

Chapter 1

SVD, PCA & Pre- processing

Part 3: Interpreting and computing the SVD



Interpreting SVD

Factor interpretation

- Let \mathbf{A} be objects-by-attributes and $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$ its SVD
- If two columns have similar values in a row of \mathbf{V}^T , these attributes are similar (have strong correlation)
- If two rows have similar values in a column of \mathbf{U} , these objects are similar

Example

- Data: people's ratings on different wines
- Scatterplot of first two LSV
 - SVD doesn't know what the data is
- Conclusion: winelovers like red and white alike, others are more biased





Figure 3.2. The first two factors for a dataset ranking wines.


Geometric interpretation

- Let $\mathbf{M} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$
- Any linear mapping $\mathbf{y}=\mathbf{M}\mathbf{x}$ can be expressed as a rotation, stretching, and rotation operation
 - $\mathbf{y}_1 = \mathbf{V}^T\mathbf{x}$ is the first rotation
 - $\mathbf{y}_2 = \mathbf{\Sigma}\mathbf{y}_1$ is the stretching
 - $\mathbf{y} = \mathbf{U}\mathbf{y}_2$ is the final rotation



Direction of largest variances


- The singular vectors give the directions of the largest variances
 - First singular vector points to the direction of the largest variance
 - Second to the second-largest
 - Spans a hyperplane with the first
- The projection distance to these hyperplanes is minimal over all hyperplanes (Eckart–Young)



Component interpretation

- We can write $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T = \sum_i \sigma_i \mathbf{u}_i \mathbf{v}_i^T = \sum_i \mathbf{A}_i$
- This explains the data as a sum of rank-1 layers
 - First layer explains the most, the second updates that, the third updates that, ...
- Each individual layer don't have to be very intuitive


Example



Applications of SVD

Removing noise

- SVD is often used as a pre-processing step to remove noise from the data
- The rank- k truncated SVD with proper k



$$\sigma_1 = 11.73$$

$$\sigma_2 = 1.71$$

Removing dimensions

- SVD can be used to project the data to smaller-dimensional subspace
- Original dimensions can have complex correlations
- Subsequent analysis is faster
- Points seem close to each other in high-dimensional space


Curse of dimensionality

Karhunen–Loève transform

- The **Karhunen–Loève transform** (KLT) works as follows:
 - Normalize $\mathbf{A} \in \mathbb{R}^{n \times m}$ to z-scores
 - Compute the SVD $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^T = \mathbf{A}$
 - Project $\mathbf{A} \mapsto \mathbf{AV}_k \in \mathbb{R}^{n \times k}$
 - $\mathbf{V}_k =$ top- k right singular vectors
- A.k.a. the **principal component analysis** (PCA)


More on KLT

- The columns of \mathbf{V}_k show the main directions of variance in columns
- The data is expressed in a new coordinate system
- The average projection distance is minimized



Visualization

Scatter plots



2D or 3D KLT




Figure 3.2. The first two factors for a dataset ranking wines.

Latent Semantic Analysis & Indexing

- **Latent semantic analysis** (LSA) is a **latent topic model**
 - Documents-by-terms matrix **A**
 - Typically normalized (e.g. tf/idf)
- Goal is to find the “topics” doing SVD
 - **U** associates documents to topics
 - **V** associates topics to terms
- Queries can be answered by projecting the query vector **q** to $\mathbf{q}' = \mathbf{qV}\boldsymbol{\Sigma}^{-1}$ and returning rows of **U** that are similar to **q'**

And many more...

- Determining the rank, finding the least-squares solution, recommending the movies, ordering results of queries, ...

Computing the SVD

Very general idea

- SVD is unique
 - If \mathbf{U} and \mathbf{V} are orthogonal s.t. $\mathbf{U}^T \mathbf{A} \mathbf{V} = \mathbf{\Sigma}$, then $\mathbf{U} \mathbf{\Sigma} \mathbf{V}^T$ is the SVD of \mathbf{A}
- Idea: find orthogonal \mathbf{U} and \mathbf{V} s.t. $\mathbf{U}^T \mathbf{A} \mathbf{V}$ is as desired
 - Iterative process: find orthogonal $\mathbf{U}_1, \mathbf{U}_2, \dots$ and set $\mathbf{U} = \mathbf{U}_1 \mathbf{U}_2 \mathbf{U}_3 \dots$
 - Still orthogonal

