

Computational Methods in Statistics with Applications

Singular Value Decomposition

Maya Neytcheva, Lars Eldén

SeSe, September 2020

Matrix factorizations/decompositions

1/58



2/58



Matrix factorizations I

- ▶ LU , LDU , Cholesky LDL^T
- ▶ Tridiagonalization $Q^T A Q = T$, A - symmetric, T - tridiagonal
Aasen's algorithm: $A = L T L^T$
- ▶ Bidiagonalization $Q^T A V = B$, $A(m, n)$, B - upper bi-diagonal
- ▶ ...
- ▶ QR

Matrix factorizations

- **Eigenvalue decomposition**

A - square, normal matrix. If all its eigenvectors are linearly independent, then $A = Q D Q^T$, where Q is orthogonal and D is diagonal, containing the eigenvalues of A .

Golub, Van Loan, *Matrix Computations*, many editions.

Note: Some of the algorithms are not numerically stable.



Matrix factorizations

- **Eigenvalue decomposition**

A - square, normal matrix. If all its eigenvectors are linearly independent, then $A = QDQ^T$, where Q is orthogonal and D is diagonal, containing the eigenvalues of A .

- **Schur decomposition**: $A = QTQ^*$

Any square matrix A is unitary similar to an upper triangular matrix T , which has the eigenvalues of A on its diagonal.

Note: Eigenvalue-revealing factorization

Matrix factorizations

- **Eigenvalue decomposition**

A - square, normal matrix. If all its eigenvectors are linearly independent, then $A = QDQ^T$, where Q is orthogonal and D is diagonal, containing the eigenvalues of A .

- **Schur decomposition**: $A = QTQ^*$

Any square matrix A is unitary similar to an upper triangular matrix T , which has the eigenvalues of A on its diagonal.

Note: Eigenvalue-revealing factorization

- Question: Can we diagonalize a general $(m \times n)$ matrix using unitary matrices?

Singular value decomposition SVD

$$Q_1 A Q_2^* = \Sigma$$

4/58



4/58



SVD

Singular value decomposition

Let $A(m, n)$, $n \leq m$ or $n \geq m$, $\text{rank}(A) = \text{rank}(A^*)$.

Definition

If there exist $\mu \neq 0$ and vectors \mathbf{u} and \mathbf{v} , such that

$$A\mathbf{v} = \mu\mathbf{u} \quad \text{and} \quad A^*\mathbf{u} = \mu\mathbf{v}$$

then μ is called a singular value of A , and \mathbf{u}, \mathbf{v} are a pair of singular vectors, corresponding to μ .

5/58

6/58



Historical notes

SVD has many different names:

- ▶ First derivation of the SVD by Eugenio Beltrami (1873)
- ▶ Full proof by Camille Jordan (1874)
- ▶ James Joseph Sylvester (1889), independently discovers SVD
- ▶ Erhard Schmidt (1907), first to derive an optimal, low-rank approximation of a larger problem
- ▶ Hermann Weyl (1912) - determination of the rank in the presence of errors
- ▶ Eckart-Young decomposition and optimality properties of SVD (1936), psychometrics
- ▶ Numerically efficient algorithms to compute the SVD - works by Gene Golub 1970 (Golub-Kahan)

The existence of singular values and vectors ...

... is shown via the following construction:

$$A\mathbf{v} = \mu\mathbf{u}, \quad A^*\mathbf{u} = \mu\mathbf{v}$$

can be written as

$$\tilde{A} \begin{bmatrix} \mathbf{v} \\ \mathbf{u} \end{bmatrix} = \begin{bmatrix} 0 & A^* \\ A & 0 \end{bmatrix} \begin{bmatrix} \mathbf{v} \\ \mathbf{u} \end{bmatrix} = \mu \begin{bmatrix} \mathbf{v} \\ \mathbf{u} \end{bmatrix}$$

The matrix \tilde{A} is selfadjoint, has real eigenvalues and a complete eigenvector space.

