# Advanced Data Analysis: Principal Component Analysis 

## Masashi Sugiyama (Computer Science)

W8E-406, sugi@cs.titech.ac.jp
http://sugiyama-www.cs.titech.ac.jp/~sugi

## Curse of Dimensionality

$$
\left\{\boldsymbol{x}_{i}\right\}_{i=1}^{n}, \quad \boldsymbol{x}_{i} \in \mathbb{R}^{d}, d \gg 1
$$

- If your data samples are high-dimensional, they are often too complex to directly analyze.
- Usual geometric intuitions are often only applicable to low-dimensional spaces; such intuitions could be even misleading in high-dimensional spaces.


## Curse of Dimensionality (cont.) ${ }^{4}$

$\square$ When the dimensionality increases,

- Volume of unit hyper-cube $V_{c}$ is always 1.
- Volume of inscribed hyper-sphere $V_{s}$ goes to 0 .
- Relative size of hyper-sphere gets small!

$$
d=1 \quad \begin{array}{lll} 
\\
d=2 & \frac{V_{s}}{V_{c}} \rightarrow 0 \\
& d=3 & \cdots \\
& & d=\infty
\end{array}
$$

## Curse of Dimensionality (cont.) ${ }^{5}$

$\square$ Grid sampling requires an exponentially large number.


- Unless you have an exponentially large number of samples, your high-dimensional samples are never dense.


## Dimensionality Reduction

- We want to reduce the dimensionality of the data while preserving the intrinsic "information" in the data.
- Dimensionality reduction is also called embedding; if the dimension is reduced up to 3 , it is also called data visualization.
- Basic assumption (or belief) behind dimensionality reduction: your highdimensional data is redundant in some sense.


## Notation: Linear Embedding

- Data samples:

$$
\left\{\boldsymbol{x}_{i}\right\}_{i=1}^{n}, \quad \boldsymbol{x}_{i} \in \mathbb{R}^{d}, d \gg 1
$$

$\square$ Embedding matrix:

$$
\boldsymbol{B} \in \mathbb{R}^{m \times d}, \quad 1 \leq m \ll d
$$

■ Embedded data samples:

$$
\left\{\boldsymbol{z}_{i}\right\}_{i=1}^{n}, \quad \boldsymbol{z}_{i}=\boldsymbol{B} \boldsymbol{x}_{i} \in \mathbb{R}^{m}
$$



## Principal Component Analysis (PCA)

- Idea: We want to get rid of a redundant dimension of the data samples

$$
\binom{1}{0},\binom{2}{0.1},\binom{3}{-0.1}
$$

$\square$ This could be achieved by minimizing the distance between embedded samples and original samples.


## Data Centering

$\square$ We center the data samples by

- In matrix,

$$
\overline{\boldsymbol{x}}_{i}=\boldsymbol{x}_{i}-\frac{1}{n} \sum_{j=1}^{n} \boldsymbol{x}_{j}
$$

$$
\frac{1}{n} \sum_{i=1}^{n} \bar{x}_{i}=0
$$

$$
\bar{X}=\boldsymbol{X} \boldsymbol{H}
$$

$$
\begin{aligned}
& \overline{\boldsymbol{X}}=\left(\overline{\boldsymbol{x}}_{1}\left|\overline{\boldsymbol{x}}_{2}\right| \cdots \mid \overline{\boldsymbol{x}}_{n}\right) \\
& \boldsymbol{X}=\left(\boldsymbol{x}_{1}\left|\boldsymbol{x}_{2}\right| \cdots \mid \boldsymbol{x}_{n}\right) \\
& \boldsymbol{H}=\boldsymbol{I}_{n}-\frac{1}{n} \mathbf{1}_{n \times n}
\end{aligned}
$$

$$
\boldsymbol{X}=\left(\boldsymbol{x}_{1}\left|\boldsymbol{x}_{2}\right| \cdots \mid \boldsymbol{x}_{n}\right) \quad \boldsymbol{I}_{n}: n \text {-dimensional identity matrix }
$$

$\mathbf{1}_{n \times n}: n \times n$ matrix with all ones

## Orthogonal Projection

$\square\left\{\boldsymbol{b}_{i}\left(\in \mathbb{R}^{d}\right)\right\}_{i=1}^{m}$ : Orthonormal basis in $m$-dimensional embedding subspace

$$
\left\langle\boldsymbol{b}_{i}, \boldsymbol{b}_{j}\right\rangle=\delta_{i, j}= \begin{cases}1 & (i=j) \\ 0 & (i \neq j)\end{cases}
$$

