

CS540 Introduction to Artificial Intelligence

Lecture 12

Young Wu

Based on lecture slides by Jerry Zhu, Yingyu Liang, and Charles Dyer

July 13, 2020


Low Dimension Representation

Motivation

- Unsupervised learning techniques are used to find low dimensional representation.
- ① Visualization.
- ② Efficient storage.
- ③ Better generalization.
- ④ Noise removal.

Dimension Reduction Diagram

Motivation




max variance
along PC₁

Projection

Definition

- The projection of x_i onto a unit vector u_k is the vector in the direction of u_k that is the closest to x_i .
 length

$$\text{proj}_{u_k} x_i = \left(\frac{u_k^T x_i}{u_k^T u_k} \right) u_k = u_k^T x_i u_k$$

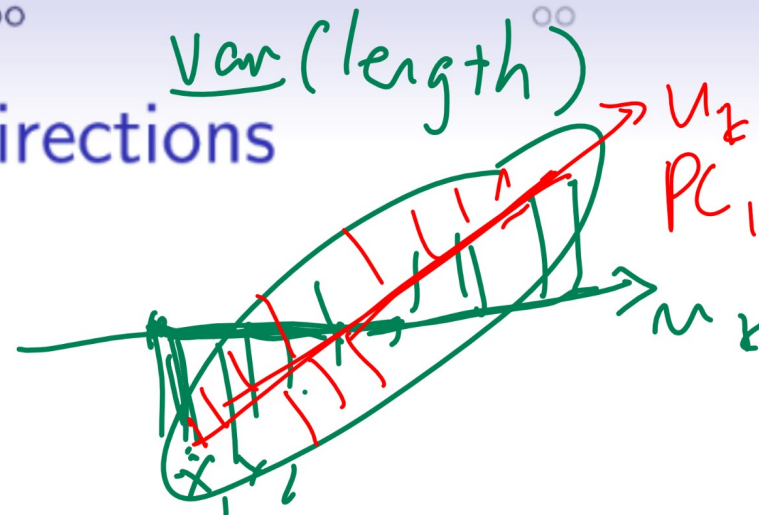


- The length of the projection of x_i onto a unit vector u_k is $u_k^T x_i$.
 proj

$$\| \text{proj}_{u_k} x_i \|_2 = u_k^T x_i$$

Maximum Variance Directions

Definition



$$x^T A x \rightarrow A x$$

$$x^T x \rightarrow x$$

- The goal is to find the direction that maximizes the projected variance.

variance

$$\max_{u_k} u_k^T \hat{\Sigma} u_k \text{ such that } u_k^T u_k = 1$$

$$\Rightarrow \max_{u_k} u_k^T \hat{\Sigma} u_k - \lambda u_k^T u_k$$

$$\Rightarrow \hat{\Sigma} u_k = \lambda u_k$$

Lagrange

derivative

max

eigen value


$$u_k^T \hat{\Sigma} u_k = \lambda$$

Projection Example 2

Quiz

- What is the projection of $\begin{bmatrix} 1 \\ 0 \end{bmatrix}$ onto $\begin{bmatrix} 1 \\ 1 \end{bmatrix}$?

$$\begin{pmatrix} \frac{1}{2} \\ \frac{1}{2} \end{pmatrix}$$



$$\frac{u^T x}{u^T u} u = \frac{1}{2} \cdot \begin{pmatrix} 1 \\ 1 \end{pmatrix} = \begin{pmatrix} \frac{1}{2} \\ \frac{1}{2} \end{pmatrix}$$

Projection Example 3

Quiz

Q8

- What is the projection of

$$\begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} \text{ onto } \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \stackrel{u}{=} ?$$