First attempt

- Recall: \mathbf{U} are the eigenvectors of \mathbf{AA}^T and σ_i^2 are the associated eigenvalues
- Idea: Compute the eigenvectors and values of \mathbf{AA}^T and $\mathbf{A}^T\mathbf{A}$ to get the SVD of \mathbf{A}
 - Not the most optimal idea because it requires \mathbf{AA}^T and $\mathbf{A}^T\mathbf{A}$
- We need a way to build orthogonal matrices that make matrices more diagonal

Rotations and reflections

2D rotation

$$\begin{pmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{pmatrix}$$

Rotates counterclockwise through an angle θ

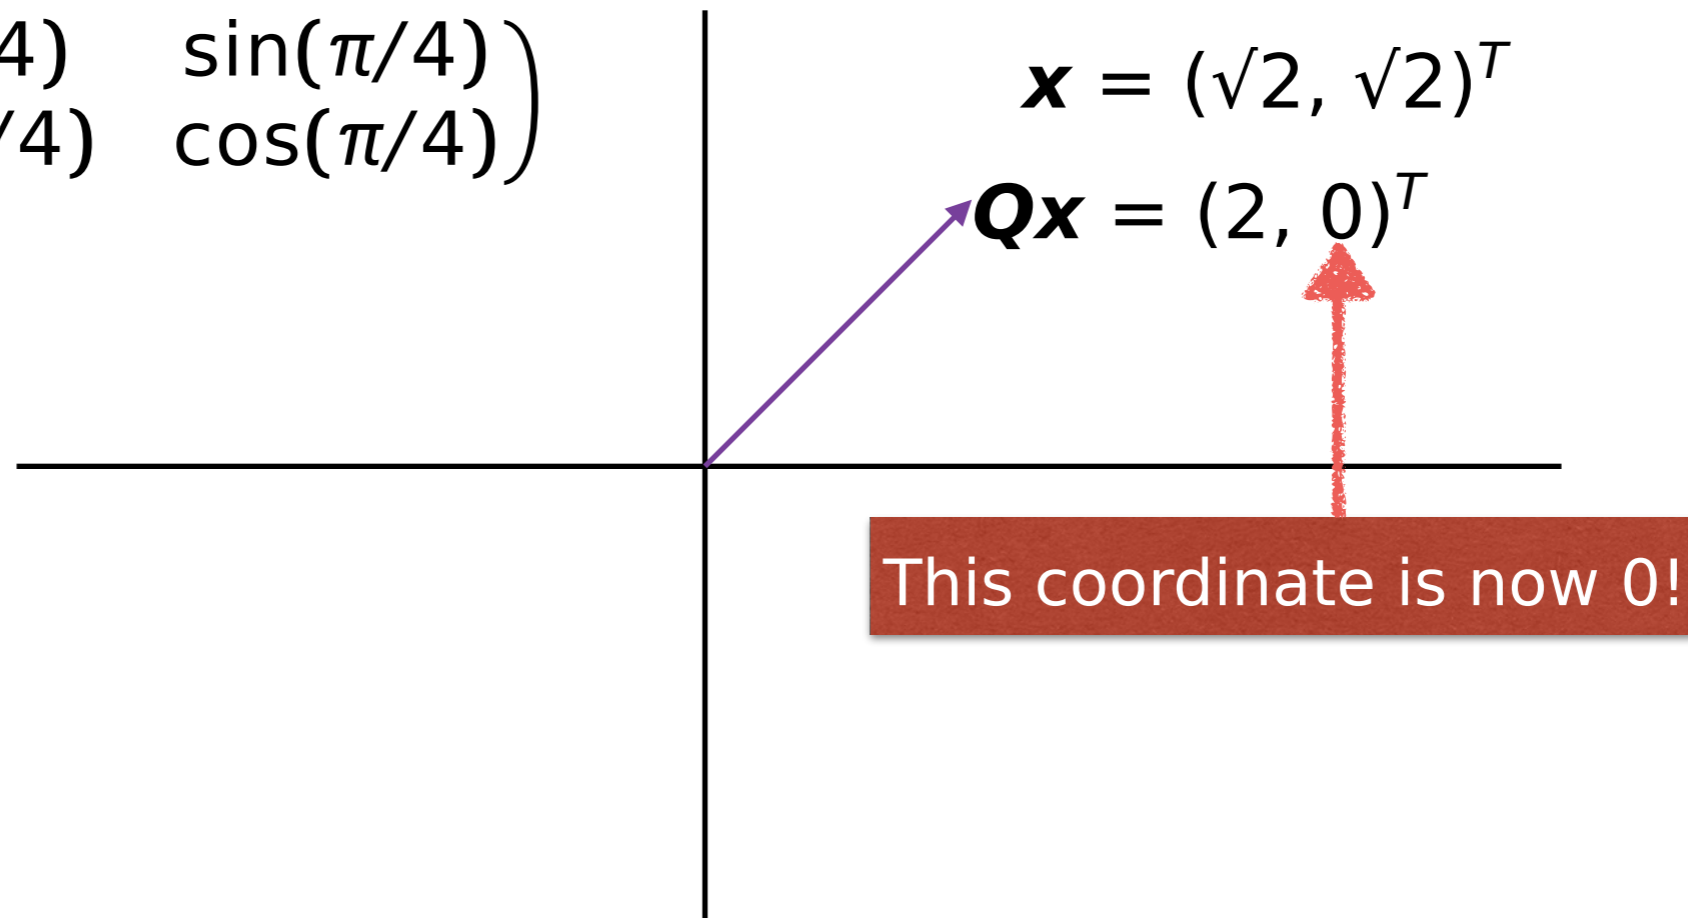
2D reflection

$$\begin{pmatrix} \cos(\theta) & \sin(\theta) \\ \sin(\theta) & -\cos(\theta) \end{pmatrix}$$

Reflects across the line spanned by $(\cos(\theta/2), \sin(\theta/2))^T$

Example

$$Q = \begin{pmatrix} \cos(\pi/4) & \sin(\pi/4) \\ -\sin(\pi/4) & \cos(\pi/4) \end{pmatrix}$$



