Mnohorozměrná analýza (NMST539)

## Z. Hlávka

Univerzita Karlova v Praze
Matematicko-fyzikální fakulta
Katedra pravděpodobnosti a matematické statistiky www.karlin.mff.cuni.cz/~hlavka

## Použitá literatura:

Mardia, K. V., Kent, J. T., \& Bibby, J. M. (1979). Multivariate analysis. Academic press.

Härdle, W. K., \& Simar, L. (2014). Applied multivariate statistical analysis, 4th edition. Springer Science \& Business Media.

Anděl, J. (1985). Matematická statistika. SNTL.

## Mnohorozměrná analýza (NMST539)

- Mnohorozměrná data.
- Opakování: lineární algebra (matice).
- Mnohorozměrné normální rozdělení, Wishartovo a Hotellingovo rozdělení.
- Metoda hlavních komponent, faktorová analýza.
- Mnohorozměrné škálování, shluková a diskriminační analýza.
- Kanonické korelace, korespondenční analýza.
- Další metody (hloubka dat, SIR, projection pursuit).


## Týden 1

Předpokládané znalosti: základní maticové operace (sčítání, násobení apod.)

Mnohorozměrná data:

- grafické znázornění,
- matice dat a popisné statistiky.


## Example

Example: The authorities have measured
$X_{1}=$ length of the bill
$X_{2}=$ height of the bill (left)
$X_{3}=$ height of the bill (right)
$X_{4}=$ distance of the inner frame to the lower border
$X_{5}=$ distance of the inner frame to the upper border
$X_{6}=$ length of the diagonal of the central picture.

Figure: An old Swiss 1000-franc bill.

## Example: Swiss bank data

The dataset consists of 200 measurements on Swiss bank notes. One half of these bank notes are genuine, the other half are forged bank notes.

It is important to be able to decide whether a given banknote is genuine.
We want to derive a good rule that separates the genuine and forged banknotes.

Which measurement is most informative? We have to visualize the difference.
Graphics Week 1 Graphics

Computers allow easy construction of informative plots:
1D Boxplot, histogram, kernel density estimator (KDE), dotplot, ....
2D Histogram, KDE, scatterplot.
3D 3D scatterplot.
4+ Scatterplot matrix, parallel coordinates, Chernoff-Flury faces, Andrew's curves.

One typically needs static graphics (PDF) for reports and interactive graphics for data exploration.


Figure: Variables $X_{6}$ (diagonal) of bank notes, the genuine at the left. $\rightarrow$ MVAboxbank6

## Mil Boxplots

* Median and mean bar indicate the central locations.
* The relative location of median (and mean) in the box is a measure of skewness.
* The length of box and whiskers is a measure of spread.
* The length of whiskers indicate the tail length of the distribution.
* Outlying points are marked as " $\star$ " or "•" outside the outside bars.
* Boxplots do not indicate multi modality or clusters.
* If we compare the relative size and location of the boxes we are comparing distributions.

Figure: Variables $X_{1}$ (length) of bank notes, the genuine at the left. $\rightarrow$ MVAboxbank1

## Histograms



Week 1
Graphics

The histogram counts relative frequencies of observations $x_{i}$ falling into predefined bins:

$$
\widehat{f}_{h}(x)=n^{-1} h^{-1} \sum_{j \in \mathbb{Z}} \sum_{i=1}^{n} \mathrm{I}\left\{x_{i} \in B_{j}\left(x_{0}, h\right)\right\} \mathrm{I}\left\{x \in B_{j}\left(x_{0}, h\right)\right\}
$$

- the histogram is a simple estimator of a probability density,
- $h$ is a smoothing parameter and controls the width of the histogram bins.


Figure: Diagonal of forged bank notes. Histograms with $x_{0}=137.8$ and $h=0.1$ (upper left), $h=0.2$ (lower left), $h=0.3$ (upper right), $h=0.4$ (lower right).

For $x \in B_{j}$ (assuming that the density $f(x)$ is 'reasonable'), it is easy to calculcate the bias $E \widehat{f}_{h}(x)-f(x) \doteq f^{\prime}\left(m_{j}\right)\left(m_{j}-x\right)$ and variance $\operatorname{var} \widehat{f}_{h}(x) \doteq \frac{1}{n h} f(x)$

It follows that the Mean Squared Error is

$$
\operatorname{MSE}\left\{\widehat{f}_{h}(x)\right\}=\frac{1}{m n} f(x)+\left\{f^{\prime}\left(m_{j}\right)\right\}^{2}\left(m_{j}-x\right)^{2}+o(h)+o(1 / n h)
$$

By integrating MSE and taking limits, we easily obtain

$$
\operatorname{AMISE}\left\{\widehat{f}_{h}(x)\right\}=\frac{1}{n h}+\frac{h^{2}}{12}\left\|f^{\prime}\right\|_{2}^{2}
$$

leading the asymptotically optimal bandwidth $h_{0}=\left\{6 /\left(n\left\|f^{\prime}\right\|_{2}^{2}\right)\right\}^{1 / 3}$.


Figure: Diagonal of forged bank notes. Histogram with $h=0.4$ and origins $x_{0}=137.65$ (upper left), $x_{0}=137.75$ (lower left), $x_{0}=137.85$ (upper right), $x_{0}=137.95$ (lower right).

## Histograms

* Modes of the density correspond to strong peaks in the histogram.
* Histograms with the same $h$ need not be identical because they also depend on the origin $x_{0}$ of the grid.
* The consequence of a too large $h$ is a too flat, unstructured histogram (large bias).
* A too small binwidth $h$ results in a wiggly histogram (large variance).
* It is recommended to use averaged histograms (so-called kernel density estimators).


## Kernel density estimators (KDEs)

Kernel density estimator is a natural generalization of a histogram (by shifting the "bin", we obtain smooth estimator of the underlying probability density).

Histogram (at the center of a bin) can be written as

$$
\widehat{f}_{h}(x)=n^{-1}(2 h)^{-1} \sum_{i=1}^{n} I\left(\left|x-x_{i}\right| \leq h\right)
$$

$K(u)=I(|u| \leq 1) / 2$

$$
\widehat{f}_{h}(x)=n^{-1} h^{-1} \sum_{i=1}^{n} K\left(\frac{x-x_{i}}{h}\right)
$$

$K$ is the so-called kernel.


Figure: Densities of diagonals of genuine and forged bank notes. Automatic density estimates. $\rightarrow$ MVAdenbank

| $K(u)=\frac{1}{2} I(\|u\| \leq 1)$ |  |
| :--- | :--- |
| $K(u)=(1-\|u\|)(\|u\| \leq 1)$ | Uniform |
| $K(u)=\frac{3}{4}\left(1-u^{2}\right) l(\|u\| \leq 1)$ | Triangle |
| $K(u)=\frac{15}{16}\left(1-u^{2}\right)^{2} u^{2}(\|u\| \leq 1)$ | Quartrichnikov (Biweight) |
| $K(u)=\frac{1}{\sqrt{2 \pi}} \exp \left(-\frac{u^{2}}{2}\right)=\varphi(u)$ | Gaussian |



## Week 1 Graphics

The bias of KDE

$$
\operatorname{Bias} \widehat{f}_{h}(x)=\frac{h^{2}}{2} f^{\prime \prime}(x) \mu_{2}(K)+o\left(h^{2}\right)
$$

is of a smaller order than the bias of histogram.
Proceeding similarly, it is straghtforward that the Asymptotic Mean Integrated Squred Error is

$$
\operatorname{AMISE}\left(\widehat{f}_{h}\right)=\frac{1}{n h}\|K\|_{2}^{2}+\frac{h^{4}}{4}\left\{\mu_{2}(K)\right\}^{2}\left\|f^{\prime \prime}\right\|_{2}^{2}
$$

leading the asymptotically optimal bandwidth

$$
h_{0}=\left(\frac{\|K\|_{2}^{2}}{\left.n\left\|f^{\prime \prime}\right\|_{2}^{2} \mu_{2}(K)\right\}^{2}}\right)^{1 / 5}
$$

## Choice of the bandwidth

Assuming normality and using Gaussian kernel $K(u)=\frac{1}{\sqrt{2 \pi}} \exp \left(-\frac{u^{2}}{2}\right)$, the unknown constants can be calculated and we obtain the so-called

## Silverman's rule of thumb

$$
h_{G}=1.06 \hat{\sigma} n^{-\frac{1}{5}},
$$

where $\hat{\sigma}^{2}=\frac{1}{n} \sum_{x=1}^{n}\left(x_{i}-\bar{x}\right)^{2}$
Using Quartic kernel, the constants are somewhat different: $h_{Q}=2.62 h_{G}$.
In practice, one must be very careful because statistical software may assume another standardization of the kernel function (i.e., the bandwidth parameter may be multiplied by some constant).

Cross-validation is a popular bandwidth selection method (producing somewhat unstable results).



Figure: Contours of the density of $X_{4}, X_{5}, X_{6}$ of genuine and forged bank notes. $\rightarrow$ MVAcontbank3

## KDEs in R

## Libraries: KernSmooth, ks, sm.

## See $R$ task-views on CRAN:

1D density (),
bkde(KernSmooth),
locpoly(KernSmooth),
3D sm.density (sm)
6D kde(ks)
Unfortunately, multivariate KDEs have slow rates of convergence (so-called curse of dimensionality) - see Modern Statistical Methods (NMST434) for more details.

## Kernel densities

* Kernel densities estimate distribution densities by the kernel method.
* The bandwidth $h$ determines the degree of smoothness of the estimate $\widehat{f}$.
* A simple (but not necessarily correct) way to find a good bandwidth is to compute the rule of thumb bandwidth $h_{G}=1.06 \widehat{\sigma} n^{-1 / 5}$. This bandwidth is to be used only in combination with a Gaussian kernel $\varphi$.
* Kernel density estimates are a good descriptive tool for seeing modes, location, skewness, tails, asymmetry etc.


## Scatterplots

- Rotation of data
- Separation lines
- Draftman plot
- Brushing
- Parallel coordinate plots

| Swiss bank notes |  |
| :---: | :---: |
|  |  |

Figure: 3D Scatterplot for $\left(X_{4}, X_{5}, X_{6}\right)$ of the bank notes. Genuine notes are circles, forged are stars. $\rightarrow$ MVAscabank456


Figure: 2D scatterplot for $X_{5}$ vs. $X_{6}$ of the bank notes. Genuine notes are circles, forged are stars. $\rightarrow$ MVAscabank56


Figure: Draftman plot of the bank notes. The pictures in the left column show $\left(X_{3}, X_{4}\right),\left(X_{3}, X_{5}\right)$ and $\left(X_{3}, X_{6}\right)$, in the middle we have $\left(X_{4}, X_{5}\right)$ and $\left(X_{4}, X_{6}\right)$, and in the lower right is $\left(X_{5}, X_{6}\right) . \rightarrow$ MVAdrafbank1


Figure: Stereo plot of the bank notes $-\left(X_{4}, X_{5}, X_{6}\right)$.
Z. Hlávka (KPMS) NMST539 4 ロ (4.

## Week 1

 Graphics

Figure: Stereo plot of the bank notes - all variables.

## Parallel coordinate plots

- based on a orthogonal coordinate system
- allows to see more than four dimensions


## Idea:

Instead of plotting observations in an orthogonal coordinate system one draws their coordinates in a system of parallel axes. This way of representation is however sensitive to the order of the variables.


Figure: Parallel coordinate plot of observations $96-105 \rightarrow$ MVAparcoo1
Figure: The full bank Data set. Genuine banknotes displayed as solid lines. The forged bank notes are shown as dashed lines. $\rightarrow$ MVAparcoo2

## Faces



Figure: Flury faces for observations 91 to 110 of the bank notes. $\rightarrow$ MVAfacebank

Six variables to the following face elements
$X_{1}=1,19$ (eye sizes)
$X_{2}=2,20$ (pupil sizes)
$X_{3}=4,22$ (eye slants)
$X_{4}=11,29$ (upper hair lines)
$X_{5}=12,29$ (lower hair lines)
$X_{6}=13,14,31,32$ (face lines and darkness of hair)

## library(aplpack)

faces(bank2)
faces(bank2[91:110])

## Summary Statistics

$\mathcal{X}(n \times p)$ data matrix

$$
\mathcal{X}=\left(\begin{array}{ccc}
x_{11} & \cdots & x_{1 p} \\
\vdots & & \vdots \\
\vdots & & \vdots \\
x_{n 1} & \cdots & x_{n p}
\end{array}\right)
$$

mean

$$
\bar{x}=\left(\begin{array}{c}
\bar{x}_{1} \\
\vdots \\
\bar{x}_{p}
\end{array}\right)=n^{-1} \mathcal{X}^{\top} 1_{n}
$$

## 

## Faces

* faces can be used to detect subgroups in multivariate data
* subgroups are characterized by similar looking faces
* outliers are identified by extreme faces (e.g. dark hair)
* if one element of $X$ is unusual the corresponding face element changes a lot in shape


## Covariance matrix

$$
\begin{gathered}
\mathcal{S}=n^{-1} \mathcal{X}^{\top} \mathcal{X}-\bar{x} \bar{x}^{\top} \\
=n^{-1}\left(\mathcal{X}^{\top} \mathcal{X}-n^{-1} \mathcal{X}^{\top} 1_{n} 1_{n}^{\top} \mathcal{X}\right)=n^{-1} \mathcal{X}^{\top} \mathcal{H X}
\end{gathered}
$$

## Centering matrix

$$
\mathcal{H}=\mathcal{I}_{n}-n^{-1} 1_{n} 1_{n}^{\top}
$$

centered data: $S=n^{-1} \mathcal{X}^{\top} \mathcal{X}$
$\mathcal{D}=\operatorname{diag}\left(s_{X_{j} X_{j}}\right)$, where $X_{j}, j=1, \ldots, p$ are the columns of $\mathcal{X}$
Correlation matrix $\mathcal{R}=\mathcal{D}^{-1 / 2} \mathcal{S D}^{-1 / 2}$

## Linear transformations

$\mathcal{A}(q \times p)$ matrix

$$
\begin{gathered}
\mathcal{Y}=\mathcal{X} \mathcal{A}^{\top}=\left(y_{1}, \ldots, y_{n}\right)^{\top} \\
\bar{y}=\mathcal{A} \bar{x} \\
\mathcal{S}_{\mathcal{Y}}=\mathcal{A} S_{\mathcal{X}} \mathcal{A}^{\top}
\end{gathered}
$$

Example: $\bar{x}=(1,2)^{\top}$
$y=4 x, x \in \mathbb{R}^{2}$
$\bar{y}=4 \bar{x}=(4,8)^{\top}$

## Mahalanobis Transformation

$$
\begin{gathered}
z_{i}=\mathcal{S}^{-1 / 2}\left(x_{i}-\bar{x}\right), \quad i=1, \ldots, n, \\
\mathcal{S}_{\mathcal{Z}}=n^{-1} \mathcal{Z}^{\top} \mathcal{H Z}=\mathcal{I}_{p}, \\
\overline{\mathcal{Z}}=0
\end{gathered}
$$

where $\mathcal{H}$ is the centering matrix.
Mahalanobis transformation leads to standardized uncorrelated zero mean data matrix $\mathcal{Z}$.

## Summary Statistics

* The center of gravity of a data matrix is given by its mean vector $\bar{x}=n^{-1} \mathcal{X}^{\top} 1_{n}$.
* The dispersion of the observations in a data matrix is given by the empirical covariance matrix $\mathcal{S}=n^{-1} \mathcal{X}^{\top} \mathcal{H} \mathcal{X}$.
* The empirical correlation matrix is given by $\mathcal{R}=\mathcal{D}^{-1 / 2} \mathcal{S D}^{-1 / 2}$
* A linear transformation $\mathcal{Y}=\mathcal{X} \mathcal{A}^{\top}$ of a data matrix $\mathcal{X}$ has mean $\mathcal{A} \bar{x}$ and empirical covariance $\mathcal{A} \mathcal{S}_{\mathcal{X}} \mathcal{A}^{\top}$.
* The Mahalanobis transformation is a linear transformation $z_{i}=\mathcal{S}^{-1 / 2}\left(x_{i}-\bar{x}\right)$ which gives a standardized, uncorrelated data matrix $\mathcal{Z}$.


## Week 2

## Týden 2

Opakování základní maticové algebry:

- spektrální rozklad matice a kvadratické formy.

Náhodné vektory:

- mnohorozměrná distribuční funkce a hustota,
- podmíněná a marginální rozdělení,
- momenty,
- mnohorozměrné normální rozdělení.

A short Excursion into Matrix Algebra

$$
\mathcal{A}_{(n \times p)}=\left(\begin{array}{ccc}
a_{11} & \cdots & a_{1 p} \\
\vdots & \ddots & \vdots \\
a_{n 1} & \cdots & a_{n p}
\end{array}\right)
$$

| Definition | Notation |
| :--- | :--- |
| Transpose | $\mathcal{A}$ |
| Sum | $\mathcal{A}+\mathcal{B}$ |
| Difference | $\mathcal{A}-\mathcal{B}$ |
| Scalar product | $c \cdot \mathcal{A}$ |
| Product | $\mathcal{A} \cdot \mathcal{B}$ |
| Rank | $\operatorname{rank}(\mathcal{A})$ |
| Trace | $\operatorname{tr}(\mathcal{A})$ |
| Determinant | $\operatorname{det}(\mathcal{A})=\|\mathcal{A}\|$ |
| Inverse | $\mathcal{A}^{-1}$ |
| Generalised Inverse | $\mathcal{A}^{-}: \mathcal{A} \mathcal{A}^{-} \mathcal{A}=\mathcal{A}$ |


| Name | Definition | Notation | Example |
| :---: | :---: | :---: | :---: |
| diagonal matrix | $a_{i j}=0, i \neq j, n=p$ | $\operatorname{diag}\left(a_{i i}\right)$ | $\left(\begin{array}{ll}1 & 0 \\ 0 & 2\end{array}\right)$ |
| identity matrix | $\operatorname{diag}(\underbrace{1, \ldots, 1}_{p})$ | $\mathcal{I}_{p}$ | $\left(\begin{array}{ll}1 & 0 \\ 0 & 1\end{array}\right)$ |
| unit matrix | $a_{i j} \equiv 1, n=p$ | $1_{n} 1_{n}^{\top}$ | $\left(\begin{array}{ll}1 & 1 \\ 1 & 1\end{array}\right)$ |
| symmetric matrix | $a_{i j}=a_{j i}$ |  | $\left(\begin{array}{ll}1 & 2 \\ 2 & 3\end{array}\right)$ |


| Name | Definition | Notation | Example |
| :---: | :---: | :---: | :---: |
| scalar | $p=n=1$ | $a$ | 3 |
| column vector | $p=1$ | $a$ | $\binom{1}{3}$ |
| row vector | $n=1$ | $a^{\top}$ | $\binom{1}{3}$ |
| vector of ones | $(\underbrace{1, \ldots, 1}_{n})^{\top}$ | $1_{n}$ | $\binom{1}{1}$ |
| vector of zeros | $(\underbrace{0, \ldots, 0}_{n})^{\top}$ | $0_{n}$ | $\binom{0}{0}$ |
| square matrix | $n=p$ | $\mathcal{A}(p \times p)$ | $\left(\begin{array}{ll}2 & 0 \\ 0 & 2\end{array}\right)$ |


| Name | Definition | Example |
| :---: | :---: | :---: |
| null matrix | $a_{i j}=0$ | $\left(\begin{array}{cc}0 & 0 \\ 0 & 0\end{array}\right)$ |
| upper triangular matrix | $a_{i j}=0, i<j$ | $\left(\begin{array}{ccc}1 & 2 & 4 \\ 0 & 1 & 3 \\ 0 & 0 & 1\end{array}\right)$ |
| idempotent matrix | $\mathcal{A}^{2}=\mathcal{A}$ | $\left(\begin{array}{cc}\frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2}\end{array}\right)$ |
| orthogonal matrix | $\mathcal{A}^{\top} \mathcal{A}=I=\mathcal{A A}^{\top}$ | $\left(\begin{array}{cc}\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}}\end{array}\right)$ |

## Properties of a Square Matrix

For any $\mathcal{A}(n \times n)$ and $\mathcal{B}(n \times n)$ and any scalar $c$

$$
\begin{aligned}
\operatorname{tr}(\mathcal{A}+\mathcal{B}) & =\operatorname{tr}(\mathcal{A})+\operatorname{tr}(\mathcal{B}) \\
\operatorname{tr}(c \mathcal{A}) & =c \operatorname{tr}(\mathcal{A}) \\
|c \mathcal{A}| & =c^{n}|\mathcal{A}| \\
\operatorname{tr}(\mathcal{A B}) & =\operatorname{tr}(\mathcal{B A}) \\
|\mathcal{A B}| & =|\mathcal{B A}| \\
|\mathcal{A B}| & =|\mathcal{A}||\mathcal{B}| \\
\left|\mathcal{A}^{-1}\right| & =|\mathcal{A}|^{-1}
\end{aligned}
$$

## $1 \mathrm{MES} \cdot \mathrm{d}$

## Matrix Algebra

* The determinant $|\mathcal{A}|$ is a product of the eigenvalues of $\mathcal{A}$.
* The inverse of a matrix $\mathcal{A}$ exists if $|\mathcal{A}| \neq 0$.
* The $\operatorname{trace} \operatorname{tr}(\mathcal{A})$ is the sum of the eigenvalues of $\mathcal{A}$.
* The sum of the traces of two matrices equals the trace of the sum of the two matrices.
* The trace $\operatorname{tr}(\mathcal{A B})$ equals $\operatorname{tr}(\mathcal{B A})$.


## Eigenvalues and Eigenvectors

## Square matrix $\mathcal{A}(n \times n)$

eigenvalue $\lambda=\operatorname{Eval}(\mathcal{A})$
eigenvector $\gamma=\operatorname{Evec}(\mathcal{A})$

$$
\mathcal{A} \gamma=\lambda \gamma
$$

Eigenvalues describe the 'size' of the matrix $\mathcal{A}$ :

$$
\begin{aligned}
|\mathcal{A}| & =\prod_{j=1}^{n} \lambda_{j} \\
\operatorname{tr}(\mathcal{A}) & =\sum_{j=1}^{n} \lambda_{j}
\end{aligned}
$$

## Spectral Decomposition

Every real symmetric matrix $\mathcal{A}(p \times p)$ can be written as:

$$
\begin{aligned}
\mathcal{A} & =\Gamma \wedge \Gamma^{\top} \\
& =\sum_{j=1}^{p} \lambda_{j} \gamma_{j} \gamma_{j}^{\top} \\
\Lambda & =\operatorname{diag}\left(\lambda_{1}, \cdots, \lambda_{p}\right) \\
\Gamma & =\left(\gamma_{1}, \cdots, \gamma_{p}\right),
\end{aligned}
$$

where the matrix $\Gamma$ is orthogonal (i.e. $\Gamma^{\top} \Gamma=\mathcal{I}_{p}$ ).
Spectral decomposition allows easier calculation of powers of the matrix $\mathcal{A}$ (very useful is the inverse $\mathcal{A}^{-1}$ and 'inverse square root' $\mathcal{A}^{-1 / 2}$ ).

## Quadratic forms

$\mathcal{A}(p \times p)$ symmetric matrix

$$
Q(x)=x^{T} \mathcal{A} x=\sum_{i=1}^{p} \sum_{j=1}^{p} a_{i j} x_{i} x_{j}
$$

## Definiteness

$$
\begin{array}{ll}
Q(x)>0 \text { for all } x \neq 0 & \text { positive definite }(\mathrm{pd}) \\
Q(x) \geq 0 \text { for all } x \neq 0 & \text { positive semidefinite }(\mathrm{psd}) .
\end{array}
$$

$\mathcal{A}$ is $\mathrm{pd}(\mathrm{psd})$ iff $Q(x)=x^{T} \mathcal{A} x$ is $\mathrm{pd}(\mathrm{psd})$

## Week 2 Matrix algebra

Example:
$Q(x)=x^{\top} \mathcal{A} x=x_{1}^{2}+x_{2}^{2}, \mathcal{A}=\left(\begin{array}{ll}1 & 0 \\ 0 & 1\end{array}\right)$
eigenvalues: $\lambda_{1}=\lambda_{2}=1$ positive definite
$Q(x)=\left(x_{1}-x_{2}\right)^{2}, \mathcal{A}=\left(\begin{array}{rr}1 & 1 \\ -1 & 1\end{array}\right)$
eigenvalues $\lambda_{1}=2, \lambda_{2}=0$ positive semidefinite
$Q(x)=x_{1}^{2}-x_{2}^{2}$
eigenvalues $\lambda_{1}=1, \lambda_{2}=-1$ indefinite.
$\mathcal{A}>0$ if and only if all $\lambda_{i}>0, i=1, \ldots, p$

## 표여 Quadratic forms

* A quadratic form can be described by a symmetric quadratic matrix $\mathcal{A}$.
* Quadratic forms can always be diagonalized.
* Positive definiteness of a quadratic form is equivalent to positiveness of the eigenvalues of the matrix $\mathcal{A}$.
* Maximum and minimum of a quadratic form under constraints can be expressed in terms of eigenvalues.

Geometrical aspects

Distance function $d: \mathbb{R}^{2 p} \rightarrow \mathbb{R}_{+} d^{2}(x, y)=(x-y)^{T} \mathcal{A}(x-y), \quad \mathcal{A}>0$ $\mathcal{A}=\mathcal{I}_{p}$, Euclidean distance
$E_{d}=\left\{x \in \mathbb{R}^{p} \mid\left(x-x_{0}\right)^{\top}\left(x-x_{0}\right)=d^{2}\right\}$
Example: $x \in \mathbb{R}^{2}, x_{0}=0, x_{1}^{2}+x_{2}^{2}=1$

## Norm of a vector

$$
\|x\|=d(0, x)=\sqrt{x^{T} x}
$$

## Week 2 Matrix algebra



Figure: Iso-distance ellipsoid.

$$
E_{d}=\left\{x:\left(x-x_{0}\right)^{\top} \mathcal{A}\left(x-x_{0}\right)=d^{2}\right\}, \gamma_{j}=\operatorname{Evec}(\mathcal{A}), \mathcal{A}>0
$$

Angle between two Vectors

Angle of vectors $x$ and $y$ can be calculated as

$$
\cos \theta=\frac{x^{\top} y}{\|x\|\|y\|}
$$

Norm of a vector

$$
\|x\|=d(0, x)=\sqrt{x^{\top} x}
$$

Unit vectors

$$
\{x:\|x\|=1\}
$$



Figure: Angle between vectors.

Angle between two Vectors

Example: Angle $=$ Correlation
Observations $\left\{x_{i}\right\}_{i=1}^{n},\left\{y_{i}\right\}_{i=1}^{n}$
$\bar{x}=\bar{y}=0$

$$
\rho_{X Y}=\frac{\sum x_{i} y_{i}}{\sqrt{\sum x_{i}^{2} \sum y_{i}^{2}}}=\cos \theta
$$

Correlation corresponds to angle between $x, y \in \mathbb{R}^{p}$.


