

$$\begin{aligned} \sigma_i(X)^2 &= \lambda_i(X^T X) \\ &= \lambda_i(X X^T) \end{aligned}$$

$$\begin{aligned} X &= U \Sigma V^T \\ X^T X &= V \Sigma^2 V^T \\ X X^T &= U \Sigma^2 U^T \end{aligned}$$

COMPSCI 514: ALGORITHMS FOR DATA SCIENCE

Cameron Musco

University of Massachusetts Amherst. Fall 2020.

Lecture 17

LOGISTICS

- ① eigenvalues of $X^T X =$ squared sing. values of X
- ② $=$ singular values of $X^T X$
- ③ any sym. A
singular values of $A =$ absolute values of eigenvals of A .

- Problem Set 3 deadline extended until Monday 10/26, 8pm.
- Week 9 Quiz will now be due Tuesday 10/27, 8pm. ?

symmetric
 $A = W \Lambda W^T$
 eig. val are
 of A (???)
 = 1, 2, 3, 4, 5

symmetric

$$A = W \Lambda W^T$$

where W is orthogonal

$$A = U \Sigma V^T$$

$$U = W \cdot S$$

$$S = \text{sign}(\Lambda)$$

$$\Sigma = |\Lambda|$$

$$V = W$$

$$A^2 = W \Lambda W^T W \Lambda W^T$$

$$= W \Lambda^2 W^T$$

SUMMARY

Last Few Classes: Low-Rank Approximation and PCA

$$A = V \Lambda V^T = V \Sigma \underbrace{S}_{\substack{\text{where } \Sigma = |\Lambda| \\ S = \text{sign}(\Lambda)}} V^T$$

- Compress data that lies close to a k -dimensional subspace.
- Equivalent to finding a low-rank approximation of the data matrix X : $X \approx \underbrace{X V}_{\text{rank } k} V^T$ for orthonormal $V \in \mathbb{R}^{d \times k}$.
- Optimal solution via PCA (eigendecomposition of $X^T X$ or equivalently, SVD of X).
- Singular vectors of X are the eigenvectors of XX^T and $X^T X$. Singular values squared are the eigenvalues.

$$\Lambda \begin{bmatrix} -6 & \\ & 2 & \\ & & 3 \end{bmatrix} \begin{bmatrix} 6 & \\ & 2 & \\ & & 3 \end{bmatrix}$$

2
1

$$X = U \Sigma V^T \quad \begin{matrix} U, V \text{ are orthonormal} \\ \Sigma \text{ is positive diagonal} \end{matrix}$$

$$XX^T = U \Sigma^2 U^T \quad \text{eigencomp.}$$

$$X^T X = V \Sigma^2 V^T$$

SVD of $X^T X$?

if I have any symmetric A with eigendecomposition

SVD of $X^T X$, positive semi-definite matrix
all eigens are +

$$V \Lambda V^T = \text{SVD of } A?$$

Last Few Classes: Low-Rank Approximation and PCA

- Compress data that lies close to a k -dimensional subspace.
- Equivalent to finding a low-rank approximation of the data matrix \mathbf{X} : $\mathbf{X} \approx \mathbf{X}\mathbf{V}\mathbf{V}^T$ for orthonormal $\mathbf{V} \in \mathbb{R}^{d \times k}$.
- Optimal solution via PCA (eigendecomposition of $\mathbf{X}^T\mathbf{X}$ or equivalently, SVD of \mathbf{X}).
- Singular vectors of \mathbf{X} are the eigenvectors of $\mathbf{X}\mathbf{X}^T$ and $\mathbf{X}^T\mathbf{X}$. Singular values squared are the eigenvalues.

This Class: Applications of low-rank approx. beyond compression.