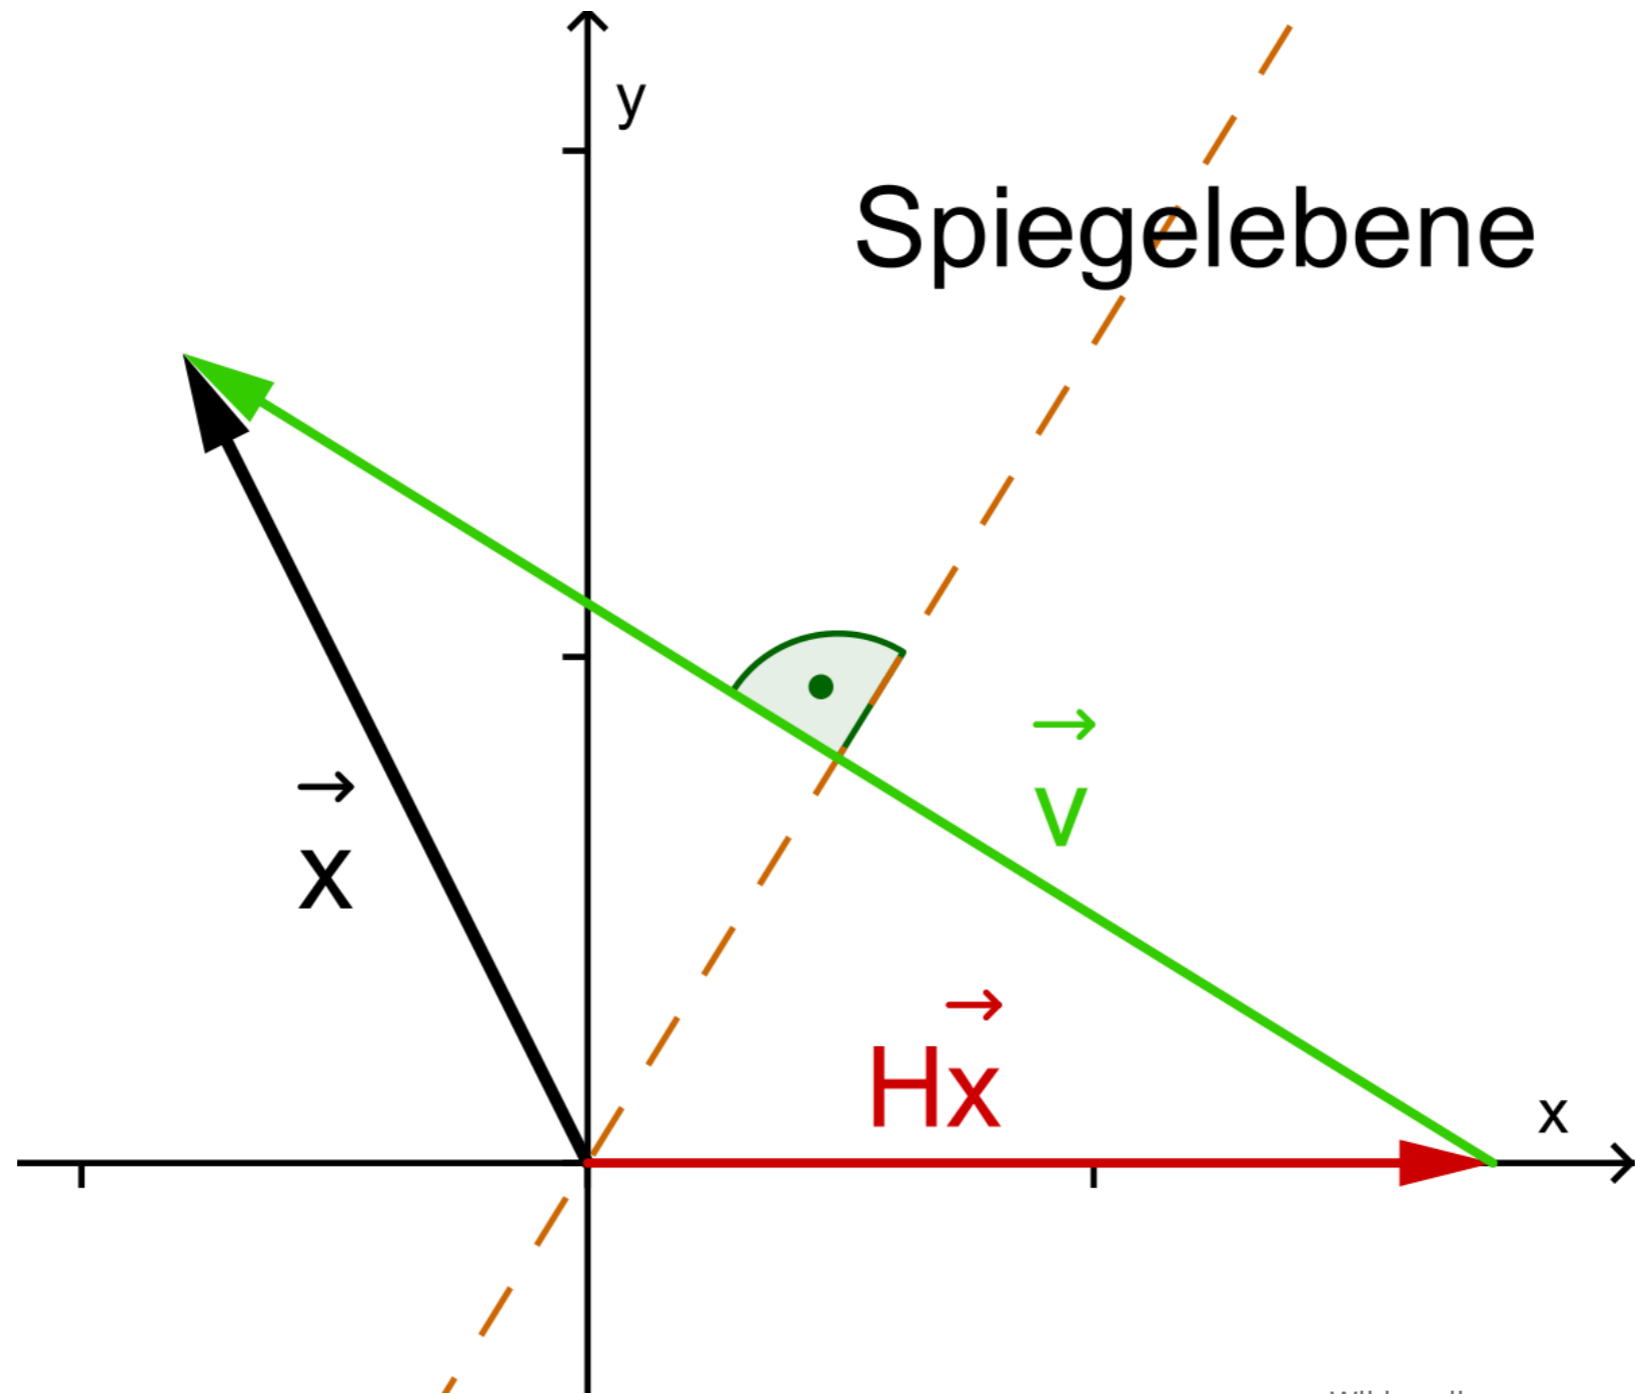
Householder reflections

- A **Householder reflection** is n -by- n matrix

$$\mathbf{P} = \mathbf{I} - \beta \mathbf{v} \mathbf{v}^T \quad \text{where} \quad \beta = \frac{2}{\mathbf{v}^T \mathbf{v}}$$

