 \lambda_1^t \vec{v}_1 + c_2 \lambda_2^t \vec{v}_2 + \dots + c_d \lambda_d^t \vec{v}_d$$

Power Method Slow Convergence

Slow Case: A has eigenvalues: $\lambda_1 = 1, \lambda_2 = .99, \lambda_3 = .9, \lambda_4 = .8, \dots$


$$\vec{z}^{(0)} = c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_d\vec{v}_d \implies \vec{z}^{(t)} = c_1\lambda_1^t\vec{v}_1 + c_2\lambda_2^t\vec{v}_2 + \dots + c_d\lambda_d^t\vec{v}_d$$



Power Method Slow Convergence

Slow Case: A has eigenvalues: $\lambda_1 = 1, \lambda_2 = .99, \lambda_3 = .9, \lambda_4 = .8, \dots$


$$\vec{z}^{(0)} = c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_d\vec{v}_d \implies \vec{z}^{(t)} = c_1\lambda_1^t\vec{v}_1 + c_2\lambda_2^t\vec{v}_2 + \dots + c_d\lambda_d^t\vec{v}_d$$



Power Method Slow Convergence

Slow Case: A has eigenvalues: $\lambda_1 = 1, \lambda_2 = .99, \lambda_3 = .9, \lambda_4 = .8, \dots$


$$\vec{z}^{(0)} = c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_d\vec{v}_d \implies \vec{z}^{(t)} = c_1\lambda_1^t\vec{v}_1 + c_2\lambda_2^t\vec{v}_2 + \dots + c_d\lambda_d^t\vec{v}_d$$



Power Method Slow Convergence

Slow Case: A has eigenvalues: $\lambda_1 = 1, \lambda_2 = .99, \lambda_3 = .9, \lambda_4 = .8, \dots$


$$\vec{z}^{(0)} = c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_d\vec{v}_d \implies \vec{z}^{(t)} = c_1\lambda_1^t\vec{v}_1 + c_2\lambda_2^t\vec{v}_2 + \dots + c_d\lambda_d^t\vec{v}_d$$



Power Method Slow Convergence

Slow Case: A has eigenvalues: $\lambda_1 = 1, \lambda_2 = .99, \lambda_3 = .9, \lambda_4 = .8, \dots$


$$\vec{z}^{(0)} = c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_d\vec{v}_d \implies \vec{z}^{(t)} = c_1\lambda_1^t\vec{v}_1 + c_2\lambda_2^t\vec{v}_2 + \dots + c_d\lambda_d^t\vec{v}_d$$



Power Method Slow Convergence

Slow Case: A has eigenvalues: $\lambda_1 = 1, \lambda_2 = .99, \lambda_3 = .9, \lambda_4 = .8, \dots$


$$\vec{z}^{(0)} = c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_d\vec{v}_d \implies \vec{z}^{(t)} = c_1\lambda_1^t\vec{v}_1 + c_2\lambda_2^t\vec{v}_2 + \dots + c_d\lambda_d^t\vec{v}_d$$



Power Method Slow Convergence

Slow Case: A has eigenvalues: $\lambda_1 = 1, \lambda_2 = .99, \lambda_3 = .9, \lambda_4 = .8, \dots$


$$\vec{z}^{(0)} = c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_d\vec{v}_d \implies \vec{z}^{(t)} = c_1\lambda_1^t\vec{v}_1 + c_2\lambda_2^t\vec{v}_2 + \dots + c_d\lambda_d^t\vec{v}_d$$



Power Method Slow Convergence

Slow Case: A has eigenvalues: $\lambda_1 = 1, \lambda_2 = .99, \lambda_3 = .9, \lambda_4 = .8, \dots$


$$\vec{z}^{(0)} = c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_d\vec{v}_d \implies \vec{z}^{(t)} = c_1\lambda_1^t\vec{v}_1 + c_2\lambda_2^t\vec{v}_2 + \dots + c_d\lambda_d^t\vec{v}_d$$



Power Method Slow Convergence

Slow Case: A has eigenvalues: $\lambda_1 = 1, \lambda_2 = .99, \lambda_3 = .9, \lambda_4 = .8, \dots$


$$\vec{z}^{(0)} = c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_d\vec{v}_d \implies \vec{z}^{(t)} = c_1\lambda_1^t\vec{v}_1 + c_2\lambda_2^t\vec{v}_2 + \dots + c_d\lambda_d^t\vec{v}_d$$



Power Method Slow Convergence

Slow Case: A has eigenvalues: $\lambda_1 = 1, \lambda_2 = .99, \lambda_3 = .9, \lambda_4 = .8, \dots$


$$\vec{z}^{(0)} = c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_d\vec{v}_d \implies \vec{z}^{(t)} = c_1\lambda_1^t\vec{v}_1 + c_2\lambda_2^t\vec{v}_2 + \dots + c_d\lambda_d^t\vec{v}_d$$