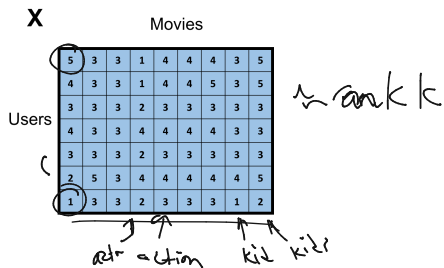
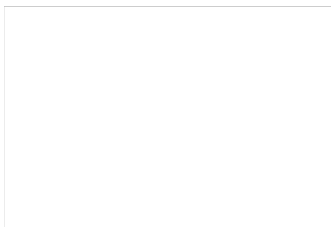
- Matrix completion and collaborative filtering
- Entity embeddings (word embeddings, node embeddings, etc.)
- Low-rank approximation for non-linear dimensionality reduction.
- Spectral graph theory, spectral clustering.

Consider a matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$ which we cannot fully observe but believe is close to rank- k (i.e., well approximated by a rank k matrix).

MATRIX COMPLETION

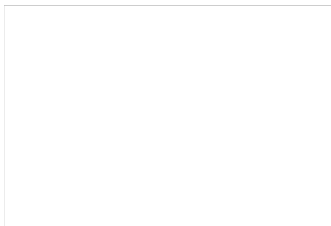
Consider a matrix $X \in \mathbb{R}^{n \times d}$ which we cannot fully observe but believe is close to rank- k (i.e., well approximated by a rank k matrix).
Classic example: the Netflix prize problem.

1, 2, 3, 4, 5



MATRIX COMPLETION

Consider a matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$ which we cannot fully observe but believe is close to rank- k (i.e., well approximated by a rank k matrix).
Classic example: the Netflix prize problem.



X

Users

Movies

5		1	4						
	3					5			
			4						
	5								5
1		2							

MATRIX COMPLETION

Consider a matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$ which we cannot fully observe but believe is close to rank- k (i.e., well approximated by a rank k matrix).
Classic example: the Netflix prize problem.

LRA: solve \mathbf{X} by PCA

$$\hat{\mathbf{Y}} = \underset{\text{rank-}k \mathbf{B}}{\operatorname{argmin}} \|\mathbf{X} - \mathbf{B}\|_F^2 = \sum_{i,j} (x_{i,j} - \hat{b}_{i,j})^2$$

$$\text{Solve: } \mathbf{Y} = \underset{\text{rank-}k \mathbf{B}}{\operatorname{argmin}} \sum_{\text{observed } (j,k)} [\mathbf{X}_{j,k} - \mathbf{B}_{j,k}]^2$$

X

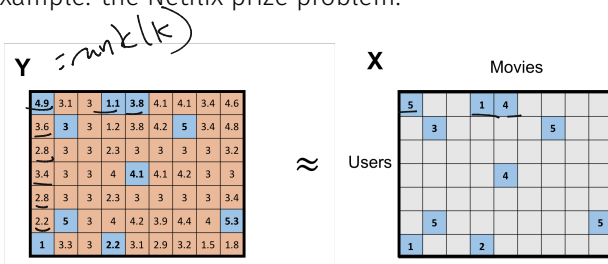
Users

Movies

5		1	4						
	3					5			
			4						
	5							5	
1		2							

MATRIX COMPLETION

Consider a matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$ which we cannot fully observe but believe is close to rank- k (i.e., well approximated by a rank k matrix).
Classic example: the Netflix prize problem.



Solve: $\mathbf{Y} = \arg \min_{\text{rank}-k \mathbf{B}} \sum_{\text{observed } (j,k)} [\mathbf{X}_{j,k} - \mathbf{B}_{j,k}]^2$

MATRIX COMPLETION

Consider a matrix $X \in \mathbb{R}^{n \times d}$ which we cannot fully observe but believe is close to rank- k (i.e., well approximated by a rank k matrix).
 Classic example: the Netflix prize problem.