Figure: Projection.

$$
\cos \theta=\frac{x^{\top} y}{\|x\|\|y\|}=\frac{x_{1} y_{1}+x_{2} y_{2}}{\|x\|\|y\|}=\cos \theta_{1} \cos \theta_{2}+\sin \theta_{1} \sin \theta_{2}
$$

## Projection on $C(\mathcal{X})$

$$
\mathcal{X}(n \times p), \quad \mathcal{P}=\mathcal{X}\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1} \mathcal{X}^{\top}
$$

## Column space

$\mathcal{X}(n \times p)$ data matrix

$$
C(\mathcal{X})=\left\{x \in \mathbb{R}^{n} \mid \exists a \in \mathbb{R}^{p} \text { so that } \mathcal{X} a=x\right\}
$$

## projection matrix

$\mathcal{P}(n \times n), \quad \mathcal{P}=\mathcal{P}^{\top}=\mathcal{P}^{2}$ ( $\mathcal{P}$ is idempotent)
let $b \in \mathbb{R}^{n}, a=\mathcal{P} b$ is the projection of $b$ on $C(\mathcal{P})$

$$
\mathcal{P} \mathcal{X}=\mathcal{X}, \mathcal{P} \text { is a projector, } \mathcal{P} \mathcal{P}=\mathcal{P} .
$$

$$
\mathcal{Q}=\mathcal{I}_{n}-\mathcal{P}, \mathcal{Q}^{2}=\mathcal{Q}
$$

$$
p_{x}=\frac{y^{\top} x}{\|y\|^{2}} y
$$

$$
\mathcal{P} \mathcal{X}=\mathcal{X}
$$

$$
\mathcal{Q X}=0
$$

## mand Geometrical aspects

* For the Euclidean distance with $\mathcal{A}=\mathcal{I}$ the correlation between two centered data vectors $x$ and $y$ is given by the cosine of the angle between them, i.e. $\cos \theta=\rho_{X Y}$.
* The projection $\mathcal{P}=\mathcal{X}\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1} \mathcal{X}^{\top}$ is the projection in the column space $C(\mathcal{X})$ of $\mathcal{X}$.
* The projection of $x \in \mathbb{R}^{n}$ on $y \in \mathbb{R}^{n}$ is given by $p_{x}=\frac{y^{\top} x}{\|y\|^{2}} y$.


## Random vector

Let us assume that random variables $X_{1}, \ldots, X_{p}$ are defined on the probability space $(\Omega, \mathcal{A}, P)$. In this setup, the vector $\left(X_{1}, \ldots, X_{p}\right)^{\top}$ is called random vector.

Theorem: The $p$-dimensional random vector $X=\left(X_{1}, \ldots, X_{p}\right)^{\top}$ is a measurable function from $(\Omega, \mathcal{A}, P)$ to $\left(\mathbb{R}^{p}, \mathcal{B}_{p}\right)$

Proof: See Theorem II.1.1 in Anděl (1985).

The function

$$
F\left(x_{1}, \ldots, x_{p}\right)=P\left(X_{1}<x, \ldots X_{p}<x_{p}\right)
$$

is the multivariate (joint) cummulative distribution function of the random vector $X=\left(X_{1}, \ldots, X_{p}\right)^{\top}$.

## Example:

$$
f\left(x_{1}, x_{2}\right)= \begin{cases}\frac{1}{2} x_{1}+\frac{3}{2} x_{2} & 0 \leq x_{1}, x_{2} \leq 1 \\ 0 & \text { otherwise }\end{cases}
$$

$f\left(x_{1}, x_{2}\right)$ is a density since

$$
\int f\left(x_{1}, x_{2}\right) d x_{1} x_{2}=\frac{1}{2}\left[\frac{x_{1}^{2}}{2}\right]_{0}^{1}+\frac{3}{2}\left[\frac{x_{2}^{2}}{2}\right]_{0}^{1}=\frac{1}{4}+\frac{3}{4}=1
$$

Example: The marginal densities

$$
\begin{aligned}
f_{X_{1}}\left(x_{1}\right) & =\int f\left(x_{1}, x_{2}\right) d x_{2}=\int_{0}^{1}\left(\frac{1}{2} x_{1}+\frac{3}{2} x_{2}\right) d x_{2}=\frac{1}{2} x_{1}+\frac{3}{4} ; \\
f_{X_{2}}\left(x_{2}\right) & =\int f\left(x_{1}, x_{2}\right) d x_{1}=\int_{0}^{1}\left(\frac{1}{2} x_{1}+\frac{3}{2} x_{2}\right) d x_{1}=\frac{3}{2} x_{2}+\frac{1}{4} .
\end{aligned}
$$

The conditional densities

$$
f\left(x_{2} \mid x_{1}\right)=\frac{\frac{1}{2} x_{1}+\frac{3}{2} x_{2}}{\frac{1}{2} x_{1}+\frac{3}{4}} \quad \text { and } \quad f\left(x_{1} \mid x_{2}\right)=\frac{\frac{1}{2} x_{1}+\frac{3}{2} x_{2}}{\frac{3}{2} x_{2}+\frac{1}{4}}
$$

Example:

$$
\begin{gathered}
f\left(x_{1}, x_{2}\right)=1, \quad 0<x_{1}, x_{2}<1, \\
f\left(x_{1}, x_{2}\right)=1+\alpha\left(2 x_{1}-1\right)\left(2 x_{2}-1\right), \quad 0<x_{1}, x_{2}<1, \quad-1 \leq \alpha \leq 1 . \\
f_{X_{1}}\left(x_{1}\right)=1, \quad f_{X_{2}}\left(x_{2}\right)=1 . \\
\int_{0}^{1} 1+\alpha\left(2 x_{1}-1\right)\left(2 x_{2}-1\right) d x_{2}=1+\alpha\left(2 x_{1}-1\right)\left[x_{2}^{2}-x_{2}\right]_{0}^{1}=1 .
\end{gathered}
$$

Definition of (statistical) independence

Absolutely continuous random vectors $X_{1}, X_{2}$ are independent iff $f(x)=f\left(x_{1}, x_{2}\right)=f_{X_{1}}\left(x_{1}\right) f_{X_{2}}\left(x_{2}\right)$.

$1!$
Two random variables may have identical marginals but different joint distribution.

## Moments

$E X \in \mathbb{R}^{p}$ denotes the $p$-dimensional vector of expected values of the random vector $X$

$$
E X=\left(\begin{array}{c}
E X_{1} \\
\vdots \\
E X_{p}
\end{array}\right)=\int x f(x) d x=\left(\begin{array}{c}
\int x_{1} f(x) d x \\
\vdots \\
\int x_{p} f(x) d x
\end{array}\right)=\mu
$$

The properties of expected value follow from the properties of the integral:

$$
E(\alpha X+\beta Y)=\alpha E X+\beta E Y
$$

If $X$ and $Y$ are independent then

$$
\begin{gathered}
E\left(X Y^{\top}\right)=\int x y^{\top} f(x, y) d x d y \\
=\int x f(x) d x \int y^{\top} f(y) d y=E X E Y^{\top}
\end{gathered}
$$

## Definition of variance matrix $(\Sigma)$

$$
\Sigma=\operatorname{Var}(X)=E(X-\mu)(X-\mu)^{\top}
$$

We say that random vector $X$ has a distribution with the vector of expected values $\mu$ and the covariance matrix $\Sigma$,

$$
X \sim(\mu, \Sigma)
$$

## Week 2 Random vector

Properties of Variances and Covariances

$$
\begin{aligned}
& \operatorname{var}\left(a^{\top} X\right)=a^{\top} \operatorname{Var}(X) a=\sum_{i, j} a_{i} a_{j} \sigma_{X_{i} X_{j}} \\
& \operatorname{Var}(\mathcal{A} X+b)=\mathcal{A} \operatorname{Var}(X) \mathcal{A}^{\top} \\
& \operatorname{Cov}(X+Y, Z)=\operatorname{Cov}(X, Z)+\operatorname{Cov}(Y, Z) \\
& \operatorname{Var}(X+Y)=\operatorname{Var}(X)+\operatorname{Cov}(X, Y)+\operatorname{Cov}(Y, X)+\operatorname{Var}(Y) \\
& \operatorname{Cov}(\mathcal{A} X, \mathcal{B} Y)=\mathcal{A} \operatorname{Cov}(X, Y) \mathcal{B}^{\top} .
\end{aligned}
$$

## Properties of the Covariance Matrix

Elements of $\Sigma$ are variances and covariances of the components of the random vector $X$ :

$$
\begin{aligned}
& \Sigma=\left(\sigma_{X_{i} X_{j}}\right) \\
& \sigma_{X_{i} X_{j}}=\operatorname{cov}\left(X_{i}, X_{j}\right) \\
& \sigma_{X_{i} X_{i}}=\operatorname{var}\left(X_{i}\right)
\end{aligned}
$$

Computational formula:

$$
\Sigma=E\left(X X^{\top}\right)-\mu \mu^{\top}
$$

Variance matrix is positive semidefinite:

## $\Sigma \geq 0$

(variance $a^{\top} \Sigma a$ of any linear combination $a^{\top} X$ cannot be negative).

## Week 2 Random vectors

## Conditional Expectations

(Absolutely continuous) random vector $X=\left(X_{1}, X_{2}\right)$
Conditional expectation of $X_{2}$, given $X_{1}=x_{1}$ :

$$
E\left(X_{2} \mid x_{1}\right)=\int x_{2} f\left(x_{2} \mid x_{1}\right) d x_{2}
$$

and conditional expectation of $X_{1}$, given $X_{2}=x_{2}$ :

$$
E\left(X_{1} \mid x_{2}\right)=\int x_{1} f\left(x_{1} \mid x_{2}\right) d x_{1}
$$

The conditional expectation $E\left(X_{2} \mid x_{1}\right)$ is a function of $x_{1}$ (it is the expected value of $X_{2}$ if we know that corresponding $X_{1}=x_{1}$-typical example of this setup is simple linear regression, where $E(Y \mid X=x)=x \beta$.

## Moments

* The expectation of a random vector $X$ is $\mu=\int x f(x) d x$, the covariance matrix $\Sigma=\operatorname{Var}(X)=E(X-\mu)(X-\mu)^{\top}$. We denote $X \sim(\mu, \Sigma)$.
* Expectations are linear, i.e., $E(\alpha X+\beta Y)=\alpha E X+\beta E Y$. If $X, Y$ are independent then $E\left(X Y^{\top}\right)=E X E Y^{\top}$.
* The covariance between two random vectors $X, Y$ is
$\Sigma_{X Y}=\operatorname{Cov}(X, Y)=E(X-E X)(Y-E Y)^{\top}=E\left(X Y^{\top}\right)-E X E Y^{\top}$. If $X, Y$ are independent then $\operatorname{Cov}(X, Y)=0$.
* The Conditional Expectation $E\left(X_{2} \mid X_{1}\right)$ is the MSE best approximation of $X_{2}$ by a function of $X_{1}$.


## Week 2 Random vectors

## Geometry of the $N_{p}(\mu, \Sigma)$ Distribution

Density of $N_{p}(\mu, \Sigma)$ is constant on ellipsoids of the form

$$
(x-\mu)^{\top} \Sigma^{-1}(x-\mu)=d^{2}
$$

If $X \sim N_{p}(\mu, \Sigma)$, then the variable $Y=(X-\mu)^{\top} \Sigma^{-1}(X-\mu)$ is $\chi_{p}^{2}$ distributed, since the Mahalonobis transformation
$Z=\Sigma^{-1 / 2}(X-\mu) \sim N_{p}\left(0, \mathcal{I}_{p}\right)$ and $Y=Z^{T} Z=\sum_{j=1}^{p} Z_{j}^{2}$.

Multivariate Normal (Multinormal) Distribution

The pdf of a multinormal is (assuming that $\Sigma$ has full rank):

$$
f(x)=|2 \pi \Sigma|^{-1 / 2} \exp \left\{-\frac{1}{2}(x-\mu)^{\top} \Sigma^{-1}(x-\mu)\right\} .
$$

$X \sim N_{p}(\mu, \Sigma)$
Expected value is $E X=\mu$,
Variance matrix of $X$ is $\operatorname{Var}\{X\}=\Sigma>0$
(what is the meaning of the quadratic form $(x-\mu)^{T} \Sigma^{-1}(x-\mu)$ in the formula for density?)

contour ellipses


Figure: Scatterplot of normal sample and contour ellipses for $\mu=\binom{3}{2}$ and $\Sigma=\left(\begin{array}{cc}1 & -1.5 \\ -1.5 & 4\end{array}\right) \rightarrow$ SMScontnorm

## Singular Normal Distribution

Definition of "Normal" distribution in case that the matrix $\Sigma$ is singular-we use its eigenvalues $\lambda_{i}$ and the generalized inverse $\Sigma^{-}$:
$\operatorname{rank}(\Sigma)=k<p, \quad \lambda_{1} \cdots \lambda_{k}$

$$
\frac{(2 \pi)^{-k / 2}}{\left(\lambda_{1} \cdots \lambda_{k}\right)^{1 / 2}} \exp \left\{-\frac{1}{2}(x-\mu)^{\top} \Sigma^{-}(x-\mu)\right\}
$$

$\Sigma^{-}=$G-inverse
The contour curves of a multinormal are ellipsoids with half-lengths proportional to $\sqrt{\lambda_{i}}$, where $\lambda_{i}$ denote the eigenvalues of $\Sigma$.

* The Mahalanobis transformation transforms $X \sim N_{p}(\mu, \Sigma)$ to $Y=\Sigma^{-1 / 2}(X-\mu) \sim N_{p}\left(0, \mathcal{I}_{p}\right)$. Vice versa, one can create a $X \sim N_{p}(\mu, \Sigma)$ from $Y \sim N_{p}\left(0, \mathcal{I}_{p}\right)$ via $X=\Sigma^{1 / 2} Y+\mu$.


## Týden 3

## Mnohorozměrné normální rozdělení:

- hustota transformovaného náhodného vektoru.
- centrální limitní věta a transformace,
- vlastnosti mnohorozměrného normálního rozdělení.


## Transformations

Density of a linear transformation

$$
\begin{gathered}
Y=\mathcal{A} X+b, \quad \mathcal{A} \text { nonsingular } \\
X=\mathcal{A}^{-1}(Y-b) \\
\mathcal{J}=\mathcal{A}^{-1} \\
f_{Y}(y)=a b s\left(|\mathcal{A}|^{-1}\right) f_{X}\left\{\mathcal{A}^{-1}(y-b)\right\}
\end{gathered}
$$

Starting from $X \sim N_{p}\left(0_{p}, \mathcal{I}_{p}\right)$, it is now easy to calculate the $p$-dimensional density of $Y=\Sigma^{1 / 2} X+\mu \sim N_{p}(\mu, \Sigma)$ (assuming that $\Sigma$ has full rank).

Note that the multivariate standard normal density (or characteristic function) can be defined as a product of univariate standard normal densities $\exp \left(-x^{2} / 2\right) / \sqrt{2 \pi}$ (or characteristic functions $\exp \left(-t^{2} / 2\right)$ ).

## Alternative definition (Cramér-Wold characterization)

We have defined $N_{p}(\mu, \Sigma)$ by writing down its density. Unfortunately, this approach has some disadvantages because we assumed that $\Sigma$ has full rank.

Definition: $X$ has multivariate Normal distribution if and only if $a^{\top} X$ is univariate normal for all $a \in \mathbb{R}^{p}$.

It easily follows that $Y=A x+c$ has multivariate normal distribution even when $A$ is not square and it does not have full rank.

The random variable $Z=t^{\top} X \sim N\left(t^{\top} \mu, t^{\top} \Sigma t\right)$ has the characteristic function

$$
\phi_{Z}(s)=E(\exp \text { is } Z)=\exp \left(i s t^{\top} \mu-s^{2} t^{\top} \Sigma t / 2\right)
$$

Therefore, the characteristic function of $X$ is

$$
\phi_{X}(t)=E\left(\exp i t^{\top} X\right)=E(\exp i Z)=\phi_{Z}(1)=\exp \left(i t^{\top} \mu-t^{\top} \Sigma t / 2\right)
$$

## Partitioned Matrices

$\mathcal{A}(n \times p)$

$$
\mathcal{A}=\left(\begin{array}{cc}
\mathcal{A}_{11} & \mathcal{A}_{12} \\
\mathcal{A}_{21} & \mathcal{A}_{22}
\end{array}\right)
$$

$\mathcal{A}_{i j}\left(n_{i} \times p_{j}\right)$

$$
\begin{aligned}
\mathcal{A}+\mathcal{B} & =\left(\begin{array}{ll}
\mathcal{A}_{11}+\mathcal{B}_{11} & \mathcal{A}_{12}+\mathcal{B}_{12} \\
\mathcal{A}_{21}+\mathcal{B}_{21} & \mathcal{A}_{22}+\mathcal{B}_{22}
\end{array}\right) \\
\mathcal{B}^{\top} & =\left(\begin{array}{ll}
\mathcal{B}_{11}^{\top} & \mathcal{B}_{21}^{\top} \\
\mathcal{B}_{12}^{1} & \mathcal{B}_{22}^{2}
\end{array}\right) \\
\mathcal{A B}^{\top} & =\left(\begin{array}{ll}
\mathcal{A}_{11} \mathcal{B}_{11}^{\top}+\mathcal{A}_{12} \mathcal{B}_{12}^{\top} & \mathcal{A}_{11} \mathcal{B}_{21}^{\top}+\mathcal{A}_{12} \mathcal{B}_{22}^{\top} \\
\mathcal{A}_{21} \mathcal{B}_{11}^{\top}+\mathcal{A}_{22} \mathcal{B}_{12}^{\top} & \mathcal{A}_{21} \mathcal{B}_{21}^{\top}+\mathcal{A}_{22} \mathcal{B}_{22}^{\top}
\end{array}\right)
\end{aligned}
$$

## Partioned Matrices

* For $\mathcal{A}$ nonsingular, $\mathcal{A}_{11}, \mathcal{A}_{22}$ square matrices,

$$
\mathcal{A}^{-1}=\left(\begin{array}{ll}
\mathcal{A}^{11} & \mathcal{A}^{12} \\
\mathcal{A}^{21} & \mathcal{A}^{22}
\end{array}\right)
$$

$$
\left\{\begin{array}{l}
\mathcal{A}^{11}=\left(\mathcal{A}_{11}-\mathcal{A}_{12} \mathcal{A}_{22}^{-1} \mathcal{A}_{21}\right)^{-1}=\left(\mathcal{A}_{11 \cdot 2}\right)^{-1} \\
\mathcal{A}^{12}=-\left(\mathcal{A}_{11 \cdot 2}\right)^{-1} \mathcal{A}_{12} \mathcal{A}_{22}^{-1} \\
\mathcal{A}^{21}=-\mathcal{A}_{22}^{-1} \mathcal{A}_{21}\left(\mathcal{A}_{11 \cdot 2}\right)^{-1} \\
\mathcal{A}^{22}=\mathcal{A}_{22}^{-1}+\mathcal{A}_{22}^{-1} \mathcal{A}_{21}\left(\mathcal{A}_{11 \cdot 2}\right)^{-1} \mathcal{A}_{12} \mathcal{A}_{22}^{-1}
\end{array}\right.
$$

* For $\mathcal{B}=\left(\begin{array}{ll}1 & b^{\top} \\ a & \mathcal{A}\end{array}\right)$ we have $|\mathcal{B}|=\left|\mathcal{A}-a b^{\top}\right|=|\mathcal{A}|\left|1-b^{\top} \mathcal{A}^{-1} a\right|$.
* $\left(\mathcal{A}-a b^{\top}\right)^{-1}=\mathcal{A}^{-1}+\frac{\mathcal{A}^{-1} a b^{\top} \mathcal{A}^{-1}}{1-b^{\top} \mathcal{A}^{-1} a}$


## Inverse of Partitioned Matrix

$\mathcal{A}$ nonsingular, $\mathcal{A}_{11}, \mathcal{A}_{22}$ square matrices

$$
\mathcal{A}^{-1}=\left(\begin{array}{ll}
\mathcal{A}^{11} & \mathcal{A}^{12} \\
\mathcal{A}^{21} & \mathcal{A}^{22}
\end{array}\right)
$$

where

$$
\left\{\begin{array}{l}
\mathcal{A}^{11}=\left(\mathcal{A}_{11}-\mathcal{A}_{12} \mathcal{A}_{22}^{-1} \mathcal{A}_{21}\right)^{-1}=\left(\mathcal{A}_{11 \cdot 2}\right)^{-1} \\
\mathcal{A}^{12}=-\left(\mathcal{A}_{11 \cdot 2}\right)^{-1} \mathcal{A}_{12} \mathcal{A}_{22}^{-1} \\
\mathcal{A}^{21}=-\mathcal{A}_{22}^{-1} \mathcal{A}_{21}\left(\mathcal{A}_{11 \cdot 2}\right)^{-1} \\
\mathcal{A}^{22}=\mathcal{A}_{22}^{-1}+\mathcal{A}_{22}^{-1} \mathcal{A}_{21}\left(\mathcal{A}_{11 \cdot 2}\right)^{-1} \mathcal{A}_{12} \mathcal{A}_{22}^{-1}
\end{array}\right.
$$

Determinant:

$$
|\mathcal{A}|=\left|\mathcal{A}_{11}\right|\left|\mathcal{A}_{22}-\mathcal{A}_{21} \mathcal{A}_{11}^{-1} \mathcal{A}_{12}\right|=\left|\mathcal{A}_{22}\right|\left|\mathcal{A}_{11}-\mathcal{A}_{12} \mathcal{A}_{22}^{-1} \mathcal{A}_{21}\right|
$$

Correlations and independence
Corollary: Let $X=\binom{X_{1}}{X_{2}} \sim N_{p}(\mu, \Sigma)$, then $X_{1}$ is independent of $X_{2}$ if and only if $X_{1}$ and $X_{2}$ are uncorrelated.

Proof: factorization of density $(\Sigma \geq 0)$ or characteristic function.

Interestingly, for two jointly multivariate Normal vectors (i.e., $X=\binom{X_{1}}{X_{2}} \sim N_{p}(\mu, \Sigma)$ ), pair-wise independence of their components implies complete independence.

The independence of two linear transforms of a multinormal $X$ can be shown via the following corollary.

Corollary: If $X \sim N_{p}(\mu, \Sigma), \mathcal{A}$ and $\mathcal{B}$ matrices, then $\mathcal{A} X$ and $\mathcal{B} X$ are independent if and only if $\mathcal{A} \Sigma \mathcal{B}^{\top}=0$.

Marginal and conditional distributions

Marginal distribution is just a special case of linear transform:

$$
X_{1}=\left(\begin{array}{ll}
\mathcal{I}_{q} & 0_{q} \times 0_{p}^{\top}
\end{array}\right)\binom{X_{1}}{X_{2}}
$$

For conditional distribution $X_{2} \mid X_{1}=x_{1}$ we have the following:
Theorem: The conditional distribution of $X_{2}$ given $X_{1}=x_{1}$ is normal with mean $\mu_{2}+\Sigma_{21} \Sigma_{11}^{-1}\left(x_{1}-\mu_{1}\right)$ and covariance $\Sigma_{22.1}=\Sigma_{22}-\Sigma_{21} \Sigma_{11}^{-1} \Sigma_{12}$, i.e.,

$$
\left(X_{2} \mid X_{1}=x_{1}\right) \sim N_{p-r}\left(\mu_{2}+\Sigma_{21} \Sigma_{11}^{-1}\left(x_{1}-\mu_{1}\right), \Sigma_{22.1}\right)
$$

Proof: e.g. via the following lemma or by factorizing the density (using the formula for inverse of partitioned matrix).

## Example:

$$
\begin{aligned}
& p=2, r=1, \mu=\binom{0}{0}, \Sigma=\left(\begin{array}{cc}
1 & -0.8 \\
-0.8 & 2
\end{array}\right) \\
& \Sigma_{11}=1, \Sigma_{21}=-0.8, \Sigma_{22.1}=2-(0.8)^{2}=1.36 \\
& \Rightarrow f_{X_{1}}\left(x_{1}\right) \quad=\frac{1}{\sqrt{2 \pi}} \exp \left(-\frac{x_{1}^{2}}{2}\right) \\
& \Rightarrow f\left(x_{2} \mid x_{1}\right)=\frac{1}{\sqrt{2 \pi(1.36)}} \exp \left\{-\frac{\left(x_{2}-0.8 x_{1}\right)^{2}}{2 \cdot(1.36)}\right\} .
\end{aligned}
$$

Decomposition of Normal Random Vector

## Lemma:

$$
\begin{aligned}
X= & \binom{X_{1}}{X_{2}}, \quad X_{1} \in \mathbb{R}^{r} \\
X_{2.1}= & X_{2}-\Sigma_{21} \Sigma_{11}^{-1} X_{1} \\
\Sigma \quad= & \left(\begin{array}{ll}
\Sigma_{11} & \Sigma_{12} \\
\Sigma_{21} & \Sigma_{22}
\end{array}\right) . \\
& \quad \Rightarrow \quad X_{1} \sim N_{r}\left(\mu_{1}, \Sigma_{11}\right), \\
& \quad \begin{array}{c}
\text { independent }
\end{array} \\
& \quad \Rightarrow \\
\mu_{2.1}= & \mu_{22}-\Sigma_{21} \Sigma_{11}^{-1} \mu_{1} \\
\mu_{22.1}= & \Sigma_{22}-\Sigma_{21} \Sigma_{11}^{-1} \Sigma_{12}
\end{aligned}
$$

Theorem: If $X_{1} \sim N_{r}\left(\mu_{1}, \Sigma_{11}\right)$ and $\left(X_{2} \mid X_{1}=x_{1}\right) \sim N_{p-r}\left(\mathcal{A} x_{1}+b, \Omega\right)$ where $\Omega$ does not depend on $x_{1}$, then

$$
X=\binom{X_{1}}{X_{2}} \sim N_{p}(\mu, \Sigma)
$$

where

$$
\mu=\binom{\mu_{1}}{\mathcal{A} \mu+b}
$$

and

$$
\Sigma=\left(\begin{array}{cc}
\Sigma_{11} & \Sigma_{11} \mathcal{A}^{\top} \\
\mathcal{A} \Sigma_{11} & \Omega+\mathcal{A} \Sigma_{11} \mathcal{A}^{\top}
\end{array}\right)
$$

Example: $X_{2} \in \mathbb{R}, X_{1} \in \mathbb{R}^{r}$
$E\left(X_{2} \mid X_{1}\right)=\mu_{2}+\Sigma_{21} \Sigma_{11}^{-1}\left(X_{1}-\mu_{1}\right)$
linear approximation!

$$
\begin{aligned}
X_{2} & =E\left(X_{2} \mid X_{1}\right)+U, U \sim N_{p-r}\left(0, \Sigma_{22.1}\right) \\
& =\beta_{0}+\beta^{\top} X_{1}+U \\
\sigma_{22} & =\operatorname{var}\left(X_{2}\right) \Sigma=\left(\begin{array}{cc}
\Sigma_{11} & \sigma_{12} \\
\sigma_{21} & \sigma_{22}
\end{array}\right) \\
& =\beta^{\top} \Sigma_{11} \beta+\sigma_{22.1}=\sigma_{21} \Sigma_{11}^{-1} \sigma_{12}+\sigma_{22.1}
\end{aligned}
$$

Consider the case where $r=p-1$.
Now $X_{2} \in \mathbb{R}$ and $B$ is a row vector $\beta^{\top}$ of dimension $(1 \times r)$

$$
X_{2}=\beta_{0}+\beta^{\top} X_{1}+U .
$$

This means that the best MSE approximation of $X_{2}$ by a function of $X_{1}$ is a straight line.

The marginal variance of $X_{2}$ can be decomposed as

$$
\begin{gathered}
\sigma_{22}=\beta^{\top} \Sigma_{11} \beta+\sigma_{22.1}=\sigma_{21} \Sigma_{11}^{-1} \sigma_{12}+\sigma_{22.1} \\
\rho_{2.1 \ldots r}^{2}=\frac{\sigma_{21} \Sigma_{11}^{-1} \sigma_{12}}{\sigma_{22}}
\end{gathered}
$$

is the square of the multiple correlation between $X_{2}$ and the $r$ variables $X_{1}$.

## 7 il Elementary Properties

* If $X \sim N_{p}(\mu, \Sigma)$ then a linear transformation $\mathcal{A} X+c, \mathcal{A}(q \times p), c \in \mathbb{R}^{q}$ has distribution $N_{q}\left(\mathcal{A} \mu+c, \mathcal{A} \Sigma \mathcal{A}^{\top}\right)$.
* Two linear transformations $\mathcal{A} X$ and $\mathcal{B} X$ of $X \sim N_{p}(\mu, \Sigma)$ are independent if and only if $\mathcal{A} \Sigma \mathcal{B}^{\top}=0$.
* If $X_{1}$ and $X_{2}$ are partitions of $X \sim N_{p}(\mu, \Sigma)$ then the conditional distribution of $X_{2}$ given $X_{1}=x_{1}$ is normal again and $X_{1}$ is independent of $X_{2}$ if and only if $\Sigma_{12}=0$.
* The conditional expectation of $\left(X_{2} \mid X_{1}\right)$ is a linear function for $\binom{x_{1}}{x_{2}} \sim N_{p}(\mu, \Sigma)$.
* The multiple correlation coefficient is the percentage of the variance of $X_{2}$ explained by the linear approximation $\beta_{0}+\beta^{\top} X_{1}$.