- In matrix, $\boldsymbol{B} \boldsymbol{B}^{\top}=\boldsymbol{I}_{m}$

$$
\boldsymbol{B}=\left(\boldsymbol{b}_{1}\left|\boldsymbol{b}_{2}\right| \cdots \mid \boldsymbol{b}_{m}\right)^{\top}
$$

$\square$ Orthogonal projection of $\bar{x}_{i}$ is expressed by

$$
\sum_{j=1}^{m}\left\langle\boldsymbol{b}_{j}, \overline{\boldsymbol{x}}_{i}\right\rangle \boldsymbol{b}_{j} \quad\left(=\boldsymbol{B}^{\top} \boldsymbol{B} \overline{\boldsymbol{x}}_{i}\right)
$$

## PCA Criterion

$\square$ Minimize the sum of squared distances.

$$
\begin{array}{r}
\sum_{i=1}^{n}\left\|\boldsymbol{B}^{\top} \boldsymbol{B} \overline{\boldsymbol{x}}_{i}-\overline{\boldsymbol{x}}_{i}\right\|^{2} \quad\left(=-\operatorname{tr}\left(\boldsymbol{B} \overline{\boldsymbol{C}} \boldsymbol{B}^{\top}\right)+\operatorname{tr}(\overline{\boldsymbol{C}})\right) \\
\overline{\boldsymbol{C}}=\sum_{i=1}^{n} \overline{\boldsymbol{x}}_{i} \overline{\boldsymbol{x}}_{i}^{\top}=\overline{\boldsymbol{X}} \overline{\boldsymbol{X}}^{\top}
\end{array}
$$

- PCA criterion:

$$
\begin{aligned}
\boldsymbol{B}_{P C A}= & \underset{\boldsymbol{B} \in \mathbb{R}^{m \times d}}{\operatorname{argmax}} \operatorname{tr}\left(\boldsymbol{B} \overline{\boldsymbol{C}} \boldsymbol{B}^{\top}\right) \\
& \text { subject to } \boldsymbol{B} \boldsymbol{B}^{\top}=\boldsymbol{I}_{m}
\end{aligned}
$$

- A PCA solution:

$$
\boldsymbol{B}_{P C A}=\left(\boldsymbol{\psi}_{1}\left|\boldsymbol{\psi}_{2}\right| \cdots \mid \boldsymbol{\psi}_{m}\right)^{\top}
$$

$\left\{\lambda_{i}, \boldsymbol{\psi}_{i}\right\}_{i=1}^{m}:$ Sorted eigenvalues and normalized eigenvectors of $\bar{C} \psi=\lambda \psi$
$\lambda_{1} \geq \lambda_{2} \geq \cdots \geq \lambda_{d} \quad\left\langle\boldsymbol{\psi}_{i}, \boldsymbol{\psi}_{j}\right\rangle=\delta_{i, j}$
$\square$ PCA embedding of a sample $\boldsymbol{x}$ :

$$
\overline{\boldsymbol{z}}=\boldsymbol{B}_{P C A}\left(\boldsymbol{x}-\frac{1}{n} \boldsymbol{X} \mathbf{1}_{n}\right)
$$

$\mathbf{1}_{n}$ : $n$-dimensional vector with all ones

- Lagrangian:


## Proof

$$
L(\boldsymbol{B}, \boldsymbol{\Delta})=\operatorname{tr}\left(\boldsymbol{B} \overline{\boldsymbol{C}} \boldsymbol{B}^{\top}\right)-\operatorname{tr}\left(\left(\boldsymbol{B} \boldsymbol{B}^{\top}-\boldsymbol{I}_{m}\right) \boldsymbol{\Delta}\right)
$$

$\Delta:$ Lagrange multipliers (symmetric)

