

COMPSCI 514: Algorithms for Data Science

Cameron Musco

University of Massachusetts Amherst. Fall 2024.

Lecture 21

- Problem Set 4 due Monday.
- No class or office hours next week. No quiz due.
- Office hours tomorrow 10am-11am in CS 142.
- Practice final exams have been posted in Canvas. I will release a more complete study guide with additional practice questions *soon*.

Summary

Last Class: Fast computation of the SVD/eigendecomposition.

- Power method for approximating the top eigenvector of a matrix.
- Start on analysis of convergence.

This Class (+ Rest of Semester):

- Finish up power method analysis.
- General iterative algorithms for optimization, specifically **gradient descent** and its variants.
- What are these methods, when are they applied, and how do you analyze their performance?
- Small taste of what you can find in COMPSCI 651.

kernel methods
- connections to random walks/
markov chains


Power Method Wrap Up

Power Method

Basic Power Method:

$$A \in \mathbb{R}^{d \times d}$$

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




Power Method Convergence Rate

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

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 $\delta > 0$?



$$\begin{aligned} (1 - \gamma)^t |\lambda_1|^t &\leq \delta |\lambda_1|^t & t &\geq \frac{\log(\delta)}{\log(1/\delta)} \\ (1 - \gamma)^t &\leq \delta \\ t \cdot \log(1 - \gamma) &\leq \log(\delta) \\ t \cdot \log\left(\frac{1}{1 - \gamma}\right) &\geq \log(1/\delta) \end{aligned}$$

\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 $\delta > 0$?

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

\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 $\delta > 0$?

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

\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 $\delta > 0$?

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

\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 $\delta > 0$?

non # know γ

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

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

\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

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

Handwritten notes: A circle around λ_1^t and a circle around λ_2^t with an arrow pointing to it from the label $\delta|\lambda_1|^t$ above.

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 :

$$\underbrace{c_1}_{N(0,1)} \vec{v}_1 + \underbrace{c_2}_{N(0,1)} \vec{v}_2 + \dots + c_d \vec{v}_d, \quad \mathbb{E} c_i = 0$$
$$\underbrace{O(1/d^2)}_{\text{anti-concentration bound}} \leq |c_i| \leq \underbrace{O(\log d)}_{\text{concentration bound}}$$

Corollary:

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



$A \in \mathbb{R}^{d \times d}$: 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 .

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 .

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

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



$A \in \mathbb{R}^{d \times d}$: 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 .

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

$A \in \mathbb{R}^{d \times d}$: 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 .

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

$A \in \mathbb{R}^{d \times d}$: 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 .

$$\begin{aligned} \bar{z}^{(t)} &:= \frac{c_1 \lambda_1^t \vec{v}_1 + \dots + c_d \lambda_d^t \vec{v}_d}{\|c_1 \lambda_1^t \vec{v}_1 + \dots + c_d \lambda_d^t \vec{v}_d\|_2} \implies \\ \|\bar{z}^{(t)} - \vec{v}_1\|_2 &\leq \left\| \frac{c_1 \lambda_1^t \vec{v}_1 + \dots + c_d \lambda_d^t \vec{v}_d}{\|c_1 \lambda_1^t \vec{v}_1\|_2} - \vec{v}_1 \right\|_2 \\ &= \left\| \frac{c_2 \lambda_2^t}{c_1 \lambda_1^t} \vec{v}_2 + \dots + \frac{c_d \lambda_d^t}{c_1 \lambda_1^t} \vec{v}_d \right\|_2 \ll \underbrace{\left(\frac{c_2 \lambda_2^t}{c_1 \lambda_1^t} \right)}_{O(d^2 \log d)} \delta + \dots + \left| \frac{c_d \lambda_d^t}{c_1 \lambda_1^t} \right| \leq \delta \cdot O(d^2 \log d) \cdot d. \\ &\quad \underbrace{\delta \cdot O(d^3 \log d)} \end{aligned}$$

$$\sqrt{\sum \left| \frac{c_i \lambda_i^t}{c_1 \lambda_1^t} \right|^2}$$

$A \in \mathbb{R}^{d \times d}$: input with eigenvalues $\lambda_1, \dots, \lambda_d$ and eigenvectors $\vec{v}_1, \dots, \vec{v}_d$. $\bar{z}^{(i)}$: 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 $t = \frac{\ln(1/\delta)}{\gamma}$, $\left| \frac{\lambda_j^t}{\lambda_1^t} \right| \leq \delta$ for all j .