Power Method Slow Convergence

Slow Case: A has eigenvalues: $\lambda_1 = 1, \lambda_2 = .99, \lambda_3 = .9, \lambda_4 = .8, \dots$


$$\vec{z}^{(0)} = c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_d\vec{v}_d \implies \vec{z}^{(t)} = c_1\lambda_1^t\vec{v}_1 + c_2\lambda_2^t\vec{v}_2 + \dots + c_d\lambda_d^t\vec{v}_d$$



Power Method Slow Convergence

Slow Case: A has eigenvalues: $\lambda_1 = 1, \lambda_2 = .99, \lambda_3 = .9, \lambda_4 = .8, \dots$


$$\vec{z}^{(0)} = c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_d\vec{v}_d \implies \vec{z}^{(t)} = c_1\lambda_1^t\vec{v}_1 + c_2\lambda_2^t\vec{v}_2 + \dots + c_d\lambda_d^t\vec{v}_d$$



Power Method Slow Convergence

Slow Case: A has eigenvalues: $\lambda_1 = 1, \lambda_2 = .99, \lambda_3 = .9, \lambda_4 = .8, \dots$

$$\vec{z}^{(0)} = c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_d\vec{v}_d \implies \vec{z}^{(t)} = c_1\lambda_1^t\vec{v}_1 + c_2\lambda_2^t\vec{v}_2 + \dots + c_d\lambda_d^t\vec{v}_d$$



Power Method Slow Convergence

Slow Case: A has eigenvalues: $\lambda_1 = 1, \lambda_2 = .99, \lambda_3 = .9, \lambda_4 = .8, \dots$

$$\vec{z}^{(0)} = c_1 \vec{v}_1 + c_2 \vec{v}_2 + \dots + c_d \vec{v}_d \implies \vec{z}^{(t)} = \underbrace{c_1 \lambda_1^t \vec{v}_1} + \underbrace{c_2 \lambda_2^t \vec{v}_2} + \dots + c_d \lambda_d^t \vec{v}_d$$

$$\lambda_1 = \lambda_2$$


$$\vec{z} = c_1 \vec{v}_1 + c_2 \vec{v}_2$$

$$A\vec{z} = A c_1 \vec{v}_1 + A c_2 \vec{v}_2$$

$$= c_1 \lambda_1 \vec{v}_1 + c_2 \lambda_2 \vec{v}_2$$

$$= \lambda_1 (c_1 \vec{v}_1 + c_2 \vec{v}_2)$$

$$= \underline{\underline{\lambda_1 \vec{z}}}$$



$$\lambda_1 > \lambda_2 > \dots > \lambda_d$$

Power Method Convergence Rate

$$\lambda_2 = 0.99\lambda_1 \quad \gamma = .01 \quad \lambda_2 = \frac{\lambda_1}{3} \quad \gamma = 2/3$$
$$\vec{z}^{(0)} = c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_d\vec{v}_d \implies \vec{z}^{(t)} = c_1\lambda_1^t\vec{v}_1 + c_2\lambda_2^t\vec{v}_2 + \dots + c_d\lambda_d^t\vec{v}_d$$

Write $|\lambda_2| = (1 - \gamma)|\lambda_1|$ for 'gap' $\gamma = \frac{|\lambda_1| - |\lambda_2|}{|\lambda_1|}$. $\gamma \in (0, 1)$

How many iterations t does it take to have $|\lambda_2|^t \leq \delta \cdot |\lambda_1|^t$ for $\delta > 0$?

$$\lambda_2^t = [(1 - \gamma)\lambda_1]^t = (1 - \gamma)^t \cdot \lambda_1^t \leq \delta \cdot \lambda_1^t$$
$$(1 - \gamma)^t \leq \delta$$

\vec{v}_1 : top eigenvector, being computed, $\vec{z}^{(i)}$: iterate at step i , converging to \vec{v}_1 .
 $\lambda_1, \lambda_2, \dots, \lambda_n$: eigenvalues of \mathbf{A} , $\gamma = \frac{|\lambda_1| - |\lambda_2|}{|\lambda_1|}$: eigengap controlling convergence rate

Power Method Convergence Rate

$$\vec{z}^{(0)} = c_1 \vec{v}_1 + c_2 \vec{v}_2 + \dots + c_d \vec{v}_d \implies \vec{z}^{(t)} = c_1 \lambda_1^t \vec{v}_1 + c_2 \lambda_2^t \vec{v}_2 + \dots + c_d \lambda_d^t \vec{v}_d$$

Write $|\lambda_2| = (1 - \gamma)|\lambda_1|$ for 'gap' $\gamma = \frac{|\lambda_1| - |\lambda_2|}{|\lambda_1|}$.

