




CPS C540



Linear dimensionality reduction



Nando de Freitas
September, 2011
University of British Columbia

Outline

We introduce the Singular Value Decomposition (SVD). This is a matrix factorization that has many applications, including:

- Information retrieval,
- Least-squares problems,
- Image processing,
- Dimensionality reduction.

A video can be treated as a matrix that can be decomposed



(a) Original frames

(b) Low-rank \hat{L}

(c) Sparse \hat{S}



[Candes et al, 2009]

Eigenvalue decomposition

Let \mathbf{A} be an $m \times m$ matrix of reals; that is $\mathbf{A} \in \mathbb{R}^{m \times m}$. If we place the eigenvalues of \mathbf{A} into a diagonal matrix $\mathbf{\Lambda}$ and gather the eigenvectors into a matrix \mathbf{Q} , then the eigenvalue decomposition of \mathbf{A} is given by

$$\mathbf{A} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^{-1}.$$



SVD decomposition

$$\mathbf{A} \in \mathbb{R}^{m \times n}$$

$$\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$$

$\mathbf{\Sigma} \in \mathbb{R}^{n \times n}$ is diagonal with positive entries (singular values in the diagonal).

$\mathbf{U} \in \mathbb{R}^{m \times n}$ has orthonormal columns.

$\mathbf{V} \in \mathbb{R}^{n \times n}$ has orthonormal columns and rows.

That is, \mathbf{V} is an orthogonal matrix, so $\mathbf{V}^{-1} = \mathbf{V}^T$.

SVD

The equations relating the right singular values $\{\mathbf{v}_j\}$ and the left singular vectors $\{\mathbf{u}_j\}$ are

$$\mathbf{A}\mathbf{v}_j = \sigma_j\mathbf{u}_j \quad j = 1, 2, \dots, n$$

$$\mathbf{A}\mathbf{V} = \mathbf{U}\mathbf{\Sigma}$$

$$\mathbf{A} \begin{bmatrix} \mathbf{v}_1 & \mathbf{v}_2 & \dots & \mathbf{v}_n \end{bmatrix} = \begin{bmatrix} \mathbf{u}_1 & \mathbf{u}_2 & \dots & \mathbf{u}_n \end{bmatrix} \begin{bmatrix} \sigma_1 & & & \\ & \sigma_2 & & \\ & & \ddots & \\ & & & \sigma_n \end{bmatrix}$$

SVD properties

1. There is no assumption that $m \geq n$ or that \mathbf{A} has full rank.
2. All diagonal elements of $\mathbf{\Sigma}$ are non-negative and in non-increasing order:

$$\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_p \geq 0$$

where $p = \min(m, n)$

SVD in terms of eigenvalues

Theorem 4 *The nonzero singular values of \mathbf{A} are the (positive) square roots of the nonzero eigenvalues of $\mathbf{A}^T \mathbf{A}$ or $\mathbf{A} \mathbf{A}^T$ (these matrices have the same nonzero eigenvalues).*





Image compression example in python

```
from scipy import *
from pylab import *


img = imread("clown.png")[:, :, 0]
gray()
figure(1)
imshow(img)

m,n = img.shape
U,S,Vt = svd(img)
S = resize(S,[m,1])*eye(m,n)

k = 20
figure(2)
imshow(dot(U[:,1:k],dot(S[1:k,1:k],Vt[1:k,:])))
show()
```



The truncated SVD


 $A = U \Sigma V^T$

$m = 5$
 $n = 4$
 $k = 2$

$A = u_1 \sigma_1 v_1^T + u_2 \sigma_2 v_2^T + u_3 \sigma_3 v_3^T + u_4 \sigma_4 v_4^T$

$A = \begin{bmatrix} \overset{u_1}{x} & \overset{u_2}{x} & \overset{u_3}{x} & \overset{u_4}{x} \\ x & x & x & x \\ x & x & x & x \\ x & x & x & x \\ x & x & x & x \end{bmatrix} \begin{bmatrix} \sigma_1 & 0 & 0 & 0 \\ 0 & \sigma_2 & 0 & 0 \\ 0 & 0 & \sigma_3 & 0 \\ 0 & 0 & 0 & \sigma_4 \end{bmatrix} \begin{bmatrix} x & x & x & x \\ x & x & x & x \\ x & x & x & x \\ x & x & x & x \end{bmatrix} \begin{matrix} v_1^T \\ v_2^T \\ v_3^T \\ v_4^T \end{matrix}$

$A_k = u_1 \sigma_1 v_1^T + u_2 \sigma_2 v_2^T ; k=2$



Smaller eigenvectors capture high frequency variations (small brush-strokes).

Image compression example

The code:

- loads a clown image into a 200 by 320 array A,
- displays the image in one figure,
- performs a singular value decomposition on A,
- displays the image obtained from a rank-20 SVD approximation of A in another figure.



The original storage requirements for A are:

The compressed representation requires:

Text retrieval:

Latent semantic indexing (LSI)

The SVD can be used to cluster documents and carry out information retrieval by using concepts as opposed to exact word-matching.

This enables us to surmount the problems of synonymy (car, auto) and polysemy (money bank, river bank).



The data is available in a term-frequency (TF) matrix:

$$A = \begin{matrix} & \begin{matrix} \text{doc1} & \text{doc2} & \text{doc3} \end{matrix} \\ \begin{matrix} \text{word1} \\ \text{word2} \\ \vdots \end{matrix} & \begin{bmatrix} 1 & 3 & 2 \\ 0 & 1 & 0 \end{bmatrix} \end{matrix}$$

a_{ij} = # times word i appears in document j

LSI example

$$m = 5$$

$$n = 3$$



$$d_1 = \{\text{I, eat, chips}\}$$

$$d_2 = \{\text{computer, chips, chips}\}$$

$$d_3 = \{\text{intel, computer, chips}\}$$

$$A = \begin{matrix} & d_1 & d_2 & d_3 \\ \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 1 & 2 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix} & \begin{matrix} \text{I} \\ \text{eat} \\ \text{chips} \\ \text{computer} \\ \text{intel} \end{matrix} \end{matrix}$$

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

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

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

$$\Sigma^{-1} U^T A = \Sigma^{-1} \Sigma V^T$$

$$V^T = \Sigma^{-1} U^T A$$

Truncated SVD for LSI

If we truncate the approximation to the k -largest singular values, we have


$$A = U_k \Sigma_k V_k^T$$

So

$$V_k^T = \Sigma_k^{-1} U_k^T A$$

In English, A is projected to a lower-dimensional space spanned by the k singular vectors U_k (eigenvectors of AA^T).

Part I: Building the search engine




Part II: Querying the search engine

To carry out **retrieval**, a **query** $\mathbf{q} \in \mathbb{R}^n$ is first projected to the low-dimensional space:


$$\hat{\mathbf{q}}_k = \Sigma_k^{-1} \mathbf{U}_k^T \mathbf{q}$$

2×1 2×5 5×1

And then we measure the angle between $\hat{\mathbf{q}}_k$ and the \mathbf{v}_k .



Part II: Querying the search engine



TF-IDF

The *term frequency* of word w in document d is equal to the total number of times w appears in d , divided by the total number of words in d .

The *inverse document frequency* of word w is based on the logarithm of the inverse frequency of the word in the corpus. If the number of documents in the corpus is D and the number of documents the word appears in is D_w , then

$$idf_w = \log \frac{D}{1 + D_w}$$

and for w and d , we can compute

$$tfidf_{w,d} = tf_{w,d} idf_w$$