Big Data Fundamentals and Applications

Principal Component Analysis – PCA

Asst. Prof. Chan, Chun-Hsiang

Master program in Intelligent Computing and Big Data, Chung Yuan Christian University, Taoyuan, Taiwan
Undergraduate program in Intelligent Computing and Big Data, Chung Yuan Christian University, Taoyuan, Taiwan
Undergraduate program in Applied Artificial Intelligence, Chung Yuan Christian University, Taoyuan, Taiwan

Outlines

1. Review
2. Why Do We Need Dimension Reduction?
3. PCA – Mathematical Explanation
4. PCA – Programming in Python
5. Part VIII Basic numerical methods (Scipy)
   Part IX Advanced numerical methods (Scipy)
6. Notice

Review

- Before we explain PCA, we need to review the mathematical meaning of three basic descriptive statistics, including expectation, variance, and covariance.

$$\mu = \frac{1}{n}\sum_{i=1}^{n} x_i$$

$$var(x) = \sigma^2 = \left(\sqrt{\frac{\sum_{i=1}^{n}(x_i - \mu)^2}{n}}\right)^2 = \frac{1}{n}\sum_{i=1}^{n}(x_i - \mu)^2$$

$$Cov(x, y) = \frac{1}{n}\sum_{i=1}^{n}(x_i - \mu_x)(y_i - \mu_y)$$

- In previous courses or your understanding, these three parameters usually perform as above equations.

Review

$$\mu = \frac{1}{n}\sum_{i=1}^{n} x_i \;\mid\; var(x) = \sigma^2 = \frac{1}{n}\sum_{i=1}^{n}(x_i - \mu)^2 \;\mid\; Cov(x,y) = \frac{1}{n}\sum_{i=1}^{n}(x_i - \mu_x)(y_i - \mu_y)$$