How many iterations t does it take to have $|\lambda_2|^t \leq \delta \cdot |\lambda_1|^t$ for $\delta > 0$?

$$|\lambda_2|^t = (1 - \gamma)^t \cdot |\lambda_1|^t$$

\vec{v}_1 : top eigenvector, being computed, $\vec{z}^{(i)}$: iterate at step i , converging to \vec{v}_1 .
 $\lambda_1, \lambda_2, \dots, \lambda_n$: eigenvalues of \mathbf{A} , $\gamma = \frac{|\lambda_1| - |\lambda_2|}{|\lambda_1|}$: eigengap controlling convergence rate

Power Method Convergence Rate

$$\vec{z}^{(0)} = c_1 \vec{v}_1 + c_2 \vec{v}_2 + \dots + c_d \vec{v}_d \implies \vec{z}^{(t)} = c_1 \lambda_1^t \vec{v}_1 + c_2 \lambda_2^t \vec{v}_2 + \dots + c_d \lambda_d^t \vec{v}_d$$

Write $|\lambda_2| = (1 - \gamma)|\lambda_1|$ for 'gap' $\gamma = \frac{|\lambda_1| - |\lambda_2|}{|\lambda_1|}$.

How many iterations t does it take to have $|\lambda_2|^t \leq \delta \cdot |\lambda_1|^t$ for $\delta > 0$?

$$\begin{aligned} |\lambda_2|^t &= (1 - \gamma)^t \cdot |\lambda_1|^t \\ &= \underbrace{(1 - \gamma)^{1/\gamma}}_{\frac{1}{e}} \cdot \gamma^t \cdot |\lambda_1|^t \end{aligned}$$

\vec{v}_1 : top eigenvector, being computed, $\vec{z}^{(i)}$: iterate at step i , converging to \vec{v}_1 .
 $\lambda_1, \lambda_2, \dots, \lambda_n$: eigenvalues of \mathbf{A} , $\gamma = \frac{|\lambda_1| - |\lambda_2|}{|\lambda_1|}$: eigengap controlling convergence rate

Power Method Convergence Rate

$$\vec{z}^{(0)} = c_1 \vec{v}_1 + c_2 \vec{v}_2 + \dots + c_d \vec{v}_d \implies \vec{z}^{(t)} = c_1 \lambda_1^t \vec{v}_1 + c_2 \lambda_2^t \vec{v}_2 + \dots + c_d \lambda_d^t \vec{v}_d$$

Write $|\lambda_2| = (1 - \gamma)|\lambda_1|$ for 'gap' $\gamma = \frac{|\lambda_1| - |\lambda_2|}{|\lambda_1|}$.

How many iterations t does it take to have $|\lambda_2|^t \leq \delta \cdot |\lambda_1|^t$ for $\delta > 0$?

$$\begin{aligned} |\lambda_2|^t &= (1 - \gamma)^t \cdot \underbrace{|\lambda_1|^t}_1 \\ &= (1 - \gamma)^{1/\gamma \cdot \gamma t} \cdot |\lambda_1|^t = e^{-\gamma t} \\ &\leq e^{-\gamma t} \cdot |\lambda_1|^t \end{aligned}$$

$$\begin{aligned} e^{-\gamma t} &= \delta \\ -\gamma t &= \log(\delta) \end{aligned}$$

$$t = \frac{-\log(\delta)}{\gamma} = \boxed{\frac{\log(1/\delta)}{\gamma}}$$

\vec{v}_1 : top eigenvector, being computed, $\vec{z}^{(i)}$: iterate at step i , converging to \vec{v}_1 .
 $\lambda_1, \lambda_2, \dots, \lambda_n$: eigenvalues of \mathbf{A} , $\gamma = \frac{|\lambda_1| - |\lambda_2|}{|\lambda_1|}$: eigengap controlling convergence rate

Power Method Convergence Rate

$$\vec{z}^{(0)} = c_1 \vec{v}_1 + c_2 \vec{v}_2 + \dots + c_d \vec{v}_d \implies \vec{z}^{(t)} = c_1 \lambda_1^t \vec{v}_1 + c_2 \lambda_2^t \vec{v}_2 + \dots + c_d \lambda_d^t \vec{v}_d$$

Write $|\lambda_2| = (1 - \gamma)|\lambda_1|$ for 'gap' $\gamma = \frac{|\lambda_1| - |\lambda_2|}{|\lambda_1|}$.

How many iterations t does it take to have $|\lambda_2|^t \leq \delta \cdot |\lambda_1|^t$ for $\delta > 0$?

$$\begin{aligned} |\lambda_2|^t &= (1 - \gamma)^t \cdot |\lambda_1|^t \\ &= (1 - \gamma)^{1/\gamma \cdot \gamma t} \cdot |\lambda_1|^t \\ &\leq e^{-\gamma t} \cdot |\lambda_1|^t \end{aligned}$$

So it suffices to set $\gamma t = \ln(1/\delta)$. Or $t = \frac{\ln(1/\delta)}{\gamma}$.

