

Multivariate Statistics

Lecture 12

Fudan University

1 Principal Components

Outline

1 Principal Components

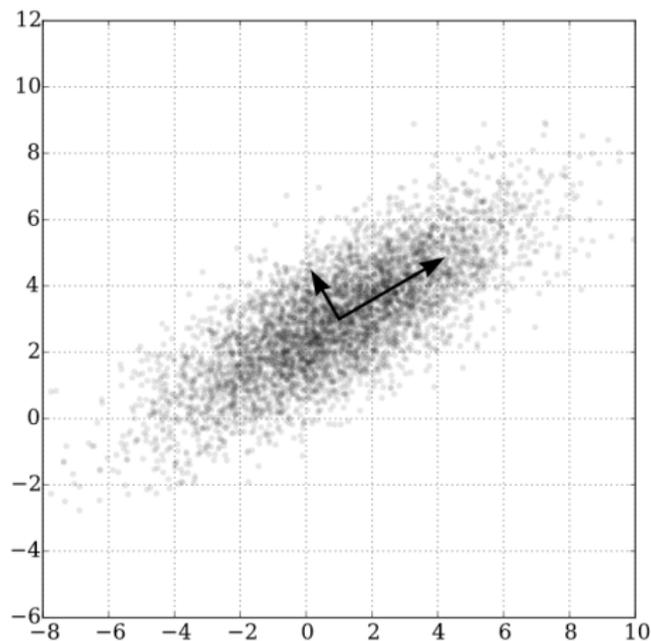
2 Canonical Correlation

1 Principal Components

2 Canonical Correlation

Principal Components

In statistical practice, the method of principal components is used to find the linear combinations with large variance.



Principal Components

Let random vector \mathbf{x} of p component has mean $\mathbf{0}$ and covariance matrix $\mathbf{\Sigma}$.

Let β be a p -component column vector such that $\|\beta\|_2 = 1$.

- ① The variance of $\beta^\top \mathbf{x}$ is

$$\mathbb{E}[(\beta^\top \mathbf{x})^2] = \beta^\top \mathbb{E}[\mathbf{x}\mathbf{x}^\top] \beta = \beta^\top \mathbf{\Sigma} \beta.$$

- ② Maximizing $\beta^\top \mathbf{\Sigma} \beta$ must satisfy

$$(\mathbf{\Sigma} - \lambda_1 \mathbf{I}) \beta = \mathbf{0},$$

where λ_1 is the largest root of

$$\det(\mathbf{\Sigma} - \lambda \mathbf{I}) = 0.$$

- ③ Let $\beta^{(1)} = \arg \max_{\|\beta\|_2=1} \beta^\top \mathbf{\Sigma} \beta$.

Principal Components

Find β such that $\beta^\top \mathbf{x}$ has maximum variance and is uncorrelated with $u_1 = \beta^{(1)\top} \mathbf{x}$.

- 1 Lack of correlation means

$$\beta^\top \beta^{(1)} = 0.$$

- 2 The vector β must satisfy

$$(\mathbf{\Sigma} - \lambda_2 \mathbf{I})\beta = \mathbf{0},$$

where λ_2 is the second largest root of

$$\det(\mathbf{\Sigma} - \lambda \mathbf{I}) = 0.$$

Principal Components

At the $(r + 1)$ -th step, we want to find a vector such that $\beta^\top \mathbf{x}$ has maximum variance and lacks correlation with $u_1 \dots, u_r$, that is

$$0 = \mathbb{E}[\beta^\top \mathbf{x} u_i] = \mathbb{E}[\beta^\top \mathbf{x} \mathbf{x}^\top \beta^{(i)}] = \beta^\top \boldsymbol{\Sigma} \beta^{(i)} = \lambda \beta^\top \beta^{(i)}$$

for $i = 1, \dots, r$, where $u_i = \beta^{(i)\top} \mathbf{x}$

Finally, we obtain $\beta^{(1)}, \dots, \beta^{(p)}$ and $\lambda_1 \geq \dots \geq \lambda_p$ such that

$$\boldsymbol{\Sigma} \mathbf{B} = \mathbf{B} \boldsymbol{\Lambda}$$

where $\mathbf{B} = [\beta^{(1)}, \dots, \beta^{(p)}]$ satisfying $\mathbf{B}^\top \mathbf{B} = \mathbf{I}$ and

$$\boldsymbol{\Lambda} = \begin{bmatrix} \lambda_1 & 0 & \dots & 0 \\ 0 & \lambda_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \lambda_p \end{bmatrix}.$$