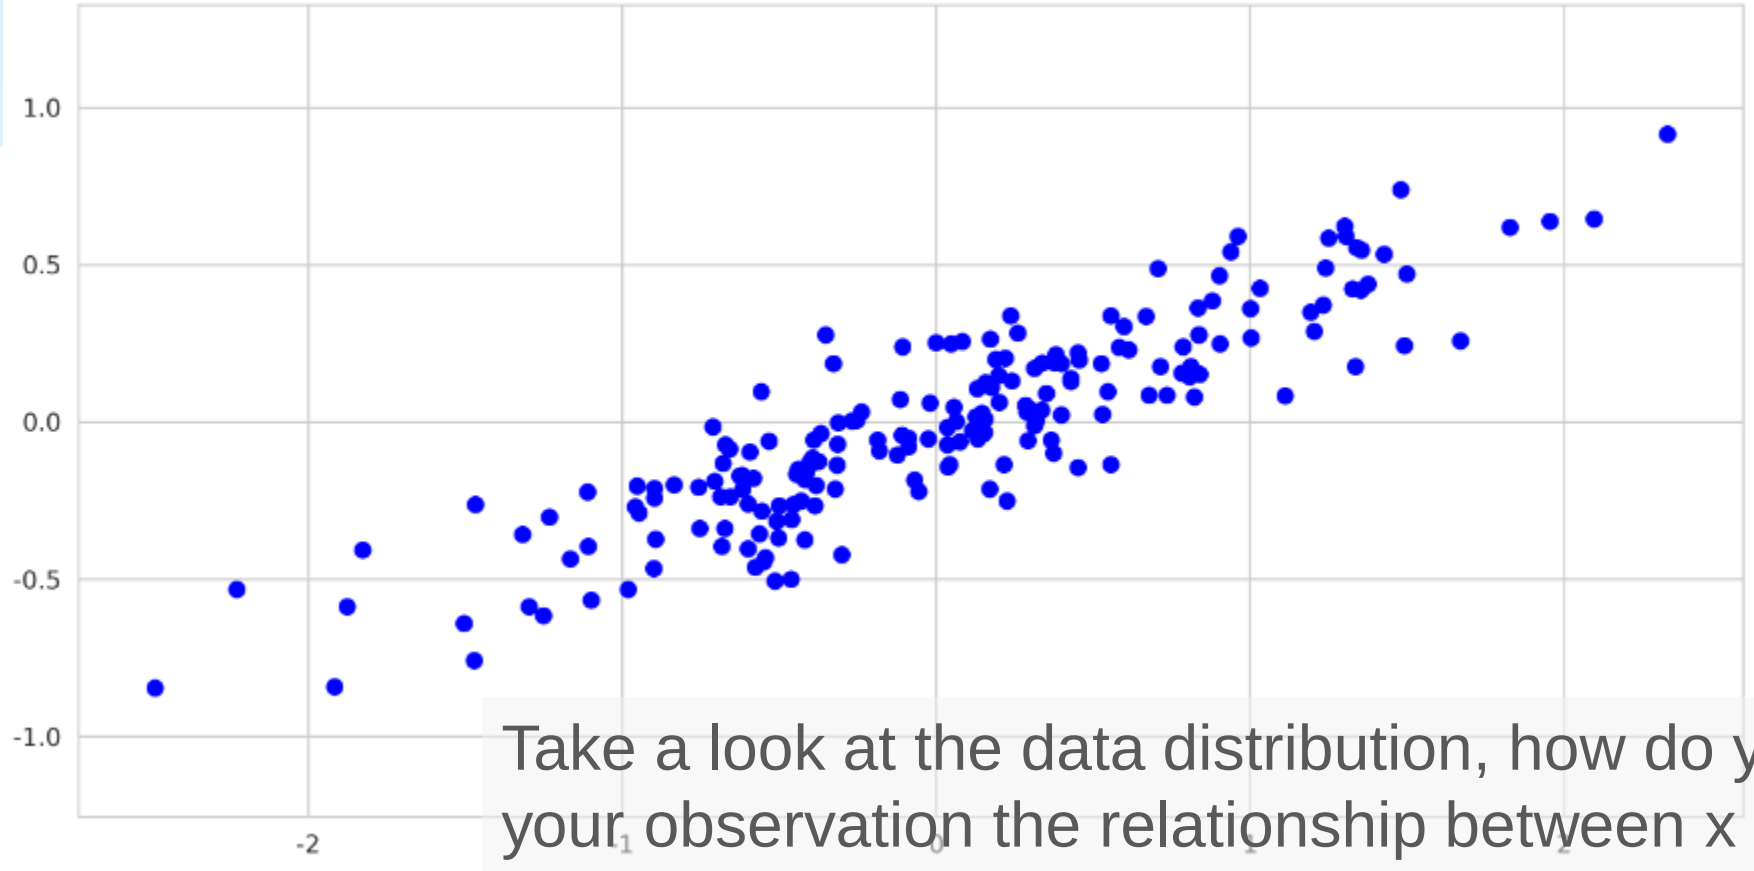

Take a look at the data distribution, how do you explain your observation the relationship between x and y?

Review

$$\mu = \frac{1}{n}\sum_{i=1}^{n} x_i \;\mid\; var(x) = \sigma^2 = \frac{1}{n}\sum_{i=1}^{n}(x_i - \mu)^2 \;\mid\; Cov(x,y) = \frac{1}{n}\sum_{i=1}^{n}(x_i - \mu_x)(y_i - \mu_y)$$