Furthermore, μ^2 is an eigenvalue of A^*A with eigenvector \mathbf{u} and of AA^* with eigenvector \mathbf{v} , because

$$\begin{aligned} A\mathbf{v} = \mu\mathbf{u}, & \quad \rightarrow \quad A^*A\mathbf{v} = \mu A^*\mathbf{u} = \mu^2\mathbf{v} \\ A^*\mathbf{u} = \mu\mathbf{v}, & \quad \rightarrow \quad AA^*\mathbf{u} = \mu A\mathbf{v} = \mu^2\mathbf{u} \end{aligned}$$

7/58



8/58



Singular Value Decomposition

SVD, $A(m \times n)$



Theorem (SVD)

Any $m \times n$ matrix A with dimensions, say, $m \geq n$, can be factorized as

$$A = U \begin{bmatrix} \Sigma \\ 0 \end{bmatrix} V^*,$$

where $U \in \mathbb{R}^{m \times m}$ and $V \in \mathbb{R}^{n \times n}$ are orthogonal, and $\Sigma \in \mathbb{R}^{m \times n}$ is diagonal,

$$\begin{aligned} \Sigma &= \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_n), \\ \sigma_1 &\geq \sigma_2 \geq \dots \geq \sigma_n \geq 0. \end{aligned}$$




9/58

10/58




SVD, observation

Thin SVD



Partition $U = (U_1 \ U_2)$, where $U_1 \in \mathbb{R}^{m \times n}$,

$$A = U_1 \Sigma V^T,$$




11/58

12/58



Thin SVD

Fundamental Subspaces I




The *range* of the matrix A :

$$\mathcal{R}(A) = \{y \mid y = Ax, \text{ for arbitrary } x\}.$$

Assume that A has rank r :

$$\sigma_1 \geq \dots \geq \sigma_r > \sigma_{r+1} = \dots = \sigma_n = 0.$$

Outer product form:



$$y = Ax = \sum_{i=1}^r \sigma_i u_i v_i^T x = \sum_{i=1}^r (\sigma_i v_i^T x) u_i = \sum_{i=1}^r \alpha_i u_i.$$

13/58

14/58



Fundamental Subspaces II

The null-space of the matrix A :

$$\mathcal{N}(A) = \{x \mid Ax = 0\}.$$

$$Ax = \sum_{i=1}^r \sigma_i u_i v_i^T x$$

Any vector $z = \sum_{i=r+1}^n \beta_i v_i$ is in the null-space:

$$Az = \left(\sum_{i=1}^r \sigma_i u_i v_i^T \right) \left(\sum_{i=r+1}^n \beta_i v_i \right) = 0.$$

15/58



Fundamental Subspaces

Theorem (Fundamental subspaces)

1. The singular vectors u_1, u_2, \dots, u_r are an orthonormal basis in $\mathcal{R}(A)$ and

$$\text{rank}(A) = \dim(\mathcal{R}(A)) = r.$$

2. The singular vectors $v_{r+1}, v_{r+2}, \dots, v_n$ are an orthonormal basis in $\mathcal{N}(A)$ and

$$\dim(\mathcal{N}(A)) = n - r.$$

3. The singular vectors v_1, v_2, \dots, v_r are an orthonormal basis in $\mathcal{R}(A^T)$.
4. The singular vectors $u_{r+1}, u_{r+2}, \dots, u_m$ are an orthonormal basis in $\mathcal{N}(A^T)$.

16/58



SVD matrix expansion

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

$$A = \sum_{i=1}^n \sigma_i u_i v_i^T = \begin{array}{|c} \text{---} \\ | \\ \text{---} \end{array} + \begin{array}{|c} \text{---} \\ | \\ \text{---} \end{array} + \dots$$

A is represented as a sum of rank-one matrices.

17/58

Rank deficient matrix I

```
>> A(:,3) = A(:,1) + 0.5 * A(:,2)
```

```
A = 1.0000    1.0000    1.5000
     1.0000    2.0000    2.0000
     1.0000    3.0000    2.5000
     1.0000    4.0000    3.0000
```

```
>> [U,S,V] = svd(A,0)
```

```
U = 0.2612   -0.7948   -0.5000
     0.4032   -0.3708    0.8333
     0.5451    0.0533   -0.1667
     0.6871    0.4774   -0.1667
```

18/58



$$S = \begin{bmatrix} 7.3944 & 0 & 0 \\ 0 & 0.9072 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

$$V = \begin{bmatrix} 0.2565 & -0.6998 & 0.6667 \\ 0.7372 & 0.5877 & 0.3333 \\ 0.6251 & -0.4060 & -0.6667 \end{bmatrix}$$

SVD is rank-revealing!

