

Machine Learning for Data Science (CS4786)

Lecture 2

Dimensionality Reduction
&
Principal Component Analysis

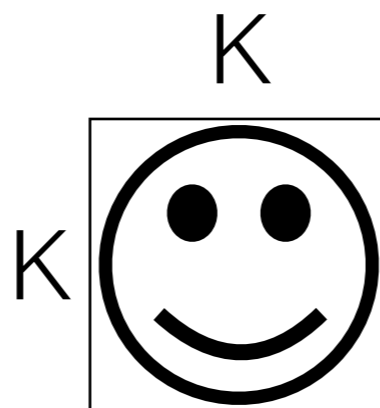
Course Webpage :

<http://www.cs.cornell.edu/Courses/cs4786/2016fa/>


REPRESENTING DATA AS FEATURE VECTORS

- How do we represent data?
- Each data-point often represented as vector referred to as feature vector

EXAMPLE: IMAGES



vectorize



$$d = K^2$$

EXAMPLE: TEXT (BAG OF WORDS)

Documents:

car
engine
hood
tires
truck
trunk


car
emissions
hood
make
model
trunk

Chomsky
corpus
noun
parsing
tagging
wonderful



car	Chomsky	corpus	emissions	engine	hood	make	model	noun	parsing	tagging	tires	truck	trunk	wonderful
1	0	0	0	1	1	0	0	0	0	0	1	1	1	0
1	0	0	1	0	1	1	1	0	0	0	0	0	1	0
0	1	1	0	0	0	0	0	1	1	1	0	0	0	1

DIMENSIONALITY REDUCTION



DIMENSIONALITY REDUCTION

Given feature vectors $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{R}^d$, compress the data points into low dimensional representation $\mathbf{y}_1, \dots, \mathbf{y}_n \in \mathbb{R}^K$ where $K \ll d$

Flowers



Iris-Setosa



Iris-versicolor



Iris-virginica

PRINCIPAL COMPONENT ANALYSIS: DEMO

WHY DIMENSIONALITY REDUCTION?

- For computational ease
 - As input to supervised learning algorithm
 - Before clustering to remove redundant information and noise
- Data compression & Noise reduction
- Data visualization

DIMENSIONALITY REDUCTION


Desired properties:

- ① Original data can be (approximately) reconstructed
- ② Preserve distances between data points
- ③ “Relevant” information is preserved
- ④ Noise is reduced

DIM REDUCTION: LINEAR TRANSFORMATION

- Pick a low dimensional subspace
- Project linearly to this subspace
- Subspace retains as much information

DIM REDUCTION: LINEAR TRANSFORMATION




$$\mathbf{y}_i^\top = \mathbf{x}_i^\top W$$


DIM REDUCTION: LINEAR TRANSFORMATION

Prelude: reducing to 1 dimension


$$y_1 = \mathbf{w}^T \mathbf{x}_1 = \|\mathbf{x}_1\| \cos(\angle \mathbf{w} \mathbf{x}_1)$$




PCA: VARIANCE MAXIMIZATION



PCA: VARIANCE MAXIMIZATION



PCA: VARIANCE MAXIMIZATION



PCA: VARIANCE MAXIMIZATION

- Pick directions along which data varies the most
- First principal component:

$$\begin{aligned}\mathbf{w}_1 &= \arg \max_{\mathbf{w}: \|\mathbf{w}\|_2=1} \frac{1}{n} \sum_{t=1}^n \left(y_t - \frac{1}{n} \sum_{t=1}^n y_t \right)^2 \\ &= \arg \max_{\mathbf{w}: \|\mathbf{w}\|_2=1} \frac{1}{n} \sum_{t=1}^n \left(\mathbf{w}^\top \mathbf{x}_t - \frac{1}{n} \sum_{t=1}^n \mathbf{w}^\top \mathbf{x}_t \right)^2 \\ &= \arg \max_{\mathbf{w}: \|\mathbf{w}\|_2=1} \frac{1}{n} \sum_{t=1}^n \left(\mathbf{w}^\top (\mathbf{x}_t - \boldsymbol{\mu}) \right)^2 \\ &= \arg \max_{\mathbf{w}: \|\mathbf{w}\|_2=1} \frac{1}{n} \sum_{t=1}^n \mathbf{w}^\top (\mathbf{x}_t - \boldsymbol{\mu}) (\mathbf{x}_t - \boldsymbol{\mu})^\top \mathbf{w} \\ &= \arg \max_{\mathbf{w}: \|\mathbf{w}\|_2=1} \mathbf{w}^\top \boldsymbol{\Sigma} \mathbf{w}\end{aligned}$$