## Central Limit Theorems

Central Limit Theorem describes the (asymptotic) behaviour of sample mean
$X_{1}, X_{2}, \ldots, X_{n}$, i.i.d with $X_{i} \sim(\mu, \Sigma)$

$$
\sqrt{n}(\bar{x}-\mu) \xrightarrow{\mathcal{L}} N_{p}(0, \Sigma) \quad \text { for } \quad n \longrightarrow \infty .
$$

The CLT can be easily applied for testing.
Normal distribution plays a central role in statistics.


Figure: The CLT in the two-dimensional case. Sample size $n=5$ (left) and $n=500$ (right). $\rightarrow$ SMScltbern3


Figure: The CLT for Bernoulli distributed random variables. Sample size $n=5$ (left) and $n=35$ (right). $\rightarrow$ SMScltbern

Transformation of statistics

If $\sqrt{n}(t-\mu) \xrightarrow{\mathcal{L}} N_{p}(0, \Sigma)$ and if $f=\left(f_{1}, \ldots, f_{q}\right)^{\top}: \mathbb{R}^{p} \rightarrow \mathbb{R}^{q}$
are real valued functions which are differentiable at $\mu \in \mathbb{R}^{p}$, then $f(t)$ is asymptotically normal with mean $f(\mu)$ and covariance $\mathcal{D}^{\top} \Sigma \mathcal{D}$, i.e.,

$$
\sqrt{n}\{f(t)-f(\mu)\} \xrightarrow{\mathcal{L}} N_{q}\left(0, \mathcal{D}^{\top} \Sigma \mathcal{D}\right) \quad \text { for } \quad n \longrightarrow \infty
$$

where

$$
\mathcal{D}=\left.\left(\frac{\partial f_{j}}{\partial t_{i}}\right)(t)\right|_{t=\mu}
$$

$(p \times q)$ matrix of all partial derivatives.
This theorem can be applied, e.g., to find the "variance stabilizing" transformation.

Example:
Suppose

$$
\left\{X_{i}\right\}_{i=1}^{n} \sim(\mu, \Sigma) ; \quad \mu=\binom{0}{0}, \quad \Sigma=\left(\begin{array}{cc}
1 & 0.5 \\
0.5 & 1
\end{array}\right), \quad p=2 .
$$

We have by CLT for $n \rightarrow \infty$

$$
\sqrt{n}(\bar{x}-\mu) \xrightarrow{\mathcal{L}} N(0, \Sigma) .
$$

The distribution of $\binom{\bar{x}_{1}^{2}-\bar{x}_{2}}{\bar{x}_{1}+3 \bar{x}_{2}}$ ?
This means to consider $f=\left(f_{1}, f_{2}\right)^{\top}$ with

$$
f_{1}\left(x_{1}, x_{2}\right)=x_{1}^{2}-x_{2}, \quad f_{2}\left(x_{1}, x_{2}\right)=x_{1}+3 x_{2}, \quad q=2
$$

## 1

## Limit Theorems

* If $X_{1}, \ldots, X_{n}$ are i.i.d. random vectors with $X_{i} \sim(\mu, \Sigma)$ then the distribution of $\sqrt{n}(\bar{x}-\mu)$ is asymptotically $N(0, \Sigma)$ (Central Limit Theorem).
* If $X_{1}, \ldots, X_{n}$ are i.i.d. random variables with $X_{i} \sim(\mu, \sigma)$ then an asymptotic confidence interval can be constructed by the CLT: $\bar{x} \pm \frac{\widehat{\sigma}}{\sqrt{n}} u_{1-\alpha / 2}$.
* For small sample sizes the Bootstrap improves the precision of this confidence interval.
* If $t$ is a statistic that is asymptotically normal, i.e., $\sqrt{n}(t-\mu) \xrightarrow{\mathcal{L}} N_{p}(0, \Sigma)$, then this holds also for a function $f(t)$, i.e., $\sqrt{n}\{f(t)-f(\mu)\}$ is asymptotically normal.

Then $f(\mu)=\binom{0}{0}$ and

$$
\mathcal{D}=\left(d_{i j}\right), \quad d_{i j}=\left.\left(\frac{\partial f_{j}}{\partial x_{i}}\right)\right|_{x=\mu}=\left.\left(\begin{array}{cc}
2 x_{1} & 1 \\
-1 & 3
\end{array}\right)\right|_{x=0}=\left(\begin{array}{rr}
0 & 1 \\
-1 & 3
\end{array}\right) .
$$

We have the covariance

$$
\begin{array}{cc}
\left(\begin{array}{rr}
0 & -1 \\
1 & 3
\end{array}\right) & \left(\begin{array}{cc}
1 & \frac{1}{2} \\
\frac{1}{2} & 1
\end{array}\right)
\end{array} \underset{\mathcal{D}^{\top}}{\Sigma} \quad\left(\begin{array}{cc}
0 & 1 \\
-1 & 3
\end{array}\right)=\left(\begin{array}{cc}
1 & -\frac{7}{2} \\
-\frac{7}{2} & 13
\end{array}\right)
$$

This yields

$$
\sqrt{n}\binom{\bar{x}_{1}^{2}-\bar{x}_{2}}{\bar{x}_{1}+3 \bar{x}_{2}} \xrightarrow{\mathcal{L}} N_{2}\left(\binom{0}{0},\left(\begin{array}{cc}
1 & -\frac{7}{2} \\
-\frac{7}{2} & 13
\end{array}\right)\right) .
$$

Datové matice, testování mnohorozměrné střední hodnoty:

- Wishartovo a Hotellingovo rozdělení,
- testy vícerozměrné střední hodnoty.

Further matrix algebra: Kronecker product
Let $\mathcal{A} \otimes \mathcal{B}$ denote the Kronecker product of matrices $\mathcal{A}$ and $\mathcal{B}$ and $\operatorname{vec}(\mathcal{A})$ denote the vector obtained by stacking the columns of $\mathcal{A}$. Kronecker product and vectorization are useful tools for working with (random) matrices:

- $\alpha(\mathcal{A} \otimes \mathcal{B})=(\alpha \mathcal{A}) \otimes \mathcal{B}=\mathcal{A} \otimes(\alpha \mathcal{B})$
- $\mathcal{A} \otimes(\mathcal{B} \otimes \mathcal{C})=(\mathcal{A} \otimes \mathcal{B}) \otimes \mathcal{C}$
- $(\mathcal{A} \otimes \mathcal{B})^{\top}=\mathcal{A}^{\top} \otimes \mathcal{B}^{\top}$
- $(\mathcal{A} \otimes \mathcal{B})(\mathcal{C} \otimes \mathcal{D})=(\mathcal{A C}) \otimes(\mathcal{B D})$
- $(\mathcal{A} \otimes \mathcal{B})^{-1}=\left(\mathcal{A}^{-1} \otimes \mathcal{B}^{-1}\right)$
- $(\mathcal{A}+\mathcal{B}) \otimes \mathcal{C}=\mathcal{A} \otimes \mathcal{C}+\mathcal{B} \otimes \mathcal{C}$
- $\mathcal{A} \otimes(\mathcal{B}+\mathcal{C})=\mathcal{A} \otimes \mathcal{B}+\mathcal{A} \otimes \mathcal{C}$
- $\operatorname{vec}(\mathcal{A X B})=\left(\mathcal{B}^{\top} \otimes \mathcal{A}\right) \operatorname{vec}(\mathcal{X})$
- $\operatorname{tr}(\mathcal{A} \otimes \mathcal{B})=\operatorname{tr}(\mathcal{A}) \operatorname{tr}(\mathcal{B})$


## Week 4 Data matrices, Wishart, Hotelling

Normal data matrices: independence

Theorem: If $\mathcal{X}$ is a data matrix from $N_{p}(\mu, \Sigma)$ then $\mathcal{Y}=\mathcal{A X B}$ and $\mathcal{Z}=\mathcal{C X D}$ are independent if and only if $\mathcal{B}^{\top} \Sigma \mathcal{D}=0$ or $\mathcal{A C}^{\top}=0$.

Proof: assume $(\mathrm{WLOG})$ that $\mu=0$, then $\operatorname{vec}(\mathcal{Y})=\left(\mathcal{B}^{\top} \otimes \mathcal{A}\right) \operatorname{vec}(\mathcal{X})$, and the covariance matrix between $\operatorname{vec}(\mathcal{Y})$ and $\operatorname{vec}(\mathcal{Z})$ is

$$
\begin{aligned}
\operatorname{Evec}(\mathcal{Y}) \operatorname{vec}(\mathcal{Z})^{\top} & =\left(\mathcal{B}^{\top} \otimes \mathcal{A}\right) E \operatorname{vec}(\mathcal{X}) \operatorname{vec}(\mathcal{X})^{\top}\left(\mathcal{D}^{\top} \otimes \mathcal{C}\right)^{\top} \\
& =\left(\mathcal{B}^{\top} \otimes \mathcal{A}\right)\left(\Sigma \otimes \mathcal{I}_{n}\right)\left(\mathcal{D} \otimes \mathcal{C}^{\top}\right) \\
& =\mathcal{B}^{\top} \Sigma \mathcal{D} \otimes \mathcal{A} \mathcal{C}^{\top} .
\end{aligned}
$$

## Normal data matrices

Definition: Let $X_{1}, \ldots, X_{n}$ be a random sample from $N_{p}(\mu, \Sigma)$. Then $\mathcal{X}=\left(X_{1}, \ldots, X_{n}\right)^{\top}$ is called a data matrix from $N_{p}(\mu, \Sigma)$.

Clearly, if $\mathcal{X}$ is a data matrix from $N_{p}(\mu, \Sigma)$ then $\bar{x}_{n}=\mathcal{X}^{\top} 1_{n} / n \sim N_{p}(\mu, \Sigma / n)$.

Theorem: If $\mathcal{X}$ is a data matrix from $N_{p}(\mu, \Sigma)$ then $\mathcal{Y}=\mathcal{A X B} \sim N_{q}\left(\alpha \mathcal{B}^{\top} \mu, \beta \mathcal{B}^{\top} \Sigma \mathcal{B}\right)$ if and only if:

- $\mathcal{A} 1_{n}=\alpha 1_{m}$ for some $\alpha \in \mathbb{R}$, or $\mathcal{B}^{\top} \mu=0$, and
- $\mathcal{A} \mathcal{A}^{\top}=\beta \mathcal{I}_{m}$ for some $\beta \in \mathbb{R}$, or $\mathcal{B}^{\top} \Sigma \mathcal{B}=0_{q} 0_{q}^{\top}$.

The proof is based on vectorization of $\mathcal{X}$ (i.e., column stacking): $\operatorname{vec}(\mathcal{X}) \sim N_{n p}\left(\mu \otimes 1_{n}, \Sigma \otimes \mathcal{I}_{n}\right)$ and $\operatorname{vec}(\mathcal{A X B})=\left(\mathcal{B}^{\top} \otimes \mathcal{A}\right) \operatorname{vec}(\mathcal{X})$.

## Wishart distribution

Definition: Assuming that $\mathcal{X}$ is a data matrix from $N_{p}\left(0_{p}, \Sigma\right)$, the random matrix

$$
\mathcal{M}(p \times p)=\mathcal{X}^{\top} \mathcal{X} \sim W_{p}(\Sigma, n)
$$

where $W_{p}(\Sigma, n)$ denotes Wishart distribution with parameters $\Sigma$ and $n$.
Example:
$p=1, \quad X \quad \sim \quad N_{1}\left(0, \sigma^{2}\right)$
$\mathcal{X}=\left(\begin{array}{c}x_{1} \\ \vdots \\ x_{n}\end{array}\right) \quad \mathcal{M}=\mathcal{X}^{\top} \mathcal{X}=\sum_{i=1}^{n} x_{i}^{2} \sim \sigma^{2} \chi_{n}^{2}$
It follows that Wishart distribution is generalisation of $\chi_{n}^{2}$

Theorem:

$$
\mathcal{M} \sim W_{p}(\Sigma, n) \text { and } \mathcal{B}(p \times q) \Rightarrow \mathcal{B}^{\top} \mathcal{M B} \sim W_{q}\left(\mathcal{B}^{\top} \Sigma \mathcal{B}, n\right)
$$

Theorem:(Cochran) $\mathcal{X}(n \times p)$ is data matrix with $N_{p}(\mu, \Sigma)$. Then:

- $\mathcal{X}^{\top} \mathcal{C X}$, where $\mathcal{C}$ is symmetric, has the same distribution as a weighted sum of independent $W_{p}(\Sigma, 1)$ matrices, where the weights are eigenvalues of $\mathcal{C}$.
- $n \mathcal{S}=\mathcal{X}^{\top} \mathcal{H} \mathcal{X} \sim W_{p}(\Sigma, n-1)$,
- $\bar{x}$ and $\mathcal{S}$ are independent.

Proof: see Theorem 3.4.4. (page 68) in MKB (using spectral decomposition of $\mathcal{C}$ or $\mathcal{H}$ ).
Z. Hlávka (KPMS)

## Week 4 Data matrices, Wishart, Hotelling

## Hotelling's $T^{2}$-distribution

Definition: Assume that random vector $Y \sim N_{p}(0, \mathcal{I})$ is independent of random matrix $\mathcal{M} \sim W_{p}(\mathcal{I}, n)$. Then

$$
n Y^{\top} \mathcal{M}^{-1} Y \sim T^{2}(p, n)
$$

where $T^{2}(p, n)$ denotes Hotelling's distribution with parameters $p$ and $n$.
Hotelling's $T^{2}$ generalizes Student's $t$-distribution
The critical values of Hotelling's $T^{2}$ can be calculated using $F$-distribution:

$$
T^{2}(p, n)=\frac{n p}{n-p+1} \quad F_{p, n-p+1}
$$

## Wilks' $\Lambda$-distribution

Definition: Assume that $\mathcal{A} \sim W_{p}(\mathcal{I}, m)$ and $\mathcal{B} \sim W_{p}(\mathcal{I}, n)$ are indepependent, $m \geq p$, we say that the random variable

$$
\Lambda=|\mathcal{A}| /|\mathcal{A}+\mathcal{B}|
$$

has Wilks' lambda distribution with parameters $p, m$, and $n$, i.e., $\Lambda \sim \Lambda(p, m, n)$

This distribution occurs frequently in likelihood ratio tests.
The random variable $\Lambda$ is basically a ratio of two 'generalized variances'-therefore, Wilks' $\Lambda$ distribution can be seen as a multivariate generalization of $F$ distribution.

## Till Distributions related to multinormal

* The Wishart distribution is a generalization of the $\chi^{2}$-distribution.
* Assuming normality, the empirical covariance matrix $\mathcal{S}$ has a $\frac{1}{n} W_{p}(\Sigma, n-1)$ distribution.
* In the normal case, $\bar{x}$ and $\mathcal{S}$ are independent.
* Hotelling's $T^{2}$-distribution is a generalization of the $t$-distribution.
* $(n-1)(\bar{x}-\mu)^{\top} \mathcal{S}^{-1}(\bar{x}-\mu)$ has a $T^{2}(p, n-1)$ distribution.
* The relation between Hotelling's $T^{2}-$ and Fisher's $F$-distribution is given by $T^{2}(p, n)=\frac{n p}{n-p+1} \quad F_{p, n-p+1}$.
* Wilks' $\Lambda$-distribution can be seen as a multivariate generalization of $F$ distribution (ratio of two variances).


## Testing the multivariate mean

$X_{i} \sim N_{p}(\mu, \Sigma)$ i.i.d.

$$
H_{0}: \mu=\mu_{0}, \quad \Sigma \text { unknown, } \quad H_{1}: \quad \text { no constraints. }
$$

Under $H_{0}:(n-1)\left(\bar{x}-\mu_{0}\right)^{\top} \mathcal{S}^{-1}\left(\bar{x}-\mu_{0}\right) \sim T^{2}(p, n-1)$.
Equivalently:

$$
\left(\frac{n-p}{p}\right)\left(\bar{x}-\mu_{0}\right)^{\top} \mathcal{S}^{-1}\left(\bar{x}-\mu_{0}\right) \sim F_{p, n-p}
$$

The rejection region may be defined as

$$
\left(\frac{n-p}{p}\right)\left(\bar{x}-\mu_{0}\right)^{\top} \mathcal{S}^{-1}\left(\bar{x}-\mu_{0}\right)>F_{1-\alpha ; p, n-p}
$$

library(mvtnorm)
$\mathrm{s}=$ matrix $(\mathrm{c}(1,-0.5,-0.5,1), 2) ; \mathrm{x}=\mathrm{seq}(-3,3, \mathrm{by}=0.015)$
contour (x, x, outer (x, x,
function $(x, y)\{d m v n o r m(\operatorname{cbind}(x, y)$, sigma=s)\}))
$\mathrm{n}=20$; X=rmvnorm(n, sigma=s); m=apply (X, 2, mean) ; $\mathrm{S}=\operatorname{cov}(\mathrm{X})$
points(m[1],m[2],pch=8, col="red", cex=2)
\#contour ( $\mathrm{x}, \mathrm{x}$, outer ( $\mathrm{x}, \mathrm{x}$, function $(\mathrm{x}, \mathrm{y})\{(\mathrm{n}-2)$ *
\# $\operatorname{diag}(t(t(\operatorname{cbind}(x, y))-m) \% * \% \operatorname{solve}(S) \% * \%(t(\operatorname{cbind}(x, y))-m))<$
\# 2*qf(0.95,2,n-2)\}), col="red",add=TRUE)
S1=solve(S)
contour ( $\mathrm{x}, \mathrm{x}$, outer ( $\mathrm{x}, \mathrm{x}$, function $(\mathrm{x}, \mathrm{y})\{(\mathrm{n}-2)$ *
apply(t (t (cbind (x,y))-m), 1 , function(x) $\{t(x) \% * \% S 1 \% * \% x\})<$ $2 * q f(0.95,2, n-2)\})$, col="red", add=TRUE)

[^0]R: library(DescTools); help(HotellingsT2Test)

## Confidence region for $\mu$

$\left(\frac{n-p}{p}\right)(\bar{x}-\mu)^{\top} \mathcal{S}^{-1}(\bar{x}-\mu) \sim F_{p, n-p}$

$$
\left\{\mu \in \mathbb{R}^{p} \left\lvert\,(\mu-\bar{x})^{\top} \mathcal{S}^{-1}(\mu-\bar{x}) \leq \frac{p}{n-p} F_{1-\alpha ; p, n-p}\right.\right\}
$$

is a confidence region at level (1- $\alpha$ ) for $\mu$; it is the interior of an iso-distance ellipsoid in $\mathbb{R}^{p}$.

When $p$ is large, ellipsoids are not easy to handle for practical purposes. One is thus interested in finding confidence intervals for $\mu_{1}, \mu_{2}, \ldots, \mu_{p}$ so that simultaneous confidence on all the intervals reaches the desired level say, $1-\alpha$.

## Week 4 Data matrices, Wishart, Hotelling

## Simultaneous Confidence Intervals for $a^{\top} \mu$

Obvious confidence interval for certain $a^{\top} \mu$ is given by:

$$
\left|\frac{\sqrt{n-1}\left(a^{\top} \mu-a^{\top} \bar{x}\right)}{\sqrt{a^{\top} \mathcal{S} a}}\right| \leq t_{1-\frac{\alpha}{2} ; n-1}
$$

or equivalently

$$
t^{2}(a)=\frac{(n-1)\left\{a^{\top}(\mu-\bar{x})\right\}^{2}}{a^{\top} \mathcal{S} a} \leq F_{1-\alpha ; 1, n-1}
$$

which provides the $(1-\alpha)$ confidence interval for $a^{\top} \mu$ :

$$
\left(a^{\top} \bar{x}-\sqrt{F_{1-\alpha ; 1, n-1} \frac{a^{\top} \mathcal{S} a}{n-1}} \leq a^{\top} \mu \leq a^{\top} \bar{x}+\sqrt{F_{1-\alpha ; 1, n-1} \frac{a^{\top} \mathcal{S} a}{n-1}}\right)
$$

Using Theorem on maximum of quadratic forms we see that:

$$
\max _{a} t^{2}(a)=(n-1)(\bar{x}-\mu)^{\top} \mathcal{S}^{-1}(\bar{x}-\mu) \sim T^{2}(p, n-1)
$$

$$
\max _{a} t^{2}(a)=(n-1)(\bar{x}-\mu)^{\top} \mathcal{S}^{-1}(\bar{x}-\mu) \sim T^{2}(p, n-1)
$$

implies that the simultaneous confidence intervals for all possible linear combinations $a^{\top} \mu, a \in \mathbb{R}^{p}$ of the elements of $\mu$ is given by:

$$
\left(a^{\top} \bar{x}-\sqrt{K_{\alpha} a^{\top} \mathcal{S} a}, a^{\top} \bar{x}+\sqrt{K_{\alpha} a^{\top} \mathcal{S} a}\right)
$$

where $K_{\alpha}=\frac{p}{n-p} F_{1-\alpha ; p, n-p}$.
Example:
$95 \%$ confidence region for $\mu_{f}$, the mean of the forged banknotes, is given by the ellipsoid:

$$
\left\{\mu \in \mathbb{R}^{6} \left\lvert\,\left(\mu-\bar{x}_{f}\right)^{\top} S_{f}^{-1}\left(\mu-\bar{x}_{f}\right) \leq \frac{6}{94} F_{0.95 ; 6,94}\right.\right\}
$$

## Week 4 <br> Data matrices, Wishart, Hotelling

Testing the difference of two multivariate means

Suppose $X_{i 1} \sim N_{p}\left(\mu_{1}, \Sigma\right), i=1, \cdots, n_{1}$ and $X_{j 2} \sim N_{p}\left(\mu_{2}, \Sigma\right), j=1, \cdots, n_{2}$, all the random vectors being independent.

$$
H_{0}: \mu_{1}=\mu_{2}, \quad H_{1}: \quad \text { no constraints. }
$$

Both samples provide the statistics $\bar{x}_{k}$ and $\mathcal{S}_{k}, \mathrm{k}=1,2$.
Let $\delta=\mu_{1}-\mu_{2}$, we have

$$
\begin{aligned}
\left(\bar{x}_{1}-\bar{x}_{2}\right) & \sim N_{p}\left(\delta, \frac{n_{1}+n_{2}}{n_{1} n_{2}} \Sigma\right) \\
n_{1} S_{1}+n_{2} S_{2} & \sim W_{p}\left(\Sigma, n_{1}+n_{2}-2\right) .
\end{aligned}
$$

$95 \%$ simultaneous c.i. are given by (using $F_{0.95 ; 6,94}=2.1966$ )

| 214.692 | $\leq \mu_{1} \leq 214.954$ |
| ---: | :--- |
| 130.205 | $\leq \mu_{2} \leq 130.395$ |
| 130.082 | $\leq \mu_{3} \leq 130.304$ |
| 10.108 | $\leq \mu_{4} \leq 10.952$ |
| 10.896 | $\leq \mu_{5} \leq 11.370$ |
| 139.242 | $\leq \mu_{6} \leq 139.658$ |

Comparison with $\mu_{0}=(214.9,129.9,129.7,8.3,10.1,141.5)^{\top}$ shows that almost all components (except the first one) are responsible for the rejection of $\mu_{0}$.

In addition, choosing e.g. $a^{\top}=(0,0,0,1,-1,0)$ gives c.i.
$-1.211 \leq \mu_{4}-\mu_{5} \leq 0.005$ shows that for the forged bills, the lower border is essentially smaller than the upper border.

The rejection region is:

$$
\begin{gathered}
\frac{n_{1} n_{2}\left(n_{1}+n_{2}-p-1\right)}{p\left(n_{1}+n_{2}\right)^{2}}\left(\left(\bar{x}_{1}-\bar{x}_{2}\right)\right)^{\top} \mathcal{S}^{-1}\left(\left(\bar{x}_{1}-\bar{x}_{2}\right)\right) \\
\geq F_{1-\alpha ; p, n_{1}+n_{2}-p-1}
\end{gathered}
$$

A $(1-\alpha) * 100 \%$ confidence region for $\delta$ is given by the ellipsoid centered at $\left(\bar{x}_{1}-\bar{x}_{2}\right)$

$$
\begin{aligned}
& \left(\delta-\left(\bar{x}_{1}-\bar{x}_{2}\right)\right)^{\top} \mathcal{S}^{-1}\left(\delta-\left(\bar{x}_{1}-\bar{x}_{2}\right)\right) \\
\leq & \frac{p\left(n_{1}+n_{2}\right)^{2}}{\left(n_{1}+n_{2}-p-1\right)\left(n_{1} n_{2}\right)} F_{1-\alpha ; p, n_{1}+n_{2}-p-1}
\end{aligned}
$$

and the simultaneous confidence intervals for all linear combinations of the elements of $\delta: a^{\top} \delta$ are given by

$$
a^{\top} \delta \in a^{\top}\left(\bar{x}_{1}-\bar{x}_{2}\right) \pm \sqrt{\frac{p\left(n_{1}+n_{2}\right)^{2}}{\left(n_{1}+n_{2}-p-1\right)\left(n_{1} n_{2}\right)} F_{1-\alpha ; p, n_{1}+n_{2}-p-1} a^{\top} \mathcal{S} a}
$$

Example: We want to compare the mean of the assets $\left(X_{1}\right)$ and of the sales $\left(X_{2}\right)$ of the two sectors energy (group 1) and manufacturing (group 2).

We have the following statistics $n_{1}=15, n_{2}=10, p=2$,
$\bar{x}_{1}=\binom{4084}{2580.5}, \bar{x}_{1}=\binom{4307.2}{4925.2}$,
$\mathcal{S}_{1}=10^{7} *\left(\begin{array}{ll}1.6635 & 1.2410 \\ 1.2410 & 1.3747\end{array}\right)$,
$\mathcal{S}_{2}=10^{7} *\left(\begin{array}{ll}1.2248 & 1.1425 \\ 1.1425 & 1.5112\end{array}\right)$,
$\mathcal{S}=10^{7} *\left(\begin{array}{ll}1.4880 & 1.2016 \\ 1.2016 & 1.4293\end{array}\right)$.

Testing means with unequal covariance matrices I
Suppose $X_{i 1} \sim N_{p}\left(\mu_{1}, \Sigma_{1}\right), i=1, \cdots, n_{1}$ and
$X_{j 2} \sim N_{p}\left(\mu_{2}, \Sigma_{2}\right), j=1, \cdots, n_{2}$, all the variables being independent.

$$
H_{0}: \mu_{1}=\mu_{2}, \quad H_{1}: \quad \text { no constraints. }
$$

$$
\left(\bar{x}_{1}-\bar{x}_{2}\right) \sim N_{p}\left(\delta, \frac{\Sigma_{1}}{n_{1}}+\frac{\Sigma_{2}}{n_{2}}\right) .
$$

Therefore,

$$
\left(\bar{x}_{1}-\bar{x}_{2}\right)^{\top}\left(\frac{\Sigma_{1}}{n_{1}}+\frac{\Sigma_{2}}{n_{2}}\right)^{-1}\left(\bar{x}_{1}-\bar{x}_{2}\right) \sim \chi_{p}^{2}
$$

Since $\mathcal{S}_{i}$ is a consistent estimator of $\Sigma_{i}, i=1,2$ we have

$$
\left(\bar{x}_{1}-\bar{x}_{2}\right)^{\top}\left(\frac{\mathcal{S}_{1}}{n_{1}}+\frac{\mathcal{S}_{2}}{n_{2}}\right)^{-1}\left(\bar{x}_{1}-\bar{x}_{2}\right) \rightarrow \chi_{p}^{2}
$$

The observed value of the test statistic is $F_{\text {obs }}=2.7036$.
Since $F_{0.95 ; 2,22}=3.4434$ the hypothesis of equal means of the two groups is not rejected although it would be rejected at a less severe level ( $p-$ value $=0.0892$ ).