- Stationary point (necessary condition):
- $\frac{\partial L}{\partial \boldsymbol{B}}=2 \boldsymbol{B} \overline{\boldsymbol{C}}-2 \boldsymbol{\Delta} \boldsymbol{B}=0$
$\begin{array}{rl}\bullet \frac{\partial L}{\partial \boldsymbol{\Delta}}=\boldsymbol{B} \boldsymbol{B}^{\top}-\boldsymbol{I}_{m}=0 & \boldsymbol{C} \boldsymbol{B}^{\top}=\boldsymbol{B}^{\top} \boldsymbol{\Delta} \\ & \longrightarrow \boldsymbol{B} \boldsymbol{B}^{\top}=\boldsymbol{I}_{m}(2)\end{array}$
$\square$ Eigendecomposition:

$$
\boldsymbol{\Delta}=\boldsymbol{T} \boldsymbol{\Gamma} \boldsymbol{T}^{\top} \text { (3) }
$$

$\boldsymbol{T}$ : orthogonal matrix $\boldsymbol{\Gamma}$ : diagonal matrix

$$
\boldsymbol{T}^{-1}=\boldsymbol{T}^{\top}
$$

## Proof (cont.)

$\square(1) \&(3)$

$$
\begin{aligned}
& \overline{\boldsymbol{C}} \boldsymbol{B}^{\top}=\boldsymbol{B}^{\top} \boldsymbol{T} \boldsymbol{\Gamma} \boldsymbol{T}^{\top}(4) \\
& \overline{\boldsymbol{C}} \boldsymbol{B}^{\top} \boldsymbol{T}=\boldsymbol{B}^{\top} \boldsymbol{T} \boldsymbol{\Gamma} \\
& \overline{\boldsymbol{C}} \boldsymbol{F}=\boldsymbol{F} \boldsymbol{\Gamma}
\end{aligned}
$$

-(5) is an eigensystem

$$
\begin{align*}
\longmapsto & \mathcal{R}(\boldsymbol{F})=\operatorname{span}\left(\left\{\boldsymbol{\psi}_{k_{i}}\right\}_{i=1}^{m}\right) \\
& \boldsymbol{\Gamma}=\operatorname{diag}\left(\lambda_{k_{1}}, \lambda_{k_{2}}, \ldots, \lambda_{k_{m}}\right) \tag{7}
\end{align*}
$$

$$
k_{i} \in\{1,2, \ldots, d\}
$$

$\square \mathcal{R}(\boldsymbol{F})=\mathcal{R}\left(\boldsymbol{B}^{\top} \boldsymbol{T}\right)=\mathcal{R}\left(\boldsymbol{B}^{\top}\right)$ (8)
$\square(6) \&(8) \quad \mathcal{R}\left(\boldsymbol{B}^{\top}\right)=\operatorname{span}\left(\left\{\boldsymbol{\psi}_{k_{i}}\right\}_{i=1}^{m}\right)$

## Proof (cont.)

$\square(2) \quad \operatorname{rank}(\boldsymbol{B})=m$
all $\left\{k_{i}\right\}_{i=1}^{m}$ are distinct
$\square$ We should choose the best $\left\{k_{i}\right\}_{i=1}^{m}$ that maximizes $\operatorname{tr}\left(\boldsymbol{B} \overline{\boldsymbol{C}} \boldsymbol{B}^{\top}\right)$.
$\square(4) \&(7) \longmapsto \operatorname{tr}\left(\boldsymbol{B} \overline{\boldsymbol{C}} \boldsymbol{B}^{\top}\right)=\operatorname{tr}\left(\boldsymbol{B} \boldsymbol{B}^{\top} \boldsymbol{T} \boldsymbol{\Gamma} \boldsymbol{T}^{\top}\right)$ $=\operatorname{tr}\left(\boldsymbol{T} \boldsymbol{\Gamma} \boldsymbol{T}^{\boldsymbol{\top}}\right)$
$=\operatorname{tr}\left(\boldsymbol{\Gamma} \boldsymbol{T}^{\top} \boldsymbol{T}\right)$
$=\sum_{i=1}^{m} \lambda_{k_{i}}$
${ }_{\square} \geq \lambda_{2} \geq \cdots \geq \lambda_{d}$
$k_{i}=i$ gives a solution.

- (9)

$$
\boldsymbol{B}=\left(\boldsymbol{\psi}_{1}\left|\boldsymbol{\psi}_{2}\right| \cdots \mid \boldsymbol{\psi}_{m}\right)^{\top}
$$

(Q.E.D.)