where $\boldsymbol{\Sigma}$ is the covariance matrix and $\boldsymbol{\mu} = \frac{1}{n} \sum_{t=1}^n \mathbf{x}_t$

PCA: VARIANCE MAXIMIZATION

Covariance matrix:

$$\Sigma = \frac{1}{n} \sum_{t=1}^n (\mathbf{x}_t - \boldsymbol{\mu})(\mathbf{x}_t - \boldsymbol{\mu})^\top$$

- Its a $d \times d$ matrix, $\Sigma[i, j]$ measures “covariance” of features i and j

$$\Sigma[i, j] = \frac{1}{n} \sum_{t=1}^n (\mathbf{x}_t[i] - \mu[i])(\mathbf{x}_t[j] - \mu[j])$$

PCA: VARIANCE MAXIMIZATION

- First principal component:

$$\mathbf{w}_1 = \arg \max_{\mathbf{w}: \|\mathbf{w}\|_2=1} \mathbf{w}^\top \Sigma \mathbf{w}$$


- The solution to the above optimization problem is \mathbf{w}_1 is the top Eigen vector of matrix Σ
- Hence in “matlab”,

$$S = \text{Cov}(X)$$

$$[W, E] = \text{eigs}(S, 1)$$

$$Y = W * X$$

PRINCIPAL COMPONENT ANALYSIS: DEMO




PRINCIPAL COMPONENT ANALYSIS ($K > 1$)

What do we do when $K > 1$?

DIM REDUCTION: LINEAR TRANSFORMATION

Prelude: reducing to 1 dimension



ORTHONORMAL PROJECTIONS

- Think of $\mathbf{w}_1, \dots, \mathbf{w}_K$ as coordinate system for PCA (in a K dimensional subspace)
- \mathbf{y} values provide coefficients in this system
- Without loss of generality, $\mathbf{w}_1, \dots, \mathbf{w}_K$ can be orthonormal, i.e. $\mathbf{w}_i \perp \mathbf{w}_j$ & $\|\mathbf{w}_i\| = 1$.

- Reconstruction:

$$\hat{\mathbf{x}}_t = \sum_{j=1}^K \mathbf{y}_t[j] \mathbf{w}_j$$

- If we take all $\mathbf{w}_1, \dots, \mathbf{w}_d$, then $\mathbf{x}_t = \sum_{j=1}^d \mathbf{y}_t[j] \mathbf{w}_j$. To reduce dimensionality we only consider first K vectors of the basis

PCA: VARIANCE MAXIMIZATION

- How do we find the K components?
- We are looking for orthogonal directions that maximize total spread in each direction
- Find orthonormal W that maximizes

$$\begin{aligned} \sum_{j=1}^K \frac{1}{n} \sum_{t=1}^n \left(\mathbf{y}_t[j] - \frac{1}{n} \sum_{t=1}^n \mathbf{y}_t[j] \right)^2 &= \sum_{j=1}^K \frac{1}{n} \sum_{t=1}^n \left(\mathbf{w}_j^\top \left(\mathbf{x}_t - \frac{1}{n} \sum_{t=1}^n \mathbf{x}_t \right) \right)^2 \\ &= \sum_{j=1}^K \mathbf{w}_j^\top \Sigma \mathbf{w}_j \end{aligned}$$

- This solutions is given by $W =$ Top K eigenvectors of Σ

PRINCIPAL COMPONENT ANALYSIS

1. $\Sigma = \text{COV}(X)$

2. $W = \text{eigs}(\Sigma, K)$

3. $Y = X \times W$

PRINCIPAL COMPONENT ANALYSIS: DEMO