$\frac{u^T x u}{u^T u}$

not sure

- A: $[1 \ 1 \ 1]^T$

- B: $[2 \ 2 \ 2]^T$

- C: $[3 \ 3 \ 3]^T$

- D: $[4 \ 4 \ 4]^T$

- E: $[6 \ 6 \ 6]^T$

$$\frac{u^T x}{u^T u} u$$

$$= \frac{6}{3} \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 2 \\ 2 \\ 2 \end{pmatrix}$$

Projection Example 4

Quiz

Q9

- What is the projection of $\begin{bmatrix} 3 \\ 2 \\ 1 \end{bmatrix}$ onto $\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$?
- A: $\begin{bmatrix} 1 & 1 & 1 \end{bmatrix}^T$
- B: $\begin{bmatrix} 2 & 2 & 2 \end{bmatrix}^T$
- C: $\begin{bmatrix} 3 & 3 & 3 \end{bmatrix}^T$
- D: $\begin{bmatrix} 4 & 4 & 4 \end{bmatrix}^T$
- E: $\begin{bmatrix} 6 & 6 & 6 \end{bmatrix}^T$

Spectral Decomposition Example 1

Quiz

- Given the following spectral decomposition of $\hat{\Sigma}$, what is the first principal component?

$$\hat{\Sigma} = \underbrace{\begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ -1 & 0 & 1 \end{bmatrix}}_{\text{eigen vectors}} \underbrace{\begin{bmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 3 \end{bmatrix}}_{\text{eigen values}} \underbrace{\begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ -1 & 0 & 1 \end{bmatrix}}^{-1}$$

$PC_1 \rightarrow$ largest e.v. = 3 \Rightarrow normalized $\begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix} = \begin{pmatrix} \frac{1}{\sqrt{2}} \\ 0 \\ \frac{1}{\sqrt{2}} \end{pmatrix}$
 $PC_2 \rightarrow \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}$ $PC_3 = \begin{pmatrix} \frac{1}{\sqrt{2}} \\ 0 \\ -\frac{1}{\sqrt{2}} \end{pmatrix}$

Spectral Decomposition Example 2

Quiz

Q10

- Given the following spectral decomposition of $\hat{\Sigma}$, what is the first principal component?

$$\hat{\Sigma} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 2 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}^{-1} = \begin{pmatrix} 2 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

- A: $\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$, B: $\begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}$, C: $\begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$, D: $\begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}$, E: $\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$

Spectral Decomposition Example 3

Quiz

Q11

• $\hat{\Sigma} = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 1 \end{bmatrix}$, what is the second principal component?

• A: $\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$, B: $\begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}$, C: $\begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$, D: $\begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}$, E: $\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$

Principal Component Analysis

Algorithm

- Input: instances: $\{x_i\}_{i=1}^n$, the number of dimensions after reduction $K < m$.
- Output: K principal components.
- Find the largest K eigenvalues $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_K$.
- Return the corresponding unit orthogonal eigenvectors $u_1, u_2 \dots u_K$.

Number of Dimensions

Discussion

- There are a few ways to choose the number of principal components K .
- K can be selected given prior knowledge or requirement.
- K can be the number of non-zero eigenvalues.
- K can be the number of eigenvalues that are large (larger than some threshold).

> 1

Eigenface

Discussion

in m dim

- Eigenfaces are eigenvectors of face images (pixel intensities or HOG features).
- Every face can be written as a linear combination of eigenfaces. The coefficients determine specific faces.

$$x_i = \sum_{k=1}^m (u_k^T x_i) u_k \approx \sum_{k=1}^K (u_k^T x_i) u_k$$

Handwritten annotations: A green box surrounds the left side of the equation. A green circle surrounds the term $(u_k^T x_i)$ in the right side. A green arrow points from the box to the circle. A green line is drawn under the right side of the equation.