The $95 \%$ simultaneous confidence intervals for the differences are given by

$$
\begin{aligned}
-4628.6 & \leq \mu_{1 a}-\mu_{2 a} \leq 4182.2 \\
-6662.4 & \leq \mu_{1 s}-\mu_{2 s} \leq 1973.0
\end{aligned}
$$ ( $n_{1}$ and $n_{2}$ are both large). The test statistic turns out to be 2436.8 which is again highly significant. The $95 \%$ simultaneous confidence intervals are now:

| -0.0389 | $\leq \delta_{1} \leq$ | 0.3309 |
| ---: | :--- | ---: |
| -0.5140 | $\leq \delta_{2} \leq-0.2000$ |  |
| -0.6368 | $\leq \delta_{3} \leq-0.3092$ |  |
| -2.6846 | $\leq \delta_{4} \leq-1.7654$ |  |
| -1.2858 | $\leq \delta_{5} \leq-0.6442$ |  |
| 1.8146 | $\leq \delta_{6} \leq \quad 2.3194$ |  |

showing that all the components except the first are different from zero, the larger difference coming from $X_{6}$ (length of the diagonal) and $X_{4}$ (lower border).

## Testing means with unequal covariance matrices II

Clearly, the $\chi^{2}$ approximation to the distribution of the test statistic

$$
\left(\bar{x}_{1}-\bar{x}_{2}\right)^{\top}\left(\frac{\mathcal{S}_{1}}{n_{1}}+\frac{\mathcal{S}_{2}}{n_{2}}\right)^{-1}\left(\bar{x}_{1}-\bar{x}_{2}\right)
$$

is usable only for sufficiently large sample sizes.
For smaller sample sizes, one can use approximate likelihood ratio tests (Mardia et al, Section 5.4.1) or Welch approximation to degrees of freedom (Mardia et al, Section 5.4.2).

Note: the problem of testing equality of means without equality of variances is known as the Behrens-Fisher problem (at least in the univariate case).

## Estimation

The aim is to estimate vector of parameters $\theta$ from a sample $\mathcal{X}$ through estimators $\widehat{\theta}(\mathcal{X})$.

Most common approaches:

- maximum likelihood,
- Bayesian approach,
- robust methods (M-estimation).

In the following, we shortly discuss maximum likelihood theory.

Týden 5

Odhadování a testování:

- odhady metodou maximální věrohodnosti,
- testování poměrem věrohodností,
- príklady.

The Likelihood Function
$X \sim f(x, \theta) \quad$ pdf. parameter $\theta$
Likelihood function

$$
L(\mathcal{X} ; \theta)=\prod_{i=1}^{n} f\left(x_{i} ; \theta\right)
$$

MLE

$$
\widehat{\theta}=\arg \max _{\theta} L(\mathcal{X} ; \theta)
$$

log-likelihood

$$
\ell(\mathcal{X} ; \theta)=\log L(\mathcal{X} ; \theta)
$$

## Derivatives

## Function $f: \mathbb{R}^{p} \rightarrow \mathbb{R}$

$\frac{\partial f(x)}{\partial x}$ is the gradient, i.e., column vector of partial derivatives
$\left\{\frac{\partial f(x)}{\partial x_{j}}\right\}, j=1, \ldots, p$
$\frac{\partial f(x)}{\partial x^{\top}}$ row vector of the same derivative
$\frac{\partial^{2} f(x)}{\partial x \partial x^{\top}}$ is the $(p \times p)$ Hessian matrix of second derivatives
$\frac{\partial^{2} f(x)}{\partial x_{i} \partial x_{j}}, i=1, \ldots, p, j=1, \ldots, p$.

## Week 5

Testing and estimation
Derivative of trace and determinant

This is useful for derivation of MLEs for multivariate normal distribution:

$$
\begin{gathered}
\frac{\partial \operatorname{tr} \mathcal{X} \mathcal{A}}{\partial \mathcal{X}}= \begin{cases}\mathcal{A}^{\top} & \text { if elements of } \mathcal{A} \text { are distinct, } \\
\mathcal{A}+\mathcal{A}^{\top}-\operatorname{diag}(\mathcal{A}) & \text { for } \mathcal{A} \text { symmetric. }\end{cases} \\
\frac{\partial|\mathcal{X}|}{\partial x_{i j}}=x_{i j} \text { if elements of } \mathcal{X} \text { are distinct. } \\
\frac{\partial|\mathcal{X}|}{\partial x_{i j}}=\begin{array}{ll}
x_{i j} & \text { for } i=j \\
2 x_{i j} & \text { for } i \neq j
\end{array} \quad \text { for } \mathcal{X} \text { symmetric. }
\end{gathered}
$$

For $\mathcal{V}=\Sigma^{-1}$ symmetric it follows that:

$$
\frac{\partial \log |\mathcal{V}|}{\partial \mathcal{V}}=2 \Sigma-\operatorname{diag}(\Sigma)
$$

Example: $\left\{x_{i}\right\}_{i=1}^{n}$ is a sample from a normal distribution $N_{p}(\mu, \Sigma)$
Due to the symmetry of $\Sigma$, the unknown parameter $\theta$ is in fact $\left\{p+\frac{1}{2} p(p+1)\right\}$-dimensional.

$$
\begin{gathered}
L(\mathcal{X} ; \theta)=|2 \pi \Sigma|^{-n / 2} \exp \left\{-\frac{1}{2} \sum_{i=1}^{n}\left(x_{i}-\mu\right)^{\top} \Sigma^{-1}\left(x_{i}-\mu\right)\right\} \\
\ell(\mathcal{X} ; \theta)=-\frac{n}{2} \log |2 \pi \Sigma|-\frac{1}{2} \sum_{i=1}^{n}\left(x_{i}-\mu\right)^{\top} \Sigma^{-1}\left(x_{i}-\mu\right)
\end{gathered}
$$

After some calculations, the log-likelihood function for $N_{p}(\mu, \Sigma)$ is:

$$
\ell(\mathcal{X} ; \theta)=-\frac{n}{2} \log |2 \pi \Sigma|-\frac{n}{2} \operatorname{tr}\left\{\Sigma^{-1} \mathcal{S}\right\}-\frac{n}{2}(\bar{x}-\mu)^{\top} \Sigma^{-1}(\bar{x}-\mu)
$$

Testing and estimation

## Score and Fisher information

The score function $s(\mathcal{X} ; \theta)$ is the derivative of the log-likelihood function w.r.t. $\theta \in \mathbb{R}^{k}$

$$
s(\mathcal{X} ; \theta)=\frac{\partial}{\partial \theta} \ell(\mathcal{X} ; \theta)=\frac{1}{L(\mathcal{X} ; \theta)} \frac{\partial}{\partial \theta} L(\mathcal{X} ; \theta) .
$$

The covariance matrix

$$
\mathcal{F}_{n}=E\left\{s(\mathcal{X} ; \theta) s(\mathcal{X} ; \theta)^{\top}\right\}=\operatorname{Var}\{s(\mathcal{X} ; \theta)\}=-E\left\{\frac{\partial^{2}}{\partial \theta \partial \theta^{\top}} \ell(\mathcal{X} ; \theta)\right\}
$$

is called the Fisher information matrix.

## Asymptotic normality of MLEs

Another important result says that the MLE is asymptotically unbiased, efficient (minimum variance), and normally distributed.

Theorem: Suppose that the sample $\left\{x_{i}\right\}_{i=1}^{n}$ is i.i.d. If $\widehat{\theta}$ is the MLE for $\theta \in \mathbb{R}^{k}$ then under some regularity conditions, as $n \rightarrow \infty$ :

$$
\sqrt{n}(\widehat{\theta}-\theta) \xrightarrow{\mathcal{L}} N_{k}\left(0, \mathcal{F}_{1}^{-1}\right),
$$

where $\mathcal{F}_{1}$ denotes the Fisher information for sample size $n=1$.
This result gives us a very useful and simple approximation whenever we are not able to calculate the exact distribution of the MLE $\widehat{\theta}$.

Even in very complicated situations, the Fisher information matrix can be approximated numerically.
Z. Hlávka (KPMS)

NMST539


## Week 5 Likelihood ratio tests

## Likelihood ratio tests (LRTs)

Consider hypotheses:

$$
\begin{aligned}
& H_{0}: \theta \in \Omega_{0}, \\
& H_{1}:
\end{aligned} \quad \theta \in \Omega_{1},
$$

where $\theta$ is a parameter of the distribution of $\left\{x_{i}\right\}_{i=1}^{n}, x_{i} \in \mathbb{R}^{p}$.
Wilks' Theorem says:
Theorem: If $\Omega_{1} \subset \mathbb{R}^{q}$ is a $q$-dimensional space and if $\Omega_{0} \subset \Omega_{1}$ is an $r$-dimensional subspace, then under regularity conditions:

$$
\forall \theta \in \Omega_{0}:-2 \log \lambda=2\left(\ell_{1}^{*}-\ell_{0}^{*}\right) \xrightarrow{\mathcal{L}} \chi_{q-r}^{2} \quad \text { as } \quad n \rightarrow \infty,
$$

where $\ell_{j}^{*}, j=1,2$ are the maxima of the log-likelihood for each hypothesis.

## Hypothesis Testing

* Maximum likelihood estimators are easy to calculate but we have to know the true distribution.
* MLEs have asymptotically normal distribution.
* The asymptotic normality of transformed MLEs can be derived by using Delta theorem.
* MLEs are asymptotically optimal.


## Week 5 Likelihood ratio tests

Testing the multivariate mean

$$
\begin{aligned}
& X_{i} \sim N_{p}(\mu, \Sigma) \text { i.i.d. } \\
& \quad H_{0}: \mu=\mu_{0}, \quad \Sigma \text { unknown }, \quad H_{1}: \text { no constraints. }
\end{aligned}
$$

Under $H_{0}$ it can be shown that

$$
\ell_{0}^{*}=\ell\left(\mu_{0}, \mathcal{S}+d d^{\top}\right), \quad d=\left(\bar{x}-\mu_{0}\right)
$$

and under $H_{1}$ we have

$$
\ell_{1}^{*}=\ell(\bar{x}, \mathcal{S}) .
$$

This leads to

$$
-2 \log \lambda=2\left(\ell_{1}^{*}-\ell_{0}^{*}\right)=n \log \left(1+d^{\top} \mathcal{S}^{-1} d\right)
$$

Note that this statistic depends on $(n-1) d^{\top} \mathcal{S}^{-1} d$ which has, under $H_{0}$, a Hotelling's $T^{2}$-distribution.
Z. Hlávka (KPMS)

## Testing the variance matrix

```
\(X_{i} \sim N_{p}(\mu, \Sigma)\) i.i.d.
\[
H_{0}: \Sigma=\Sigma_{0}, \quad \mu \text { unknown, } \quad H_{1}: \quad \text { no constraints. }
\]
```

Looking at the log-likelihood function, we observe that the MLEs are $\bar{x}$ and $\Sigma_{0}$ under $H_{0}$ and $\bar{x}$ and $\mathcal{S}$ under the alternative.

Therefore:

$$
I_{0}^{*}=-\frac{1}{2} n \log \left|2 \pi \Sigma_{0}\right|-\frac{1}{2} n \operatorname{tr} \Sigma_{0}^{-1} \mathcal{S}, \text { and } I_{1}^{*}=-\frac{1}{2} n \log |2 \pi \mathcal{S}|-\frac{1}{2} n p .
$$

This leads to

$$
-2 \log \lambda=n \operatorname{tr} \Sigma_{0}^{-1} \mathcal{S}-n \log \left|\Sigma_{0}^{-1} \mathcal{S}\right|-n p=n p(a-\log g-1)
$$

where $a$ and $g$ denote respectively the arithmetic and geometric mean of eigenvalues of $\Sigma_{0}^{-1} \mathcal{S}$.

## Test of independence

$$
\text { Let } X_{i} \sim N_{p}(\mu, \Sigma), i=1 \cdots, n \text { be independent, }
$$

$$
H_{0}: \Sigma_{12}=0, \quad H_{1}: \quad \text { no constraints. }
$$

We partition the variables into two sets with dimensions $p_{1}$ and $p_{2}$. The estimators under $H_{0}$ is

$$
\hat{\mu}=\bar{x}_{n}
$$

and

$$
\hat{\Sigma}=\left(\begin{array}{cc}
S_{11} & 0 \\
0 & S_{22}
\end{array}\right)
$$

## Test of homogeneity of covariances

$$
\text { Let } X_{i h} \sim N_{p}\left(\mu_{h}, \Sigma_{h}\right), i=1 \cdots, n_{h} ; h=1, \cdots, k
$$

all variables being independent,

$$
H_{0}: \Sigma_{1}=\Sigma_{2}=\cdots=\Sigma_{k}, \quad H_{1}: \quad \text { no constraints. }
$$

$\mathcal{S}_{h}$ is the MLE estimator of $\Sigma_{h}$ under the alternative and the weighted average $\mathcal{S}=\frac{n_{1} \mathcal{S}_{1}+\cdots+n_{k} \mathcal{S}_{k}}{n}$ is the MLE of $\Sigma$ under the null $\left(H_{0}\right)$.

The likelihood ratio test leads to the statistic

$$
-2 \log \lambda=n \log |\mathcal{S}|-\sum_{h=1}^{k} n_{h} \log \left|\mathcal{S}_{h}\right|
$$

which under $H_{0}$ is approximately distributed as a $\mathcal{X}_{m}^{2}$ where $m=\frac{1}{2}(k-1) p(p+1)$.

It follows that the LRT test statistics for $H_{0}: \Sigma_{12}=0$ is:

$$
\begin{aligned}
-2 \log \lambda & =-n \log \left|\hat{\Sigma}^{-1} S\right|=-n \log \left|S_{22}-S_{21} S_{11}^{-1} S_{12}\right| /\left|S_{22}\right| \\
& =-n \log \left|\mathcal{I}-S_{22}^{-1} S_{21} S_{11}^{-1} S_{12}\right|=-n \log \prod_{i=1}^{k}\left(1-\lambda_{i}\right)
\end{aligned}
$$

where $\lambda_{i}$ are non-zero eigenvalues of $S_{22}^{-1} S_{21} S_{11}^{-1} S_{12}$.
It can be shown that the test statistics follows the so-called Wilks' lambda distribution (distribution of a ratio of determinants of independent Wishart matrices).

This test is applicable in canonical correlation analysis (investigating correlations between two sets of variables).

For $p_{1}=1$, the LRT test statistics simplifies to a function of multiple correlation coefficient.

## Tests of multivariate normality

Multivariate skewness:

$$
\beta_{1, p}=E\left\{(X-\mu)^{\top} \Sigma^{-1}(Y-\mu)\right\}^{3},
$$

where $X$ and $Y$ are iid.
Multivariate kurtosis:

$$
\beta_{2, p}=E\left\{(X-\mu)^{\top} \Sigma^{-1}(X-\mu)\right\}^{2}
$$

It can be shown that $\beta_{1, p}=0$ and $\beta_{2, p}=p(p+2)$ for $X \sim N_{p}(\mu, \Sigma)$. This easily follows from the symmetry of $V=(X-\mu)^{\top} \Sigma^{-1}(Y-\mu)$ and $(X-\mu)^{\top} \Sigma^{-1}(X-\mu) \sim \chi_{p}^{2}$.

## Týden 6-7

## Metoda hlavních komponent:

- definice a interpretace,
- standardizace,
- asymptotické vlastnosti,
- použití.


## Tests of multivariate normality

Assuming normality, the distribution of $b_{1, p}$ and $b_{2, p}$ (the sample counterparts of $\beta_{1, p}$ and $\beta_{2, p}$ ) is:

$$
\frac{1}{6} n b_{1, p} \sim \chi_{p(p+1)(p+1) / 6}^{2}
$$

and

$$
\sqrt{n} \frac{b_{2, p}-p(p+2)}{\sqrt{8 p(p+2)}} \sim N(0,1)
$$

Note: QQ diagram can be plotted using quantiles of $\chi_{p}^{2}$ distribution and ordered values of $n\left(X_{i}-\bar{x}\right)^{\top} S^{-1}\left(X_{i}-\bar{x}\right)$.

## Principal Components

Principal components are (orthogonal) linear combinations maximizing the variance of standardized linear combinations (SLC):

$$
\delta^{\top} X=\sum_{j=1}^{p} \delta_{j} X_{j} \text { such that }\|\delta\|=\delta^{\top} \delta=1
$$

Maximizing:

$$
\max _{\{\delta:\|\delta\|=1\}} \operatorname{var}\left(\delta^{\top} X\right)=\max _{\{\delta:\|\delta\|=1\}} \delta^{\top} \operatorname{Var}(X) \delta
$$

is easy using the spectral decomposition $\operatorname{Var}(X)=\Gamma \Lambda \Gamma^{\top}$.

## Example:

Bivariate normal distribution $N(0, \Sigma), \Sigma=\left(\begin{array}{l}1 \\ \rho \\ \rho\end{array}\right), \rho>0$.
Eigenvalues of this matrix are $\lambda_{1}=1+\rho$ and $\lambda_{2}=1-\rho$ with corresponding eigenvectors

$$
\gamma_{1}=\frac{1}{\sqrt{2}}\binom{1}{1}, \quad \gamma_{2}=\frac{1}{\sqrt{2}}\binom{1}{-1} .
$$

The PC transformation is thus
or

$$
Y=\Gamma^{\top}(X-\mu)=\frac{1}{\sqrt{2}}\left(\begin{array}{rr}
1 & 1 \\
1 & -1
\end{array}\right) X
$$

$$
\binom{Y_{1}}{Y_{2}}=\frac{1}{\sqrt{2}}\binom{X_{1}+X_{2}}{X_{1}-X_{2}} .
$$

## Properties of PCs

Let $X \sim(\mu, \Sigma)$ and let $Y$ be the PC transformation $Y=\Gamma^{\top}(X-\mu)$.
Then

$$
\begin{aligned}
E Y & =0_{p} \\
\operatorname{var}(Y) & =\Lambda \\
\operatorname{var}\left(Y_{1}\right) & \geq \cdots \geq \operatorname{var}\left(Y_{p}\right) \geq 0 \\
\sum_{j} \operatorname{var}\left(Y_{j}\right) & =\sum_{j} \lambda_{j}=\operatorname{tr}(\Sigma)=\sum_{j} \operatorname{var}\left(X_{j}\right) \\
\prod \operatorname{var}\left(Y_{j}\right) & =|\Sigma| .
\end{aligned}
$$

Note: $|\Sigma|$ is called the (population) generalized variance and $\operatorname{tr}(\Sigma)$ the (population) total variation.

The first principal component is

$$
Y_{1}=\frac{1}{\sqrt{2}}\left(X_{1}+X_{2}\right)
$$

and the second is

$$
Y_{2}=\frac{1}{\sqrt{2}}\left(X_{1}-X_{2}\right) .
$$

Let us compute the variances of these PCs:

$$
\begin{aligned}
\operatorname{var}\left(Y_{1}\right) & =\operatorname{var}\left\{\frac{1}{\sqrt{2}}\left(X_{1}+X_{2}\right)\right\}=\frac{1}{2} \operatorname{var}\left(X_{1}+X_{2}\right) \\
& =\frac{1}{2}\left\{\operatorname{var}\left(X_{1}\right)+\operatorname{var}\left(X_{2}\right)+2 \operatorname{cov}\left(X_{1}, X_{2}\right)\right\} \\
& =\frac{1}{2}(1+1+2 \rho)=1+\rho \\
& =\lambda_{1}
\end{aligned}
$$

Similarly we find that: $\operatorname{var}\left(Y_{2}\right)=\lambda_{2}$.


## Week 6-7

 Principal componentsExample: In practice, the sample principal components are calculated from the sample variance matrix:

$$
\begin{aligned}
\mathcal{S} & =\mathcal{G} \mathcal{L G} \mathcal{G}^{\top} \\
\mathcal{Y} & =\left(\mathcal{X}-1_{n} \bar{X}^{\top}\right) \mathcal{G}
\end{aligned}
$$

```
data(bank2)
eigen(var(bank2))
pcb=prcomp(bank2)
pcb
plot(pcb)
pcb$x
```


## PCA stopping rules

For dimension reduction, the number of PCs is usually chosen by simple ad-hoc rules:

- scree-plot (of eigenvalues),
- log-eigenvalue diagram (LEV),
- percentage of total variation (explain 80 or $90 \%$ of total variation),
- Kaiser criterion (choose PCs with higher than the "average variance").

The interpretation of the Kaiser criterion simplifies for standardized data sets: $\operatorname{tr} \Sigma=p$ implies that the average variance is 1 .

In practice, one should consider standardization of variables before running PCA.

## prcomp(bank2,scale.=TRUE)

Example: Some examples:

- bank2,
- athletic records,
- geopol,
- timebudget.

The proof uses transformation $m \Gamma^{\top} \mathcal{U} \Gamma \sim W_{p}(\Lambda, m)$, see MKB, p. 231.

## Example:

Assuming normality:
$n \mathcal{S} \sim W_{p}(\Sigma, n-1)$
$\sqrt{n-1}\left(\ell_{j}-\lambda_{j}\right) \xrightarrow{\mathcal{L}} N\left(0,2 \lambda_{j}^{2}\right), \quad j=1, \ldots, p$,
using log transformation:

$$
\sqrt{\frac{n-1}{2}}\left(\log \ell_{j}-\log \lambda_{j}\right) \xrightarrow{\mathcal{L}} N(0,1)
$$

Example: The first PC for Swiss bank notes resolves 67\% of the variation. Let us test whether the true proportion could be $75 \%$.

The 95\% confidence interval for the true proportion is

$$
0.668 \pm 1.96 \sqrt{\frac{0.142}{199}}=(0.615,0.720)
$$

We reject the hypothesis that $\psi=75 \%$ !

Clearly, the estimator of variance explained by first $q$ PCs $\widehat{\psi}=\left(\ell_{1}+\cdots+\ell_{q}\right) / \sum_{j=1}^{p} \ell_{j}$ is a nonlinear transformation of $\ell$.

Therefore,

$$
\sqrt{n-1}(\widehat{\psi}-\psi) \xrightarrow{\mathcal{L}} N\left(0, \omega^{2}\right),
$$

where

$$
\begin{aligned}
\omega^{2} & =\frac{2}{\{\operatorname{tr}(\Sigma)\}^{2}}\left\{(1-\psi)^{2}\left(\lambda_{1}^{2}+\cdots+\lambda_{q}^{2}\right)+\psi^{2}\left(\lambda_{q+1}^{2}+\cdots+\lambda_{p}^{2}\right)\right\} \\
& =\frac{2 \operatorname{tr}\left(\Sigma^{2}\right)}{\{\operatorname{tr}(\Sigma)\}^{2}}\left(\psi^{2}-2 \beta \psi+\beta\right),
\end{aligned}
$$

where $\beta=\left(\lambda_{1}^{2}+\cdots+\lambda_{q}^{2}\right) /\left(\lambda_{1}^{2}+\cdots+\lambda_{p}^{2}\right)$.
Remark: use $\operatorname{tr}(\Lambda)=\operatorname{tr}(\Sigma)$ and $\operatorname{tr}\left(\Lambda^{2}\right)=\operatorname{tr}\left(\Sigma^{2}\right)$ to simplify the calculation!

## Application of PCA

The usual flow of PCA:
(1) Is it necessary to standardize the data set?
(2) How many PCs?
(3) Interpretation!

Usual applications:

- dimension reduction,
- visualization (plotting) of high-dimensional datasets,
- regression on PCs (removes multicollinearity).

Týden 7-8

Faktorová analýza:

- model faktorové analýzy,
- odhadování a rotace faktorů,
- interpretace.


## Week 7-8 Factor analysis

## Factor Analysis Model

We want to explain $p$ components of $X$ by smaller number of common factors.

$$
\begin{aligned}
& X=\mathcal{Q F + U + \mu} \\
& \mathcal{Q}=(p \times k) \text { loadings } \\
& F=(k \times 1) \text { common factors } \\
& U=(p \times 1) \text { specific factors }
\end{aligned}
$$

where $F$ and $U$ are centered, $\operatorname{Var}(F)=\mathcal{I}_{k}$,
$\operatorname{Var}(U)=\Psi=\operatorname{diag}\left(\psi_{11}, \ldots, \psi_{p p}\right)$, and $\operatorname{Cov}(F, U)=0$
Estimates of the loadings $\mathcal{Q}$ and specific variances $\Psi$ are deduced from $\operatorname{var} X$ (using var $X=\Sigma=\mathcal{Q} \mathcal{Q}^{\top}+\Psi$ ).

## Factor analysis

Factor analysis has provoked rather turbulent controversy throughout its history.
... each application of the technique must be examined on its own merits to determine its success.

The essential purpose of factor analysis is to describe, if possible, the covariance relationships among many variables in terms of a few underlying, but unobservable, random quantities called factors.

Factor analysis can be considered as an extension of principal component analysis ... the approximation based on the factor analysis model is more elaborate.
(Johnson and Wichern, Applied Multivariate Statistical Analysis, Prentice Hall, 1992)

## Week 7-8

Factor analysis
Example: Perfect FA is PCA with only $k$ positive eigenvalues:

$$
\begin{array}{ll}
\Sigma & =\sum_{\ell=1}^{k} \lambda_{\ell} \gamma_{\ell} \gamma_{\ell}^{\top} \\
X & =\mathcal{Q} F+\mu \\
\mathcal{Q} & =\left(\sqrt{\lambda_{1}} \gamma_{1}, \ldots, \sqrt{\lambda_{k}} \gamma_{k}\right) \\
F & =k-\operatorname{dim} \text { vector (random) } \\
E F & =0 \\
\operatorname{Var}(F) & =\mathcal{I}_{k}
\end{array}
$$

Clearly, the matrix $Q$ is not unique (because rotation leads to equivalent solution).

Communality and specific variance

Define

$$
\begin{array}{ll}
h_{j}^{2}=\sum_{\ell=1}^{k} q_{j \ell}^{2} & \text { communality } \\
\psi_{j j} & \text { specific variance }
\end{array}
$$

Notice that $\operatorname{var} X_{j}=h_{j}^{2}+\psi_{j j}$, i.e., the communality is the part of variance of $X_{j}$ explained by the common factors. The specific variance is the unexplained part.

Two important properties of FA model are invariance of scale and non-uniqueness (with respect to rotations).

## Week 7-8 <br> Factor analysis

Non-Uniqueness of Factor Loadings

For orthogonal matrix $\mathcal{G}$ we get:

$$
X=(\mathcal{Q G})\left(\mathcal{G}^{\top} F\right)+U+\mu
$$

We get a $k$-factor model with factor loadings $\mathcal{Q G}$ and common factors $\mathcal{G}^{\top} F$. In practical analysis, we will choose the rotation which gives "desirable" interpretation.