Principal Components

The transformation

$$\mathbf{u} = \mathbf{B}^T \mathbf{x}$$

leads to the r -th component of \mathbf{u} has maximum variance of all normalized linear combinations uncorrelated with u_1, \dots, u_{r-1} .

The vector \mathbf{u} is defined as the vector of principal components of \mathbf{x} .

Principal Components

Theorem 1

An orthogonal transformation $\mathbf{v} = \mathbf{C}\mathbf{x}$ of a random vector \mathbf{x} with $\mathbb{E}[\mathbf{x}] = \mathbf{0}$ leaves invariant the generalized variance and the sum of the variances of the components.

Corollary 1

The generalized variance of the vector of principal components is the generalized variance of the original vector, and the sum of the variances of the principal components is the sum of the variances of the original variates.

Principal Components

Another approach is based on the density of the normal distribution.

- 1 Let $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma})$, whose density is

$$\frac{1}{(2\pi)^{\frac{p}{2}} (\det(\mathbf{\Sigma}))^{\frac{1}{2}}} \exp\left(-\frac{1}{2} \mathbf{x}^{\top} \mathbf{\Sigma}^{-1} \mathbf{x}\right).$$

- 2 Surfaces of constant density are ellipsoids

$$\{\mathbf{x} : \mathbf{x}^{\top} \mathbf{\Sigma}^{-1} \mathbf{x} = C\}$$

- 3 A principal axis of this ellipsoid is defined as the line from $-\mathbf{y}$ to \mathbf{y} , where \mathbf{y} is a point on the ellipsoid where its squared distance has a stationary point.

Maximum Likelihood Estimators of Principal Components

Let $\mathbf{x}_1, \dots, \mathbf{x}_N$ be N observations from $\mathcal{N}_p(\mathbf{0}, \mathbf{\Sigma})$, where $\mathbf{\Sigma}$ has p different characteristic roots and $N > p$. Then a set of maximum likelihood estimators of $\lambda_1, \dots, \lambda_p$ and $\beta^{(1)}, \dots, \beta^{(p)}$ consists of the roots $\lambda_1 > \dots > \lambda_p$ of

$$\det(\hat{\mathbf{\Sigma}} - \lambda \mathbf{I}) = 0$$

and a set of vectors $\hat{\beta}^{(1)}, \dots, \hat{\beta}^{(p)}$ satisfying $\|\hat{\beta}^{(i)}\|_2 = 1$ and

$$(\hat{\mathbf{\Sigma}} - \lambda_i \mathbf{I})\hat{\beta}^{(i)} = \mathbf{0}$$

for $i = 1, \dots, p$, where $\hat{\mathbf{\Sigma}}$ is the the maximum likelihood estimate of $\mathbf{\Sigma}$.

Asymptotic Distributions

Let the characteristic roots of Σ be $\lambda_1 > \dots > \lambda_p$ and the corresponding characteristic vectors be $\beta^{(1)}, \dots, \beta^{(p)}$ with $\|\beta^{(i)}\|_2 = 1$ and $\beta_{1i} \geq 0$.

Let the characteristic roots of S be $l_1 > \dots > l_p$ and the corresponding characteristic vectors be $\mathbf{b}^{(1)}, \dots, \mathbf{b}^{(p)}$ with $\|\mathbf{b}^{(i)}\|_2 = 1$ and $b_{1i} \geq 0$.

