Lecture 3
The Singular Value Decomposition

MIT 18.335J / 6.337J

Introduction to Numerical Methods

Per-Olof Persson

September 14, 2006

The SVD - The Main Idea

- Motivation:

  The image of the unit sphere under any $m \times n$ matrix is a hyperellipse

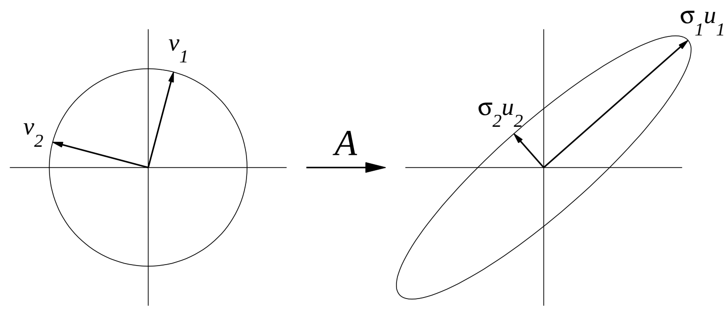

The SVD - Brief Description

- Suppose (for the moment) that A is $m \times n$ with $m \geq n$ and full rank n
- Choose orthonormal bases

$$v_1, \ldots, v_n \text{ for the row space}$$
$$u_1, \ldots, u_n \text{ for the column space}$$

such that Av_i is in the direction of u_i:

$$Av_i = \sigma_i u_i$$

- In MATLAB: `eigshow`
- The *singular values* $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_n > 0$

The SVD - Brief Description

- In Matrix form, $Av_i = \sigma_i u_i$ becomes

$$AV = \hat{U}\hat{\Sigma}, \text{ that is, } \boxed{A = \hat{U}\hat{\Sigma}V^*}$$

where $\hat{\Sigma} = \mathrm{diag}(\sigma_1, \sigma_2, \ldots, \sigma_n)$

- This is the *reduced singular value decomposition*
- Add orthogonal extension to $\hat{U}$ and add rows to $\hat{\Sigma}$ to obtain the *full singular value decomposition*

$$\boxed{A = U\Sigma V^*}$$

The Full Singular Value Decomposition

- Let A be an $m \times n$ matrix. The *singular value decomposition* of A is the factorization $A = U\Sigma V^*$ where

U is $m \times m$ unitary (the left singular vectors of A)

V is $n \times n$ unitary (the right singular vectors of A)

Σ is $m \times n$ diagonal (the singular values of A)

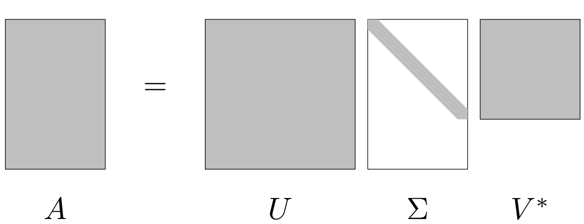

The Reduced Singular Value Decomposition

- A more compact representation is the *Reduced SVD*, for $m \geq n$:

$$A = \hat{U}\hat{\Sigma}V^*$$

where

$$\hat{U} \text{ is } m \times n, \quad V \text{ is } n \times n, \text{ and } \hat{\Sigma} \text{ is } n \times n$$

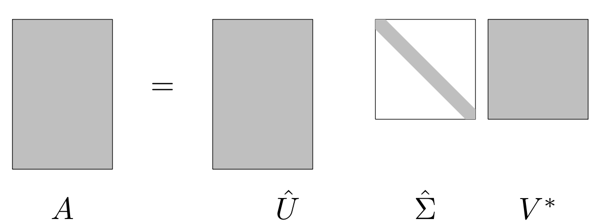

The SVD and The Eigenvalue Decomposition

- The *eigenvalue decomposition* $A = X\Lambda X^{-1}$
  - uses the same basis X for row and column space, but the SVD uses two different bases V, U
  - generally does not use an orthonormal basis, but the SVD does
  - is only defined for square matrices, but the SVD exists for all matrices
- For *symmetric positive definite* matrices A, the eigenvalue decomposition and the SVD are equal

Matrix Properties

1. The rank of A is r, the number of nonzero singular values
2. $\mathrm{range(A)} = \langle u_1, \dots, u_r \rangle$ and $\mathrm{null(A)} = \langle v_{r+1}, \dots, v_n \rangle$
3. $\|A\|_2 = \sigma_1$ and $\|A\|_F = \sqrt{\sigma_1^2 + \sigma_2^2 + \ldots + \sigma_r^2}$
4. Nonzero eigenvalues of A^*A are nonzero σ_i^2, eigenvectors are v_i
   Nonzero eigenvalues of AA^* are nonzero σ_i^2, eigenvectors are u_i
5. If $A = A^*$, $\sigma_i = |\lambda_i|$ where λ_i are eigenvalues of A
6. For square A, $|\det(A)| = \prod_{i=1}^m \sigma_i$

Existence and Uniqueness

- Every matrix has a singular value decomposition
- The singular values σ_j are uniquely determined
- If A square and σ_j distinct, left/right singular vectors u_j, v_j are uniquely determined up to complex signs
- *Proof.* Textbook / Black board

Low-Rank Approximations

- The SVD can be written as a sum of rank-one matrices

$$A = \sum_{j=1}^{r} \sigma_j u_j v_j^*$$

- The best rank ν approximation of A in the 2-norm is

$$A_\nu = \sum_{j=1}^{\nu} \sigma_j u_j v_j^*$$

  with $\|A - A_\nu\|_2 = \sigma_{\nu+1}$
- Also true in the Frobenius norm, with $\|A - A_\nu\|_2 = \sqrt{\sigma_{\nu+1}^2 + \cdots + \sigma_r^2}$

Applications of the SVD

- Calculation of matrix properties:
  - Rank of matrix (counting σ_j's $>$ tolerance)
  - Bases for range and nullspace (in U and V)
  - Induced matrix norm $\|\cdot\|_2$ ($= \sigma_1$)
- Low-rank approximations (optimal in $\|\cdot\|_2$ and $\|\cdot\|_F$)
- Least squares fitting (more later, another option is QR)
- Signal and image processing
  - Compression (see next slide)
  - Noise removal (noise tends to have low σ_j)

Application: Image Compression

- View $m \times n$ image as a (real) matrix A, find best rank ν approx. by SVD
- Storage $\nu(m+n)$ instead of mn