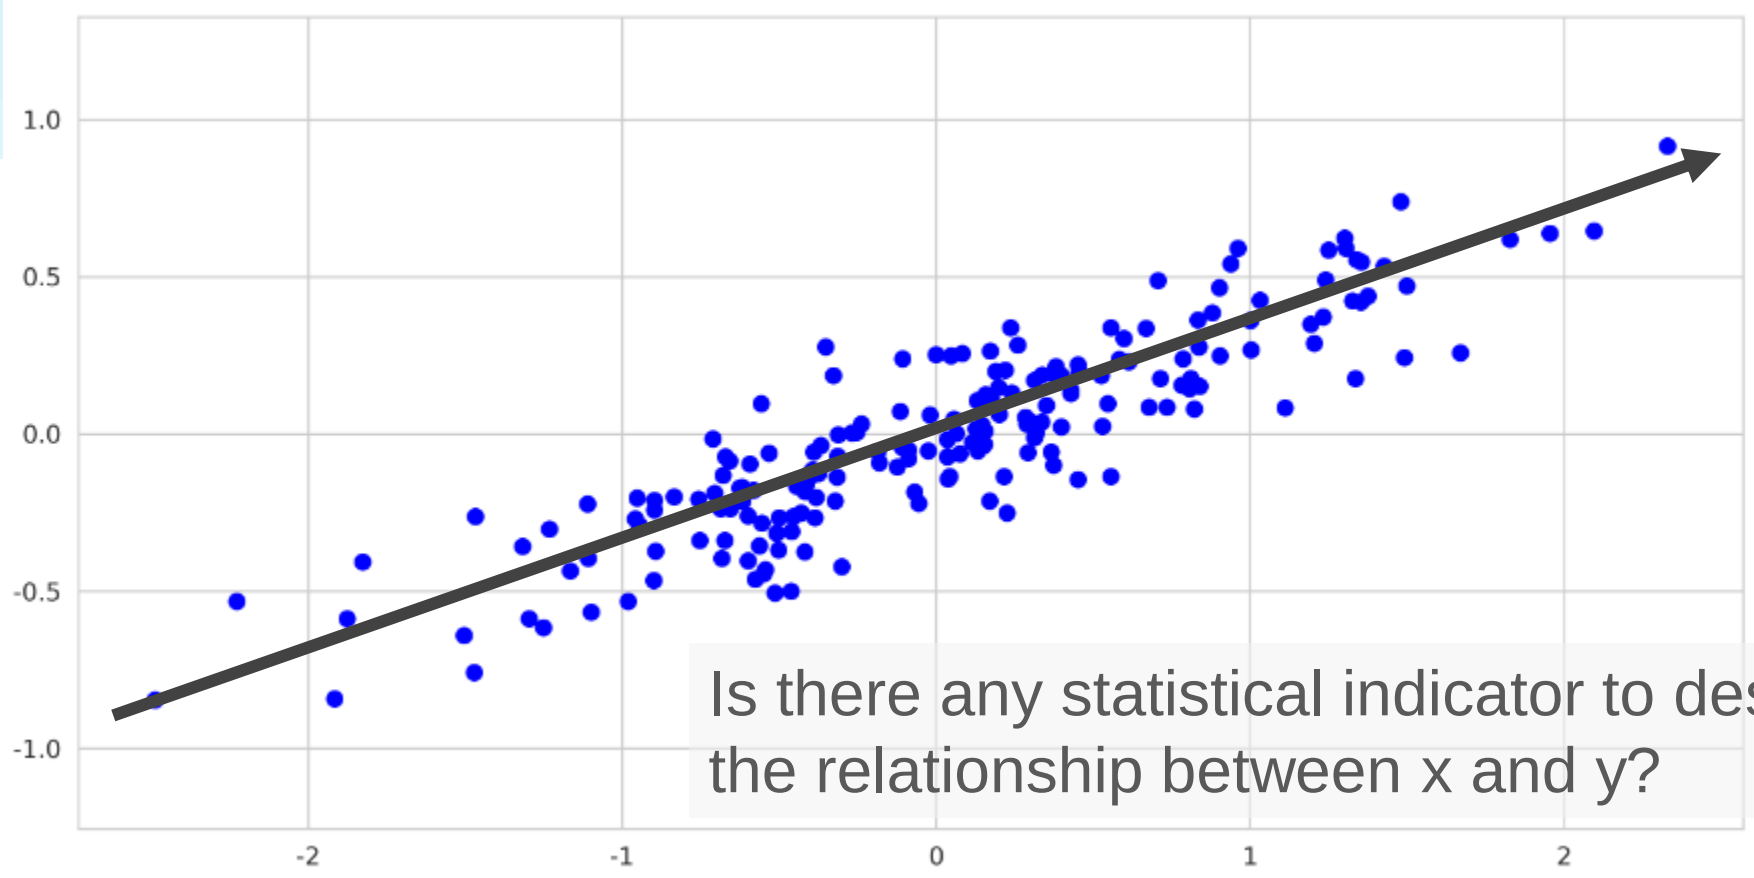

Is there any statistical indicator to describe the relationship between x and y?

Review

$$\mu = \frac{1}{n}\sum_{i=1}^{n} x_i \;\mid\; var(x) = \sigma^2 = \frac{1}{n}\sum_{i=1}^{n}(x_i - \mu)^2 \;\mid\; Cov(x,y) = \frac{1}{n}\sum_{i=1}^{n}(x_i - \mu_x)(y_i - \mu_y)$$

- Pearson‘s correlation coefficient
- Given two parameters x_i and y_i, where i ranges from 1 to n. Then, Pearson’s correlation coefficient could be defined as follows.