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

$$\begin{aligned} \|\vec{z}^{(t)} - \vec{v}_1\|_2 &\leq \left\| \frac{c_1 \lambda_1^t \vec{v}_1 + \dots + c_d \lambda_d^t \vec{v}_d}{\|c_1 \lambda_1^t \vec{v}_1\|_2} - \vec{v}_1 \right\|_2 \\ &= \left\| \frac{c_2 \lambda_2^t}{c_1 \lambda_1^t} \vec{v}_2 + \dots + \frac{c_d \lambda_d^t}{\lambda_1^t} \vec{v}_d \right\|_2 = \left| \frac{c_2 \lambda_2^t}{c_1 \lambda_1^t} \right| + \dots + \left| \frac{c_d \lambda_d^t}{\lambda_1^t} \right| \leq \underbrace{\delta \cdot O(d^2 \log d) \cdot d}_{\frac{\log(1/\delta)}{\gamma} = \frac{\log(d^2 \frac{\log d}{\delta})}{\gamma} = \frac{0(2 \log d)}{\gamma}} \end{aligned}$$

Setting $\delta = \underbrace{O\left(\frac{\epsilon}{d^3 \log d}\right)}_{\frac{\log(1/\delta)}{\gamma} = \frac{\log(d^2 \frac{\log d}{\delta})}{\gamma} = \frac{0(2 \log d)}{\gamma}}$ gives $\|\vec{z}^{(t)} - \vec{v}_1\|_2 \leq \epsilon$.

$A \in \mathbb{R}^{d \times d}$: 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:

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

$$X^T(Xz)$$

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

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

$$\underline{O(nd)}$$

$O(nd^2)$ for full eigen decomposition.

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 $\mathbf{A} = \mathbf{X}^T \mathbf{X}$:

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

How is ϵ dependence? very good

How is γ dependence? not great

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

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

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 **slds/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 . $\sim f(x) \approx \underline{x^t}$
- 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.)

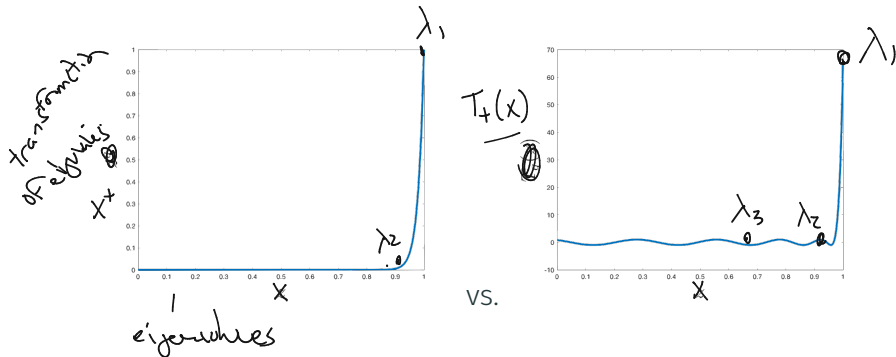
- How **svds/eigs** are actually implemented. Only need $t = O\left(\frac{\ln(d/\epsilon)}{\sqrt{\gamma}}\right)$ steps for the same guarantee.

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?**

$$\begin{aligned} & c_1 X + c_2 X^2 + \dots + c_t X^t \\ & c_1 A^0 v_0 + c_2 A^2 v_0 + \dots + c_t A^t v_0 \end{aligned}$$

krylov subspace methods



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)
- Block Krylov methods

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

to accurately compute the top k singular vectors.

Generalizations to Larger k

- Block Power Method (a.k.a. Simultaneous Iteration, Subspace Iteration, or Orthogonal Iteration)
- Block Krylov methods

$$A \overset{l \times k}{Z} \rightarrow Z = \text{orth}(AZ)$$

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.

Connection Between Random Walks, Eigenvectors, and Power Method


Connection to Random Walks

Consider a random walk on a graph G with adjacency matrix A .