\vec{v}_1 : top eigenvector, being computed, $\vec{z}^{(i)}$: iterate at step i , converging to \vec{v}_1 .
 $\lambda_1, \lambda_2, \dots, \lambda_n$: eigenvalues of \mathbf{A} , $\gamma = \frac{|\lambda_1| - |\lambda_2|}{|\lambda_1|}$: eigengap controlling convergence rate

Power Method Convergence Rate

$$\vec{z}^{(0)} = c_1 \vec{v}_1 + c_2 \vec{v}_2 + \dots + c_d \vec{v}_d \implies \vec{z}^{(t)} = \boxed{c_1 \lambda_1^t} \vec{v}_1 + \boxed{c_2 \lambda_2^t} \vec{v}_2 + \dots + c_d \lambda_d^t \vec{v}_d$$

Write $|\lambda_2| = (1 - \gamma)|\lambda_1|$ for 'gap' $\gamma = \frac{|\lambda_1| - |\lambda_2|}{|\lambda_1|}$.

How many iterations t does it take to have $|\lambda_2|^t \leq \delta \cdot |\lambda_1|^t$ for $\delta > 0$?

$$\begin{aligned} |\lambda_2|^t &= (1 - \gamma)^t \cdot |\lambda_1|^t \\ &= (1 - \gamma)^{1/\gamma \cdot \gamma t} \cdot |\lambda_1|^t \\ &\leq e^{-\gamma t} \cdot |\lambda_1|^t \end{aligned}$$

So it suffices to set $\gamma t = \ln(1/\delta)$. Or $t = \frac{\ln(1/\delta)}{\gamma}$.

How small must we set δ to ensure that $c_1 \lambda_1^t$ dominates all other components and so $\vec{z}^{(t)}$ is very close to \vec{v}_1 ?

\vec{v}_1 : top eigenvector, being computed, $\vec{z}^{(i)}$: iterate at step i , converging to \vec{v}_1 .
 $\lambda_1, \lambda_2, \dots, \lambda_n$: eigenvalues of \mathbf{A} , $\gamma = \frac{|\lambda_1| - |\lambda_2|}{|\lambda_1|}$: eigengap controlling convergence rate

Random Initialization

Claim: When $z^{(0)}$ is chosen with random Gaussian entries, writing $z^{(0)} = c_1 \vec{v}_1 + c_2 \vec{v}_2 + \dots + c_d \vec{v}_d$, with very high probability, for all i :

$$O(1/d^2) \leq |c_i| \leq O(\log d)$$

Corollary:

$$\max_j \left| \frac{c_j}{c_1} \right| \leq O(d^2 \log d).$$

$\mathbf{A} \in \mathbb{R}^{d \times d}$: input matrix with eigendecomposition $\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^T$. \vec{v}_1 : top eigenvector, being computed, $\vec{z}^{(i)}$: iterate at step i , converging to \vec{v}_1 .

Random Initialization

Claim 1: When $z^{(0)}$ is chosen with random Gaussian entries, writing $z^{(0)} = c_1 \vec{v}_1 + c_2 \vec{v}_2 + \dots + c_d \vec{v}_d$, with very high probability,
 $\max_j \left| \frac{c_j}{c_1} \right| \leq O(d^2 \log d)$.

Claim 2: For gap $\gamma = \frac{|\lambda_1| - |\lambda_2|}{|\lambda_1|}$, and $t = \frac{\ln(1/\delta)}{\gamma}$, $\left| \frac{\lambda_i^t}{\lambda_1^t} \right| \leq \delta$ for all i .

$A \in \mathbb{R}^{d \times d}$: input with eigenvalues $\lambda_1 \dots, \lambda_d$ and eigenvectors $\vec{v}_1, \dots, \vec{v}_d$. $\vec{z}^{(i)}$:
iterate at step i . c_1, \dots, c_d : coefficients of $\vec{z}^{(0)}$ in the eigenvector basis.

Random Initialization

Claim 1: When $z^{(0)}$ is chosen with random Gaussian entries, writing $z^{(0)} = c_1 \vec{v}_1 + c_2 \vec{v}_2 + \dots + c_d \vec{v}_d$, with very high probability, $\max_j \left| \frac{c_j}{c_1} \right| \leq O(d^2 \log d)$.