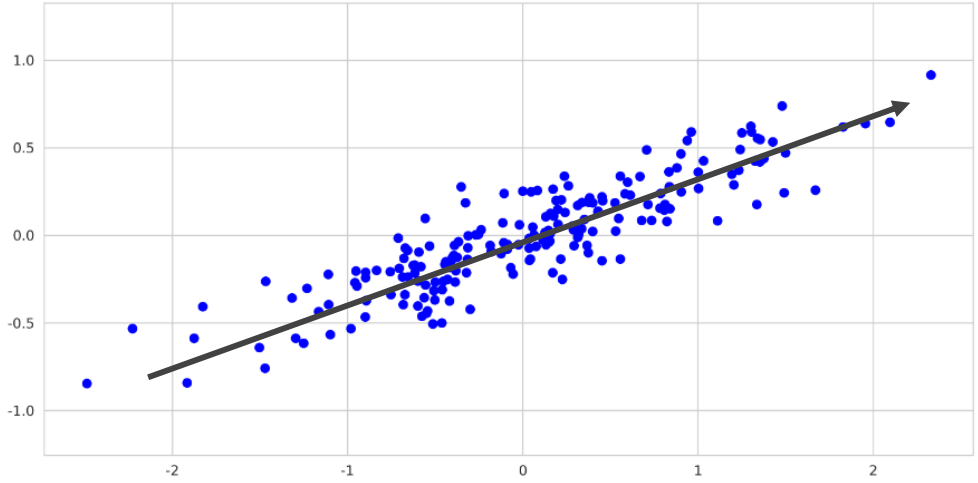

$$\rho = \frac{\sum_{i=1}^{n}(x_i - \mu_x)(y_i - \mu_y)}{\sqrt{\sum_{i=1}^{n}(x_i - \mu_x)^2 \sum_{i=1}^{n}(y_i - \mu_y)^2}}$$

Question 1

If x is highly correlated with y, and then what do you expect from their covariance and standard deviations?

Review

- **Expectation**

$$E(X) = \sum_{x} xP(X = x) = \mu$$

- **Variance**

$$var(X) = E([X - \mu]^2)$$
$$= E(X^2 - 2\mu X + \mu^2)$$
$$= E(X^2) - 2\mu E(X) + \mu^2$$
$$= E(X^2) - 2\mu^2 + \mu^2$$
$$= E(X^2) - \mu^2$$
$$= E(X^2) - E(X)^2$$

Characteristics of Expectation

$E(aX + bY) = aE(X) + bE(Y), \mathrm{a}, \mathrm{b} \in \mathbb{R}$

X and Y can be independent or dependent.

$E(XY) = E(X)E(Y)$

Where $cov(X, Y) = 0$

Review

$E(aX + bY) = aE(X) + bE(Y), \mathrm{a}, \mathrm{b} \in \mathbb{R}$
X and Y can be independent or dependent.
$E(XY) = E(X)E(Y)$
Where $cov(X, Y) = 0$

- Covariance
- If x and y are independent…

$$var(X + Y) = var(X) + var(Y)$$

- If x and y are dependent…

$$var(X + Y) = E([(X + Y) - E(X + Y)]^2)$$
$$= E\left(\left[(X + Y) - \left(E(X) + E(Y)\right)\right]^2\right)$$
$$= E\left(\left[\left(X - E(X)\right) + \left(Y - E(Y)\right)\right]^2\right)$$
$$= E\left(\left(X - E(X)\right)^2 + 2\left(X - E(X)\right)\left(Y - E(Y)\right) + \left(Y - E(Y)\right)^2\right)$$
$$= E\left[\left(X - E(X)\right)^2\right] + E\left[\left(Y - E(Y)\right)^2\right] + 2E\left[\left(X - E(X)\right)\left(Y - E(Y)\right)\right]$$
$$= var(X) + var(Y) + 2cov(X, Y)$$

Why Do We Need Dimension Reduction?

- Here comes the first question into your mind.
  - Why do we need dimension reduction?
  - What's the importance of dimension reduction?
  - Can we directly import all datasets into your model without dimension reduction?
- Statistical models (e.g., linear regression) have several assumptions when you adopt them. One of them is "all variables have to be linearly independent," indicating no collinearity.
- To achieve this goal, various methods were developed for orthogonalizing parameters and reducing the dimension of the dataset, such as PCA, LDA, LLE, and Laplacian Eigenmaps.