$$B = \underbrace{[U]_k} \underbrace{[W^T]_k}^k$$

Y

4.9	3.1	3	1.1	3.8	4.1	4.1	3.4	4.6
3.6	3	3	1.2	3.8	4.2	5	3.4	4.8
2.8	3	3	2.3	3	3	3	3	3.2
3.4	3	3	4	4.1	4.1	4.2	3	3
2.8	3	3	2.3	3	3	3	3	3.4
2.2	5	3	4	4.2	3.9	4.4	4	5.3
1	3.3	3	2.2	3.1	2.9	3.2	1.5	1.8

\approx

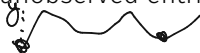
X Movies

Users	1	2	3	4	5	6	7	8	9
1	5		1	4					
2		3				5			
3									
4		1		4					
5							2		
6									
7		5							5
8	1			2					

Solve: $Y = \arg \min_{\text{rank-}k \text{ } B} \sum_{\text{observed } (j,k)} [X_{j,k} - B_{j,k}]^2$

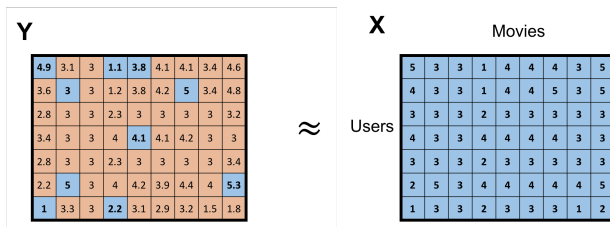
Under certain assumptions, can show that Y well approximates X on both the observed and (most importantly) unobserved entries.

$$y = \arg \min_z f(z)$$



MATRIX COMPLETION

Consider a matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$ which we cannot fully observe but believe is close to rank- k (i.e., well approximated by a rank k matrix). Classic example: the Netflix prize problem.



Solve:
$$\mathbf{Y} = \arg \min_{\text{rank}-k \mathbf{B}} \sum_{\text{observed } (j,h)} [X_{j,h} - B_{j,h}]^2$$

Under certain assumptions, can show that \mathbf{Y} well approximates \mathbf{X} on both the observed and (most importantly) unobserved entries.

Dimensionality reduction embeds d -dimensional vectors into k dimensions. But what about when you want to embed objects other than vectors?

Dimensionality reduction embeds d -dimensional vectors into k dimensions. But what about when you want to embed objects other than vectors?

- Documents (for topic-based search and classification)
- Words (to identify synonyms, translations, etc.)
- Nodes in a social network

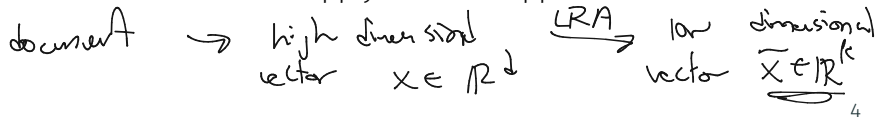
Dimensionality reduction embeds d -dimensional vectors into k dimensions. But what about when you want to embed objects other than vectors?

- Documents (for topic-based search and classification)
- Words (to identify synonyms, translations, etc.)
- Nodes in a social network

here with a neural net



Usual Approach: Convert each item into a high-dimensional feature vector and then apply low-rank approximation.



EXAMPLE: LATENT SEMANTIC ANALYSIS



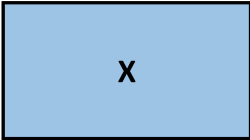
Term Document Matrix X

	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
...	0	0	0	0	0	0	0	1	1
doc_n	1	0	0	0	0	0	0	1	1