Connection to Random Walks


Consider a random walk on a graph G with adjacency matrix A .



At each step, move to a random vertex, chosen uniformly at random from the neighbors of the current vertex.


Connection to Random Walks

Consider a random walk on a graph G with adjacency matrix A .



Connection to Random Walks

Consider a random walk on a graph G with adjacency matrix A .




Connection to Random Walks

Consider a random walk on a graph G with adjacency matrix A .

$$P^{(0)} = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}$$

$$P^{(1)} = \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}$$

$$P_2 = \begin{pmatrix} 1/3 \\ 0 \\ 1/3 \\ 1/3 \\ 1 \\ 1/3 \\ 0 \end{pmatrix}$$



Connection to Random Walks

Let $\vec{p}^{(t)} \in \mathbb{R}^n$ have i^{th} entry $\vec{p}_i^{(t)} = \Pr(\text{walk at node } i \text{ at step } t)$.

Connection to Random Walks

Let $\vec{p}^{(t)} \in \mathbb{R}^n$ have i^{th} entry $\vec{p}_i^{(t)} = \Pr(\text{walk at node } i \text{ at step } t)$.


- **Initialize:** $\vec{p}^{(0)} = [1, 0, 0, \dots, 0]$.

Connection to Random Walks

Let $\vec{p}^{(t)} \in \mathbb{R}^n$ have i^{th} entry $\vec{p}_i^{(t)} = \Pr(\text{walk at node } i \text{ at step } t)$.

- **Initialize:** $\vec{p}^{(0)} = [1, 0, 0, \dots, 0]$.
- **Update:**

$$\Pr(\text{walk at } i \text{ at step } t) = \sum_{j \in \text{neigh}(i)} \Pr(\text{walk at } j \text{ at step } t-1) \cdot \frac{1}{\text{degree}(j)}$$



Connection to Random Walks

Let $\vec{p}^{(t)} \in \mathbb{R}^n$ have i^{th} entry $\vec{p}_i^{(t)} = \Pr(\text{walk at node } i \text{ at step } t)$.

- **Initialize:** $\vec{p}^{(0)} = [1, 0, 0, \dots, 0]$.
- **Update:**

$$\Pr(\text{walk at } i \text{ at step } t) = \sum_{j \in \text{neigh}(i)} \Pr(\text{walk at } j \text{ at step } t-1) \cdot \frac{1}{\text{degree}(j)}$$

$$\left[\begin{array}{cccc} 0 & \frac{1}{2} & 0 & 0 \\ \frac{1}{2} & 0 & 0 & \frac{1}{2} \\ 0 & 0 & \frac{1}{5} & 0 \\ 0 & \frac{1}{2} & 0 & \frac{1}{2} \end{array} \right] = \vec{z}^T \vec{p}^{(t-1)}$$

where $\vec{z}(j) = \frac{1}{\text{degree}(j)}$ for all $j \in \text{neigh}(i)$, $\vec{z}(j) = 0$ for all $j \notin \text{neigh}(i)$.

Connection to Random Walks

Let $\vec{p}^{(t)} \in \mathbb{R}^n$ have i^{th} entry $\vec{p}_i^{(t)} = \Pr(\text{walk at node } i \text{ at step } t)$.

- **Initialize:** $\vec{p}^{(0)} = [1, 0, 0, \dots, 0]$.
- **Update:**

$$\begin{aligned}\Pr(\text{walk at } i \text{ at step } t) &= \sum_{j \in \text{neigh}(i)} \Pr(\text{walk at } j \text{ at step } t-1) \cdot \frac{1}{\text{degree}(j)} \\ &= \vec{z}^T \vec{p}^{(t-1)}\end{aligned}$$

where $\vec{z}(j) = \frac{1}{\text{degree}(j)}$ for all $j \in \text{neigh}(i)$, $\vec{z}(j) = 0$ for all $j \notin \text{neigh}(i)$.