PCA – Math

$$\mu = \frac{1}{n}\sum_{i=1}^{n} x_i \;\mid\; var(x) = \sigma^2 = \frac{1}{n}\sum_{i=1}^{n}(x_i - \mu)^2 \;\mid\; Cov(x,y) = \frac{1}{n}\sum_{i=1}^{n}(x_i - \mu_x)(y_i - \mu_y)$$

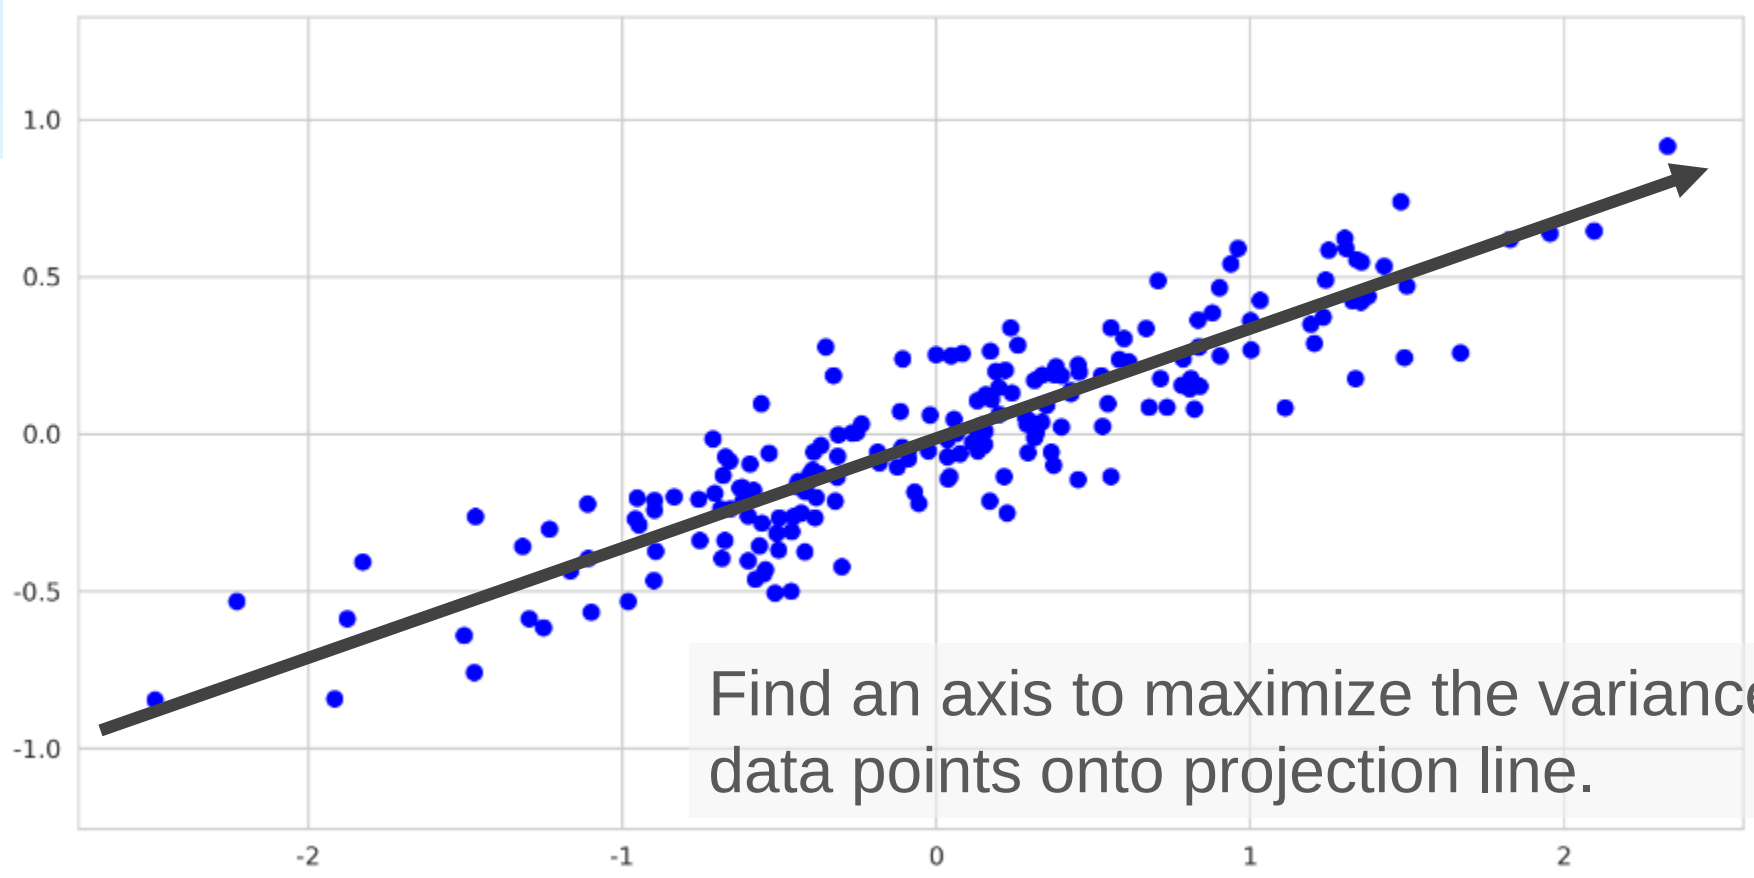

Find an axis to maximize the variance of all data points onto projection line.