bag of words
[] → { the, dog, cat }



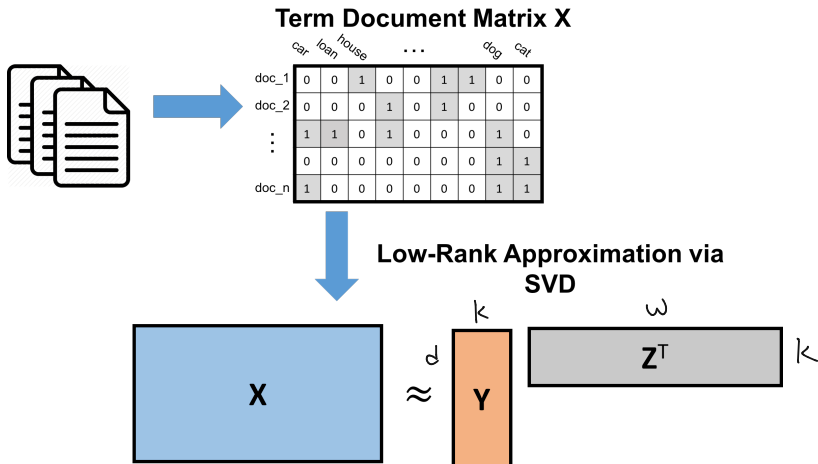
Low-Rank Approximation via SVD



≈



EXAMPLE: LATENT SEMANTIC ANALYSIS



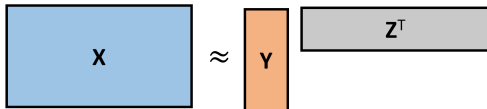
EXAMPLE: LATENT SEMANTIC ANALYSIS

Term Document Matrix X

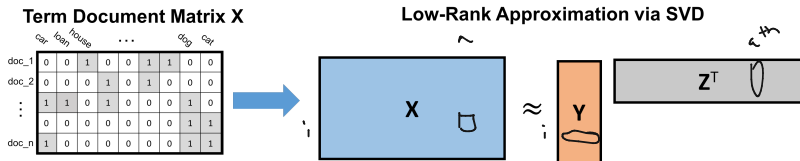
	car	loan	house	...	dog	cat			
doc_1	0	0	1	0	0	1	1	0	0
doc_2	0	0	0	1	0	1	0	0	0
⋮	1	1	0	1	0	0	0	1	0
⋮	0	0	0	0	0	0	0	1	1
doc_n	1	0	0	0	0	0	0	1	1



Low-Rank Approximation via SVD



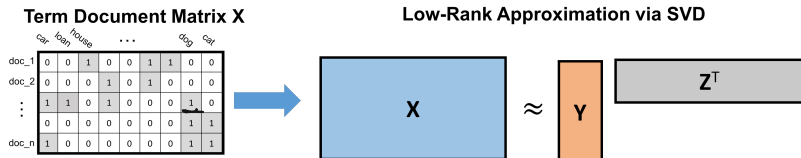
EXAMPLE: LATENT SEMANTIC ANALYSIS



- If the error $\|X - YZ^T\|_F$ is small, then on average,

$$\underline{X_{i,a}} \approx \underline{(YZ^T)_{i,a}} = \langle \vec{y}_i, \vec{z}_a \rangle.$$

EXAMPLE: LATENT SEMANTIC ANALYSIS

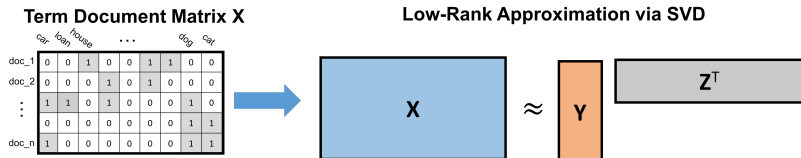


- If the error $\|X - YZ^T\|_F$ is small, then on average,

$$\underline{X_{i,a}} \approx \underline{(YZ^T)_{i,a}} = \langle \underline{\vec{y}_i}, \underline{\vec{z}_a} \rangle.$$

- I.e., $\langle \underline{\vec{y}_i}, \underline{\vec{z}_a} \rangle \approx 1$ when doc_i contains $word_a$.