For the purpose of evaluation, the non-uniqueness can be solved by imposing additional constraints, e.g.,

$$
\mathcal{Q}^{\top} \Psi^{-1} \mathcal{Q} \text { is diagonal. }
$$

## Invariance of scale

Assume that we have the following FA model for $X: \operatorname{var} X=\mathcal{Q}_{X} \mathcal{Q}_{X}^{\top}+\Psi_{X}$.
What happens if we change the scale of $X$ ?

$$
\begin{aligned}
Y & =\mathcal{C} X, \mathcal{C}=\operatorname{diag}\left(c_{1}, \ldots, c_{p}\right) \\
\operatorname{Var}(Y) & =\mathcal{C} \Sigma \mathcal{C}^{\top} \\
& =\mathcal{C} \mathcal{Q}_{X} \mathcal{Q}_{X}^{\top} \mathcal{C}^{\top}+\mathcal{C} \Psi_{X} \mathcal{C}^{\top}
\end{aligned}
$$

Hence the $k$-factor model is also true for $Y$ with

$$
\begin{aligned}
\mathcal{Q}_{Y} & =\mathcal{C} \mathcal{Q}_{X} \\
\Psi_{Y} & =\mathcal{C} \Psi_{X} \mathcal{C}^{\top}
\end{aligned}
$$

nterpretation of unobserved latent factors $F$ is based on covariances and correlations:

$$
\begin{aligned}
& \Sigma_{X F}=E\left\{(\mathcal{Q F}+U) F^{\top}\right\}=\mathcal{Q} \\
& \mathcal{P}_{X F}=D^{-1 / 2} \mathcal{Q}
\end{aligned}
$$

where $D=\operatorname{diag}\left(\sigma_{X_{1} X_{1}}, \ldots \sigma_{X_{p} X_{p}}\right)$.
Correlations $\mathcal{P}_{\text {XF }}$ show the relationship between the original variables $X_{1}, \ldots, X_{p}$ and the common factors $F_{1}, \ldots, F_{k}$.

Number of parameters in the model

We have $p(p+1) / 2$ equations and $p k+p$ parameters ( $p k$ parameters from $\mathcal{Q}$ and $p$ parameters from $\Psi$ ) with $\frac{1}{2}\{(k(k-1)\}$ constraints (e.g. $\mathcal{Q}^{\top} \Psi^{-1} \mathcal{Q}$ is diagonal):

$$
\begin{aligned}
\Rightarrow d= & \# \text { pars for } \sum \text { unconstrained } \\
& -\# \text { pars for } \sum \text { constrained } \\
= & \frac{1}{2}(p-k)^{2}-\frac{1}{2}(p+k) .
\end{aligned}
$$

$d<0$ infinity of exact solutions
$d>0$ look for approximate solutions

Example: $p=3, k=1 \Rightarrow d=0$

$$
\Sigma=\left(\begin{array}{lll}
q_{1}^{2}+\psi_{11} & & \\
q_{1} q_{2} & q_{2}^{2}+\psi_{22} & \\
q_{1} q_{3} & q_{2} q_{3} & q_{3}^{2}+\psi_{33}
\end{array}\right)
$$

1$d=0$ yields only a unique numerical solution! It need not be consistent with statistical thinking

Example: Suppose now $p=2$ and $k=1$, then $d<0$.

$$
\Sigma=\left(\begin{array}{ll}
1 & \\
\rho & 1
\end{array}\right)=\left(\begin{array}{cc}
q_{1}^{2}+\psi_{1} & \\
q_{1} q_{2} & q_{2}^{2}+\psi_{2}
\end{array}\right)
$$

We have an infinity of solutions: for any $\alpha(\rho<\alpha<1)$ a solution is provided by:

$$
q_{1}=\alpha ; \quad q_{2}=\rho / \alpha ; \Psi_{1}=1-\alpha^{2} ; \Psi_{2}=1-(\rho / \alpha)^{2}
$$

* The factor analysis model aims to describe the dependencies between the $p$ variables in a data set by a lower number $k<p$ of latent factors, i.e. it assumes $X=\mathcal{Q} F+U+\mu$. The random vector $F$ ( $k$-dimensional) contains the common factors, $U$ ( $p$-dimensional) the specific factors, $\mathcal{Q}(p \times k)$ the loadings matrix.
* It is supposed that $F$ and $U$ are uncorrelated and have mean zero and uncorrelated components, i.e., $F \sim(0, \mathcal{I}), U \sim(0, \Psi)$ with a diagonal $\Psi, \operatorname{Cov}(F, U)=0$
This leads to the covariance structure $\Sigma=\mathcal{Q} \mathcal{Q}^{\top}+\Psi$.
* The interpretation of the factor $F$ is obtained through the correlation $\mathcal{P}_{\text {XF }}=D^{-1 / 2} \mathcal{Q}$.


## Weak 7-8

Factor analysis

Example: Data set carmean2 consists of the averaged marks (from 1 low to 7 high) for 31 car types.

We consider price, security and easy handling

$$
\mathcal{R}=\left(\begin{array}{ccc}
1 & 0.975 & 0.613 \\
& 1 & 0.620 \\
& & 1
\end{array}\right)
$$

We look for one factor, i.e. $k=1$. (\# number of parameters of $\Sigma$ unconstrained $-\#$ parameters of $\Sigma$ constrained) equals here $\frac{1}{2}(p-k)^{2}-\frac{1}{2}(p+k)=\frac{1}{2}(3-1)^{2}-\frac{1}{2}(3+1)=0$.

So there is an exact solution!

The equation

$$
\left(\begin{array}{ccc}
1 & r_{X_{1} X_{2}} & r_{X_{1} X_{2}} \\
& 1 & r_{X_{1} X_{3}} \\
& & 1
\end{array}\right)=\mathcal{R}=\left(\begin{array}{ccc}
\widehat{q}_{1}^{2}+\widehat{\psi}_{11} & \widehat{q}_{1} \widehat{q}_{2} & \widehat{q}_{1} \widehat{q}_{3} \\
& \widehat{q}_{2}^{2}+\widehat{\psi}_{22} & \widehat{q}_{2} \widehat{q}_{3} \\
& & \widehat{q}_{3}^{2}+\widehat{\psi}_{33}
\end{array}\right)
$$

yields the communalities $\widehat{h}_{i}^{2}=\widehat{q}_{i}^{2}$

$$
\widehat{q}_{1}^{2}=\frac{r x_{1} x_{2} r x_{1} x_{3}}{r_{X_{2} x_{3}}} \quad \hat{q}_{2}^{2}=\frac{r x_{1} x_{2} r x_{2} x_{3}}{r_{X_{1} x_{3}}} \quad \widehat{q}_{3}^{2}=\frac{r_{X_{1} x_{3}} r_{x_{2} x_{3}}}{r_{X_{1} x_{2}}} .
$$

Together with $\widehat{\psi}_{11}=1-\widehat{q}_{1}^{2}, \widehat{\psi}_{22}=1-\widehat{q}_{2}^{2}$ and $\widehat{\psi}_{33}=1-\widehat{q}_{3}^{2}$ we get the solution

$$
\begin{aligned}
\widehat{q}_{1} & =0.982 & \widehat{q}_{2} & =0.993 \\
\widehat{\psi}_{11} & =0.035 & \widehat{\psi}_{22} & =0.014
\end{aligned}
$$

Since the first two communalities are close to one, we conclude that the first two variables, namely price and security, are explained by the factor very well.

This factor might be interpreted as a "price+security" factor.

## Week 7-8 Factor analysis

The Principal Component Method

Decompose $\operatorname{var}(X)=\mathcal{S}=\mathcal{G} \mathcal{L G}^{\top}$.
Retain the first $k$ eigenvectors to build

$$
\hat{\mathcal{Q}}=\left[\sqrt{\ell_{1}} g_{1}, \ldots, \sqrt{\ell_{k}} g_{k}\right]
$$

Omitting $p-k$ eigenvectors shouldn't cause big error if the corresponding eigenvalues $\lambda_{i}, i=k+1, \ldots, p$ are small.

Specific variance are estimated by diagonal elements of

$$
\mathcal{S}-\hat{\mathcal{Q}} \hat{\mathcal{Q}}^{\top}
$$

This gives $\hat{\psi}$.

Method of Principal Factors

We start with an estimate of the communality:

1) $\widetilde{h}_{j}^{2}=$ the square of the multiple correlation coefficient, i.e. $\rho^{2}(V, W \widehat{\beta})$ with $V=X_{j}$
$W=\left(X_{\ell}\right)_{\ell \neq j}$
$\widehat{\beta}=$ OLS of regression of $V$ on $W$
2) $\widetilde{h}_{j}^{2}=\max _{\ell \neq j}\left|r_{\chi_{j} x_{\ell}}\right|$
$\mathcal{R}=$ correlation matrix

The Maximum Likelihood Method

Log-likelihood function $\ell$ for a data matrix $\mathcal{X}$ of observations for $X \sim N_{p}(\mu, \Sigma):$

$$
\begin{aligned}
\ell(\mathcal{X} ; \mu, \Sigma) & =-\frac{n}{2} \log |2 \pi \Sigma|-\frac{1}{2} \sum_{i=1}^{n}\left(x_{i}-\mu\right) \Sigma^{-1}\left(x_{i}-\mu\right)^{\top} \\
& =-\frac{n}{2} \log |2 \pi \Sigma|-\frac{n}{2} \operatorname{tr}\left(\Sigma^{-1} \mathcal{S}\right)-\frac{n}{2}(\bar{x}-\mu) \Sigma^{-1}(\bar{x}-\mu)^{\top}
\end{aligned}
$$

Evaluated at its maximum $\widehat{\mu}=\bar{x}$ :

$$
\ell(\mathcal{X} ; \widehat{\mu}, \Sigma)=-\frac{n}{2}\left\{\log (|2 \pi \Sigma|)-\operatorname{tr}\left(\Sigma^{-1} \mathcal{S}\right)\right\}
$$

## Algorithm of Principal Factors Method

$$
\begin{array}{ll}
\tilde{\psi}_{j j} & =1-\tilde{h}_{j}^{2} \\
\text { Construct } & \mathcal{R}-\tilde{\Psi} \\
\mathcal{R}-\widetilde{\Psi} & =\sum_{\ell=1}^{p} \lambda_{\ell} \gamma_{\ell} \gamma_{\ell}^{\top} \\
\widehat{q_{\ell}} & =\sqrt{\lambda_{\ell}} \gamma_{\ell}, \quad \ell=1, \ldots k \\
\widehat{\mathcal{Q}} & =\Gamma_{1} \Lambda_{1}^{1 / 2} \\
\Gamma_{1} & =\left(\gamma_{1}, \ldots, \gamma_{k}\right) \\
\Lambda_{1} & =\operatorname{diag}\left(\lambda_{1}, \ldots, \lambda_{k}\right) \\
\widehat{\psi_{j j}} & =1-\sum_{\ell=1}^{k} \widehat{q_{j \ell}^{2}}
\end{array}
$$

Z. Hlávka (KPMS)

## Week 7-8

## Factor analysis

By substituting $\Sigma=\mathcal{Q} \mathcal{Q}^{\top}+\Psi$

$$
\ell(\mathcal{X} ; \widehat{\mu}, \mathcal{Q}, \Psi)=-\frac{n}{2}\left[\log \left\{\left|2 \pi\left(\mathcal{Q} \mathcal{Q}^{\top}+\Psi\right)\right|\right\}-\operatorname{tr}\left\{\left(\mathcal{Q} \mathcal{Q}^{\top}+\Psi\right)^{-1} \mathcal{S}\right\}\right]
$$

This model is not well defined.
Therefore, we require that $\mathcal{Q}^{\top} \Psi^{-1} \mathcal{Q}$ is diagonal matrix.
The maximum likelihood estimates of $\mathcal{Q}$ and $\Psi$ are obtained using an iterative numerical algorithm (function factanal() in R library MASS).

## LR test for the Number of Common Factors

The test follows directly from the assumption of normality. We test
$H_{0}: \Sigma=\mathcal{Q} \mathcal{Q}^{\top}+\Psi$
$H_{1}: \Sigma$ arbitrary (positive definite) matrix
See the chapter on Likelihood Ratio tests
The likelihood ratio statistic is

$$
\begin{gathered}
-2 \Lambda=-2 \log \left[\frac{\text { maximized likelihood under } H_{0}}{\text { maximized likelihood }}\right] \\
=-2 \log \left(\frac{\left|\hat{\mathcal{Q}} \hat{\mathcal{Q}}^{\top}+\hat{\Psi}\right|}{\left|S_{n}\right|}\right)^{-n / 2}+n\left\{\operatorname{tr}\left[\left(\hat{\mathcal{Q}} \hat{\mathcal{Q}}^{\top}+\hat{\Psi}\right)^{-1} S_{n}\right]-p\right\}
\end{gathered}
$$

with $(1 / 2)\left[(p-m)^{2}-p-m\right]$ degrees of freedom.

## Week 7-8

Factor analysis

## Varimax

The varimax method tries to find "reasonable rotation" automatically.
The interpretation of the loadings would be simple if the variables split into disjoint sets, each of which is associated with one factor. A well known analytical algorithm which tries to rotate the loadings in this way is the varimax rotation method.

Varimax method tries to find the rotation which maximizes the sum of the variances of the squared loadings $\hat{q}_{i j}^{*}$ within each column of $\hat{\mathcal{Q}}^{*}$ (this should lead to $q_{i j} s$ close to 0 or 1 ):

$$
\max _{\text {rotations }} Q^{*} \sum_{j=1}^{k}\left\{\frac{1}{p} \sum_{i}\left(q_{i j}^{*}\right)^{4}-\left[\frac{1}{p} \sum_{i}\left(q_{i j}^{*}\right)^{2}\right]^{2}\right\}
$$

## Rotation

The factor analysis model is not uniquely defined and the factors can be rotated without any loss of information.

We are free to rotate the estimated factors rather arbitrary. This feature of factor analysis is rather controversial.

Usually, we rotate the factors in a way which provides reasonable interpretation which is consistent with the measured variables.

In the most simple case of $k=2$ factors a rotation matrix $\mathcal{G}$ is given by

$$
\mathcal{G}(\theta)=\left(\begin{array}{rr}
\cos \theta & \sin \theta \\
-\sin \theta & \cos \theta
\end{array}\right)
$$

which represents a clockwise rotation of the coordinate axes by the angle $\theta$ (then $\hat{\mathcal{Q}}^{*}=\hat{\mathcal{Q}} \mathcal{G}(\theta)$ ).

## Promax

The promax rotation is similar to varimax but it works without the condition of orthogonality (so-called oblique rotation).

The resulting correlated(!) factors are typically easier to interpret.

Strategy for Factor Analysis
(1) Perform a principal component factor analysis, look for suspicious observations, try varimax rotation
(2) Perform maximum likelihood factor analysis including varimax rotation
(3) Compare the factor analyses: do the loadings group in the same manner?
(1) Repeat the previous steps for other number of common factors
(5) For large data sets, split them in half and perform a factor analysis on each part. Compare the solutions.

## Week 7-8 <br> Factor analysis

Assuming joint normality, the conditional distribution of $F \mid X$ is multinormal with $E(F \mid X=x)=\mathcal{Q}^{\top} \Sigma^{-1}(X-\mu)$ and the covariance matrix $\operatorname{var}(F \mid X=x)=\mathcal{I}_{k}-\mathcal{Q}^{\top} \Sigma^{-1} \mathcal{Q}$.

In practice, we replace the unknown $\mathcal{Q}, \Sigma$ and $\mu$ by corresponding estimators, leading to the estimated individual factor scores:

$$
\widehat{f}_{i}=\widehat{\mathcal{Q}}^{\top} \mathcal{S}^{-1}\left(x_{i}-\bar{x}\right) .
$$

We prefer to use the original sample covariance matrix $\mathcal{S}$ as an estimator of $\Sigma$, instead of the factor analysis approximation $\widehat{\mathcal{Q}} \widehat{\mathcal{Q}}^{\top}+\widehat{\Psi}$, in order to be more robust against incorrect determination of the number of factors.

Note that the upper left entry of this matrix equals $\Sigma$ and that the matrix has size $(p+k) \times(p+k)$.

## Notes

(1) The same rule can be followed when using $\mathcal{R}$ instead of $\mathcal{S}$. In this case the factors are given by

$$
\widehat{f}_{i}=\widehat{\mathcal{Q}}^{\top} \mathcal{R}^{-1}\left(z_{i}\right)
$$

where $z_{i}=\mathcal{D}_{S}^{-1 / 2}\left(x_{i}-\bar{x}\right), \widehat{\mathcal{Q}}$ is the loading obtained with the matrix $\mathcal{R}$, and $\mathcal{D}_{S}=\operatorname{diag}\left(s_{11}, \ldots, s_{p p}\right)$.
(2) Using MLE (treating $F$ as unknown parameters), one arrives to Bartlett's scores.
(3) Clearly, if the factors are rotated by the orthogonal matrix $\mathcal{G}$, the factor scores have to be rotated accordingly, that is

$$
\widehat{f}_{i}^{*}=\mathcal{G}^{\top} \widehat{f}_{i} .
$$

## Týden 9

Mnohorozměrné škálování:

- matice vzdáleností,
- metrické řešení,
- nemetrické řešení (PAVA a STRESS)

In practice, this technique is also called Exploratory Factor Analysis.
After exploring the factors, one can perform the so-called Confirmatory Factor Analysis allowing more detailed investigation of the underlying factors (one can imagine that oblique factors could be explained by another factor analysis leading to a hieararchical model).

Relationship between the unobserved factors can be investigated using Structural Equation Models (R library sem, M-plus, LISREL).

Factor analysis models are popular mainly in psychology and behavioral science.

## Euclidean distance



$$
d(x, y)=\sqrt{(x-y)^{T}(x-y)}
$$

## Distance matrix

$\mathcal{X}(n \times p)$ with $n$ measurements (objects) of $p$ variables.
The distance matrix $\mathcal{D}(n \times n)$ is a matrix of all distances between all pairs of observations:

$$
\mathcal{D}=\left(\begin{array}{cccccc}
d_{11} & d_{12} & \ldots & \ldots & \ldots & d_{1 n} \\
\vdots & d_{22} & & & & \vdots \\
\vdots & \vdots & \ddots & & & \vdots \\
\vdots & \vdots & & \ddots & & \vdots \\
\vdots & \vdots & & & \ddots & \vdots \\
d_{n 1} & d_{n 2} & \ldots & \ldots & \ldots & d_{n n}
\end{array}\right)
$$

Example: $L_{2}$-norm: $d_{i j}=\left\|x_{i}-x_{j}\right\|_{2}$, where $x_{i}$ and $x_{j}$ denote the rows of the data matrix $\mathcal{X}$.
Z. Hlávka (KPMS)

NMST539


## Week 9

 Distance matrixExample: Let us consider data set on songs in 20 medieval songbooks.
spev P U K E G Kl SMF Pa S M W Lc B F SG To C Me Q R

 ELEVANS AUTEM x x x x x x ROGO ERGO $\mathrm{x} x \mathrm{x} \mathrm{x} \mathrm{x} \mathrm{x}$ $\mathrm{x} \mathrm{x} \mathrm{x}-\mathrm{x} \mathrm{x}-\mathrm{x}_{\mathrm{x}} \mathrm{x}-\mathrm{x}-\mathrm{x} \mathrm{x} \mathrm{x}$ DIVES ILLE x x x x x x x _ _ _ _ _ _ x x x x _ _ DEUS CARITAS $\mathrm{x} x \mathrm{x}_{\mathrm{z}}$ _ x _ _ _ _ _ _ _ _ x _ x .







Distance and similarity

Distance can be easily calculated for numerical measurements (Euclidean distance, $L_{1}$ distance, Mahalanobis distance, etc.)

## Example:

data(bank2)
dist(bank2,method="euclidean")
dist(scale(bank2),method="euclidean")

For nominal (or binary) variables it is easier to define a measure of similarity (e.g., various ratios of "number of concordances" such as Jaccard, Tanimoto, Simple Matching, etc.)

Considering observations $x_{i}$ and $x_{j}$ and denoting

$$
\begin{aligned}
& a_{1}=\sum_{k=1}^{p} I\left(x_{i k}=x_{j k}=" x^{\prime \prime}\right), \\
& a_{2}=\sum_{k=1}^{p} I\left(x_{i k}="-", x_{j k}=" x^{\prime \prime}\right), \\
& a_{3}=\sum_{k=1}^{p} I\left(x_{i k}=" x^{\prime \prime}, x_{j k}="-\prime\right), \\
& a_{4}=\sum_{k=1}^{p} I\left(x_{i k}=x_{j k}="{ }_{-}^{\prime \prime}\right) .
\end{aligned}
$$

we can define a proximity measure as

$$
s_{i j}=\frac{a_{1}+\delta a_{4}}{a_{1}+\delta a_{4}+\lambda\left(a_{2}+a_{3}\right)}
$$

by choosing some $\delta$ and $\lambda$.

| Name | $\delta$ | $\lambda$ | Definition |
| :--- | :---: | :---: | :---: |
| Jaccard | 0 | 1 | $\frac{a_{1}}{a_{1}+a_{2}+a_{3}}$ |
| Tanimoto | 1 | 2 | $\frac{a_{1}+a_{4}}{a_{1}+2\left(a_{2}+a_{3}\right)+a_{4}}$ |
| Simple Matching (M) | 1 | 1 | $\frac{a_{1}+a_{4}}{p}$ |
| Dice | 0 | 0.5 | $\frac{2 a_{1}}{2 a_{1}+\left(a_{2}+a_{3}\right)}$ |

In the songbooks example, the Jaccard measure seems to be reasonable. Calculating the similarity measure for all pairs of songbooks, we obtain a similarity matrix $S=s_{i j}$.

## Distance and similarity

Distances and similarities are closely related. It may be useful to "transform" similarity to distance because some methods require distances

Denoting similarities as $s_{i j}$, distances can be defined as
$d_{i j}=\sqrt{s_{i j}-2 s_{i j}+s_{j j}}$ or $d_{i j}=\max _{i, j}\left\{s_{i j}\right\}-s_{i j}$ (we want symmetry and $\left.d_{i i}=0\right)$.

In the songbook example, we can define

$$
d_{i j}=\frac{a_{2}+a_{3}}{a_{1}+a_{2}+a_{3}}=1-s_{i j}
$$

as a ratio of "common songs" (from songs that are contained in songbooks $i$ and $j$ ).

## Euclidean matrix

It is easy to calculate distance matrix $\mathcal{D}$ from the data matrix $\mathcal{X}$ but is it possible to calculate the data matrix $\mathcal{X}$ from a distance matrix $\mathcal{D}$ ?

Definition: We say that $\mathcal{D}=\left(d_{i j}\right)$ is a distance matrix if $d_{i j}=d_{j i} \geq 0$ and $d_{i i}=0$, for $i, j=1, \ldots, n$.

The first step would be to verify that the matrix $\mathcal{D}$ is Euclidean (i.e., that it contains Euclidean distances).

Definition: We say that a matrix $\mathcal{D}=\left(d_{i j}\right)$ is Euclidean if for some points $x_{1}, \ldots, x_{n} \in \mathbb{R}^{p} ; d_{i j}^{2}=\left(x_{i}-x_{j}\right)^{\top}\left(x_{i}-x_{j}\right)$.

Theorem: Define $\mathcal{A}=\left(a_{i j}\right), a_{i j}=-\frac{1}{2} d_{i j}^{2}, \mathcal{B}=\mathcal{H} \mathcal{A}, \mathcal{H}$ being the centering matrix. Then the matrix $\mathcal{D}$ is Euclidean if and only if $\mathcal{B}$ is positive semidefinite.

## Idea of the proof:

1/ Assuming that $\mathcal{D}$ is Euclidean for centered data matrix $\mathcal{X}$, we have $d_{i j}^{2}=x_{i}^{\top} x_{i}+x_{j}^{\top} x_{j}-2 x_{i}^{\top} x_{j}$.

Writing $\mathcal{B}=\mathcal{H} \mathcal{A H}$ implies that $b_{i j}=a_{i j}-\bar{a}_{i .}-\bar{a}_{. j}+\bar{a}_{. .}=\cdots=x_{i}^{\top} x_{j}$. Therefore $\mathcal{B}=\mathcal{X X}^{\top} \geq 0$,

2/ Assuming that $\mathcal{B} \geq 0$ and $\operatorname{rank}(B)=p$, we can write $\mathcal{B}=\Gamma_{p} \Lambda_{p} \Gamma_{p}^{\top}$ and it follows (similarly as above) that $\mathcal{D}$ is matrix of Euclidean distances of points in $\mathcal{X}=\Gamma_{\rho} \Lambda_{\rho}^{1 / 2}$.

Note that the matrix $\mathcal{X}=\Gamma_{p} \Lambda_{p}^{1 / 2}$ is centered.

## Week 9 Multidimensional scaling <br> MDS solution (metric MDS)

Recall that $\mathcal{B}=\mathcal{H} \mathcal{A H}$, where $a_{i j}=-\frac{1}{2} d_{i j}^{2}$.
Assuming that $\operatorname{rank}(B)=p$ and writing $\mathcal{B}=\Gamma_{p} \Lambda_{p} \Gamma_{p}^{\top}$, we obtain data matrix $\mathcal{X}=\Gamma_{p} \Lambda_{p}^{1 / 2}$ that preserves the observed distances in $p$-dimensional space.y

If some of the eigenvalues are small, we can obtain a good representation (of the distances) in $k$-dimensional space by $\mathcal{X}=\Gamma_{k} \Lambda_{k}^{1 / 2}$.

Note that the final configuration of points in $\mathbb{R}^{k}$ can be arbitrarily rotated and shifted without changing the distances.

## Multidimensional scaling

MDS uses proximities (distances) between objects to produce a spatial representation of these items.

In contrast to the techniques considered so far MDS does not start from the raw multivariate data matrix, $\mathcal{X}$, but from a $(n \times n)$ dissimilarity or a distance matrix $\mathcal{D}$. Hence, the underlying dimensionality of the data under investigation is not known.

More precisely: MDS searches for a "configuration" of points in $\mathbb{R}^{k}$ that "preserves" the distances of objects in $\mathbb{R}^{p}$ (where $p$ is not known).

MDS-techniques can help to understand how people perceive and evaluate certain items:
metric MDS is based on Euclidean distances,
non-metric MDS assumes that distances are on ordinal scale.

## Week 9

Multidimensional scaling
Example: Consumers' impressions of the dissimilarity of certain cars.

|  | Audi 100 | BMW 5 | Citroen AX | Ferrari | $\ldots$ |
| :--- | :---: | :---: | :---: | :---: | :---: |
| Audi 100 | 0 | 2.232 | 3.451 | 3.689 | $\ldots$ |
| BMW 5 | 2.232 | 0 | 5.513 | 3.167 | $\ldots$ |
| Citroen AX | 3.451 | 5.513 | 0 | 6.202 | $\ldots$ |
| Ferrari | 3.689 | 3.167 | 6.202 | 0 | $\ldots$ |
| $\vdots$ | $\vdots$ | $\vdots$ | $\vdots$ | $\vdots$ | $\ddots$ |

[^1]

MDS solution.
Z. Hlávka (KPMS)

Week 9 Multidimensional scaling


Correlations between the MDS direction and the original variables.

* MDS uses distances between $n$ items to project high-dimensional data in a low-dimensional space.
* MDS (using Euclidean distances in $p$ dimensions) leads to first $k$ principal components of the original data set.
* It can be shown that the metric solution to MDS leads to optimal representation of the original data set in $k$ dimensional space (from the point of view of $\sum\left(d_{i j}^{2}-\hat{d}_{i j}^{2}\right)$, where $d_{i j}$ are the original distances and $\hat{d}_{i j}$ are the projections in $\mathbb{R}^{k}$ —note also that $\left.\hat{d}_{i j} \leq d_{i j}\right)$.

Example: Consider a small example with 4 objects based on the car marks data set

|  | j | 1 | 2 | 3 | 4 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| i |  | Mercedes | Jaguar | Ferrari | VW |
| 1 | Mercedes | - |  |  |  |
| 2 | Jaguar | 3 | - |  |  |
| 3 | Ferrari | 2 | 1 | - |  |
| 4 | VW | 5 | 4 | 6 | - |

We demonstrate each step of the algorithm using a simple example.

Our aim is to find a $p^{*}=2$ dimensional representation via MDS. Suppose that we choose as initial configuration $\mathcal{X}_{0}$ the coordinates as:


| i |  | $x_{i 1}$ | $x_{i 2}$ |
| :---: | :---: | :---: | :---: |
| 1 | Mercedes | 3 | 2 |
| 2 | Jaguar | 2 | 7 |
| 3 | Ferrari | 1 | 3 |
| 4 | VW | 10 | 4 |



A plot of the dissimilarities is not satisfactory since the ranking of the $\delta_{i j}$ did not result in a monotone relation of the corresponding distances $d_{i j}$. We apply therefore the PAV algorithm.