2-norm, Frobenius norm

$$\|A\|_2 = \sup_{x \neq 0} \frac{\|Ax\|}{\|x\|}, \quad \|A\|_F = \left(\sum_{i,j} a_{ij}^2 \right)^{1/2}$$

One can show that $\|A\|_2 = \left(\max_{1 \leq i \leq n} \lambda_i(A^T A) \right)^{1/2} = \sigma_1$

19/58



20/58



Theorem

Assume that the matrix $A \in \mathbb{R}^{m \times n}$ has rank r and choose k , such that $r > k$. The Frobenius norm matrix approximation problem

$$\min_{\text{rank}(Z)=k} \|A - Z\|_F$$

has the solution

$$Z = A_k = U_k \Sigma_k V_k^T,$$

where $U_k = (u_1, \dots, u_k)$, $V_k = (v_1, \dots, v_k)$, and $\Sigma_k = \text{diag}(\sigma_1, \dots, \sigma_k)$.

Theorem

Let $A \in \mathbb{C}^{m \times n}$ have rank r and SVD $A = U \Sigma V^* = \sum_{i=1}^r \sigma_i u_i v_i^*$.

The solution of the problem

$$\min_X \|A - X\|_2$$

where X belongs to the set of matrices in $\mathbb{C}^{m \times n}$ of rank $k < r$, is

obtained for the matrix $X = A_k = \sum_{j=1}^k \sigma_j u_j v_j^*$ and

$$\|A - A_k\|_2 = \sigma_{k+1}.$$

21/58

22/58



Best approximation / Eckart-Young Property, 1936 IV

Best approximation / Eckart-Young Property, 1936 V

Proof:

(1) Observe: $U^T A_k V = \text{diag}\{\sigma_1, \sigma_2, \dots, \sigma_k, 0, \dots, 0\}$, e.g.

$\text{rank}(A_k) = k$.

$U^T (A - A_k) V = \text{diag}\{0, \dots, 0, \sigma_{k+1}, \dots, \sigma_r\}$, then

$\|A - A_k\|_2 = \sigma_{k+1}$.

(2) Assume that there exists a matrix B of rank k , such that

$\|A - B\|_2 < \sigma_{k+1}$.

$\text{rank}(B) = k$, then there exists a basis x_1, \dots, x_{n-k} such that

$(B) = \text{span}\{x_1, \dots, x_{n-k}\}$.

From dimension arguments

$\text{span}\{x_1, \dots, x_{n-k}\} \cap \text{span}\{v_1, \dots, v_{k+1}\} \neq \emptyset$.

Let z belong to the intersection. Then

$Az = \sum_{i=1}^{k+1} \sigma_i (v_1^T z) u_i$. Thus,

$$\|A - B\|_2^2 \geq \|(A - B)z\|_2^2 = \|Az\|_2^2 = \sum_{i=1}^{k+1} \sigma_i^2 (v_1^T z)^2 \geq \sigma_{k+1}^2.$$

The latter contradicts to the assumption.

23/58





24/58



Image compression, example

SVD, geometric view

Matlab demo



$$A = U \Sigma V^* \quad AV = U \Sigma$$

25/58

26/58



Example

Test example borrowed from

Computational Statistics with Application to Bioinformatics

Prof. William H. Press Spring Term, 2008, The University of Texas at Austin

Consider some **gene expression** data, represented by the so-called 'design matrix' $X = \{X_{ij}\}$

Each column of X corresponds to a separate observation, in this case, a separate micro array experiment under a different condition. N rows are genes (1:500) and M columns are the corresponding responses.

Assumptions:

- the individual experiments (columns of X) have zero mean.
- scale data to unit standard deviation.

27/58



28/58




```

load yeastarray_t2.txt;
size(yeastarray_t2)

ans = 500 300
yclip = prctile(yeastarray_t2(:), [1,99])
yclip = [-204 244]

data = max(yclip(1), min(yclip(2), yeastarray_t2));
dmean= mean(data,1);
dstd = std(data,1);
data = (data - repmat(dmean, [size(data,1), 1])) ./ ...
       repmat(dstd, [size(data,1), 1]);
genecolormap = [min(1, (1:64)/32); 1-abs(1-(1:64)/32);
               min(1, (64-(1:64))/32)]';
figure(1), clf, colormap(genecolormap);
image(20*data+32)

```




29/58

30/58



The squares of the singular values are proportional to the portion of the total variance (L_2 norm of X) that each accounts for.

```
ssq = diag(S).^2;
semilogy(ssq, 'b')
```




31/58



We can produce fake data and compare:

```
fakedata = randn(500,300);
[Uf Sf Vf] = svd(fakedata,0);
sfsq = diag(Sf).^2;
semilogy(sfsq, 'r')
```