EXAMPLE: LATENT SEMANTIC ANALYSIS



- If the error $\|X - YZ^T\|_F$ is small, then on average,

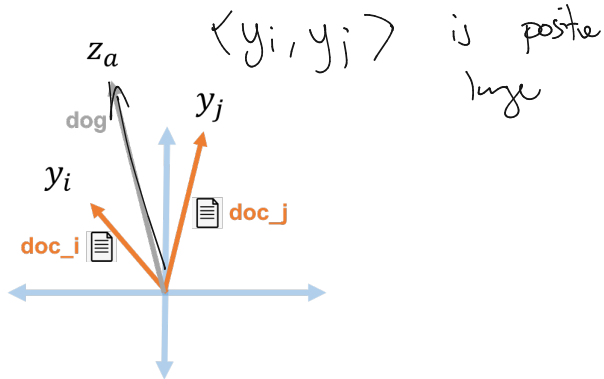
$$X_{i,a} \approx (YZ^T)_{i,a} = \langle \vec{y}_i, \vec{z}_a \rangle.$$

y_i
 y_j

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

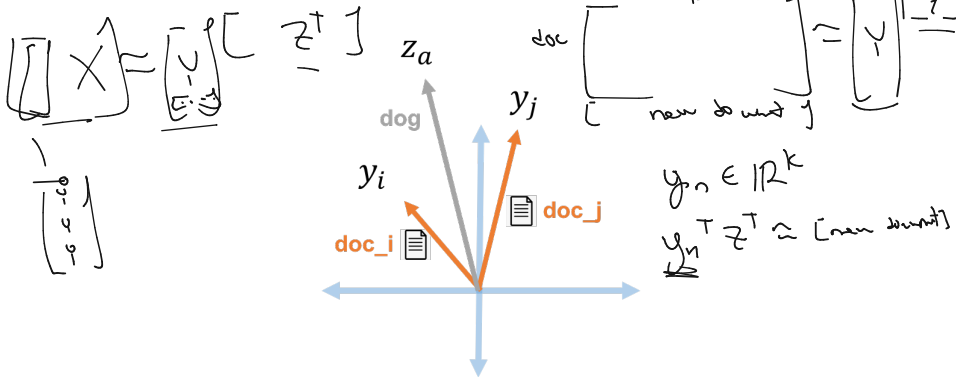
EXAMPLE: LATENT SEMANTIC ANALYSIS

If doc_i and doc_j both contain $word_a$, $\langle \vec{y}_i, \vec{z}_a \rangle \approx \langle \vec{y}_j, \vec{z}_a \rangle \approx 1$



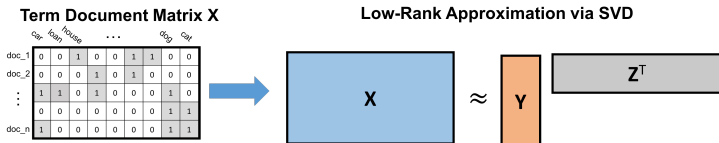
EXAMPLE: LATENT SEMANTIC ANALYSIS

If doc_i and doc_j both contain $word_a$, $\langle \vec{y}_i, \vec{z}_a \rangle \approx \langle \vec{y}_j, \vec{z}_a \rangle \approx 1$



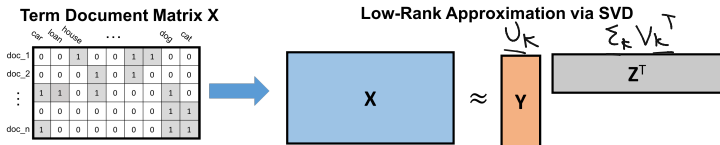
Another View: Each column of Y represents a 'topic'. $\vec{y}_i(j)$ indicates how much doc_i belongs to topic j . $\vec{z}_a(j)$ indicates how much $word_a$ associates with that topic.

EXAMPLE: LATENT SEMANTIC ANALYSIS



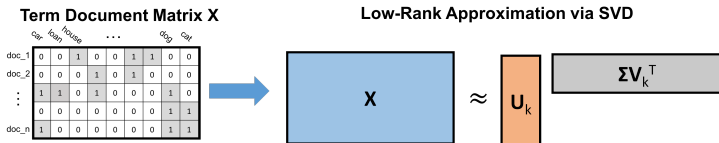
- Just like with documents, \vec{z}_a and \vec{z}_b will tend to have high dot product if $word_a$ and $word_b$ appear in many of the same documents.