- Eigenfaces and SVM can be combined to detect or recognize faces.

new features

Reduced Space Example 1

Quiz

- 2017 Fall Final Q10

If $u_1 = \begin{bmatrix} 1 \\ \sqrt{2} \\ 1 \\ \sqrt{2} \\ 0 \end{bmatrix}$ and $u_2 = \begin{bmatrix} 1 \\ -\sqrt{2} \\ 1 \\ \sqrt{2} \\ 0 \end{bmatrix}$. If one original data is

$$x = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}$$

What is the new representation?

$$\begin{pmatrix} u_1 \cdot x \\ u_2 \cdot x \end{pmatrix} = \begin{pmatrix} \frac{3}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{pmatrix}$$

$u_3 = x = 3$

$$u_3 = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$$

Reduced Space Example 1 Diagram

Quiz

$$\begin{aligned}
 X &\approx \frac{3}{\sqrt{2}} \begin{pmatrix} \frac{\sqrt{2}}{2} \\ \frac{1}{\sqrt{2}} \\ 0 \end{pmatrix} + \frac{1}{\sqrt{2}} \begin{pmatrix} -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \\ 0 \end{pmatrix} + \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \\
 &= \begin{pmatrix} \frac{3}{2} \\ \frac{3}{2} \\ 0 \end{pmatrix} + \begin{pmatrix} -\frac{1}{2} \\ \frac{1}{2} \\ 0 \end{pmatrix} = \begin{pmatrix} 1 \\ 2 \\ 0 \end{pmatrix} \approx \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix} \\
 &\quad \times \begin{pmatrix} 0 \\ 0 \\ 3 \end{pmatrix} \quad \times \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix} = X_1
 \end{aligned}$$

Reduced Space Example 2

Quiz

Q12

- $\hat{\Sigma} = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 1 \end{bmatrix}$. If one original data is $x = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}$. What is the new representation using only the first two principal components?
- $PC_1 = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}$, $PC_2 = \begin{pmatrix} 1 \\ 0 \\ 2 \end{pmatrix}$

- A: $\begin{bmatrix} 1 \\ 2 \end{bmatrix}$, B: $\begin{bmatrix} 2 \\ 1 \end{bmatrix}$, C: $\begin{bmatrix} 1 \\ 3 \end{bmatrix}$, D: $\begin{bmatrix} 3 \\ 1 \end{bmatrix}$, E: $\begin{bmatrix} 2 \\ 3 \end{bmatrix}$

Reduced Space Example 3

Quiz

Q13, (last)

- $\hat{\Sigma} = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 1 \end{bmatrix}$. If one original data is $x = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}$. What is the reconstructed vector using only the first two principal components?

- A: $\begin{bmatrix} 1 \\ 2 \\ 0 \end{bmatrix}$, B: $\begin{bmatrix} 2 \\ 1 \\ 0 \end{bmatrix}$, C: $\begin{bmatrix} 2 \\ 3 \\ 1 \end{bmatrix}$, D: $\begin{bmatrix} 2 \\ 1 \\ 3 \end{bmatrix}$, E: $\begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}$


$$2 \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} + 1 \cdot \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} = \begin{pmatrix} 1 \\ 2 \\ 0 \end{pmatrix}$$

Autoencoder

Discussion

PCA \approx no activation $a = x^T w + b$

$$a = g(x^T w + b)$$



- A multi-layer neural network with the same input and output $y_i = x_i$ is called an autoencoder.
- The hidden layers have fewer units than the dimension of the input m .
- The hidden units form an encoding of the input with reduced dimensionality.

non-linear PCA

Kernel PCA

Discussion

- A kernel can be applied before finding the principal components.

$$\hat{\Sigma} = \frac{1}{n-1} \sum_{i=1}^n \varphi(x_i) \varphi(x_i)^T$$

Handwritten notes: A red circle around $\hat{\Sigma}$ and a red arrow pointing from the $\varphi(x_i)$ term to the text $\infty \text{ dim}$.

- The principal components can be found without explicitly computing $\varphi(x_i)$, similar to the kernel trick for support vector machines.
- Kernel PCA is a non-linear dimensionality reduction method.