- If we set $\mathbf{v} = \mathbf{x} - \|\mathbf{x}\|_2 \mathbf{e}_1$, then $\mathbf{P}\mathbf{x} = \|\mathbf{x}\|_2 \mathbf{e}_1$
 - $\mathbf{e}_1 = (1, 0, 0, \dots, 0)^T$
- Note: $\mathbf{P}\mathbf{A} = \mathbf{A} - (\beta \mathbf{v})(\mathbf{v}^T \mathbf{A})$ where $\beta = 2/(\mathbf{v}^T \mathbf{v})$
 - We never have to compute matrix \mathbf{P}

Example



Wikimedia commons

Almost there: bidiagonalization

- Given n -by- m ($n \geq m$) \mathbf{A} , we can **bidiagonalize** it with Householder transformations
- Fix $\mathbf{A}[1:n,1]$, $\mathbf{A}[1,2:m]$, $\mathbf{A}[2:n,2]$, $\mathbf{A}[2,3:m]$,
 $\mathbf{A}[3:n,3]$, $\mathbf{A}[3,4:m]$...
- The results has non-zeros in main diagonal and the one above it

Example

$$\mathbf{U}_4^T \mathbf{U}_3^T \mathbf{U}_2^T \mathbf{U}_1^T \mathbf{A} \mathbf{V}_1 \mathbf{V}_2 = \begin{pmatrix} * & * & \otimes & \otimes \\ \otimes & * & * & \otimes \\ \otimes & \otimes & * & * \\ \otimes & \otimes & \otimes & * \\ \otimes & \otimes & \otimes & \otimes \end{pmatrix}$$

Givens rotations

- Householder is too crude to give identity
- Givens rotations** are rank-2 corrections to the identity of form

$$\mathbf{G}(i, k, \theta) = \begin{pmatrix} 1 & \cdots & 0 & \cdots & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & & \vdots & & \vdots \\ 0 & \cdots & \cos(\theta) & \cdots & \sin(\theta) & \cdots & 0 \\ \vdots & & \vdots & \ddots & \vdots & & \vdots \\ 0 & \cdots & -\sin(\theta) & \cdots & \cos(\theta) & \cdots & 0 \\ \vdots & & \vdots & & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & \cdots & 0 & \cdots & 1 \end{pmatrix} \begin{matrix} \\ \\ i \\ \\ k \\ \\ \end{matrix}$$

Applying Givens

- Set θ s.t.

$$\cos(\theta) = \frac{x_i}{\sqrt{x_i^2 + x_k^2}} \quad \text{and} \quad \sin(\theta) = \frac{-x_k}{\sqrt{x_i^2 + x_k^2}}$$

- Now

$$\begin{pmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{pmatrix}^T \begin{pmatrix} x_i \\ x_k \end{pmatrix} = \begin{pmatrix} r \\ 0 \end{pmatrix}$$


- N.B. $\mathbf{G}(i, k, \theta)^T \mathbf{A}$ only affects to the 2 rows $\mathbf{A}[c(i, k),]$

- Also, no inverse trig. operations are needed

Given's in SVD

- We use Given's transformations to erase the superdiagonal
 - Consider principal 2-by-2 submatrices $\mathbf{A}[k:k+1, k:k+1]$
 - Rotations can introduce unwanted non-zeros to $\mathbf{A}[k+2, k]$ (or $\mathbf{A}[k, k+2]$)
 - Fix them in the next sub-matrix

Example



Putting it all together

1. Compute the bidiagonal matrix **B** from **A** using Householder transformations
2. Apply the Givens rotations to **B** until it is fully diagonal
3. Collect the required results

Time complexity

Output	Time
Σ	$4nm^2 - 4m^3/3$
Σ, V	$4nm^2 + 8m^3$
Σ, U	$4n^2m - 8nm^2$
Σ, U_1	$14nm^2 - 2m^3$
Σ, U, V	$4n^2m + 8nm^2 + 9m^3$
Σ, U_1, V	$14nm^2 + 8m^3$

Summary of computing SVD

- Rotations and reflections allow us to selectively zero elements of a matrix with orthogonal transformations
 - Used in many, many decompositions
- Fast and accurate results require careful implementations
- Other techniques are faster for truncated SVD in large, sparse matrices

Summary of SVD

- Truly the workhorse of numerical linear algebra
 - Many useful theoretical properties
 - Rank-revealing, pseudo-inverses, scalar norm computation, ...
 - Reasonably easy to compute
- But it also has some major shortcomings in data analysis... *stay tuned!*