EXAMPLE: LATENT SEMANTIC ANALYSIS



- Just like with documents, \vec{z}_a and \vec{z}_b will tend to have high dot product if $word_a$ and $word_b$ appear in many of the same documents.
- In an SVD decomposition we set $Z^T = \sum_k V_k^T$.
- The columns of V_k are equivalently: the top k eigenvectors of $X^T X$.

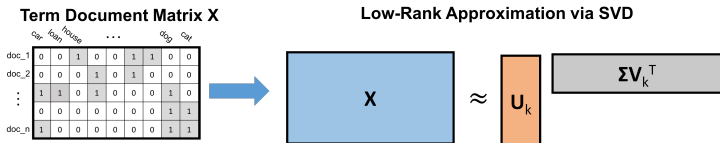
EXAMPLE: LATENT SEMANTIC ANALYSIS



- Just like with documents, \vec{z}_a and \vec{z}_b will tend to have high dot product if $word_a$ and $word_b$ appear in many of the same documents.
- In an SVD decomposition we set $Z^T = \Sigma_k V_k^T$.
- The columns of V_k are equivalently: the top k eigenvectors of $X^T X$.
The eigendecomposition of $X^T X$ is $X^T X = V \Sigma^2 V^T$.

$$X = U \Sigma V^T \quad V \Sigma U^T U \Sigma V^T = U \Sigma^2 V^T$$

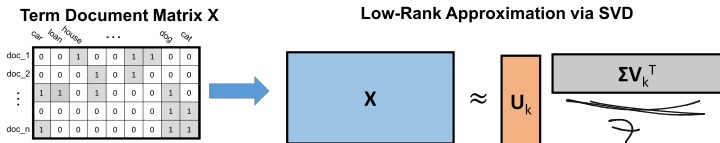
EXAMPLE: LATENT SEMANTIC ANALYSIS



- Just like with documents, \vec{z}_a and \vec{z}_b will tend to have high dot product if $word_a$ and $word_b$ appear in many of the same documents.
- In an SVD decomposition we set $Z^T = \sum_k V_k^T$.
- The columns of V_k are equivalently: the top k eigenvectors of $X^T X$.
The eigendecomposition of $X^T X$ is $X^T X = V \Sigma^2 V^T$ SVD $X^T X$
- **What is the best rank- k approximation of $X^T X$?** i.e.

$$\arg \min_{\text{rank} = k \text{ B}} \|X^T X - B\|_F = V_k \Sigma_k^2 V_k^T = Z Z^T$$

EXAMPLE: LATENT SEMANTIC ANALYSIS



- Just like with documents, \vec{z}_a and \vec{z}_b will tend to have high dot product if $word_a$ and $word_b$ appear in many of the same documents.
- In an SVD decomposition we set $Z^T = \Sigma_k V_k^T$.
- The columns of V_k are equivalently: the top k eigenvectors of $X^T X$. The eigendecomposition of $X^T X$ is $X^T X = V \Sigma^2 V^T$.
- **What is the best rank- k approximation of $X^T X$?** I.e.
$$\arg \min_{\text{rank} - k \text{ B}} \|X^T X - B\|_F$$
- $X^T X = \underbrace{V_k \Sigma_k^2 V_k^T}_{Z Z^T} = Z Z^T$.

EXAMPLE: WORD EMBEDDING

$Z \in \mathbb{R}^{w \times k}$ each row is an embedding of a word
 $Z Z^T \approx X^T X$ $w \times w$ matrix

LSA gives a way of embedding words into k -dimensional space.

- Embedding is via low-rank approximation of $X^T X$: where $(X^T X)_{a,b}$ is the number of documents that both $word_a$ and $word_b$ appear in.

$$0 \cdot 0 + 1 \cdot 1 + 1 \cdot 1 = 2$$

$(X^T X)_{a,b} = \# \text{ docs that } word_a \text{ \& } word_b \text{ both appear in}$

EXAMPLE: WORD EMBEDDING

LSA gives a way of embedding words into k -dimensional space.