PCA – Math

$$cos\theta = \frac{{x_i}^T \cdot v}{\|x_i\|\|v\|}$$

- Given a point "x" and project onto a vector "v".

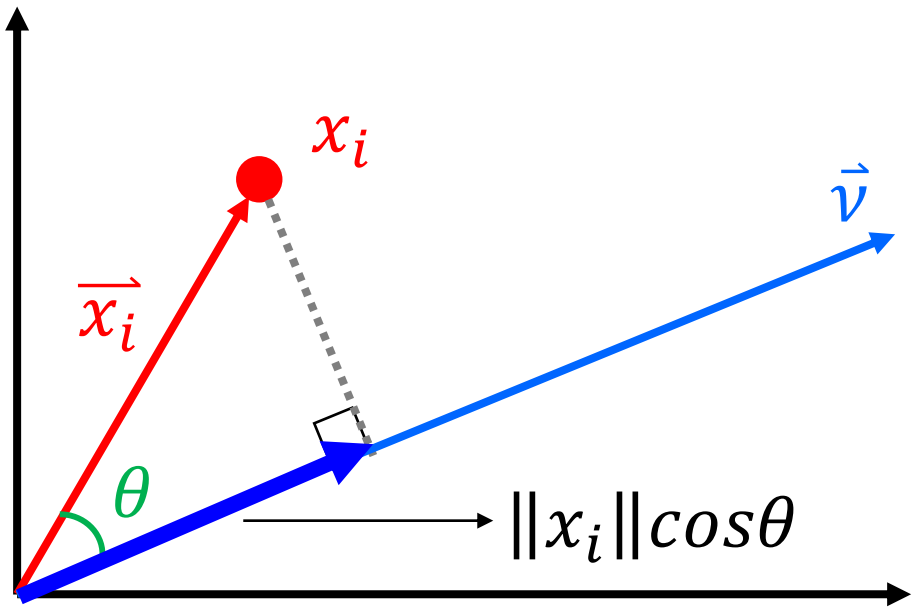

$$\|x_i\|cos\theta = \|x_i\|\frac{{x_i}^T \cdot v}{\|x_i\|\|v\|} = \frac{{x_i}^T \cdot v}{\|v\|}$$

$$if\ v\ is\ unit\ vector\ \ldots\ \|v\| = 1$$

$$= \frac{{x_i}^T \cdot v}{\|v\|} = {x_i}^T \cdot v$$

PCA – Math

$$X = \begin{bmatrix} | & | & \cdots & | \\ x_1 & x_2 & \cdots & x_n \\ | & | & \cdots & | \end{bmatrix} \rightarrow X^T = \begin{bmatrix} - & x_1 & - \\ - & x_2 & - \\ - & \vdots & - \\ - & x_n & - \end{bmatrix}$$

$$P = \begin{bmatrix} {x_1}^T \\ {x_2}^T \\ \vdots \\ {x_n}^T \end{bmatrix} = X^T u \Rightarrow solve\ P$$