PAVA = "pool adjacent violators" algorithm (= algoritmus "zprůměrování sousedních narusitelů") is used to calculate the LS estimator under assumption of monotonicity.

The first violator of monotonicity is the second point $(1,3)$ we therefore average the distances $d_{13}$ and $d_{23}$ to obtain the disparities

$$
\hat{d}_{13}=\hat{d}_{23}=\frac{d_{13}+d_{23}}{2}=\frac{2.2+4.1}{2}=3.17 .
$$

We apply the same procedure to the pair $(2,4)$ and $(1,4)$ to yield $\hat{d}_{24}=\hat{d}_{14}=7.9$. The plot od $\delta_{i j}$ versus the disparities $\hat{d}_{i j}$ represents a monotone regression relationship.

In the initial configuration, the point 3 (Ferrari) could be moved so that the distance to object 2 (Jaguar) is smaller. This procedure however also alters the distance between objects 3 and 4 . More care has therefore to be taken for an establishment of a monotone relation between $\delta_{i j}$ and $d_{i j}$.

## STRESS

In order to assess how well the derived configuration fits the given dissimilarities Kruskal suggests a measure called STRESS1 that is given by

$$
\text { STRESS1 }=\left(\frac{\sum_{i<j}\left(d_{i j}-\hat{d}_{i j}\right)^{2}}{\sum_{i<j} d_{i j}^{2}}\right)^{\frac{1}{2}} .
$$

An alternative measure of STRESS1 is given by

$$
\text { STRESS2 }=\left(\frac{\sum_{i<j}\left(d_{i j}-\hat{d}_{i j}\right)^{2}}{\sum_{i<j}\left(d_{i j}-\bar{d}\right)^{2}}\right)^{\frac{1}{2}},
$$

where $\bar{d}$ denotes the average distance

The aim is a point configuration that balances the effects STRESS and non monotonicity. This is achieved by an iterative procedure defining new position of object $i$ relative to object $j$ by

$$
x_{i l}^{N E W}=x_{i l}+\alpha\left(1-\frac{\hat{d}_{i j}}{d_{i j}}\right)\left(x_{j I}-x_{i l}\right), \quad I=1, \ldots, p^{*}
$$

Here $\alpha$ denotes the step width of the iteration.
The configuration of object $i$ is improved relative to object $j$. In order to obtain an overall improvement relative to all remaining points one uses:

$$
x_{i l}^{N E W}=x_{i l}+\frac{\alpha}{n-1} \sum_{j=1, j \neq i}^{n}\left(1-\frac{\hat{d}_{i j}}{d_{i j}}\right)\left(x_{j l}-x_{i l}\right), \quad l=1, \ldots, p^{*}
$$

The choice of step width $\alpha$ is crucial. Kruskal proposes a starting value of $\alpha=0.2$. The iteration is continued by a numerical approximation procedure.


First iteration for Ferrari

Example: Dissimilarity matrix for car marks data set:

|  | j | 1 | 2 | 3 | 4 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| i |  | Nissan | Kia | BMW | Audi |
| 1 | Nissan | - |  |  |  |
| 2 | Kia | 2 | - |  |  |
| 3 | BMW | 5 | 6 | - |  |
| 4 | Audi | 3 | 4 | 1 | - |

The dissimilarity matrix contains obviously only ranks of dissimilarity. Applying metric MDS may not be appropriate in this situation.


The outcome of the Shepard-Kruskal algorithm. It is important that both axes have the same scale, different scales could lead to wrong
interpretations.

Example: The nonmetric MDS solution for the songbooks example $(\rightarrow$ SMSclussong)


The Euclidean distances between the points are

|  | j | 1 | 2 | 3 | 4 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| i |  | Nissan | Kia | BMW | Audi |
| 1 | Nissan | - |  |  |  |
| 2 | Kia | 2.00 | - |  |  |
| 3 | BMW | 5.02 | 6.02 | - |  |
| 4 | Audi | 3.20 | 4.16 | 1.88 | - |

These distances are different from the original dissimilarities but their order is the same, i.e., the STRESS measure is equal to 0 .

## Nonmetric Multidimensional Scaling

* Nonmetric MDS is based only upon the rank order of dissimilarities.
* The object of nonmetric MDS is to create a spatial representation of the objects with low dimensionality.
* A practical algorithm is given as:
(1) Choose an initial configuration
(2) Normalize the configuration.
(3) Find $d_{i j}$ from the normalized configuration
(9) Fit $\hat{d}_{i j}$, the disparities by the PAV algorithm
(5) Find the new configuration $\mathcal{X}_{n+1}$ by using steepest descent.
(c) Go to 2 and interate until STRESS is small enough.


## Týden 10

Shluková analýza:

- shlukovací algoritmy,
- hierarchické aglomerativní algoritmy,
- dendrogram.


## Group-building algorithms

Two types of clustering methods:

- partioning algorithms (typically computationally intensive optimization of a given criterion),
- hierarchical algorithms:
- agglomerative,
- partioning

In partitioning techniques the assignment of objects into groups may change during the (iterative) algorithm.

In hierarchical clustering this assignment cannot be changed (the algorithm produces a sequence of clusters by "splitting" or "joining").

In the following, we look at agglomerative techniques.

Agglomerative techniques are computationaly simple because the distances between clusters can be easily calculated from the distance matrix $\mathcal{D}$.

If two objects or groups $P$ and $Q$ are to be united one obtains the distance to another group (object) $R$ by the following distance function
$d(R, P+Q)=\delta_{1} d(R, P)+\delta_{2} d(R, Q)+\delta_{3} d(P, Q)+\delta_{4}|d(R, P)-d(R, Q)|$
$\delta_{j}$ weighting factors

## Week 10 <br> Cluster analysis

Example: $x_{1}=(0,0), x_{2}=(1,0), x_{3}=(5,5)$ and the squared Euclidean distance matrix with single linkage weighting.

The algorithm starts with $N=3$ clusters $P=\left\{x_{1}\right\}, Q=\left\{x_{2}\right\}, R=\left\{x_{3}\right\}$.
The single linkage distance between the remaining two clusters:

$$
\begin{aligned}
d(R, P+Q) & =\frac{1}{2} d(R, P)+\frac{1}{2} d(R, Q)-\frac{1}{2}|d(R, P)-d(R, Q)| \\
& =\min (d(R, P), d(R, Q)) \\
& =\min (d(R, P), d(R, Q)) \\
& =41
\end{aligned}
$$

The reduced distance matrix is then $\left(\begin{array}{cc}0 & 41 \\ 41 & 0\end{array}\right)$.
Single linkage $=$ nearest neighbor!

|  | $\delta_{1}$ | $\delta_{2}$ | $\delta_{3}$ | $\delta_{4}$ |
| :--- | :---: | :---: | :---: | :---: |
| Single linkage | $1 / 2$ | $1 / 2$ | 0 | $-1 / 2$ |
| Complete linkage | $1 / 2$ | $1 / 2$ | 0 | $1 / 2$ |
| Average linkage | $1 / 2$ | $1 / 2$ | 0 | 0 |
| (unweighted) | $\frac{n_{P}}{n_{P}+n_{Q}}$ | $\frac{n_{Q}}{n_{P}+n_{Q}}$ | 0 | 0 |
| Average linkage |  |  |  |  |
| (weighted) | $\frac{n_{P}}{n_{P}+n_{Q}}$ | $\frac{n_{Q}}{n_{P}+n_{Q}}$ | $-\frac{n_{P} n_{Q}}{\left(n_{P}+n_{Q}\right)^{2}}$ | 0 |
| Centroid | $1 / 2$ | $1 / 2$ | $-1 / 4$ | 0 |
| Median | $\frac{n_{R}+n_{P}}{n_{.}}$ | $\frac{n_{R}+n_{Q}}{n_{.}}$ | $-\frac{n_{R}}{n_{.}}$ | 0 |
| Ward |  |  |  |  |

$n_{P}=\sum_{i=1}^{n} I\left(x_{i} \in P\right)$ denotes the number of objects in group $P$ $n_{\text {. }}=n_{R}+n_{P}+n_{Q}$

## Dendrogram

## Dendrogram:

- a graphical representation of the sequence of clustering,
- displays the observations, the sequence of clusters and the distances between the clusters.


## Construction of dendrogram

- tree displaying the progress of the agglomerative clustering algorithm
- the row name (or row number) is given on the horizontal axis.
- the vertical axis gives the distance between clusters.

Example：


The 8 points example：
eight $=$ cbind（c $(4,2,-2,-3,-2,-2,1,1), c(-3,-4,-1,0,-2,4,2,4))$


## Week 10

Cluster analysis

## Cutting the tree

If we decide to cut the tree at the level 10 we define three clusters：$\{1,2\}$ ， $\{3,4,5\}$ and $\{6,7,8\}$ ．
gr＝cutree（hclust（dist（eight）～2，method＝＂single＂），k＝3）
In practice，it is important to interpret the resulting clusters using tables of means and（multivariate）graphics：

```
sapply(data.frame(eight),tapply,gr,
    function(x)sprintf("%0.1f",mean(x)))
```

plot(eight,col=as.numeric(gr),pch=as.numeric(gr)+1)

In practice，the choice of the number of clusters is usually based on the visual inspection of the dendrogram

## Ward algorithm

The measure of heterogeneity for a group $R$ is the inertia inside the group:

$$
I_{R}=\frac{1}{n_{R}} \sum_{i=1}^{n_{R}} d^{2}\left(x_{i}, \bar{x}_{R}\right),
$$

where $\bar{x}_{R}$ is the mean (center of gravity) of the group $R$.
When two objects or groups $P$ and $Q$ will be joined, the new group $P+Q$ will have a larger inertia $I_{P+Q}$. The increase of inertia is given by

$$
\Delta(P, Q)=\frac{n_{P} n_{Q}}{n_{P}+n_{Q}} d^{2}(P, Q) .
$$

The Ward algorithm joins groups $P$ and $Q$ that give the smallest increase of $\Delta(P, Q)$.

This interpretation is correct if we are working with squared Euclidean distances (note that hclust () contains two versions of Ward algorithm) , ac

## Week 10 <br> Cluster analysis

Cluster Dendrogram


Example: US companies data set.
data (uscomp)
uscomp\$Sales=as.numeric(as.character (uscomp\$Sales))
uscomp\$Sales[65]= 1601
d=dist (scale(uscomp[,c(-7)]))
plot(cluscomp.c<-hclust(d)
plot(cluscomp.s<-hclust(d,method="single"))
plot(cluscomp.w<-hclust(d,method="ward"))
gr3=cutree(cluscomp.w,k=3)
sapply (uscomp [, -7], tapply,gr3,function(x)round (mean(x)))
parcoord(uscomp [, -7], col=as.numeric (gr3))
table (gr3, uscomp[,7])

It seems that better results could be obtained by logarithmic
transformation of the data set

## Cluster analysis

* The class of clustering algorithms can be divided into two types: hierarchical and partitioning algorithms. Partitioning algorithms start from a preliminary clustering and optimize given criterion by exchanging group elements.
* Hierarchical agglomerative techniques start from the finest possible structure, compute the distance matrix, and join clusters with the smallest distance. This step is repeated until all points are united in one cluster.
* The agglomerative procedure depends on the definition of the distance between two clusters. Often used distances are single linkage, complete linkage, Ward distance.
* The process of the unification of clusters can be graphically represented by a dendrogram.

The dendrogram for the songbooks example (Ward algorithm based on Jaccard measure): two cluster solution corresponds to the division of Francia into West Francia (more-or-less current France) and East Francia (more-or-less current Germany) in the 9th century (after the death of Charlemagne) $\rightarrow$ SMSclussong ac

## Týden 10-11

## Diskriminační analýza:

- motivace a maximální věrohodnost,
- lineární a kvadratická diskriminační analýza,
- pravděpodobnost chybné klasifikace,
- Fisherův přístup.


## Bayes rule

Suppose that observations from $\Pi_{j}$ have density $f_{j}(x)$ and that the $\pi_{j}$ is the prior probability of $\Pi_{j}$.

Using Bayes theorem:

$$
P\left(\Pi_{j} \mid X=x\right)=\frac{f_{j}(x) \pi_{j}}{\sum_{i=1}^{J} f_{i}(x) \pi_{i}}
$$

Interpreting $P\left(\Pi_{j} \mid X=x\right)$ as the posterior probability of population $\Pi_{j}$ (after observing $X=x$ ), we classify $X$ to $\Pi_{\operatorname{argmax}_{j}} P\left(\Pi_{j} \mid X\right)$.
The corresponding discriminant rule $R_{j}$ is defined as $\left\{x: f_{j}(x) \pi_{j} \geq f_{j} i(x) \pi_{i}, i \neq j\right\}$ (maximum likelihood).

## Discriminant analysis

The aim of discriminant analysis is to construct discriminant rules allowing classification of new items (subjects) into known populations $\Pi_{j}$, $j=1, \ldots, J$.

Discriminant rule is a partition of the sample space:

$$
\begin{aligned}
& \qquad \bigcup_{j=1}^{J} R_{j}=\mathbb{R}^{p} \\
& \text { partition }
\end{aligned}
$$

The new observation is classified into population $\Pi_{j}$ if it falls in $R_{j}$.

## Week 10-11

Discriminant analysis
Example: A discrimination rule based on observations of a one-dimensional variable with an exponential distribution.

The pdf is $f(x)=\lambda \exp \{-\lambda x\}$ for $x>0$. Comparing the likelihoods for two populations $\Pi_{1}: \operatorname{Exp}\left(\lambda_{1}\right)$ and $\Pi_{2}: \operatorname{Exp}\left(\lambda_{2}\right)$, we allocate the observation $x$ into population $\Pi_{1}$ if

$$
\begin{aligned}
L_{1}(x) / L_{2}(x) & \geq 1 \\
x\left(\lambda_{1}-\lambda_{2}\right) & \leq \log \frac{\lambda_{1}}{\lambda_{2}}
\end{aligned}
$$

Assuming that $\lambda_{1}<\lambda_{2}$, we obtain:

$$
R_{1}=\left\{x: x \geq \frac{\log \lambda_{1}-\log \lambda_{2}}{\lambda_{1}-\lambda_{2}}\right\}
$$

The observation $x$ is classified into $\Pi_{1}$ if it is greater than the constant $\left(\log \lambda_{1}-\log \lambda_{2}\right) /\left(\lambda_{1}-\lambda_{2}\right)$.

## Credit scoring

Example: Let $\gamma$ denote the gain of the bank from a correctly classified good client. Let $\Pi_{2}$ denote the population of good clients.
$\Pi_{1}$ represents the population of bad clients that bring the loss $C(2 \mid 1)$ if they are classified as good clients.
$C(1 \mid 2)$ denotes the cost of loosing a good client classified as bad.
The gain of the bank as a function of the discriminant rule "client is good if he falls in region $R^{\prime \prime}$ is:

$$
\begin{aligned}
G(R)= & \gamma \pi_{2} \int I(x \in R) f_{2}(x) d x-C(2 \mid 1) \pi_{1} \int I(x \in R) f_{1}(x) d x \\
& -C(1 \mid 2) \pi_{2} \int\{1-I(x \in R)\} f_{2}(x) d x
\end{aligned}
$$

## Week 10-11 <br> iscriminant analysis

One-dimensional normal distributions

Consider two normal populations $\Pi_{1}: N\left(\mu_{1}, \sigma_{1}^{2}\right)$ and $\Pi_{2}: N\left(\mu_{2}, \sigma_{2}^{2}\right)$ and $\pi_{1}=\pi_{2}$

Then

$$
L_{i}(x)=\left(2 \pi \sigma_{i}^{2}\right)^{-1 / 2} \exp \left\{-\frac{1}{2}\left(\frac{x-\mu_{i}}{\sigma_{i}}\right)^{2}\right\}
$$

and $L_{1}(x)>L_{2}(x)$ (i.e., $x \in R_{1}$ is classified to $\left.\Pi_{1}\right)$

$$
\begin{aligned}
& \Longleftrightarrow \quad \frac{\sigma_{2}}{\sigma_{1}} \exp \left\{-\frac{1}{2}\left[\left(\frac{x-\mu_{1}}{\sigma_{1}}\right)^{2}-\left(\frac{x-\mu_{2}}{\sigma_{2}}\right)^{2}\right]\right\}>1 \\
& \Longleftrightarrow x^{2}\left(\frac{1}{\sigma_{1}^{2}}-\frac{1}{\sigma_{2}^{2}}\right)-2 x\left(\frac{\mu_{1}}{\sigma_{1}^{2}}-\frac{\mu_{2}}{\sigma_{2}^{2}}\right)+\left(\frac{\mu_{1}^{2}}{\sigma_{1}^{2}}-\frac{\mu_{2}^{2}}{\sigma_{2}^{2}}\right)<2 \log \frac{\sigma_{2}}{\sigma_{1}} .
\end{aligned}
$$

This is quadratic inequality $\Longleftrightarrow \sigma_{1}^{2} \neq \sigma_{2}^{2}$.
If $\sigma_{1}=\sigma_{2}$ then (for $\mu_{1}<\mu_{2}$ ) we obtain a very simple linear discriminant rule:

$$
\begin{aligned}
R_{1} & =\left\{x: x \leq \frac{1}{2}\left(\mu_{1}+\mu_{2}\right)\right\} \\
R_{2} & =\left\{x: x>\frac{1}{2}\left(\mu_{1}+\mu_{2}\right)\right\}
\end{aligned}
$$



Maximum likelihood rule for one-dimensional normal distributions with different variances.

Multinormal distribution with common variance matrix

Rearranging terms leads to:

$$
\begin{aligned}
-2 \mu_{1}^{\top} \Sigma^{-1} x+2 \mu_{2}^{\top} \Sigma^{-1} x+\mu_{1}^{\top} \Sigma^{-1} \mu_{1}-\mu_{2}^{\top} \Sigma^{-1} \mu_{2} & <0 \\
2\left(\mu_{2}-\mu_{1}\right)^{\top} \Sigma^{-1} x+\left(\mu_{1}-\mu_{2}\right)^{\top} \Sigma^{-1}\left(\mu_{1}+\mu_{2}\right) & <0 \\
\left(\mu_{1}-\mu_{2}\right)^{\top} \Sigma^{-1}\left\{x-\frac{1}{2}\left(\mu_{1}+\mu_{2}\right)\right\} & >0 \\
\alpha^{\top}(x-\mu) & >0
\end{aligned}
$$

where $\alpha=\Sigma^{-1}\left(\mu_{1}-\mu_{2}\right)$ and $\mu=\frac{1}{2}\left(\mu_{1}+\mu_{2}\right)$.
The resulting discriminant rule is linear (see also R command lda()).

Multinormal distribution with common variance matrix

## Suppose $\Pi_{i}: N_{p}\left(\mu_{i}, \Sigma\right)$.

The Bayes rule (assuming equal prior probabilities) allocates $x$ to $\Pi_{j}$, where $j \in\{1, \ldots, J\}$ is the value that minimizes the square Mahalanobis distance between $x$ and $\mu_{i}$ :

$$
\delta^{2}\left(x, \mu_{i}\right)=\left(x-\mu_{i}\right)^{\top} \Sigma^{-1}\left(x-\mu_{i}\right), i=1, \ldots, J
$$

In the case of $J=2: x$ is allocated to $\Pi_{1}$ if

$$
\left(x-\mu_{1}\right)^{\top} \Sigma^{-1}\left(x-\mu_{1}\right)<\left(x-\mu_{2}\right)^{\top} \Sigma^{-1}\left(x-\mu_{2}\right)
$$

## Week 10-11

Discriminant analysis

## Probability of misclassification

Suppose that $\Pi_{i}: N_{p}\left(\mu_{i}, \Sigma\right)$.
Consider

$$
p_{12}=P\left(x \in R_{1} \mid \Pi_{2}\right)=P\left\{\alpha^{\top}(x-\mu)>0 \mid \Pi_{2}\right\}
$$

$\ln \Pi_{2}, \alpha^{\top}(X-\mu) \sim N\left(-\frac{1}{2} \delta^{2}, \delta^{2}\right)$ where $\delta^{2}=\left(\mu_{1}-\mu_{2}\right)^{\top} \Sigma^{-1}\left(\mu_{1}-\mu_{2}\right)$ is the squared Mahalanobis distance between the two populations, we obtain

$$
p_{12}=\Phi\left(-\frac{1}{2} \delta\right) .
$$

Similarly, we obtain the probability of misclassification into population 2 for $x$ from $\Pi_{1}$ as $p_{21}=\Phi\left(-\frac{1}{2} \delta\right)$.

Two multinormal distributions

Assuming that $\Pi_{i}: N_{p}\left(\mu_{i}, \Sigma_{i}\right)$, for $i=1,2$, the discriminant rule becomes more complicated.

$$
\begin{gathered}
R_{1}=\left\{x:-\frac{1}{2} x^{T}\left(\Sigma_{1}^{-1}-\Sigma_{2}^{-1}\right) x+\left(\mu_{1}^{T} \Sigma_{1}^{-1}-\mu_{2}^{T} \Sigma_{2}^{-1}\right) x-k\right. \\
\left.\geq \ln \left[\left\{\frac{C(1 \mid 2)}{C(2 \mid 1)}\right\}\left\{\frac{\pi_{2}}{\pi_{1}}\right\}\right]\right\}
\end{gathered}
$$

where $k=\frac{1}{2} \ln \left(\frac{\left|\Sigma_{1}\right|}{\left|\Sigma_{2}\right|}\right)+\frac{1}{2}\left(\mu_{1}^{T} \Sigma_{1}^{-1} \mu_{1}-\mu_{2}^{T} \Sigma_{2}^{-1} \mu_{2}\right)$.
This is a quadratic classification rule (notice that $\frac{1}{2} x^{T}\left(\Sigma_{1}^{-1}-\Sigma_{2}^{-1}\right) x=0$ if $\Sigma_{1}=\Sigma_{2}$ )

## Week 10-1

Discriminant analysis

## Discriminant rules in practice

The unknown parameters $\left(\mu_{j}, \Sigma_{j}\right)$ are estimated by $\left(\bar{x}_{j}, S_{j}\right)$ in each $\Pi_{j}$.
The common variance matrix $\Sigma$ can be estimated by the pooled variance matrix $\mathcal{S}_{u}=\sum_{j=1}^{J} n_{j}\left(\frac{\mathcal{S}_{j}}{n-J}\right)$, where $n=\sum_{j=1}^{J} n_{j}$.

R library MASS contains the following simple functions for discriminant analysis:

Ida(): linear discriminant analysis (assuming equal variance matrices),
qda(): quadratic discriminant analysis (with possibly different variance matrices)

## Discriminant Analysis

* Discriminant analysis is a set of methods for distinguishing between groups in data and allocating new observations into groups.
* The Bayes discriminant rule allocates an observation $x$ to the population $\Pi_{j}$ that maximizes $\max _{j} \pi_{j} f_{j}(x)$.
* For the ML rule and $J=2$ multivariate normal populations, the discriminant rule can be derived from ratio of the densities. The discriminant rule is linear for common variance matrices and quadratic if the variance matrices are different.
* For the ML rule and $J=2$ normal populations with common variance matrix, the probabilities of misclassification are given by $p_{12}=p_{21}=\Phi\left(-\frac{1}{2} \delta\right)$ where $\delta$ is the square root of the Mahalanobis distance between the 2 populations.

[^2]
## Apparent and actual error rate

The apparent error rate (APER) is defined as the percentage of misclassified observations. APER is based on the observations which were used to construct the discriminant rule and it might be too optimistic.

In order to obtain a more appropriate estimate of the misclassification probability, we may use simple leave-one-out (or cross-validation) algorithm
(1) Calculate the discrimination rule from all but one observation.Allocate the omitted observation according to the rule from step 1.Repeat steps 1 and 2 for all observations and count the number of correct and wrong classifications.

The estimate of the misclassification rate based on this procedure is called the actual error rate (AER).
Z. Hlávka (KPMS)

## NMST539



## Week 10-11

 Discriminant analysis
## Three (or more) groups

Allocation regions for $J=3$ groups:

$$
\begin{aligned}
& h_{12}(x)=\left(\bar{x}_{1}-\bar{x}_{2}\right)^{\top} \mathcal{S}_{u}^{-1}\left(x-\frac{1}{2}\left(\bar{x}_{1}+\bar{x}_{2}\right)\right) \\
& h_{13}(x)=\left(\bar{x}_{1}-\bar{x}_{3}\right)^{\top} \mathcal{S}_{u}^{-1}\left(x-\frac{1}{2}\left(\bar{x}_{1}+\bar{x}_{3}\right)\right) \\
& h_{23}(x)=\left(\bar{x}_{2}-\bar{x}_{3}\right)^{\top} \mathcal{S}_{u}^{-1}\left(x-\frac{1}{2}\left(\bar{x}_{2}+\bar{x}_{3}\right)\right) .
\end{aligned}
$$

The ML rule is to allocate $x$ to

$$
\left\{\begin{array}{lllll}
\Pi_{1} & \text { if } & h_{12}(x)>0 & \text { and } & h_{13}(x)>0 \\
\Pi_{2} & \text { if } & h_{12}(x)<0 & \text { and } & h_{23}(x)>0 \\
\Pi_{3} & \text { if } & h_{13}(x)<0 & \text { and } & h_{23}(x)<0
\end{array}\right.
$$

In R , discriminant analysis with 3 groups works differently.
\#\# cross-validation
z.cv.cl=lda(Species ~ ., iris, prior $=c(1,1,1) / 3$, CV=TRUE) \$class
z.al.cl=predict(lda(Species ~ ., iris, prior = c ( $1,1,1$ )/3), iris $[,-5])$ \$class

## Example:

```
data(iris)
```

\#\# training data set
train=sample(1:150,75);table(iris\$Species[train])
z=lda(Species ${ }^{\sim}$., iris, prior=c $(1,1,1) / 3$, subset $=$ train $)$
table(predict(z, iris[-train, -5$]) \$ c l a s s$,
iris[-train,"Species"])
(2)

Fisher's approach

Based on projections $\mathcal{Y}=\mathcal{X} a$ of the original data set $\mathcal{X}$.
Projections leading to a good separation are found by maximizing the ratio of the between-group-sum of squares to the within-group-sum of squares.

The within-sum-of-squares measures the sum of variations within each group:

$$
\sum_{j=1}^{J} \mathcal{Y}_{j}^{\top} \mathcal{H}_{j} \mathcal{Y}_{j}=\sum_{j=1}^{J} a^{\top} \mathcal{X}_{j}^{\top} \mathcal{H}_{j} \mathcal{X}_{j} a=a^{\top} \mathcal{W} a
$$

where $\mathcal{Y}_{j}$ denotes the $j$-th submatrix of $\mathcal{Y}$ corresponding to observations of group $j$ and $\mathcal{H}_{j}$ denotes the ( $n_{j} \times n_{j}$ ) centering matrix.

The total-sum-of-squares $\sum_{i=1}^{n}\left(y_{i}-\bar{y}\right)^{2}=\mathcal{Y}^{\top} \mathcal{H} \mathcal{Y}=a^{\top} \mathcal{X}^{\top} \mathcal{H} \mathcal{X} a=a^{\top} \mathcal{T} a$ can be decomposed as

$$
\begin{aligned}
\text { total } S S & =\text { within } S S \\
a^{\top} \mathcal{T} a & =a^{\top} \mathcal{W} a \\
\operatorname{l}^{\top} & a^{\top} \mathcal{B} a
\end{aligned}
$$

The idea is to select a maximizing maximizes the ratio

$$
\frac{a^{\top} \mathcal{B} a}{a^{\top} \mathcal{W} a}
$$

## Week 10-11

 Discriminant analysisExample: For two groups of sizes $n_{1}$ and $n_{2}$, we obtain:

$$
\begin{aligned}
a^{\top} \mathcal{B} a & =n_{1}\left\{a^{\top}\left(\bar{x}_{1}-\bar{x}\right)\right\}^{2}+n_{2}\left\{a^{\top}\left(\bar{x}_{2}-\bar{x}\right)\right\}^{2} \\
& =n_{1}\left\{a^{\top}\left(\bar{x}_{1}-\bar{x}_{2}\right) / 2\right\}^{2}+n_{2}\left\{a^{\top}\left(\bar{x}_{1}-\bar{x}_{2}\right) / 2\right\}^{2} \\
& =\frac{n_{1}+n_{2}}{4}\left\{a^{\top}\left(\bar{x}_{1}-\bar{x}_{2}\right)^{2}\right\}
\end{aligned}
$$

Clearly, $\mathcal{B}=\left\{\left(n_{1}+n_{2}\right) / 4\right\} d d^{\top}$, where $d=\left(\bar{x}_{1}-\bar{x}_{2}\right)$ and the largest eigenvalue of $\mathcal{W}^{-1} \mathcal{B}$ is $\left(n_{1}+n_{2} / 4\right) d^{\top} \mathcal{W}^{-1} d$.