- Embedding is via low-rank approximation of $X^T X$: where $(X^T X)_{a,b}$ is the number of documents that both $word_a$ and $word_b$ appear in.
- Think about $X^T X$ as a **similarity matrix** (gram matrix, kernel matrix) with entry (a, b) being the similarity between $word_a$ and $word_b$.

$Z Z^T$ is best LRA to this
similarity matrix.

EXAMPLE: WORD EMBEDDING

LSA gives a way of embedding words into k -dimensional space.

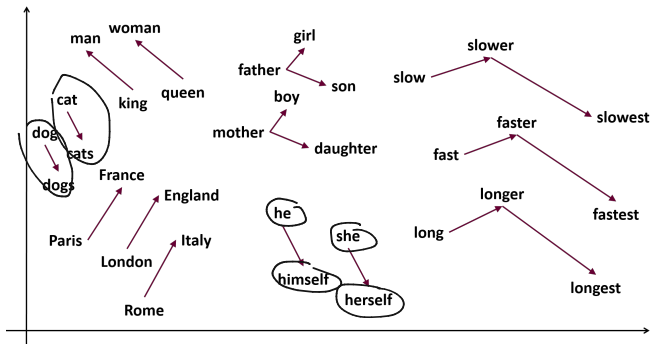
- Embedding is via low-rank approximation of $\mathbf{X}^T\mathbf{X}$: where $(\mathbf{X}^T\mathbf{X})_{a,b}$ is the number of documents that both $word_a$ and $word_b$ appear in.
- Think about $\mathbf{X}^T\mathbf{X}$ as a **similarity matrix** (gram matrix, kernel matrix) with entry (a, b) being the similarity between $word_a$ and $word_b$.
- Many ways to measure similarity: number of sentences both occur in, number of times both appear in the same window of w words, in similar positions of documents in different languages, etc.

EXAMPLE: WORD EMBEDDING

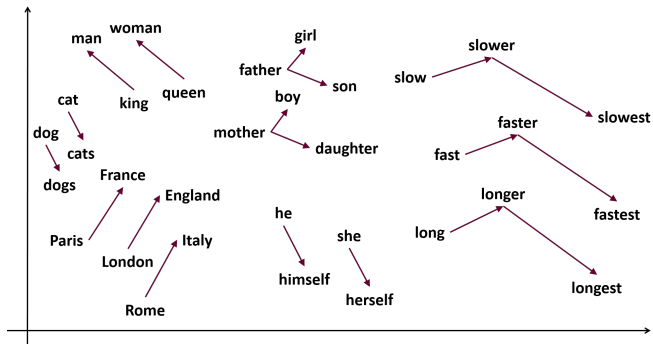
LSA gives a way of embedding words into k -dimensional space.

- Embedding is via low-rank approximation of $\mathbf{X}^T\mathbf{X}$: where $(\mathbf{X}^T\mathbf{X})_{a,b}$ is the number of documents that both $word_a$ and $word_b$ appear in.
- Think about $\mathbf{X}^T\mathbf{X}$ as a **similarity matrix** (gram matrix, kernel matrix) with entry (a, b) being the similarity between $word_a$ and $word_b$.
- Many ways to measure similarity: number of sentences both occur in, number of times both appear in the same window of w words, in similar positions of documents in different languages, etc.
- Replacing $\mathbf{X}^T\mathbf{X}$ with these different metrics (sometimes appropriately transformed) leads to popular word embedding algorithms: word2vec, GloVe, fastText, etc.

EXAMPLE: WORD EMBEDDING



EXAMPLE: WORD EMBEDDING



Note: word2vec is typically described as a neural-network method, but it is really just low-rank approximation of a specific similarity matrix. *Neural word embedding as implicit matrix factorization*, Levy and Goldberg.