## Correlation

$\square$ Correlation coefficient for $\left\{s_{i}, t_{i}\right\}_{i=1}^{n}$ :

$$
\rho=\frac{\sum_{i=1}^{n}\left(s_{i}-\bar{s}\right)\left(t_{i}-\bar{t}\right)}{\sqrt{\left(\sum_{i=1}^{n}\left(s_{i}-\bar{s}\right)^{2}\right)\left(\sum_{i=1}^{n}\left(t_{i}-\bar{t}\right)^{2}\right)}}
$$

$$
\bar{s}=\sum_{i=1}^{n} s_{i} \quad \bar{t}=\sum_{i=1}^{n} t_{i}
$$





## PCA Uncorrelates Data

$$
\boldsymbol{B}_{P C A}=\left(\boldsymbol{\psi}_{1}\left|\boldsymbol{\psi}_{2}\right| \cdots \mid \boldsymbol{\psi}_{m}\right)^{\top}
$$

Covariance matrix of the PCAembedded samples is diagonal.

$$
\frac{1}{n} \sum_{i=1}^{n} \bar{z}_{i} \bar{z}_{i}^{\top}=\operatorname{diag}\left(\lambda_{1}, \lambda_{2}, \ldots, \lambda_{m}\right)
$$

(Homework)
Each element in $\overline{\boldsymbol{z}}$ is uncorrelated!

## Examples



- Data is well described.
- PCA is intuitive, easy to implement, analytic solution available, and fast.


## Examples (cont.)

19

Iris data (4d->2d) Letter data (16d->2d)


■ Embedded samples seem informative.

## Examples (cont.)




■ However, PCA does not necessarily preserve interesting information such as clusters.

## Homework

1. Implement PCA and reproduce the 2dimensional examples shown in the class.

- Data sets 1 and 2 are available from http://sugiyama-www.cs.titech.ac.jp/~sugi/data/DataAnalysis


- Test PCA on your own (artificial or real) data and analyze the characteristics of PCA.


## Homework (cont.)

2. Let

- B : $m \times d,(1 \leq m \leq d)$
- $\boldsymbol{C}, \boldsymbol{D}: d \times d$, positive definite, symmetric
- $\left\{\lambda_{i}, \boldsymbol{\psi}_{i}\right\}_{i=1}^{m}$ : Sorted generalized eigenvalues and normalized eigenvectors of $\boldsymbol{C} \boldsymbol{\psi}=\lambda \boldsymbol{D} \boldsymbol{\psi}$

$$
\lambda_{1} \geq \lambda_{2} \geq \cdots \geq \lambda_{d} \quad\left\langle\boldsymbol{D} \boldsymbol{\psi}_{i}, \boldsymbol{\psi}_{j}\right\rangle=\delta_{i, j}
$$

Prove that a solution of

$$
\boldsymbol{B}_{\text {min }}=\underset{\boldsymbol{B} \in \mathbb{R}^{m \times d}}{\operatorname{argmin}}\left[\operatorname{tr}\left(\boldsymbol{B} \boldsymbol{C} \boldsymbol{B}^{\top}\right)\right]
$$

is given by

$$
\text { subject to } \boldsymbol{B D} \boldsymbol{B}^{\top}=\boldsymbol{I}_{m}
$$

$$
\boldsymbol{B}_{\text {min }}=\left(\boldsymbol{\psi}_{d}\left|\boldsymbol{\psi}_{d-1}\right| \cdots \mid \boldsymbol{\psi}_{d-m+1}\right)^{\top}
$$

## Homework (cont.)

23
3. Prove that PCA uncorrelates the samples; more specifically, prove that the covariance matrix of the PCA-embedded samples is the following diagonal matrix:

$$
\sum_{i=1}^{n} \bar{z}_{i} \bar{z}_{i}^{\top}=\operatorname{diag}\left(\lambda_{1}, \lambda_{2}, \ldots, \lambda_{m}\right)
$$

$$
\overline{\boldsymbol{z}}_{i}=\boldsymbol{B}_{P C A} \overline{\boldsymbol{x}}_{i}
$$

$$
\boldsymbol{B}_{P C A}=\left(\boldsymbol{\psi}_{1}\left|\boldsymbol{\psi}_{2}\right| \cdots \mid \boldsymbol{\psi}_{m}\right)^{\top}
$$

## Suggestion

$\square$ Read the following article for upcoming classes:

- X. He \& P. Niyogi: Locality preserving projections, In Advances in Neural Information Processing Systems 16, MIT Press, Cambridge, MA, 2004.
http://books.nips.cc/papers/files/nips16/NIPS2003_AA20.pdf