- \vec{z} is the i^{th} row of the right normalized adjacency matrix \underline{AD}^{-1} .

$$\begin{bmatrix} & A & \\ 0 & 1 & 0 & 1 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1/2, & & & & & \\ & 1/2 & & & & \\ & & 1/2 & & & \\ & & & \ddots & & \\ & & & & 1/2 & \\ & & & & & 1/2 \end{bmatrix} = \begin{bmatrix} 0 & 1/2 & 0 & 1/4 & 0 & 1/6 \end{bmatrix}$$

Connection to Random Walks

Let $\vec{p}^{(t)} \in \mathbb{R}^n$ have i^{th} entry $\vec{p}_i^{(t)} = \Pr(\text{walk at node } i \text{ at step } t)$.

- **Initialize:** $\vec{p}^{(0)} = [1, 0, 0, \dots, 0]$.
- **Update:**

$$\begin{aligned}\Pr(\text{walk at } i \text{ at step } t) &= \sum_{j \in \text{neigh}(i)} \Pr(\text{walk at } j \text{ at step } t-1) \cdot \frac{1}{\text{degree}(j)} \\ &= \vec{z}^T \vec{p}^{(t-1)}\end{aligned}$$

where $\vec{z}(j) = \frac{1}{\text{degree}(j)}$ for all $j \in \text{neigh}(i)$, $\vec{z}(j) = 0$ for all $j \notin \text{neigh}(i)$.

- \vec{z} is the i^{th} row of the right normalized adjacency matrix \mathbf{AD}^{-1} .
- $\vec{p}^{(t)} = \mathbf{AD}^{-1} \vec{p}^{(t-1)}$

Connection to Random Walks

Let $\vec{p}^{(t)} \in \mathbb{R}^n$ have i^{th} entry $\vec{p}_i^{(t)} = \Pr(\text{walk at node } i \text{ at step } t)$.

- **Initialize:** $\vec{p}^{(0)} = [1, 0, 0, \dots, 0]$.
- **Update:**

$$\begin{aligned}\Pr(\text{walk at } i \text{ at step } t) &= \sum_{j \in \text{neigh}(i)} \Pr(\text{walk at } j \text{ at step } t-1) \cdot \frac{1}{\text{degree}(j)} \\ &= \vec{z}^T \vec{p}^{(t-1)}\end{aligned}$$

where $\vec{z}(j) = \frac{1}{\text{degree}(j)}$ for all $j \in \text{neigh}(i)$, $\vec{z}(j) = 0$ for all $j \notin \text{neigh}(i)$.

- \vec{z} is the i^{th} row of the right normalized adjacency matrix \mathbf{AD}^{-1} .
- $\vec{p}^{(t)} = \mathbf{AD}^{-1} \vec{p}^{(t-1)} = \underbrace{\mathbf{AD}^{-1} \mathbf{AD}^{-1} \dots \mathbf{AD}^{-1}}_{t \text{ times}} \vec{p}^{(0)}$

Random Walking as Power Method

Claim: After t steps, the probability that a random walk is at node i is given by the i^{th} entry of

$$\vec{p}^{(t)} = \underbrace{AD^{-1}AD^{-1}\dots AD^{-1}}_{t \text{ times}} \vec{p}^{(0)}.$$

Random Walking as Power Method

Claim: After t steps, the probability that a random walk is at node i is given by the i^{th} entry of

$$D^{-1/2} \vec{p}^{(t)} = \underbrace{D^{-1/2} A D^{-1} A D^{-1} \dots A D^{-1}}_{t \text{ times}} D^{-1/2} \vec{p}^{(0)}.$$

$$D^{-1/2} \vec{p}^{(t)} = \underbrace{(D^{-1/2} A D^{-1/2})}_{t \text{ times}} \dots (D^{-1/2} A D^{-1/2}) D^{-1/2} \vec{p}^{(0)}.$$

t step power method

Random Walking as Power Method

Claim: After t steps, the probability that a random walk is at node i is given by the i^{th} entry of

$$\vec{p}^{(t)} = \underbrace{\mathbf{A}\mathbf{D}^{-1}\mathbf{A}\mathbf{D}^{-1} \dots \mathbf{A}\mathbf{D}^{-1}}_{t \text{ times}} \vec{p}^{(0)}.$$

$$\mathbf{D}^{-1/2} \vec{p}^{(t)} = \underbrace{(\mathbf{D}^{-1/2}\mathbf{A}\mathbf{D}^{-1/2})(\mathbf{D}^{-1/2}\mathbf{A}\mathbf{D}^{-1/2}) \dots (\mathbf{D}^{-1/2}\mathbf{A}\mathbf{D}^{-1/2})}_{t \text{ times}} (\mathbf{D}^{-1/2} \vec{p}^{(0)}).$$

- $\mathbf{D}^{-1/2} \vec{p}^{(t)}$ is exactly what would be obtained by applying t iterations of power method to $\mathbf{D}^{-1/2} \vec{p}^{(0)}$!