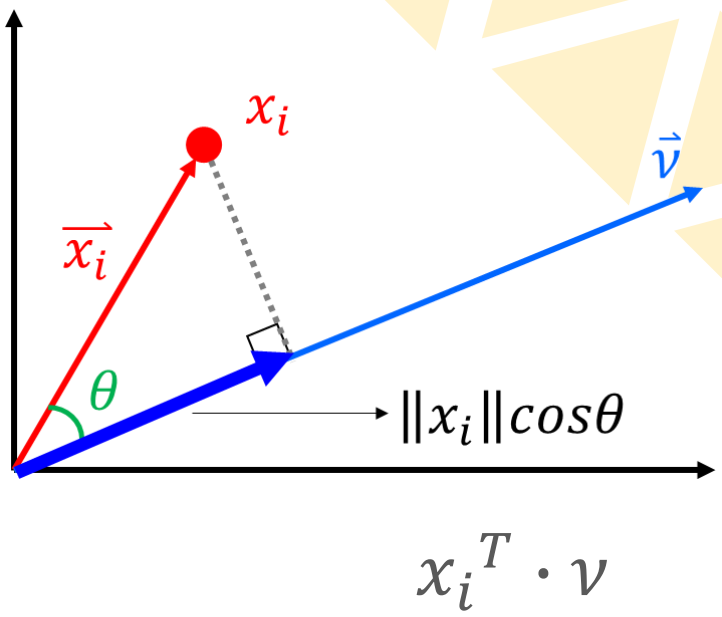

PCA – Math

$$J(u) = \|P^2\| = P^T P = (X^T u)^T (X^T u) = u^T X X^T u$$

$$\underset{u}{\operatorname{argmax}}\ J(u) = u^T X X^T u, where\ subject\ to\ u^T u = 1$$

Add Lagrange multiplier

$$\underset{u,\lambda}{\operatorname{argmax}}\ J(u,\lambda) = u^T X X^T u + \lambda(1 - u^T u)$$

$$\nabla_u J(u,\lambda) = \nabla_u \left(u^T X X^T u + \lambda(1 - u^T u)\right) = 0$$

$$\Rightarrow 2XX^T u - 2\lambda u = 0$$

$$\Rightarrow \underbrace{\boxed{XX^T}}_{cov(X)} u = \underbrace{\boxed{\lambda}}_{\text{eigenvalue}} \boxed{u} \rightarrow \text{eigenvector} \longrightarrow Au = \lambda u$$

PCA – Math

$$XX^T u = \lambda u$$

Eigenvector is a symmetry matrix

u is an unit vector

$$uu^T = u^T u = 1$$

$when\ u\ is\ the\ eigen\ vector$

$$J(u) = \|P^2\| = u^T XX^T u = u^T \lambda u = \lambda u^T u = \lambda$$

Given an ***eigenvector****, the* ***total square of projected values*** *is the* ***eigenvalue*** *= λ*

PCA – Math

$$XX^T u = \lambda u$$

Eigenvector is a symmetry matrix

u is an unit vector

$$uu^T = u^T u = 1$$

A is a square symmetric matrix has orthogonal eigenvectors with different eigenvalues.

$[x_1, \lambda_1], [x_2, \lambda_2]$

$$\begin{cases} Ax_1 = \lambda_1 x_1 \\ Ax_2 = \lambda_2 x_2 \end{cases}$$

$$x_1{}^T A x_2 = x_1{}^T \lambda_2 x_2 = \lambda_2 x_1{}^T x_2$$

Equal

$$x_1{}^T A^T x_2 = (Ax_1)^T x_2 = (\lambda_1 x_1)^T x_2 = \lambda_1 x_1{}^T x_2$$

$$(\because A \in symmetric\ matrix, \therefore A = A^T)$$

$$\lambda_2 x_1{}^T x_2 = \lambda_1 x_1{}^T x_2$$

$$\underline{x_1{}^T x_2}\,\underline{(\lambda_2 - \lambda_1)} = 0$$

Orthogonal

$x_1{}^T x_2 = 0$

All eigenvalues are different

PCA – Math

- Conversion between orthogonal bases

$$u_i \cdot u_j = u_i^T \cdot u_j = \begin{cases} 1, if\ i = j \\ 0, otherwise \end{cases}$$

$$U = \begin{bmatrix} | & | & & | \\ u_1 & u_2 & \cdots & u_d \\ | & | & & | \end{bmatrix} \Rightarrow U^T U = I = U^{-1} = U^T$$

$$x = y_1 u_1 + y_2 u_2 + \cdots + y_d u_d = \begin{bmatrix} | & | & & | \\ u_1 & u_2 & \cdots & u_d \\ | & | & & | \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_d \end{bmatrix} = Uy$$

$$\Rightarrow y = U^{-1} x = U^T x$$

PCA – Math

$$XX^T v = \lambda v \Rightarrow AV = \lambda V$$

where V is eigenvector and λ is eigenvalue.