32/58




For the data in this example, a sensible use of PCA (i.e., SVD) would be to project the data into the subspace of the first 20 SVs, where we can be sure that it is not noise.

```
% Truncate the first 20 singular values/vectors
strunc = diag(S);
strunc(21:end) = 0;
filtdata20 = U*diag(strunc)*V';
figure(2), clf, colormap(genecolormap);
image(20*filtdata20+32)
```


```
% Truncate the first 5 singular values/vectors
strunc(6:end) = 0;
filtdata5 = U*diag(strunc)*V';
figure(3), clf, colormap(genecolormap);
image(20*filtdata5+32)
```



(a) The original



(b) truncate to 20




33/58

34/58



How to interpret the singular vectors? The first three vectors u are 'eigengenes', the linear combination of genes that explain the most data.

```
plot(U(:,1:3))
```




35/58



The first three vectors v are 'eigenarrays', the linear combination of experiments that explain the most data.

```
plot(V(:,1:3))
```




36/58



Consider a toy example

```
pdata = randn(500,300);
pdata(101:200,51:100) = pdata(101:200,51:100) + 1;
pdata(301:400,201:250) = pdata(301:400,201:250) - 1;
pmean = mean(pdata,1);
pstd = std(pdata,1);
pdata = (pdata - repmat(pmean, [size(pdata,1),1]))./...
        repmat(pstd, [size(pdata,1),1]);
colormap(genecolormap)
image(20*pdata+32)
```



37/58


38/58



Consider a toy example

```
[Up Sp Vp] = svd(pdata,0);
spsq = diag(Sp).^2;
semilogy(spsq(1:50),'.b')
```

Should we expect the eigengenes/eigenarrays to show the separate main effects?




39/58




40/58



```
plot(Up(:,1:2)),
plot(Up(:,1:3))
```




(d)




(e)

```
plot(Vp(:,1:2)),
plot(Vp(:,1:3))
```



(f)



(g)

41/58

42/58



Linear dependence – SVD

Linear dependence I

Theorem

Let the singular values of A satisfy

$$\sigma_1 \geq \dots \geq \sigma_r > \sigma_{r+1} = \dots = \sigma_n = 0.$$

Then the rank of A is equal to r .

Rank = the number of linearly independent columns of A .

```
A=[1 1; 1 2; 1 3; 1 4]
singval=svd(A)
```

```
% Third col=linear combination of first two
A1=[A A(:,1)+0.5*A(:,2)]
singval1=svd(A1)
```

43/58



44/58



Linear dependence II

Almost linear dependence I

Result:

```
A = 1 1
     1 2
     1 3
     1 4
singval = 5.7794
         0.7738
```

```
A1 = 1.0000 1.0000 1.5000
     1.0000 2.0000 2.0000
     1.0000 3.0000 2.5000
     1.0000 4.0000 3.0000
```

```
singval1 = 7.3944
          0.9072
          0
```

```
A2=[A A(:,1)+0.5*A(:,2)+0.0001*randn(4,1)]
singval2=svd(A2)
```

```
-----
A2 = 1.0000 1.0000 1.4999
     1.0000 2.0000 2.0001
     1.0000 3.0000 2.5000
     1.0000 4.0000 3.0001
```

```
singval2 = 7.3944
          0.9072
          0.0001
```

45/58

46/58



Run Matlab demo

```
~/.../STAT/Labs/Lab_QR_SVD/Small_singular_values.m
```

▶ Matlab:

```
[U, S, V] = svd(A); [U1, S, V]=svd(A, 0);
```

▶ R:

```
(s <- svd(A))
s <- La.svd(A)
```

▶ python:

```
U, S, V = np.linalg.svd(A)
```

47/58



48/58



Golub-Kahan (Golub-Kahan-Lanczos) bidiagonalization

Computing the SVD in a numerically efficient way

$$A = \begin{bmatrix} x & x & x & x \\ x & x & x & x \\ x & x & x & x \\ x & x & x & x \\ x & x & x & x \end{bmatrix} \rightarrow U_1^* A = \begin{bmatrix} x & x & x & x \\ 0 & x & x & x \\ 0 & x & x & x \\ 0 & x & x & x \\ 0 & x & x & x \end{bmatrix} \rightarrow U_1^* A V_1 = \begin{bmatrix} x & x & 0 & 0 \\ 0 & x & x & x \\ 0 & x & x & x \\ 0 & x & x & x \\ 0 & x & x & x \end{bmatrix}$$



Golub-Kahan (Golub-Kahan-Lanczos) bidiagonalization

$$U_4^* U_3^* U_2^* U_1^* A V_1 V_2 = \begin{bmatrix} x & x & 0 & 0 \\ 0 & x & x & 0 \\ 0 & 0 & x & x \\ 0 & 0 & 0 & x \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

Thus, $U^*AV = B$, B - bidiagonal.