Claim 2: For gap $\gamma = \frac{|\lambda_1| - |\lambda_2|}{|\lambda_1|}$, and $t = \frac{\ln(1/\delta)}{\gamma}$, $\left| \frac{\lambda_i^t}{\lambda_1^t} \right| \leq \delta$ for all i .

$$\vec{z}^{(t)} := \frac{c_1 \lambda_1^t \vec{v}_1 + \dots + c_d \lambda_d^t \vec{v}_d}{\|c_1 \lambda_1^t \vec{v}_1 + \dots + c_d \lambda_d^t \vec{v}_d\|_2}$$

$A \in \mathbb{R}^{d \times d}$: input with eigenvalues $\lambda_1, \dots, \lambda_d$ and eigenvectors $\vec{v}_1, \dots, \vec{v}_d$. $\vec{z}^{(i)}$: iterate at step i . c_1, \dots, c_d : coefficients of $\vec{z}^{(0)}$ in the eigenvector basis.

Random Initialization


Claim 1: When $z^{(0)}$ is chosen with random Gaussian entries, writing $z^{(0)} = c_1 \vec{v}_1 + c_2 \vec{v}_2 + \dots + c_d \vec{v}_d$, with very high probability, $\max_j \left| \frac{c_j}{c_1} \right| \leq O(d^2 \log d)$.

Claim 2: For gap $\gamma = \frac{|\lambda_1| - |\lambda_2|}{|\lambda_1|}$, and $t = \frac{\ln(1/\delta)}{\gamma}$, $\left| \frac{\lambda_i^t}{\lambda_1^t} \right| \leq \delta$ for all i .

$$\vec{z}^{(t)} := \frac{c_1 \lambda_1^t \vec{v}_1 + \dots + c_d \lambda_d^t \vec{v}_d}{\|c_1 \lambda_1^t \vec{v}_1 + \dots + c_d \lambda_d^t \vec{v}_d\|_2} \quad \left. \begin{array}{l} \\ \end{array} \right\} \begin{array}{l} \text{proper} \\ \text{unit vector} \end{array}$$

$$\|\vec{z}^{(t)} - \vec{v}_1\|_2 \leq \left\| \frac{c_1 \lambda_1^t \vec{v}_1 + \dots + c_d \lambda_d^t \vec{v}_d}{\|c_1 \lambda_1^t \vec{v}_1\|_2} - \vec{v}_1 \right\|_2$$

not quite a unit vector



$A \in \mathbb{R}^{d \times d}$: input with eigenvalues $\lambda_1, \dots, \lambda_d$ and eigenvectors $\vec{v}_1, \dots, \vec{v}_d$. $\vec{z}^{(i)}$: iterate at step i . c_1, \dots, c_d : coefficients of $\vec{z}^{(0)}$ in the eigenvector basis.

Random Initialization

Claim 1: When $z^{(0)}$ is chosen with random Gaussian entries, writing $z^{(0)} = c_1 \vec{v}_1 + c_2 \vec{v}_2 + \dots + c_d \vec{v}_d$, with very high probability, $\max_j \left| \frac{c_j}{c_1} \right| \leq O(d^2 \log d)$.

for $i=1, \dots, t$
 $z^{(i)} = A z^{(i-1)}$

Claim 2: For gap $\gamma = \frac{|\lambda_1| - |\lambda_2|}{|\lambda_1|}$, and $t = \frac{\ln(1/\delta)}{\gamma}$, $\left| \frac{\lambda_i^t}{\lambda_1^t} \right| \leq \delta$ for all i .

$z^{(i)} = \frac{z^{(i)}}{\|z^{(i)}\|_2}$

$$\vec{z}^{(t)} := \frac{c_1 \lambda_1^t \vec{v}_1 + \dots + c_d \lambda_d^t \vec{v}_d}{\|c_1 \lambda_1^t \vec{v}_1 + \dots + c_d \lambda_d^t \vec{v}_d\|_2} = \frac{\vec{z}^{(t)}}{\|\vec{z}^{(t)}\|_2}$$

$$\|\vec{z}^{(t)} - \vec{v}_1\|_2 \leq \left\| \frac{c_1 \lambda_1^t \vec{v}_1 + \dots + c_d \lambda_d^t \vec{v}_d}{\|c_1 \lambda_1^t \vec{v}_1 + \dots + c_d \lambda_d^t \vec{v}_d\|_2} - \vec{v}_1 \right\|_2$$

$c_1 \lambda_1^t$

simpler but worse approx to \vec{v}_1

$$= \left\| \frac{c_2 \lambda_2^t \vec{v}_2 + \dots + c_d \lambda_d^t \vec{v}_d}{c_1 \lambda_1^t} \right\|_2$$

$A \in \mathbb{R}^{d \times d}$: input with eigenvalues $\lambda_1, \dots, \lambda_d$ and eigenvectors $\vec{v}_1, \dots, \vec{v}_d$. $\vec{z}^{(i)}$: iterate at step i . c_1, \dots, c_d : coefficients of $\vec{z}^{(0)}$ in the eigenvector basis.