Principal component PC

$$AV = \lambda V$$

From Singular Vector Decomposition (SVD)

$$A = U\Sigma V^T \Rightarrow AV = U\Sigma \Rightarrow \lambda = \frac{\Sigma^2}{N}, \therefore AV = U\Sigma$$

Variable loading L (A onto PC)

$$Standardize\ PC = \sqrt{N}U$$

$$cov(A, PC) = \frac{A^T\sqrt{N}U}{N} = \frac{V\Sigma U^T U}{\sqrt{N}} = V\frac{\Sigma}{\sqrt{N}} = V\sqrt{E} = L$$

PCA – Math

$$J(u) = \|P^2\| = P^T P = (X^T u)^T (X^T u) = u^T X X^T u$$

$$\underset{u}{\operatorname{argmax}}\ J(u) = u^T X X^T u, where\ subject\ to\ u^T u = 1$$

Add Lagrange multiplier

$$\underset{u,\lambda}{\operatorname{argmax}}\ J(u,\lambda) = u^T X X^T u - \lambda(1 - u^T u)$$

$$\nabla_u J(u,\lambda) = \nabla_u\left(u^T X X^T u - \lambda(1 - u^T u)\right) = 0$$

$$\Rightarrow 2XX^T u - 2\lambda u = 0$$

$$\Rightarrow \underbrace{\boxed{XX^T}}_{cov(X)} u = \underbrace{\boxed{\lambda}}_{\text{eigenvalue}}\boxed{u} \rightarrow \text{eigenvector} \longrightarrow Au = \lambda u$$

12

PCA – Math

The steps of PCA

1. Find the sample mean $\mu = \frac{1}{n}\sum_{i=1}^{n} x_i$
2. Subtract mean
3. Compute covariance matrix $C = \frac{1}{n}XX^T = \frac{1}{n}\sum_{i=1}^{n}(x_i - \mu)\,(x_i - \mu)$
4. Find the eigenvalues of C and arrange them into descending order

   $\lambda_1 > \lambda_2 > \cdots > \lambda_d, \{u_1, u_2, \ldots, u_d\}$
5. The transformation is $y = U^T X$.

PCA – Python

scikit learn | Install | User Guide | API | Examples | Community | More | Go

Prev | Up | Next

scikit-learn 1.1.2
Other versions

Please **cite us** if you use the software.

`sklearn.decomposition.PCA`
Examples using `sklearn.decomposition.PCA`

`sklearn.decomposition.PCA`

class `sklearn.decomposition.`**`PCA`**(*n_components=None, *, copy=True, whiten=False, svd_solver='auto', tol=0.0, iterated_power='auto', n_oversamples=10, power_iteration_normalizer='auto', random_state=None*) [source]

Principal component analysis (PCA).

Linear dimensionality reduction using Singular Value Decomposition of the data to project it to a lower dimensional space. The input data is centered but not scaled for each feature before applying the SVD.

It uses the LAPACK implementation of the full SVD or a randomized truncated SVD by the method of Halko et al. 2009, depending on the shape of the input data and the number of components to extract.

It can also use the scipy.sparse.linalg ARPACK implementation of the truncated SVD.

Notice that this class does not support sparse input. See `TruncatedSVD` for an alternative with sparse data.

Read more in the User Guide.

Parameters::

n_components : *int, float or 'mle', default=None*

Number of components to keep. if n_components is not set all components are kept:

```
n_components == min(n_samples, n_features)
```

If `n_components == 'mle'` and `svd_solver == 'full'`, Minka's MLE is used to guess the dimension. Use of `n_components == 'mle'` will interpret `svd_solver == 'auto'` as `svd_solver == 'full'`.

If `0 < n_components < 1` and `svd_solver == 'full'`, select the number of components such that the amount of variance that needs to be explained is greater than the percentage specified by n_components.

If `svd_solver == 'arpack'`, the number of components must be strictly less than the minimum of n_features and n_samples.

PCA – Python

```
# linear algebra
from sklearn.decomposition import PCA
# call PCA
pca = PCA(n_components=2) # number of preserved components
pca.fit(X)
# show results
print(pca.mean_) # mean
print(pca.explained_variance_) # eigenvalues
print(pca.components_) # eigenvectors
```

Notice

- Advantages of PCA
  - Easy to use
  - Accelerate model fitting
  - Avoid overfitting
- Disadvantage of PCA
  - Low interpretability
  - The trade-off between information loss and features (dimensions)
- Limitations of PCA
  - Linear relationship between features
  - Correlation between features
  - Sensitive to the scale of data
  - Not robust to outliers
  - Not accept missing values

Source: https://www.keboola.com/blog/pca-machine-learning

Question Time

If you have any questions, please do not hesitate to ask me.

References:

張智星 科學計算講義

Sklearn

The End

Thank you for your attention))