Random Walking as Power Method

Claim: After t steps, the probability that a random walk is at node i is given by the i^{th} entry of

$$\vec{p}^{(t)} = \underbrace{\mathbf{A}\mathbf{D}^{-1}\mathbf{A}\mathbf{D}^{-1} \dots \mathbf{A}\mathbf{D}^{-1}}_{t \text{ times}} \vec{p}^{(0)}.$$

$$\mathbf{D}^{-1/2}\vec{p}^{(t)} = \underbrace{(\mathbf{D}^{-1/2}\mathbf{A}\mathbf{D}^{-1/2})(\mathbf{D}^{-1/2}\mathbf{A}\mathbf{D}^{-1/2}) \dots (\mathbf{D}^{-1/2}\mathbf{A}\mathbf{D}^{-1/2})}_{t \text{ times}} (\mathbf{D}^{-1/2}\vec{p}^{(0)}).$$

- $\mathbf{D}^{-1/2}\vec{p}^{(t)}$ is exactly what would be obtained by applying $t/2$ iterations of power method to $\mathbf{D}^{-1/2}\vec{p}^{(0)}$!
- Will converge to the top eigenvector of the normalized adjacency matrix $\mathbf{D}^{-1/2}\mathbf{A}\mathbf{D}^{-1/2}$. **Stationary distribution.**

Random Walking as Power Method

Claim: After t steps, the probability that a random walk is at node i is given by the i^{th} entry of

$$\vec{p}^{(t)} = \underbrace{\mathbf{A}\mathbf{D}^{-1}\mathbf{A}\mathbf{D}^{-1} \dots \mathbf{A}\mathbf{D}^{-1}}_{t \text{ times}} \vec{p}^{(0)}.$$

$$\mathbf{D}^{-1/2}\vec{p}^{(t)} = \underbrace{(\mathbf{D}^{-1/2}\mathbf{A}\mathbf{D}^{-1/2})(\mathbf{D}^{-1/2}\mathbf{A}\mathbf{D}^{-1/2}) \dots (\mathbf{D}^{-1/2}\mathbf{A}\mathbf{D}^{-1/2})}_{t \text{ times}} (\mathbf{D}^{-1/2}\vec{p}^{(0)}).$$

- $\mathbf{D}^{-1/2}\vec{p}^{(t)}$ is exactly what would be obtained by applying $t/2$ iterations of power method to $\mathbf{D}^{-1/2}\vec{p}^{(0)}$!
- Will converge to the top eigenvector of the normalized adjacency matrix $\mathbf{D}^{-1/2}\mathbf{A}\mathbf{D}^{-1/2}$. **Stationary distribution.**
- Like the power method, the time a random walk takes to converge to its stationary distribution (mixing time) is dependent on the gap between the top two eigenvalues of $\mathbf{D}^{-1/2}\mathbf{A}\mathbf{D}^{-1/2}$. The **spectral gap**.

Continuous Optimization and Gradient Descent

Discrete vs. Continuous Optimization

Discrete (Combinatorial) Optimization: (traditional CS algorithms)

- Graph Problems: min-cut, max flow, shortest path, matchings, maximum independent set, traveling salesman problem
- Problems with discrete constraints or outputs: bin-packing, scheduling, sequence alignment, submodular maximization
- Generally searching over a finite but exponentially large set of possible solutions. Many of these problems are NP-Hard.

Continuous Optimization: (maybe seen in ML/advanced algorithms)

- Unconstrained convex and non-convex optimization.
- Linear programming, quadratic programming, semidefinite programming

Continuous Optimization Examples