Random Initialization

Claim 1: When $z^{(0)}$ is chosen with random Gaussian entries, writing $z^{(0)} = c_1 \vec{v}_1 + c_2 \vec{v}_2 + \dots + c_d \vec{v}_d$, with very high probability, $\max_j \left| \frac{c_j}{c_1} \right| \leq O(d^2 \log d)$.

Claim 2: For gap $\gamma = \frac{|\lambda_1| - |\lambda_2|}{|\lambda_1|}$, and $t = \frac{\ln(1/\delta)}{\gamma}$, $\left| \frac{\lambda_i^t}{\lambda_1^t} \right| \leq \delta$ for all i .

$$\begin{aligned} \bar{z}^{(t)} &:= \frac{c_1 \lambda_1^t \vec{v}_1 + \dots + c_d \lambda_d^t \vec{v}_d}{\|c_1 \lambda_1^t \vec{v}_1 + \dots + c_d \lambda_d^t \vec{v}_d\|_2} \implies \\ \|\bar{z}^{(t)} - \vec{v}_1\|_2^2 &\leq \left\| \frac{c_1 \lambda_1^t \vec{v}_1 + \dots + c_d \lambda_d^t \vec{v}_d}{\|c_1 \lambda_1^t \vec{v}_1\|_2} - \vec{v}_1 \right\|_2^2 \\ &= \left\| \frac{c_2 \lambda_2^t}{c_1 \lambda_1^t} \vec{v}_2 + \dots + \frac{c_d \lambda_d^t}{c_1 \lambda_1^t} \vec{v}_d \right\|_2^2 = \left| \frac{c_2 \lambda_2^t}{c_1 \lambda_1^t} \right|^2 + \dots + \left| \frac{c_d \lambda_d^t}{c_1 \lambda_1^t} \right|^2 \quad (\text{Pythagorean}) \end{aligned}$$



$A \in \mathbb{R}^{d \times d}$: input with eigenvalues $\lambda_1, \dots, \lambda_d$ and eigenvectors $\vec{v}_1, \dots, \vec{v}_d$. $\bar{z}^{(i)}$: iterate at step i . c_1, \dots, c_d : coefficients of $\bar{z}^{(0)}$ in the eigenvector basis.

Random Initialization

Claim 1: When $z^{(0)}$ is chosen with random Gaussian entries, writing $z^{(0)} = c_1 \vec{v}_1 + c_2 \vec{v}_2 + \dots + c_d \vec{v}_d$, with very high probability, $\max_j \left| \frac{c_j}{c_1} \right| \leq O(d^2 \log d)$.

Claim 2: For gap $\gamma = \frac{|\lambda_1| - |\lambda_2|}{|\lambda_1|}$, and $t = \frac{\ln(1/\delta)}{\gamma}$, $\left| \frac{\lambda_i^t}{\lambda_1^t} \right| \leq \delta$ for all i .

$$\begin{aligned} \bar{z}^{(t)} &:= \frac{c_1 \lambda_1^t \vec{v}_1 + \dots + c_d \lambda_d^t \vec{v}_d}{\|c_1 \lambda_1^t \vec{v}_1 + \dots + c_d \lambda_d^t \vec{v}_d\|_2} \implies \\ \|\bar{z}^{(t)} - \vec{v}_1\|_2 &\leq \left\| \frac{c_1 \lambda_1^t \vec{v}_1 + \dots + c_d \lambda_d^t \vec{v}_d}{\|c_1 \lambda_1^t \vec{v}_1\|_2} - \vec{v}_1 \right\|_2^2 \\ &= \left\| \frac{c_2 \lambda_2^t}{c_1 \lambda_1^t} \vec{v}_2 + \dots + \frac{c_d \lambda_d^t}{c_1 \lambda_1^t} \vec{v}_d \right\|_2^2 = \left| \frac{c_2 \lambda_2^t}{c_1 \lambda_1^t} \right|^2 + \dots + \left| \frac{c_d \lambda_d^t}{c_1 \lambda_1^t} \right|^2 \leq \left[\delta \cdot O(d^2 \log d) \right]^2 \cdot d. \end{aligned}$$

$A \in \mathbb{R}^{d \times d}$: input with eigenvalues $\lambda_1, \dots, \lambda_d$ and eigenvectors $\vec{v}_1, \dots, \vec{v}_d$. $\bar{z}^{(i)}$: iterate at step i . c_1, \dots, c_d : coefficients of $\bar{z}^{(0)}$ in the eigenvector basis.

Random Initialization

Claim 1: When $z^{(0)}$ is chosen with random Gaussian entries, writing $z^{(0)} = c_1 \vec{v}_1 + c_2 \vec{v}_2 + \dots + c_d \vec{v}_d$, with very high probability, $\max_j \left| \frac{c_j}{c_1} \right| \leq O(d^2 \log d)$.