Therefore, the corresponding eigenvector has to satisfy:

$$
\begin{aligned}
\mathcal{W}^{-1} \mathcal{B} \gamma & =\left\{\left(n_{1}+n_{2}\right) / 4\right\} d^{\top} \mathcal{W}^{-1} d \gamma \\
\mathcal{W}^{-1} d d^{\top} \gamma & =d^{\top} \mathcal{W}^{-1} d \gamma \\
\text { leading } \gamma=\mathcal{W}^{-1} d=\mathcal{W}^{-1}\left(\bar{x}_{1}\right. & \left.-\bar{x}_{2}\right)
\end{aligned}
$$

## Proportion of trace and more groups

In this way, we find only one direction maximizing the differences between two groups

For three $\operatorname{groups}, \operatorname{rank}(\mathcal{B})=2$, and we obtain two directions (i.e., a linear transformation of the original data set maximizing the between-group differences w.r.t. the within-group variability). The eigenvalues of $\mathcal{W}^{-1} \mathcal{B}$ correspond to the importance of these directions (its percentages can be interpreted as percentages of between-group differences explained by the corresponding directions).

For $g$ groups, $\operatorname{rank}(\mathcal{B}) \leq \min (p, g-1)$. I.e., we obtain at most $g-1$ linear discriminants.

Other usable methods
logistic regression
classification trees
k-nearest neighbors
support vector machine
neural networks

Example:

```
data(iris)
## training data set
train=sample(1:150,75);table(iris$Species[train])
z=lda(Species~.,iris,prior=c(1,1,1)/3,subset = train)
pz<-predict(z, iris[-train,-5])
table(pz$class,iris[-train,"Species"])
eqscplot(pz$x, type="n",xlab="LD1",ylab="LD2")
spec=as.numeric(iris[-train,5],1,1)
text(pz$x,labels=spec,col=spec)
z ## see "proportion of trace"
```


## 1 <br> Discrimination Rules in Practice

* Linear discriminant rule allocates $x$ to the population with smallest Mahalanobis distance

$$
\delta^{2}\left(x ; \mu_{i}\right)=\left(x-\mu_{i}\right)^{\top} \Sigma^{-1}\left(x-\mu_{i}\right) .
$$

* Classification for different covariance structures in the two populations leads to quadratic discrimination rules.
* The probability of misclassification can be estimated by cross-validation.
* Fisher's linear discrimination finds a linear combination $a^{\top} x$ that maximizes the ratio of the "between-sum-of-squares" and the "within-sum-of-squares". This rule is identical to the (linear) ML rule in the case of $J=2$ for normal populations.


## Týden 11

## Kanonické korelace:

- kanonické proměnné, kanonické vektory a kanonické korelace,
- praktické použití a příklad.

Assuming that

$$
\binom{X}{Y} \sim\left(\binom{\mu}{\nu},\binom{\Sigma_{X X} \Sigma_{X Y}}{\Sigma_{Y X} \Sigma_{Y Y}}\right)
$$


Note that $\rho(c a, b)=\rho(a, b)$ for any $c \in \mathbb{R}$. Therefore, we can maximize $a^{\top} \Sigma_{X Y} b$ under the constraints $a^{\top} \Sigma_{X X} a=b^{\top} \Sigma_{Y Y} b=1$.

And this is the same as maximizing $u^{\top} \Sigma_{X X}^{-1 / 2} \Sigma_{X Y} \Sigma_{Y Y}^{-1 / 2} v$ under the constraints $\|u\|=\|v\|=1$.

## Canonical correlations

We have random vectors $X \in \mathbb{R}^{q}$ and $Y \in \mathbb{R}^{p}$.
Consider linear combinations:

$$
a^{\top} X \quad \text { and } \quad b^{\top} Y
$$

Correlation of the linear combinations:

$$
\rho(a, b)=\rho_{a^{\top} X b^{\top} Y .}
$$

We want to find $a, b$ maximizing the correlation $\rho(a, b)$.
The linear combinations $a^{\top} X$ and $b^{\top} Y$ describe the structure of "common variability" of $X$ and $Y$.

Denoting $\mathcal{K}=\Sigma_{X X}^{-1 / 2} \Sigma_{X Y} \Sigma_{Y Y}^{-1 / 2}$, we have $u^{\top} \Sigma_{X X}^{-1 / 2} \Sigma_{X Y} \Sigma_{Y Y}^{-1 / 2} v=u^{\top} \mathcal{K} v$.
Clearly, for each $v$ fixed such that $\|v\|=1$, we have the following

$$
\begin{aligned}
\max _{u,\|u\|=1}\left(u^{\top} \mathcal{K} v\right)^{2} & \leq \max _{u,\|u\|=1} u^{\top} \mathcal{K} v v^{\top} \mathcal{K}^{\top} u \\
& =v^{\top} \mathcal{K}^{\top} \mathcal{K} v \\
& \leq \lambda_{1}
\end{aligned}
$$

where $\lambda_{1}$ is the largest eigenvalue of $\mathcal{K}^{\top} \mathcal{K}$.
The SVD decomposition $\mathcal{K}=\Gamma \Lambda^{1 / 2} \Delta^{\top}$ with $k=\operatorname{rank}(\mathcal{K})$ and $\lambda_{1} \geq \lambda_{2} \geq \ldots \geq \lambda_{k}>0$ then leads

$$
\gamma_{1}^{\top} \mathcal{K} \delta_{1}=\lambda_{1}^{1 / 2}\left(=\rho\left(\Sigma_{X X}^{-1 / 2} \gamma_{1}, \Sigma_{Y Y}^{-1 / 2} \delta_{1}\right)\right)
$$

Theorem: Define $f_{r}=\max _{a, b} a^{\top} \Sigma_{X Y} b$ under the constraints
$a^{\top} \Sigma_{X X} a=b^{\top} \Sigma_{Y Y} b=1$ and $a_{i}^{\top} \Sigma_{X X a}=b_{i}^{\top} \Sigma_{Y Y} b=0$ for $i=1, \ldots, r-1$ (for some $r \in\{1, \ldots, k\}$ fixed).

The maximum of $\rho(a, b)$ under the above constraints is given by $f_{r}$ and it is attained when $a=a_{r}=\Sigma_{X X}^{-1 / 2} \gamma_{r}$ and $b=b_{r}=\Sigma_{Y Y}^{-1 / 2} \delta_{r}$.

The correlation $\rho(a, b)$ is maximized for $a=a_{1}$ and $b=b_{1}$ and $\rho\left(a_{1}, b_{1}\right)=\lambda_{1}^{1 / 2}$ is the correlation of random variables $\eta_{1}$ and $\varphi_{1}$.

The vectors $a_{r}$ and $b_{r}$ maximize the correlation subject to the condition that $a^{\top} X$ and $b^{\top} X$ are uncorrelated with the previous canonical variables $a_{i}^{\top} X$ and $b_{i}^{\top} X$, respectively.

## Week 11-12

 Canonical correlations
## Properties

Theorem: Let $\eta$ and $\varphi$ be the canonical variables, i.e., the components of the vector $\eta$ are

$$
\eta_{i}=\left(\Sigma_{X X}^{-1 / 2} \gamma_{i}\right)^{\top} X
$$

and the components of the vector $\varphi$ are

$$
\varphi_{i}=\left(\Sigma_{Y Y}^{-1 / 2} \delta_{i}\right)^{\top} Y
$$

for $1 \leq i \leq k$. Then

$$
\operatorname{Var}\binom{\eta}{\varphi}=\left(\begin{array}{cc}
\mathcal{I} & \Lambda^{1 / 2} \\
\Lambda^{1 / 2} & \mathcal{I}
\end{array}\right)
$$

where $\Lambda^{1 / 2}=\operatorname{diag}\left(\lambda_{1}^{1 / 2}, \ldots, \lambda_{k}^{1 / 2}\right)$.

## Terminology

## Canonical correlation vectors

$$
\begin{aligned}
a_{i} & =\Sigma_{X X}^{-1 / 2} \gamma_{i} \\
b_{i} & =\Sigma_{Y Y}^{-1 / 2} \delta_{i}
\end{aligned}
$$

## Canonical variables

$$
\begin{aligned}
\eta_{i} & =a_{i}^{\top} X \\
\varphi_{i} & =b_{i}^{\top} Y
\end{aligned}
$$

## Canonical correlation coefficients

$$
\lambda_{1}^{1 / 2}, \ldots, \lambda_{k}^{1 / 2}
$$

NMST539

## Week 11-12

Canonical correlations

## Relation to principal components

Both PC and CC are calculated using eigenvalues and eigenvectors of some (covariance) matrices.

PC analysis decomposes the "total variability" of one dataset.
CC analysis decomposes the total "common variability" of two datasets.
The "common variability" is described in terms of linear combinations (it is common and therefore we get description of the common variability in terms of both datasets).

The canonical variables in both datasets are related: the first canonical variable in the first dataset describes the same part of the common variablity as the first canonical variability in the second dataset.

## Canonical correlation analysis

* Canonical correlation analysis aims to identify possible links between two (sub-)sets of variables $X \in \mathbb{R}^{q}$ and $Y \in \mathbb{R}^{p}$. The idea is to find indices $a^{\top} X$ and $b^{\top} Y$ such that the correlation $\rho(a, b)=\rho_{a^{\top} X b^{\top} Y}$ is maximal.
* The maximum correlation is found by $a_{i}=\Sigma_{X X}^{-1 / 2} \gamma_{i}$ and $b_{i}=\Sigma_{Y Y}^{-1 / 2} \delta_{i}$, where $\gamma_{i}$ and $\delta_{i}$ denote the eigenvectors of $\mathcal{K} \mathcal{K}^{\top}$ and $\mathcal{K}^{\top} \mathcal{K}$, $\mathcal{K}=\Sigma_{X X}^{-1 / 2} \Sigma_{X Y} \Sigma_{Y Y}^{-1 / 2}$.
* The vectors $a_{i}$ and $b_{i}$ are the canonical correlation vectors, $\eta_{i}=a_{i}^{\top} X$ and $\varphi_{i}=b_{i}^{\top} Y$ are the canonical variables.
* The covariance between the canonical variables is $\operatorname{cov}\left(\eta_{i}, \varphi_{i}\right)=\sqrt{\lambda}{ }_{i}$, $i=1, \ldots, k$
* Canonical correlations are invariant w.r.t. linear transformations of the original variables $X$ and $Y$.


## Test of independence

We have already seen that $-2 \log \lambda=-n \log \left|\mathcal{I}-\mathcal{S}_{22}^{-1} \mathcal{S}_{21} \mathcal{S}_{11}^{-1} \mathcal{S}_{12}\right|$.
It can be shown that this LRT test statistic is distributed as a ratio of determinants of independent Wishart matrices (this is the Wilks' lambda distribution).

For large values of $n$, the Wilks' lambda distribution can be approximated (Bartlett's approximation):

$$
-\{n-(p+q+3) / 2\} \log \left|\mathcal{I}-\mathcal{S}_{22}^{-1} \mathcal{S}_{21} \mathcal{S}_{11}^{-1} \mathcal{S}_{12}\right| \sim \chi_{p q}^{2},
$$

i.e., we reject independence of $X$ and $Y$ if

$$
-\{n-(p+q+3) / 2\} \log \left|\mathcal{I}-\mathcal{S}_{22}^{-1} \mathcal{S}_{21} \mathcal{S}_{11}^{-1} \mathcal{S}_{12}\right|>\chi_{q p}^{2}(1-\alpha)
$$

## Canonical Correlations in Practice

In practice, the covariance matrices $\Sigma_{X X}, \Sigma_{X Y}, \Sigma_{Y Y}$ are estimated by sample covariance matrices $S_{X X}, S_{X Y}, S_{Y Y}$. The canonical correlation analysis is carried out on the estimates.

Before running the analysis, one should test the hypothesis of independence between $X$ and $Y$ (using, e.g., the ML test described previously):

Let $Z_{i}=\left(X_{i}^{\top}, Y_{i}^{\top}\right)^{\top} \sim N_{q+p}(\mu, \Sigma), i=1 \cdots, n$ be independent,

$$
H_{0}: \Sigma_{X Y}=0, \quad H_{1}: \quad \text { no constraints. }
$$

## Week 11-12

Canonical correlations
Example: What is the relationship between the datasets on US crimes (murder, rape, robbery, assault, burglary, larceny, autotheft) and US health (accident, cardiovascular, cancer, pulmonar, pneumonia, diabetis, liver)?

```
data(uscrime)
x=sqrt(as.matrix(uscrime[,3:9]))
x=scale(x)
data(ushealth)
y=sqrt(as.matrix(ushealth[,3:9]))
y=scale(y)
n=nrow(x); p=ncol(x); q=ncol(y)
```

x denotes US crimes
y denotes US health
$s x x=\operatorname{cov}(x) ; s y y=\operatorname{cov}(y) ; s x y=\operatorname{cov}(x, y)$

```
t=-(n-(p+q+3)/2)*log(\operatorname{det}(\operatorname{diag}(1,q)-
    solve(syy)%*%t(sxy)%*%solve(sxx) %*%sxy))
```

format.pval(1-pchisq(t,p*q)) \#\# test of independence
e=eigen (sxx)
sxx12=e\$vectors $\% * \%$ (sqrt (diag(1/e\$values)) ) $\% * \%$ t(e\$vectors)
e=eigen(syy)
syy12=e\$vectors\%*\% (sqrt (diag(1/e\$values))) \% * \% t (e\$vectors)
kkt $=$ sxx $12 \% * \%$ sxy $\% * \%$ syy $12 \% * \%$ syy $12 \% * \%$ t (sxy) $\% * \%$ sxx 12
ktk $=$ syy $12 \% * \%$ t (sxy) $\% * \%$ sxx $12 \% * \%$ sxx $12 \% * \%$ sxy $\% * \%$ syy 12
e1=eigen(kkt)
e2=eigen(ktk)
print (cbind (e1\$values, e2\$values))
a=sxx $12 \% * \%$ e1\$vectors
b=syy $12 \% * \%$ e2\$vectors

## ${ }^{192}=4$

## Canonical Correlations in Practice

* In practice, we estimate $\Sigma_{X X}, \Sigma_{X Y}, \Sigma_{Y Y}$ by the empirical covariances and to compute estimates $\ell_{i}, g_{i}, d_{i}$ for $\lambda_{i}, \gamma_{i}, \delta_{i}$ from the SVD of $\widehat{\mathcal{K}}=\mathcal{S}_{X X}^{-1 / 2} \mathcal{S}_{X Y} \mathcal{S}_{Y Y}^{-1 / 2}$.
* The coefficients of the canonical variables (i.e., the canonical vectors) tell us the influence of these variables.
* The independence of the two random vectors can be tested by a likelihood ratio test leading to Wilks' lambda distribution.
* Barlett's test of the null hypothesis that only s population canonical correlation coefficients are non zero is based on the statistic $-\{n-(p+q+3) / 2\} \log \prod_{i=s+1}^{\min (p, q)}\left(1-r_{i}\right) \sim \chi_{(p-s)(q-s)}^{2}$, where $r_{i}$ are the sample canonical correlation coefficients.

```
## canonical variables
cvx=x%*%a
cvy=y%*%%
## plot of the first pair
plot(cvx[,1],cvy[,1],type="n")
text(cvx[,1],cvy[,1],row.names(ushealth))
## canonical correlation
cor(cvx[,1],cvy[,1])
sqrt(e1$values[1])
## R library stats
cancor(x,y)
## coefficients are divided by sqrt(dim)?
```

Korespondenční analýza:

- testy nezávislosti v kontingenční tabulce,
- reprezentace řádků a sloupců.


## Correspondence analysis

Categorical scales are pervasive in the social sciences for measuring attitudes and opinions on various issues and demographic characteristics such as gender, race, and social class.

Categorical scales (...) occur frequently in the behavioral sciences, public health, ecology, education, and marketing. They even occur in highly quantitative fields such as engineering sciences and industrial quality control. Such applications often involve subjective evaluation of some characteristic-how soft to the touch a certain fabric is, how good a particular food product tastes, or how easy a worker finds a certain task to be.
(Alan Agresti, Categorical Data Analysis, Wiley, 1990)
Z. Hlávka (KPMS)

## Week 12 Correspondence analysis

Example:

$$
\begin{gathered}
\mathcal{X}=\left(\begin{array}{rrr|r}
8 & 4 & 3 & \\
2 & 1 & 6 & \\
3 & 6 & 2 & \\
\hline 13 & 11 & 11 & 35
\end{array}\right) \quad \begin{array}{l}
\leftarrow \text { Beer } \\
\leftarrow \text { Wine } \\
\leftarrow \text { Spirit } \\
\uparrow \text { Czechia } \\
\uparrow \text { Russia } \\
\uparrow \mathrm{GB}
\end{array}
\end{gathered}
$$

Joint distribution: $\pi_{i j}=P(Z=i, Y=j)$ is the probability that $Z$ is equal to $i$ and at the same time $Y$ is $j$.

Marginal distribution of $Z: \pi_{i}$. is the probability that $Z$ is equal to $i$
Marginal distribution of $Y: \pi_{\cdot j}$ is the probability that $Y$ is equal to $j$

Two-way contingency table

Variable $Z$ has I levels
Variable $Y$ has $J$ levels
This gives IJ combinations of levels of $Z$ and $Y$
We count the responses $(Z, Y)$ in our sample and display this information in rectangular table which has I rows and $J$ columns.

In each cell we give the number of subjects in our sample having the corresponding combination of responses on $Z$ and $Y$.

The entry $x_{i j}$ in the contingency table $\mathcal{X}(n \times p)$ is the number of observations in a sample that simultaneously fall in the ith row category and the $j$ th column category.
Z. Hâvka (KPMS) NMST539

## Sampling Distributions

This is the way in which the table was created. It is important for understanding the table correctly.

The likelihoods depend on the sampling distribution.

Poisson sampling: everything is random
Multinomial sampling: total number of observed subjects is fixed,
Independent multinomial sampling: number of subject in each row or column is fixed.

Estimators and likelihood ratio tests are often identical for all types of sampling (NMST432 Advanced Regression Models).

## Maximum Likelihood Estimates

By maximizing the likelihood function we obtain the ML estimator

$$
\hat{\pi}_{i j}=p_{i j}=x_{i j} / x_{\bullet \bullet}
$$

where $x_{\bullet \bullet}=\sum_{i=1}^{n} x_{i \bullet \bullet}=\sum_{j=1}^{n} x_{\bullet j}$ is the total number of observations.
Notice that $Z$ and $Y$ are independent if for all $i$ and $j: \pi_{i \mid j}=\pi_{i j} / \pi_{. j}=\pi_{i}$. or $\pi_{j \mid i}=\pi_{j i} / \pi_{i}=\pi_{\cdot j}$ or or $\pi_{i j}=\pi_{i \cdot} \pi_{\cdot j}$.

The ML estimators of cell probabilities $\pi_{i j}$ under independence are

$$
\hat{\pi}_{i j}=p_{i \bullet} p_{\bullet j}=\left(x_{i \bullet} x_{\bullet j}\right) / x_{\bullet \bullet}^{2}
$$

$x_{i \bullet}=\sum_{j=1}^{n} x_{i j}$ is the number of observations falling into the $i$ th row category.

Example: Alcohol consumption in three countries.

```
alc=matrix(c(8, 4, 3, 2, 1, 6, 3, 6, 2),3,byrow=T)
row.names(alc)=c("Beer","Wine","Spirit")
colnames(alc)=c("Czechia","Russia","GB")
chisq.test(alc)
Pearson's Chi-squared test
data: alc
X-squared = 9.8406, df = 4, p-value = 0.0432
Warning message: Chi-squared approximation may be incorrect
```


## Decomposition of $\chi^{2}$-statistic

The SVD of $\mathcal{C}=\left(c_{i j}\right)_{i=1, \ldots, n ; j=1, \ldots, p}$ yields

$$
\mathcal{C}=\Gamma \Lambda^{1 / 2} \Delta^{\top}
$$

with $\Lambda^{1 / 2}=\operatorname{diag}\left(\lambda_{1}^{1 / 2}, \ldots, \lambda_{R}^{1 / 2}\right)$, where $\lambda_{1}, \ldots, \lambda_{R}$ are the nonzero eigenvalues of both $\mathcal{C}^{\top} \mathcal{C}$ and $\mathcal{C} \mathcal{C}^{\top}$

Now, it is easy to see that

$$
t=\sum_{i=1}^{n} \sum_{j=1}^{p}\left(x_{i j}-E_{i j}\right)^{2} / E_{i j}=\sum_{i=1}^{n} \sum_{j=1}^{p} c_{i j}^{2}=\operatorname{tr}\left(\mathcal{C C}^{\top}\right)=\sum_{k=1}^{R} \lambda_{k} .
$$

Hence, the SVD of the matrix $\mathcal{C}$ decomposes the $\chi^{2}$-statistic $t$.

## Week 12

 Correspondence analysis
## Marginal frequencies

Defining $\mathcal{A}=\operatorname{diag}\left(x_{i \bullet}\right)$ and $\mathcal{B}=\operatorname{diag}\left(x_{\bullet j}\right)$ leads the vectors of marginal row and column frequencies:

$$
a=\mathcal{A} 1_{n} \quad \text { and } \quad b=\mathcal{B} 1_{p}
$$

This allows to write $E=a b^{\top} x_{\bullet \bullet}^{-1}$ and $\mathcal{C}=\mathcal{A}^{-1 / 2}(\mathcal{X}-E) \mathcal{B}^{-1 / 2} \sqrt{x_{\bullet \bullet}}$.
It is easy to verify that

$$
\begin{array}{rlrl}
\mathcal{C} \sqrt{b}=0 & & \text { and } & \\
\mathcal{C}^{\top} \sqrt{a}=0, \\
\delta_{k}^{\top} \sqrt{b}=0 & & \text { and } & \\
\gamma_{k}^{\top} \sqrt{a}=0 .
\end{array}
$$

## Example:

decomp=svd(C)
gamma1=decomp\$u[,1]
delta1=decomp\$v[,1]
lambda=decomp\$d
sum(lambda^2) \#\# chi2 statistika

Row and column coordinates

The row coordinates $r_{k}=\mathcal{A}^{-\frac{1}{2}} \mathcal{C} \delta_{k}$ and column coordinates $s_{k}=\mathcal{B}^{-\frac{1}{2}} \mathcal{C}^{\top} \gamma_{k}$ satisfy

$$
r_{k}^{\top} a=\delta_{k}^{\top} \mathcal{C}^{\top} \mathcal{A}^{-\frac{1}{2}} a=\delta_{k}^{\top} \mathcal{C}^{\top} \sqrt{a}=\delta_{k}^{\top} 0=0
$$

and

$$
s_{k}^{\top} b=\gamma_{k}^{\top} \mathcal{C B} \mathcal{B}^{\frac{1}{2}} b=\gamma_{k}^{\top} \mathcal{C} \sqrt{b}=\gamma_{k}^{\top} 0=0 .
$$

The true meaning of relations $r_{k}^{\top}=0$ and $s_{k}^{\top} b=0$ is

$$
\bar{r}_{k}=\frac{1}{x_{\bullet \bullet}} r_{k}^{\top} a=0 \quad \text { and } \quad \bar{s}_{k}=\frac{1}{x_{\bullet \bullet}} s_{k}^{\top} b=0
$$

where means are (of course) weighted by the row and column marginal frequencies. Hence, both row and column factors are centered.

Example:

```
decomp=svd(C); gamma1=decomp$u[,1]; delta1=decomp$v[,1]
A=diag(a); B=diag(b)
r1=diag(1/sqrt(a)) %*%%%**%delta1
s1=diag(1/sqrt(b))%*%t (C)%*%gamma1
row.names(r1)=row.names (C)
row.names(s1)=colnames(C)
```


## Example:

\#\# prumery
sum(r1*a)
sum (s1*b)
\#\# rozptyly
$\operatorname{sum}((r 1 \sim 2) * a) / n$
$\operatorname{sum}\left(\left(s 1^{\sim} 2\right) * b\right) / n$
\#\#
(lambda[1] ~2) /n

## Variance of row and column factors

For the sample variances of $r_{k}$ and $s_{k}$ we have the following:

$$
\begin{aligned}
& \widehat{\operatorname{Var}}\left(r_{k}\right)=\frac{1}{x_{\bullet \bullet}} \sum_{i=1}^{n} x_{i \bullet} r_{k i}^{2}=r_{k}^{\top} \mathcal{A} r_{k} / x_{\bullet \bullet}=\delta_{k}^{\top} \mathcal{C}^{\top} \mathcal{C} \delta_{k} / x_{\bullet \bullet}=\frac{\lambda_{k}}{x_{\bullet \bullet}} \\
& \widehat{\operatorname{Var}}\left(s_{k}\right)=\frac{1}{x_{\bullet \bullet}} \sum_{j=1}^{p} x_{\bullet j} s_{k j}^{2}=s_{k}^{\top} \mathcal{B} s_{k} / x_{\bullet \bullet}=\gamma^{\top} \mathcal{C} \mathcal{C}^{\top} \gamma_{k} / x_{\bullet \bullet}=\frac{\lambda_{k}}{x_{\bullet \bullet}} .
\end{aligned}
$$

In practice, statistical software may return differently scaled values (than $r_{k}$ nad $s_{k}$ ). Functions corresp() and ca() in R libraries MASS and ca standardize row and column factors by $\rho_{k}=\left(\lambda_{k} / x_{\bullet \bullet}\right)^{1 / 2}$.

This means that the row and column factors given by standard software are standardized.

## Week 12

Correspondence analysis

## Proportion of explained variance

Hence, the proportion of the variance explained by the $k$ th factor is

$$
\widehat{\operatorname{Var}}\left(r_{k}\right) / \sum_{i=1}^{R} \widehat{\operatorname{Var}}\left(r_{k}\right)=\lambda_{k} / \sum_{i=1}^{R} \lambda_{i}
$$

The variance of the $k$ th row factor, $\widehat{\operatorname{Var}}\left(r_{k}\right)$, can be further decomposed into the absolute single row contributions defined as

$$
C_{a}\left(i, r_{k}\right)=\frac{x_{i} \bullet r_{k i}^{2}}{\lambda_{k}}, \text { for } i=1, \ldots, n, k=1, \ldots, R
$$

Similarly $C_{a}\left(j, s_{k}\right)=x_{\bullet j} s_{k j}^{2} / \lambda_{k}$ for $j=1, \ldots, p, k=1, \ldots, R$ are the absolute contributions of column $j$ to the variance of the column factor $s_{k}$.

These absolute contributions may help to interpret the row and column factors obtained by the correspondence analysis.