51/58



Golub-Kahan (Golub-Kahan-Lanczos) bidiagonalization

$$B = \begin{bmatrix} 1 & 2 & 0 & 0 \\ 0 & 3 & 4 & 0 \\ 0 & 0 & 5 & 6 \\ 0 & 0 & 0 & 7 \end{bmatrix}$$

$$H = \begin{bmatrix} 0 & B \\ B' & 0 \end{bmatrix}, \quad p = [5 \ 1 \ 6 \ 2 \ 7 \ 3 \ 8 \ 4]$$

$$H(p, p) = \begin{bmatrix} 1 & & & & & & & \\ & 1 & & & & & & \\ & & 2 & & & & & \\ & & & 3 & & & & \\ & & & & 3 & & & \\ & & & & & 4 & & \\ & & & & & & 5 & \\ & & & & & & & 6 \\ & & & & & & & & 6 & \\ & & & & & & & & & 7 & \\ & & & & & & & & & & 7 \end{bmatrix}$$

52/58



- ▶ How much does it cost to compute the SVD?
The cost for the bidiagonalization is $4mn^2 - 4/3n^3$.
The cost for SVD: $4m^2n + 8mn^2 + 9n^3$.

- ▶ How much does it cost to compute the SVD?
The cost for the bidiagonalization is $4mn^2 - 4/3n^3$.
The cost for SVD: $4m^2n + 8mn^2 + 9n^3$.
- ▶ How much does it cost to LU-factorize a full matrix?



- ▶ How much does it cost to compute the SVD?
The cost for the bidiagonalization is $4mn^2 - 4/3n^3$.
The cost for SVD: $4m^2n + 8mn^2 + 9n^3$.
- ▶ How much does it cost to LU-factorize a full matrix?
- ▶ And how much does it take to LU-factorize a sparse matrix?

Applications of SVD

53/58



Applications of SVD I

- ▶ Data compression
- ▶ Estimation/inversion
- ▶ Sensitivity of linear equations to data errors
The condition number of A is defined as $\kappa(A) = \sigma_{max}/\sigma_{min}$.
- ▶ Low rank approximation
- ▶ Image processing
- ▶ Distance to singularity Another interpretation of σ_i :

$$\sigma_i = \min\{\|A - B\|, s.t. rank(B) \leq i - 1\},$$

i.e., σ_i is the distance, measured by the matrix norm, to the nearest rank $i - 1$ matrix.

55/58

54/58



Applications of SVD II

- ▶ General pseudoinverse Let $A = U\Sigma V^T$. Then

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

is the so-called Moore-Penrose inverse of A .
Let $A(m, n)$, $m > n$. Then $A^\dagger = (A^T A)^{-1} A^T$ gives the least square solution of $Ax = y$, $x = A^\dagger y$.
Let $m < n$. Then $A^\dagger = A^T (A A^T)^{-1}$ gives the least-norm solution of $Ax = y$

56/58



"Search engines like Google use enormous matrices of cross-references, which pages link to which other pages, and what words are on each page. When you do a Google search, the higher ranks usually go to pages with your key words that have lots of links to them. But there are billions of pages out there, and storing a billion by billion matrix is trouble, not to mention searching through it.

In searching, one only cares about the main directions that the Web is taking. So the first few singular values create a very good approximation for the enormous matrix, can be searched relatively quickly (just a few billion entries) and provide compression ratios of millions to one. "

$$\begin{aligned}\|Ax - b\|_2^2 &= \|U^T(AVV^T x - b)\|_2^2 = \|\Sigma(V^T x) - U^T b\|_2^2 \\ &= \sum_{i=1}^r (\sigma_i z_i - u_i^T b)^2 + \sum_{i=r+1}^m (u_i^T b)^2.\end{aligned}$$

Thus,

$$\min \|Ax - b\|_2^2 = \underbrace{\sum_{i=1}^r (\sigma_i z_i - u_i^T b)^2}_{\text{set to 0}} + \underbrace{\sum_{i=r+1}^m (u_i^T b)^2}_{\text{does not depend on } x}$$