Claim 2: For gap $\gamma = \frac{|\lambda_1| - |\lambda_2|}{|\lambda_1|}$, and $t = \frac{\ln(1/\delta)}{\gamma}$, $\left| \frac{\lambda_i^t}{\lambda_1^t} \right| \leq \delta$ for all i .

$$\vec{z}^{(t)} := \frac{c_1 \lambda_1^t \vec{v}_1 + \dots + c_d \lambda_d^t \vec{v}_d}{\|c_1 \lambda_1^t \vec{v}_1 + \dots + c_d \lambda_d^t \vec{v}_d\|_2} \Rightarrow$$

$$\|\vec{z}^{(t)} - \vec{v}_1\|_2 \leq \left\| \frac{c_1 \lambda_1^t \vec{v}_1 + \dots + c_d \lambda_d^t \vec{v}_d}{\|c_1 \lambda_1^t \vec{v}_1\|_2} - \vec{v}_1 \right\|_2^2$$

$$= \left\| \frac{c_2 \lambda_2^t}{c_1 \lambda_1^t} \vec{v}_2 + \dots + \frac{c_d \lambda_d^t}{c_1 \lambda_1^t} \vec{v}_d \right\|_2^2 = \left| \frac{c_2 \lambda_2^t}{c_1 \lambda_1^t} \right|^2 + \dots + \left| \frac{c_d \lambda_d^t}{c_1 \lambda_1^t} \right|^2 \leq \left[\delta \cdot O(d^2 \log d) \right]^2 \cdot d$$

$$\delta = O\left(\frac{\epsilon^{1/2}}{d^{2.5} \log d}\right)$$

Setting $\delta = O\left(\frac{\epsilon^{1/2}}{d^{2.5} \log d}\right)$ gives $\|\vec{z}^{(t)} - \vec{v}_1\|_2 \leq \epsilon$.

$$\begin{aligned} \log(1/\delta) &= \log\left(\frac{d^{2.5}}{\epsilon}\right) \\ &= O\left(\frac{\log(d/\epsilon)}{\gamma}\right) \\ &\leq \left[\delta \cdot O(d^2 \log d) \right]^2 \cdot d \\ &< \epsilon \\ \delta^2 &\leq \frac{\epsilon^{1/2}}{d^{4.5} \log^2 d} \cdot d \\ &= \frac{\epsilon^{1/2}}{d^{3.5} \log^2 d} \end{aligned}$$

$A \in \mathbb{R}^{d \times d}$: input with eigenvalues $\lambda_1, \dots, \lambda_d$ and eigenvectors $\vec{v}_1, \dots, \vec{v}_d$. $\vec{z}^{(i)}$: iterate at step i . c_1, \dots, c_d : coefficients of $\vec{z}^{(0)}$ in the eigenvector basis.

Power Method Theorem

Theorem (Basic Power Method Convergence)

Let $\gamma = \frac{|\lambda_1| - |\lambda_2|}{|\lambda_1|}$ be the relative gap between the first and second eigenvalues. If Power Method is initialized with a random Gaussian vector $\vec{v}^{(0)}$ then, with high probability, after $t = O\left(\frac{\ln(d/\epsilon)}{\gamma}\right)$ steps:

$$\|\vec{z}^{(t)} - \vec{v}_1\|_2 \leq \epsilon.$$

Power Method Theorem

Theorem (Basic Power Method Convergence)

Let $\gamma = \frac{|\lambda_1| - |\lambda_2|}{|\lambda_1|}$ be the relative gap between the first and second eigenvalues. If Power Method is initialized with a random Gaussian vector $\vec{v}^{(0)}$ then, with high probability, after $t = O\left(\frac{\ln(d/\epsilon)}{\gamma}\right)$ steps:

$$\|\vec{z}^{(t)} - \vec{v}_1\|_2 \leq \epsilon.$$

Total runtime: $O(t)$ matrix-vector multiplications. If $\mathbf{A} = \mathbf{X}^T \mathbf{X}$:

$$\sum_{i=1}^t O(\text{mv}(\mathbf{X}) \cdot \frac{\ln(d/\epsilon)}{\gamma}) = O\left(nd \cdot \frac{\ln(d/\epsilon)}{\gamma}\right).$$

$$(X^T(X z^{i-1}))$$

$$O(nd^2)$$

Power Method Theorem

Theorem (Basic Power Method Convergence)

Let $\gamma = \frac{|\lambda_1| - |\lambda_2|}{|\lambda_1|}$ be the relative gap between the first and second eigenvalues. If Power Method is initialized with a random Gaussian vector $\vec{v}^{(0)}$ then, with high probability, after $t = O\left(\frac{\ln(d/\epsilon)}{\gamma}\right)$ steps:

$$\|\vec{z}^{(t)} - \vec{v}_1\|_2 \leq \epsilon.$$

Total runtime: $O(t)$ matrix-vector multiplications. If $\mathbf{A} = \mathbf{X}^T \mathbf{X}$:

$$O\left(\text{mv}(\mathbf{X}) \cdot \frac{\ln(d/\epsilon)}{\gamma}\right) = O\left(nd \cdot \frac{\ln(d/\epsilon)}{\gamma}\right).$$

How is ϵ dependence?