Let $d_i = \sqrt{n}(l_i - \lambda_i)$ and $\mathbf{g}^{(i)} = \mathbf{b}^{(i)} - \beta^{(i)}$ for $i = 1, \dots, p$.

- 1 The limiting normal distribution the sets $\{d_1, \dots, d_p\}$ and $\{\mathbf{g}^{(1)}, \dots, \mathbf{g}^{(p)}\}$ are independent and d_1, \dots, d_p are mutually independent.
- 2 The element d_i has the limiting distribution $\mathcal{N}(0, 2\lambda_i^2)$.
- 3 The random vectors $\mathbf{g}^{(1)}, \dots, \mathbf{g}^{(p)}$ has limiting normal distribution with

$$\lim_{n \rightarrow \infty} \text{Cov}[\mathbf{g}^{(i)}, \mathbf{g}^{(j)}] = \begin{cases} \sum_{k=1, k \neq i}^p \frac{\lambda_i \lambda_k}{(\lambda_i - \lambda_k)^2} \beta^{(k)} \beta^{(k)\top}, & i = j \\ -\frac{\lambda_i \lambda_j}{(\lambda_i - \lambda_j)^2} \beta^{(j)} \beta^{(i)\top}, & i \neq j \end{cases}$$

Asymptotic Distributions

One treats l_i as approximately normal with mean λ_i and variance $2\lambda_i^2/n$. Since l_i is a consistent estimate of λ_i , the limiting distribution of

$$\frac{\sqrt{n}(l_i - \lambda_i)}{\sqrt{2} l_i}$$

is $\mathcal{N}(0, 1)$.

- ① A two-tailed test of hypothesis $\lambda_i = \lambda_i^0$ has (asymptotic) acceptance region

$$-z(\epsilon) \leq \frac{\sqrt{n}(l_i - \lambda_i^0)}{\sqrt{2} \lambda_i^0} \leq z(\epsilon)$$

where the value of the $\mathcal{N}(0, 1)$ distribution beyond $z(\epsilon)$ is $\epsilon/2$.

- ② The confidence interval for λ_i with confidence $1 - \epsilon$ is

$$\frac{l_i}{1 + \sqrt{2/n} z(\epsilon)} \leq \lambda_i \leq \frac{l_i}{1 - \sqrt{2/n} z(\epsilon)}.$$

Exact Confidence Limits on the Characteristic Roots

Let $n\mathbf{S} \sim \mathcal{W}(\boldsymbol{\Sigma}, n)$, then

$$\frac{n\boldsymbol{\beta}^{(1)\top} \mathbf{S} \boldsymbol{\beta}^{(1)}}{\lambda_1} \quad \text{and} \quad \frac{n\boldsymbol{\beta}^{(p)\top} \mathbf{S} \boldsymbol{\beta}^{(p)}}{\lambda_p}.$$

are independently distributed as χ^2 -distribution with n degrees of freedom.

Let l and u be two numbers such that

$$1 - \epsilon = \Pr \{nl \leq \chi_n^2\} \Pr \{\chi_n^2 \leq nu\}.$$

Then a confidence interval for the characteristic roots of $\boldsymbol{\Sigma}$ with confidence at least $1 - \epsilon$ is

$$\frac{l_p}{u} \leq \lambda_p \leq \lambda_1 \leq \frac{l_1}{l}.$$

Outline

1 Principal Components

2 Canonical Correlation

Canonical Correlations

We still consider random vector \mathbf{x} of p components has zero means and the covariance matrix $\Sigma \succ \mathbf{0}$.

We partition \mathbf{x} into two subvectors of p_1 and p_2 components ($p_1 \leq p_2$)

$$\mathbf{x} = \begin{bmatrix} \mathbf{x}^{(1)} \\ \mathbf{x}^{(2)} \end{bmatrix}.$$

The covariance matrix is partitioned into p_1 and p_2 rows and columns

$$\Sigma = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{bmatrix}.$$

Here we shall develop a transformation of $\mathbf{x}^{(1)}$ and another transformation of $\mathbf{x}^{(2)}$ to a new system that exhibit clearly the intercorrelations between $\mathbf{x}^{(1)}$ and $\mathbf{x}^{(2)}$.

Canonical Correlations

Consider linear combinations

$$u = \alpha^\top \mathbf{x}^{(1)} \quad \text{and} \quad v = \gamma^\top \mathbf{x}^{(2)}.$$

We ask for α and γ that maximize the correlation between u and v .

- ① We require α and γ such that

$$1 = \mathbb{E}[u^2] = \mathbb{E}[\alpha^\top \mathbf{x}^{(1)} \mathbf{x}^{(1)\top} \alpha] = \alpha^\top \Sigma_{11} \alpha,$$

$$1 = \mathbb{E}[v^2] = \mathbb{E}[\gamma^\top \mathbf{x}^{(2)} \mathbf{x}^{(2)\top} \gamma] = \gamma^\top \Sigma_{22} \gamma.$$

- ② The correlation between u and v is

$$\mathbb{E}[uv] = \mathbb{E}[\alpha^\top \mathbf{x}^{(1)} \mathbf{x}^{(2)\top} \gamma] = \alpha^\top \Sigma_{12} \gamma.$$

- ③ Then the problem is

$$\begin{aligned} \max \quad & \alpha^\top \Sigma_{12} \gamma. \\ \alpha^\top \Sigma_{11} \alpha = & 1 \\ \gamma^\top \Sigma_{22} \gamma = & 1 \end{aligned}$$

Canonical Correlations

The solution of

$$\begin{aligned} \max & \quad \alpha^\top \Sigma_{12} \gamma. \\ \alpha^\top \Sigma_{11} \alpha &= 1 \\ \gamma^\top \Sigma_{22} \gamma &= 1 \end{aligned}$$

must satisfy

$$\begin{bmatrix} -\lambda \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & -\lambda \Sigma_{22} \end{bmatrix} \begin{bmatrix} \alpha \\ \gamma \end{bmatrix} = \mathbf{0},$$

where λ is the root of

$$\det \left(\begin{bmatrix} -\lambda \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & -\lambda \Sigma_{22} \end{bmatrix} \right) = 0.$$

Denote the largest root and the corresponds vectors be λ_1 , $\alpha^{(1)}$ and $\gamma^{(1)}$

Canonical Correlations

Then we consider $u = \alpha^\top \mathbf{x}^{(1)}$ and $v = \gamma^\top \mathbf{x}^{(2)}$ for $\mathbf{x}^{(2)}$ with maximum correlation, such that u is uncorrelated with $u_1 = \alpha^{(1)\top} \mathbf{x}^{(1)}$ and v is uncorrelated with $v_1 = \gamma^{(1)\top} \mathbf{x}^{(2)}$.

This procedure is continued. At r -th step, we have

$$\begin{aligned} u_1 &= \alpha^{(1)\top} \mathbf{x}^{(1)}, \dots, u_r = \alpha^{(r)\top} \mathbf{x}^{(1)} \\ v_1 &= \gamma^{(1)\top} \mathbf{x}^{(2)}, \dots, v_r = \gamma^{(r)\top} \mathbf{x}^{(2)} \end{aligned}$$

and each of them are uncorrelated. Let the correlation between u_i and v_i be λ_i .

We obtain α^{r+1} and $\gamma^{(r+1)}$ by maximizing the correlation between $u = \alpha^\top \mathbf{x}^{(1)}$ and $v = \gamma^\top \mathbf{x}^{(2)}$ such that u is uncorrelated with u_1, \dots, u_r and v is uncorrelated with v_1, \dots, v_r .