How is γ dependence?

Krylov subspace methods (Lanczos method, Arnoldi method.)

- How **svds/eigs** are actually implemented. Only need $t = O\left(\frac{\ln(d/\epsilon)}{\sqrt{\gamma}}\right)$ steps for the same guarantee.

Krylov subspace methods (Lanczos method, Arnoldi method.)

- How `svecs/eigs` are actually implemented. Only need $t = O\left(\frac{\ln(d/\epsilon)}{\sqrt{\gamma}}\right)$ steps for the same guarantee.

Main Idea: Need to separate λ_1 from λ_i for $i \geq 2$.

Krylov subspace methods (Lanczos method, Arnoldi method.)

- How `svals/eigs` are actually implemented. Only need $t = O\left(\frac{\ln(d/\epsilon)}{\sqrt{\gamma}}\right)$ steps for the same guarantee.

Main Idea: Need to separate λ_1 from λ_i for $i \geq 2$.

- Power method: power up to λ_1^t and λ_j^t .

Krylov Subspace Methods

Krylov subspace methods (Lanczos method, Arnoldi method.)

- How **svds/eigs** are actually implemented. Only need $t = O\left(\frac{\ln(d/\epsilon)}{\sqrt{\gamma}}\right)$ steps for the same guarantee.

Main Idea: Need to separate λ_1 from λ_i for $i \geq 2$.

- Power method: power up to λ_1^t and λ_i^t . $p(x) = x^t$
- Krylov methods: apply a **better** degree t polynomial $T_t(\cdot)$ to the eigenvalues to separate $T_t(\lambda_1)$ from $T_t(\lambda_i)$.

Krylov Subspace Methods

Krylov subspace methods (Lanczos method, Arnoldi method.)

- How **svds/eigs** are actually implemented. Only need $t = O\left(\frac{\ln(d/\epsilon)}{\sqrt{\gamma}}\right)$ steps for the same guarantee.


Main Idea: Need to separate λ_1 from λ_i for $i \geq 2$.

- Power method: power up to λ_1^t and λ_j^t .
- Krylov methods: apply a **better** degree t polynomial $T_t(\cdot)$ to the eigenvalues to separate $T_t(\lambda_1)$ from $T_t(\lambda_j)$.
- Still requires just t matrix vector multiplies. **Why?**


$$(A(A(Az^0)))$$

$$c_4 A^4 z^0 + \dots + c_2 A^2 z^0 + c_1 A z^0 + c_0 z^0 = p(A) \cdot z^0$$

krylov subspace methods



vs.



Optimal ‘jump’ polynomial in general is given by a degree t **Chebyshev polynomial**. Krylov methods find a polynomial tuned to the input matrix that does at least as well.

Generalizations to Larger k

- Block Power Method (a.k.a. Simultaneous Iteration, Subspace Iteration, or Orthogonal Iteration)
- Pick random block of vectors $Z^{(0)} \in \mathbb{R}^{n \times k}$. $Z^{(i)} = \text{orth}(AZ^{(i-1)})$.
- Block Krylov methods use an analogous approach.

Runtime: $O\left(\underbrace{ndk}_{\text{circled}} \cdot \frac{\ln(d/\epsilon)}{\sqrt{\gamma}}\right)$

to accurately compute the top k singular vectors.

$$A A \begin{bmatrix} z^0 \end{bmatrix} \quad A A A \begin{bmatrix} z^0 \\ \vdots \\ z^k \end{bmatrix}$$

Generalizations to Larger k

- Block Power Method (a.k.a. Simultaneous Iteration, Subspace Iteration, or Orthogonal Iteration)
- Pick random block of vectors $\mathbf{Z}^{(0)} \in \mathbb{R}^{n \times k}$. $\mathbf{Z}^{(i)} = \text{orth}(\mathbf{AZ}^{(i-1)})$.
- Block Krylov methods use an analogous approach.

Runtime: $O\left(ndk \cdot \frac{\ln(d/\epsilon)}{\sqrt{\gamma}}\right)$

to accurately compute the top k singular vectors.

'Gapless' Runtime: $O\left(ndk \cdot \frac{\ln(d/\epsilon)}{\sqrt{\epsilon}}\right)$

if you just want a set of vectors that gives an ϵ -optimal low-rank approximation when you project onto them.