Canonical Correlations

Let $\mathbf{A} = [\boldsymbol{\alpha}^{(1)}, \dots, \boldsymbol{\alpha}^{(p_1)}]$, $\boldsymbol{\Gamma} = [\boldsymbol{\Gamma}_1, \boldsymbol{\Gamma}_2] = [\boldsymbol{\gamma}^{(1)}, \dots, \boldsymbol{\gamma}^{(p_2)}]$ and

$$\boldsymbol{\Lambda} = \begin{bmatrix} \lambda_1 & 0 & \dots & 0 \\ 0 & \lambda_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \lambda_{p_1} \end{bmatrix}.$$

All of conditions can be summarized as

$$\mathbf{A}^\top \boldsymbol{\Sigma}_{11} \mathbf{A} = \mathbf{I},$$

$$\mathbf{A}^\top \boldsymbol{\Sigma}_{12} \boldsymbol{\Gamma}_1 = \boldsymbol{\Lambda},$$

$$\boldsymbol{\Gamma}_1^\top \boldsymbol{\Sigma}_{22} \boldsymbol{\Gamma}_1 = \mathbf{I},$$

$$\boldsymbol{\Gamma}_2^\top \boldsymbol{\Sigma}_{22} \boldsymbol{\Gamma}_1 = \mathbf{0},$$

$$\boldsymbol{\Gamma}_2^\top \boldsymbol{\Sigma}_{22} \boldsymbol{\Gamma}_2 = \mathbf{I}.$$

Canonical Correlations

Each $\alpha^{(i)}$, $\gamma^{(i)}$ can be obtained by solving

$$\begin{bmatrix} -\lambda_i \mathbf{\Sigma}_{11} & \mathbf{\Sigma}_{12} \\ \mathbf{\Sigma}_{21} & -\lambda_i \mathbf{\Sigma}_{22} \end{bmatrix} \begin{bmatrix} \alpha \\ \gamma \end{bmatrix} = \mathbf{0},$$

where λ_i is the i -th largest root of

$$\det \left(\begin{bmatrix} -\lambda \mathbf{\Sigma}_{11} & \mathbf{\Sigma}_{12} \\ \mathbf{\Sigma}_{21} & -\lambda \mathbf{\Sigma}_{22} \end{bmatrix} \right) = 0.$$

This can be written as generalized eigenvalue problems

$$(\mathbf{\Sigma}_{12} \mathbf{\Sigma}_{22}^{-1} \mathbf{\Sigma}_{21} - \lambda^2 \mathbf{\Sigma}_{11}) \gamma = \mathbf{0}$$

and

$$(\mathbf{\Sigma}_{21} \mathbf{\Sigma}_{11}^{-1} \mathbf{\Sigma}_{12} - \lambda^2 \mathbf{\Sigma}_{22}) \alpha = \mathbf{0}.$$

Canonical Correlations

Let

$$\mathbf{x} = \begin{bmatrix} \mathbf{x}^{(1)} \\ \mathbf{x}^{(2)} \end{bmatrix}$$

be a random vector where $\mathbf{x}^{(1)}$ has p_1 components and $\mathbf{x}^{(2)}$ has p_2 components.

In the r -th pair of canonical variates is the pair of linear combinations

$$u_r = \boldsymbol{\alpha}^{(r)\top} \mathbf{x}^{(1)} \quad \text{and} \quad v_r = \boldsymbol{\gamma}^{(r)\top} \mathbf{x}^{(2)},$$

each of unit variance and uncorrelated with the first $r - 1$ pairs of canonical variates and having maximum correlation.

The correlation between u_r and v_r is the r -th canonical correlation.

Canonical Correlations

The canonical correlations are invariant with respect to transformations

$$\begin{cases} \mathbf{x}^{*(1)} = \mathbf{C}_1 \mathbf{x}^{(1)}, \\ \mathbf{x}^{*(2)} = \mathbf{C}_2 \mathbf{x}^{(2)}, \end{cases}$$

where \mathbf{C}_1 and \mathbf{C}_2 are non-singular. Additionally, any function of $\boldsymbol{\Sigma}$ that is invariant (under any such transformation) is a function of the canonical correlations.

Estimation of Canonical Correlation

Let $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}$ be N observations from $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. Let each $\mathbf{x}^{(\alpha)}$ be partitioned into sub-vectors of p_1 and p_2 components

$$\mathbf{x} = \begin{bmatrix} \mathbf{x}^{(1)} \\ \mathbf{x}^{(2)} \end{bmatrix}.$$

The maximum likelihood of $\boldsymbol{\Sigma}$ is

$$\begin{aligned} \hat{\boldsymbol{\Sigma}} &= \begin{bmatrix} \hat{\boldsymbol{\Sigma}}_{11} & \hat{\boldsymbol{\Sigma}}_{12} \\ \hat{\boldsymbol{\Sigma}}_{21} & \hat{\boldsymbol{\Sigma}}_{22} \end{bmatrix} \\ &= \frac{1}{N} \begin{bmatrix} \sum_{\alpha=1}^N (\mathbf{x}_{\alpha}^{(1)} - \bar{\mathbf{x}}^{(1)})(\mathbf{x}_{\alpha}^{(1)} - \bar{\mathbf{x}}^{(1)})^{\top} & \sum_{\alpha=1}^N (\mathbf{x}_{\alpha}^{(1)} - \bar{\mathbf{x}}^{(1)})(\mathbf{x}_{\alpha}^{(2)} - \bar{\mathbf{x}}^{(2)})^{\top} \\ \sum_{\alpha=1}^N (\mathbf{x}_{\alpha}^{(2)} - \bar{\mathbf{x}}^{(2)})(\mathbf{x}_{\alpha}^{(1)} - \bar{\mathbf{x}}^{(1)})^{\top} & \sum_{\alpha=1}^N (\mathbf{x}_{\alpha}^{(2)} - \bar{\mathbf{x}}^{(2)})(\mathbf{x}_{\alpha}^{(2)} - \bar{\mathbf{x}}^{(2)})^{\top} \end{bmatrix}. \end{aligned}$$

The maximum likelihood estimators $\hat{\boldsymbol{\Lambda}}$, $\hat{\mathbf{A}}$ and $\hat{\boldsymbol{\Gamma}}$ of $\boldsymbol{\Lambda}$, \mathbf{A} and $\boldsymbol{\Gamma}$ involve applying the algebra of the previous slides to $\hat{\boldsymbol{\Sigma}}$.

Sample Canonical Variates and Correlations

We could define the sample canonical variates and correlations in terms of

$$\mathbf{S} = \frac{1}{N-1} \sum_{\alpha=1}^N (\mathbf{x}_\alpha - \bar{\mathbf{x}})(\mathbf{x}_\alpha - \bar{\mathbf{x}})^\top = \begin{bmatrix} \mathbf{S}_{11} & \mathbf{S}_{12} \\ \mathbf{S}_{21} & \mathbf{S}_{22} \end{bmatrix}.$$

Let $\mathbf{a}^{(j)} = \sqrt{(N-1)/N} \hat{\boldsymbol{\alpha}}^{(j)}$, $\mathbf{c}^{(j)} = \sqrt{(N-1)/N} \hat{\boldsymbol{\gamma}}^{(j)}$ and l_j satisfies

$$\mathbf{S}_{12} \mathbf{c}^{(j)} = l_j \mathbf{S}_{11} \mathbf{a}^{(j)},$$

$$\mathbf{S}_{21} \mathbf{c}^{(j)} = l_j \mathbf{S}_{22} \mathbf{c}^{(j)},$$

$$\mathbf{a}^{(j)\top} \mathbf{S}_{11} \mathbf{a}^{(j)} = 1,$$

$$\mathbf{c}^{(j)\top} \mathbf{S}_{22} \mathbf{c}^{(j)} = 1.$$

We call linear combinations $\mathbf{a}^{(j)\top} \mathbf{x}_\alpha^{(1)}$ and $\mathbf{c}^{(j)\top} \mathbf{x}_\alpha^{(2)}$ the sample canonical variates.