## Relation between row and column coordinates

From the properties of SVD we know the relationship between $\delta_{k}$ and $\gamma_{k}$ :

$$
\delta_{k}=\frac{1}{\sqrt{\lambda_{k}}} \mathcal{C}^{\top} \gamma_{k} \quad \text { and } \quad \gamma_{k}=\frac{1}{\sqrt{\lambda_{k}}} \mathcal{C} \delta_{k} .
$$

Therefore

$$
s_{k}=\mathcal{B}^{-1 / 2} \mathcal{C}^{\top} \gamma_{k}=\sqrt{\lambda_{k}} \mathcal{B}^{-1 / 2} \delta_{k} .
$$

Using the definition of $r_{k}$, we have

$$
\begin{aligned}
r_{k} & =\mathcal{A}^{-1 / 2} \mathcal{C} \delta_{k}=\sqrt{x_{\bullet \bullet}} \mathcal{A}^{-1 / 2} \mathcal{A}^{-1 / 2}(\mathcal{X}-E) \mathcal{B}^{-1 / 2} \delta_{k} \\
& =\sqrt{\frac{x_{\bullet \bullet}}{\lambda_{k}}} \mathcal{A}^{-1}(\mathcal{X}-E) s_{k}=\sqrt{\frac{x_{\bullet \bullet}}{\lambda_{k}}} \mathcal{A}^{-1}\left(\mathcal{X} s_{k}-\frac{a b^{\top} s_{k}}{x_{\bullet \bullet}}\right) \\
& =\sqrt{\frac{x_{\bullet \bullet}}{\lambda_{k}}} \mathcal{A}^{-1} \mathcal{X} s_{k} .
\end{aligned}
$$

## Covariance

$$
\begin{aligned}
\widehat{\operatorname{Cov}}\left(r_{k}, s_{k}\right) & =\frac{1}{x_{\bullet \bullet}} \sum_{i=1}^{n} \sum_{j=1}^{p} x_{i j} s_{k j} \\
& =r_{k}^{\top} \mathcal{X} s_{k} / x_{\bullet \bullet}=\sqrt{\frac{\lambda_{k}}{x_{\bullet \bullet}}} r_{k}^{\top} \mathcal{A} r_{k} / x_{\bullet \bullet} \\
& =\sqrt{\frac{\lambda_{k}}{x_{\bullet \bullet}}} \widehat{\operatorname{Var}}\left(r_{k}\right) \\
& =\sqrt{\frac{\lambda_{k}}{x_{\bullet \bullet}}} \frac{\lambda_{k}}{x_{\bullet \bullet}}
\end{aligned}
$$

$$
\widehat{\operatorname{Cov}}\left(r_{1}, r_{2}\right)=\frac{1}{x_{\bullet \bullet}} \sum_{i=1}^{n} x_{i} r_{1 i} r_{2 i}=r_{1}^{\top} \mathcal{A} r_{2} / x_{\bullet \bullet}=\delta_{1}^{\top} \mathcal{C}^{\top} \mathcal{C} \delta_{2} / x_{\bullet \bullet}=0
$$

Example: Plot of two pairs of indices

```
gamma2=decomp$u[,2]; delta2=decomp$v[,2]
r2=diag(1/sqrt(a)) %*%C%*%delta2
s2=diag(1/sqrt (b)) %*%t (C) %*%gamma2
```

cumsum(lambda^2)/sum(lambda^2)
plot (NULL, xlim=c $(-1,1), y \operatorname{lim=c}(-1,1), x l a b=" ", y l a b=" ")$
text (r1, r2, labels=row. names (C))
text (s1,s2,labels=colnames (C), col="blue")
library (ca)
calc=ca(alc)
plot(calc)

## Correlation

It follows that the sample correlation coefficient of $r_{k}$ and $s_{k}$ is:

$$
\rho_{k}=\sqrt{\frac{\lambda_{k}}{x_{\bullet \bullet}}}
$$

This means that the correlation structure of the row and column coordinates (in correspondence analysis) is similar to the structure of canonical variables (in canonical correlation analysis).

Example: Analysis of Journaux data set (newspapers in Belgium).
\# load data
data(journaux); $x=$ journaux; $a=\operatorname{rowSums}(x) ; b=\operatorname{colSums}(x)$
e = matrix(a) \%*\% b/sum(a)
\# chi-matrix
$c c=(x-e) / s q r t(e)$
\# singular value decomposition
sv = svd(cc);g = sv\$u;l = sv\$d;d = sv\$v
\# eigenvalues
$11=1 * 1$
\# cumulated percentage of the variance aux $=$ cumsum(ll)/sum(ll); perc $=$ cbind(ll, aux)

## Week 12

 Correspondence analysis\# labels for journals
types =c("va", "vb", "vc", "vd", "ve", "ff", "fg", "fh", "fi", "bj", "bk", "bl", "vm", "fn", "fo")
\# labels for regions
regions =c("brw", "bxl", "anv", "brf", "foc", "for", "hai", "lig", "lim", "lux")
\# plot
plot(rr, type="n", xlim=c(-1.1, 1.5), ylim=c(-1.1, 0.6), xlab="r_1,s_1", ylab="r_2,s_2", main="Journal Data", cex.axis=1.2, cex.lab=1.2, cex.main=1.6)
text(rr, types, cex=1.5, col="blue")
text(ss, regions, col="red"); abline(h=0, v=0, lwd=2)
\#\# library(ca); plot(ca(journaux)); plot3d.ca(ca(journaux))
r1=matrix(l, nrow=nrow(g), ncol=ncol(g), byrow=T) * g r=r1/matrix (sqrt (a), nrow=nrow (g), ncol=ncol (g) , byrow=F) s1=matrix(l, nrow=nrow(d), ncol=ncol(d), byrow=T) * d s=s1/matrix (sqrt (b) , nrow=nrow (d), ncol=ncol (d), byrow=F)
car=matrix(matrix(a), nrow=nrow(r), ncol=ncol(r), byrow=F
) * r^2/matrix(l^2,nrow=nrow(r), ncol=ncol(r), byrow=T) row.names(car)=row.names (x)
cas=matrix(matrix(b), nrow=nrow(s), ncol=ncol(s), byrow=F
) * s^2/matrix(l^2, nrow=nrow(s), ncol=ncol(s), byrow=T) row.names(cas)=colnames (x)
rr=r[, 1:2]; row.names(rr)=row.names(x)
ss=s[, 1:2]; row.names(ss)=colnames (x)

## Example:

data(food); plot(ca(food))
plot3d.ca(ca(food))

## Example:

data(carmean); plot(ca(carmean2))
plot(ca(5-carmean2))

## Example:

data(uscrime); plot3d.ca(ca(uscrime[,3:9]))
?ca \#\#

## Týden 13

* Correspondence analysis investigates dependencies in contingency tables.
* Correlations between row and column coordinates correspond to contributions to $\chi^{2}$ statistic
* The structure of row and column coordinates is similar to canonical variables in canonical correlation analysis.
* Plot of the row and column coordinates displays dependencies in the contingency table.
* The solution allows adding of additional (supplementary) variables that do not influence the calculation of the original coordinates.


## Characteristic function

The characteristic function (CF) of a random vector $X \in \mathbb{R}^{p}$ is:

$$
\varphi_{X}(t)=E\left(e^{\mathbf{i t} t^{\top} x}\right)=\int e^{i t^{\top} x} f(x) d x, \quad t \in \mathbb{R}^{p}
$$

The CF has many interesting and useful properties, e.g.:
(1) The CF always exists, $\varphi_{X}(0)=1$, and $\left|\varphi_{X}(t)\right| \leq 1$.
(2) Two random vectors have the same CF if and only if they have the same distribution. If CF $\varphi_{X}(t)$ is absolutely integrable then $f(x)=\frac{1}{(2 \pi)^{\rho}} \int_{-\infty}^{\infty} e^{-i t^{\top} x} \varphi_{X}(t) d t$.
(3) Random vectors $X_{1}$ and $X_{2}$ are independent if and only if $\varphi_{X}(t)=\varphi_{X_{1}}\left(t_{1}\right) \varphi_{X_{2}}\left(t_{2}\right)$, where $X=\left(X_{1}^{\top}, X_{2}^{\top}\right)^{\top}$.
(9) CF of the sum of two independent random vectors $X$ and $Y$ is the product $\varphi_{X}(t) \varphi_{Y}(t)=\varphi_{X+Y}(t)$.

Obecnější mnohorozměrná rozdělení:

- sférická a eliptická rozdělení,
- kopule.

Kvantily mnohorozměrných rozdělení:

- hloubka.

Směrová data.

## Spherical distributions

Definition: $\mathrm{A}(p \times 1)$ random vector $Y$ is said to have a spherical distribution $S_{p}(\phi)$ if its characteristic function $\psi_{Y}(t)$ satisfies: $\psi_{Y}(t)=\phi\left(t^{\top} t\right)$ for some scalar function $\phi($.$) (the characteristic generator$ of the spherical distribution). We will write $Y \sim S_{p}(\phi)$.

Clearly, $\varphi_{X_{1}}\left(t_{1}\right)=\varphi_{X}\left(t_{1}, 0, \ldots, 0\right)$. This implies that all marginal distributions of a spherical distribution are identical (and symmetric)

Example: The multivariate t-distribution. Let $Z \sim N_{p}\left(0, \mathcal{I}_{p}\right)$ and $S \sim \chi_{m}^{2}$ be independent. The random vector

$$
Y=\sqrt{m} \frac{Z}{S}
$$

has a multivariate $t$-distribution with $m$ degrees of freedom.

## Week 13

Elliptical distributions

## Elliptical distributions

The characteristic function of elliptically symmetric $X$ is of the form

$$
\psi(t)=e^{i t^{\top} \mu} \phi\left(t^{\top} \Sigma t\right)
$$

for a scalar function $\phi$.
Marginal distributions of elliptically distributed variables are elliptical.
The assumption that the returns on all assets available for portfolio formation are jointly elliptically distributed is used in portfolio theory (multinormal distribution of returns usually does not work).

Clearly, the contours of a spherical distribution are $p$-dimensional spheres and contours of an elliptical distribution are $p$-dimensional ellipsoids (if the density exists).

## Elliptical distributions

Definition: $\mathrm{A}(p \times 1)$ random vector $X$ has an elliptical distribution with parameters $\mu(p \times 1)$ and $\Sigma(p \times p)$ if $X$ has the same distribution as $\mu+\mathcal{A}^{\top} Y$, where $Y \sim S_{k}(\phi)$ and $\mathcal{A}$ is a $(k \times p)$ matrix such that $\mathcal{A}^{\top} \mathcal{A}=\Sigma$ with $\operatorname{rank}(\Sigma)=k$. We shall write $X \sim E C_{p}(\mu, \Sigma, \phi)$

The elliptical distribution can be seen as an extension of $N_{p}(\mu, \Sigma)$.
Example: The CF of standard multinormal distribution is
$\varphi_{Y}(t)=e^{-t^{\top} t / 2}$ and it is spherically symmetric with the characteristic generator $\exp (-x / 2)$. The CF of $X=\mu+\mathcal{A}^{\top} Y$ is
$\varphi_{X}(t)=e^{i t^{\top} \mu-t^{\top} \mathcal{A}^{\top} \mathcal{A} t / 2}$ because $t^{\top}\left(\mu+\mathcal{A}^{\top} Y\right)$ has univariate normal distribution and

$$
E e^{\mathrm{it} \boldsymbol{\top}^{\top}\left(\mu+\mathcal{A}^{\top} Y\right)}=\left.\varphi_{N\left(t^{\top} \mu, t^{\top} \mathcal{A}^{\top} \mathcal{A} t\right)}(s)\right|_{s=1}=\left.e^{i t^{\top} \mu s-t^{\top} \mathcal{A}^{\top} \mathcal{A} t s / 2}\right|_{s=1}
$$

## Spherical and elliptical distributions

* The characteristic function is always defined and it uniquely determines the probability distribution.
* An arbitrary function $\phi: \mathbb{R}^{n} \rightarrow C$ is the characteristic function of some random variable if and only if $\phi$ is positive definite, continuous at the origin, and if $\phi(0)=1$ (Bochner's theorem)
* Spherical distribution can be seen as a generalization of $N_{p}\left(0, \mathcal{I}_{p}\right)$, elliptical distributions generalize $N_{p}(\mu, \Sigma)$.
* Elliptical distributions can also be defined in terms of their density functions (if it exists): $f(x)=k \cdot g\left((x-\mu)^{\prime} \Sigma^{-1}(x-\mu)\right.$ ) for some density $g($.$) .$


## Copula

## A copula allows a generalized representation of (complicated)

 dependencies between random variables (risk factors).The basic idea is to describe the joint distribution of a random variable $X=\left(X_{1}, \ldots, X_{p}\right)^{\top}$ using a function $C:[0,1]^{p} \rightarrow[0,1]:$

$$
F\left(x_{1}, \ldots, x_{p}\right)=C\left(F_{1}\left(x_{1}\right), \ldots, F_{p}\left(x_{p}\right)\right),
$$

where $F_{1}, \ldots, F_{p}$ represent the marginal cumulative distributions function of the variables $X_{j}, j=1, \cdots, p$.

A copula C typically depends on some "tuning parameters" determining the dependence.

## Week 13

 CopulasExample: The product copula $\Pi$ : two random variables $X_{1}$ and $X_{2}$ are independent if and only if

$$
H\left(x_{1}, x_{2}\right)=F_{1}\left(x_{1}\right) \cdot F_{2}\left(x_{2}\right)
$$

Hence, the so-called product copula $C=\Pi$ is given by:

$$
\Pi\left(u_{1}, \cdots, u_{p}\right)=\prod_{j=1}^{p} u_{p}
$$

Example: Gaussian or normal copula:
$C_{\rho}\left(u_{1}, u_{2}\right)=\int_{-\infty}^{\Phi_{1}^{-1}\left(u_{1}\right)} \int_{-\infty}^{\Phi_{2}^{-1}\left(u_{2}\right)} \varphi_{\rho}\left(r_{1}, r_{2}\right) d r_{2} d r_{1}=\Phi_{\rho}\left\{\Phi_{1}^{-1}\left(u_{1}\right), \Phi_{2}^{-1}\left(u_{2}\right)\right\}$, where $\varphi_{\rho}$ denotes the bivariate normal density function with correlation $\rho$ and $\Phi_{j}, j=1,2$ represent the gaussian marginal distribution (GOOGLE: gaussian copula financial crisis).

## Two-dimensional copula

Definition: A two-dimensional copula is a function $C:[0,1]^{2} \rightarrow[0,1]$ with the following properties:

- For every $u \in[0,1] C(0, u)=C(u, 0)=0$ (grounded function).
- For every $u \in[0,1]: C(u, 1)=u$ and $C(1, u)=u$ (uniform marginals).
- For every $\left(u_{1}, u_{2}\right),\left(v_{1}, v_{2}\right) \in[0,1] \times[0,1]$ with $u_{1} \leq v_{1}$ and $u_{2} \leq v_{2}$ : $C\left(v_{1}, v_{2}\right)-C\left(v_{1}, u_{2}\right)-C\left(u_{1}, v_{2}\right)+C\left(u_{1}, u_{2}\right) \geq 0$ (two-increasing).

Theorem: (Sklar) Every joint distribution function $H($.$) with marginal$ distributions $F_{1}($.$) and F_{2}($.$) can be expressed as:$

$$
H\left(x_{1}, x_{2}\right)=C\left(F_{1}\left(x_{1}\right), F_{2}\left(x_{2}\right)\right)
$$

and the copula $C($.$) is unique when F_{1}($.$) and F_{2}($.$) are continuous.$

## Week 13

Copulas

Example: An important class of copulas is the Gumbel-Hougaard family:

$$
C_{\theta}\left(u_{1}, u_{2}\right) \stackrel{\text { def }}{=} \exp \left\{-\left[\left(-\ln u_{1}\right)^{\theta}+\left(-\ln u_{2}\right)^{\theta}\right]^{1 / \theta}\right\}
$$

For $\theta=1$ we obtain the product copula: $C_{1}\left(u_{1}, u_{2}\right)=\Pi\left(u_{1}, u_{2}\right)=u_{1} u_{2}$. For $\theta \rightarrow \infty$ we obtain the so-called minimum copula:

$$
C_{\theta}\left(u_{1}, u_{2}\right) \longrightarrow \min \left(u_{1}, u_{2}\right) \stackrel{\text { def }}{=} M\left(u_{1}, u_{2}\right)
$$

(that dominates every other copula; $M($.$) is therefore referred to as the$ Fréchet-Hoeffding upper bound.

The two-dimensional function $W\left(u_{1}, u_{2}\right) \stackrel{\text { def }}{=} \max \left(u_{1}+u_{2}-1,0\right)$ satisfies $W\left(u_{1}, u_{2}\right) \leq C\left(u_{1}, u_{2}\right)$ for all copulas and is called the Fréchet-Hoeffding lower bound.

## Copulas

＊Copulas provide a very flexible way of describing dependencies between random variables．Mathematically，a copula is a multivariate probability distribution function for which the marginal probability distribution of each variable is uniform．
＊Copulas are popular in high－dimensional statistical applications as they allow to model and estimate the distribution of random vectors by estimating marginals and copulae separately．
＊There are many parametric copula families available，which usually have parameters that control the strength of dependence． Archimedean copulas are defined by $\psi_{\theta}^{[-1]}\left(\psi_{\theta}\left(u_{1}\right)+\cdots+\psi_{\theta}\left(u_{d}\right)\right)$ ， where $\psi_{\theta}($.$) is a generator function．This allows modeling of$ dependence in high dimensions with only one parameter $(\theta)$ ．
＊More information can be found in［Nelsen，R．B．（1999）．An Introduction to Copulas，Springer，New York．］

## Quantiles in more dimensions

THE AIM IS to generalize the definition of＂one－dimensional＂quantiles to more dimensional data sets．

The definition of quantiles in 1D uses the order of observations（but observations in more dimensions are not clearly ordered）．

In order to define some kind of＂ordering＂we may define depth function－a measure of＂how deep in the dataset＂is some point．

The $p$－dimensional data set will be ordered from outside to inside instead from left to right．

## Multivariate quantiles

The contours（iso－density regions）for multinormal（and elliptical） distributions are ellipsoids that can be understood as multivariate generalization of quantiles．Unfortunately，defining multivariate quantiles in general is very complicated．

Notice that：
－Median and other quantiles are naturally defined for 1－dim random variable BUT definition of quantile（apart of multinormal or elliptical distribution）is not straigtforward，
－Median：measure of location，＂most central＂point．
－Quantiles：testing，construction of prediciton regions，boxplots．．．

Z．Hlávka（KPMS）NMST539 $338 / 413$
Motivation Week 13 Depth

## Example：boxplot（carmean2）

The standard definition of boxplot is based on sample quantiles（median and quartiles）that are not naturally defined in two and more dimensions．

Boxplot（defined without quantiles，by using the inside $\times$ outside ordering）：
Central box contains $1 / 2$ of＂most central＂observations．
Whiskers extend to most extreme observations by（at most） $1.5 \times$＂central box＂．
Outliers are＂too far away＂from the centre．

Depth
$\square$ dimensions?

Example: $X$ is one-dimensional random variable with d.f. $F_{X}($.$) .$

$$
D(x)=1-2|F(x)-1 / 2|
$$

Deepest point: $F(x)=1 / 2$ (median).
Point with min. depth: $F(x)=0$ (extremes).
Note: for a random sample $X_{1}, \ldots, X_{n}$, we define sample version of the depth function as $D(x)=1-2\left|F_{n}(x)-1 / 2\right|$.

Some desired properties of (sample) depth functions:

- Depth should not depend on the coordinate system (rotation and scale invariance).
- If a distribution is symmetric around $s$ then $s$ is the deepest point.
- Decreasing along rays from the deepest point.
- Vanishing at infinity, i.e., $D(x) \longrightarrow 0$ if $\|x\| \longrightarrow \infty$.
- Quasi-concavity (level sets of depth function are convex).

More details: Liu (1990), Serfling (2000).
Example: pairs(carmean2); How to find deepest point in more
Technically: for a given random vector $X \in \mathbb{R}^{p}$ (with distribution function $\left.F_{X}\right)$ depth is a function $D: \mathbb{R}^{p} \longrightarrow \mathbb{R}$.

Region $R(a)$ with given depth a in $\mathbb{R}^{p}$ is $\left\{x \in \mathbb{R}^{p}: D(x) \geq a\right\} \ldots$ the border of the region $R(a)$ is the a-depth contour (and this is the "multivariate contour").

## Depth functions

Popular depth functions:

- Simplicial depth (Liu depth).
- Halfspace depth (Tukey).

Implementation in R: library(depth), commands: perspdepth, isodepth, depth...

Other approaches: convex hull peeling, zonoids, L1-depth, location-scale depth, and many other.

## Simplicial depth

The simplicial depth (or Liu depth) of a data point $x$ is defined as the number of convex hulls formed from all possible selections of $p+1$ points covering $x$ (convex hull of $p+1$ points $=$ simplex)

The multivariate median (the deepest point) may be defined as the point with the largest simplicial depth, i.e.

$$
x_{\text {med }}=\arg \max _{i} \#\left\{k_{0}, \ldots, k_{p} \in\{1, \ldots, n\}: x_{i} \in \operatorname{hull}\left(x_{k_{0}}, \ldots, x_{k_{p}}\right)\right\} .
$$

in 1D: closed intervals given by 2 points $\left[x_{i}, x_{j}\right.$ ],
in 2D: triangles given by 3 points,
in 3D: "pyramids" given by 4 points etc

Liu depth in $R$

## library(depth)

perspdepth(carmean2[,1:2], method="Liu")
d=perspdepth(carmean2[,1:2], method="Liu", output=TRUE)
contour (d)
text(carmean2[,1:2], rownames(carmean2))
$\rightarrow$ MVAsimdep1
Simplicial Depth Example


Depth

## Halfspace depth

The (sample) halfspace depth of point $x$ (with respect to sample points $\left.x_{1}, \ldots, x_{n}\right)$ is defined as the minimum number of sample points on one side of a hyperplane through the point $x$.

In other words, minimum number of sample points in a halfspace containing the point $x$.

Example: 1D, points are lying on real line...
Example: 2D
isodepth(carmean2[,1:2],mustdith=TRUE)
text(carmean2[,1:2], rownames(carmean2))

## Halfspace depth

## Example: 3D

```
for (i in 1:nrow(carmean2)) {
    print(rownames(carmean2) [i])
    print(depth(carmean2[i,1:3], carmean2[,1:3]))
}
```

Example: 8D, almost all point are "outside" (see also Ggobi).

NMST539 $349 / 413$

## Week 13 Depth

## Mid Depth

* Depth can be seen as a multivariate generalization of (empirical) quantile.
* The most popular depth functions are simplicial (Liu) depth and halfspace (Tukey) depth but many other depth functions have been proposed [D. Hlubinka: Výpravy do hlubin dat, Robust 2008, http://www.karlin.mff.cuni.cz/~hlubinka/soubory/robust08.pdf; D. Hlubinka: O kvantilech ve více rozměrech, Robust 2002, http://www.statspol.cz/oldstat/robust/2002_hlubinka.pdf].
* Most depth functions are computationally intensive.
* Using depth, it is possible to define bagplot as a two-dimensional generalization of the boxplot.


## Distribution function and density

Let $\theta$ be a random angle. Its distribution function $F$ is given by

$$
F(\theta)=P(0<\theta \leq \theta), \quad 0<\theta \leq 2 \pi
$$

and

$$
F(\theta+2 \pi)-F(\theta)=1, \quad-\infty<\theta \leq \infty
$$

Let the distribution function $F$ of random angle $\theta$ be absolutely continuous. Then for the probability density function $f$ of random angle $\theta$, the following holds:
(1) $f(\theta) \geq 0$ almost everywhere on $(-\infty, \infty)$,
(2) $f(\theta+2 \pi)=f(\theta)$ almost everywhere on $(-\infty, \infty)$,
(3) $\int_{0}^{0+2 \pi} f(\theta) d \theta=1$ and $\int_{x}^{x+2 \pi} f(\theta) d \theta=1$.

## Week 13 <br> Directional data

Denoting $(\bar{R}, \bar{\theta})$ the polar coordinates of $\bar{X}$, we obtain:
Definition: The sample mean resultant length $\bar{R} \geq 0$ is given by

$$
\bar{R}=\|V\|=\sqrt{\bar{C}^{2}+\bar{S}^{2}}
$$

If $\bar{R}>0$, the sample mean direction $\bar{\theta}$ is defined as follows:

$$
\bar{\theta}=\arctan ^{*}(\bar{S} / \bar{C})= \begin{cases}\arctan (\bar{S} / \bar{C}), & \text { if } \bar{C}>0, \bar{S} \geq 0, \\ \pi / 2 & \text { if } \bar{C}=0, \bar{S}>0, \\ \arctan (\bar{S} / \bar{C})+\pi, & \text { if } \bar{C}<0, \\ \arctan (\bar{S} / \bar{C})+2 \pi, & \text { if } \bar{C} \geq 0, \bar{S}<0, \\ \text { undefined } & \text { if } \bar{C}=0, \bar{S}=0 .\end{cases}
$$

Definition: The sample circular variance $V$ is defined as

$$
V=1-\bar{R}, \quad 0 \leq V \leq 1 .
$$

von Mises distribution (for angle $\theta$ )
von Mises distribution (on the unit circle)


## 1 <br> Directional statistics

* See [Mardia, K. V. and Jupp, P. E. (2000). Directional statistics. 2nd ed., Wiley Series in Probability and Statistics. Wiley, Chichester] or [Malá, O. C. (2012) Fisherovo-Binghamovo rozdělení, bakalăřská práce, MFF UK] for more information.
* Directional (or axial) data are encountered in various fields: geology, meteorology, astronomy, geography, medicine and others.
* von Mises distribution can be generalized to more dimensions (Fisher-Bingham distribution, von Mises-Fisher distribution, etc.)
* One can also consider distributions defined on more general 'surfaces' (manifolds).

Z. Hlávka (KPMS)

NMST539
358/413

## Týden 13

- jádrové odhady hustoty,
- projection pursuit.

Histograms

The histogram counts relative frequencies of observations $x_{i}$ falling into predefined bins:

$$
\widehat{f}_{h}(x)=n^{-1} h^{-1} \sum_{j \in \mathbb{Z}} \sum_{i=1}^{n} \mathrm{I}\left\{x_{i} \in B_{j}\left(x_{0}, h\right)\right\} \mathrm{I}\left\{x \in B_{j}\left(x_{0}, h\right)\right\}
$$

- the histogram is a simple estimator of a probability density,
- $h$ is a smoothing parameter and controls the width of the histogram bins.

Kernel density estimators

Kernel density estimator is a natural generalization of a histogram (by shifting the "bin", we obtain smooth estimator of the underlying probability density).

Assume we have $n$ independent observations $x_{1}, \ldots, x_{n}$ from the random variable $X$. The kernel density estimator $\widehat{f}_{h}(x)$ for the estimation of the density value $f(x)$ at point $x$ is defined as

$$
\widehat{f}_{h}(x)=\frac{1}{n h} \sum_{i=1}^{n} K\left(\frac{x_{i}-x}{h}\right),
$$

where $K($.$) denotes a kernel function and h$ the bandwidth.
library(sm); library(MSES); data(athletic)
\# univariate kernel density estimator
plot(density(athletic[,"100m"]))
plot(density(athletic[,"Marathon"]))
\#\# bivariate kernel density estimator
library(MASS)
plot(athletic[,"Marathon"], athletic[,"100m"])
d1=kde2d(athletic[,"Marathon"], athletic[,"100m"])

```
image(d1, zlim = c(0, 0.13))
```

persp(d1, phi $=30$, theta $=20, \mathrm{~d}=5$ )
contour (d1)
\# add original points
points(athletic[,"Marathon"], athletic[,"100m"])
identify(athletic[,"Marathon"],athletic[,"100m"]
label=row.names (athletic))
. Hâvka (KPMS) NMST539


Projection pursuit

## Projection pursuit

Projection pursuit searches for interesting directions in a $p$-dimensional data set by maximizing a chosen index.

Exploratory projection pursuit: look for interesting linear combinations-"interestingness" is usually defined by some measure (index) of non-normality.

Projection pursuit regression: the goal is to estimate regression function $m(x)=E(Y \mid x)$ using approximating function $\widehat{f}(x)=\sum \widehat{g}_{k}\left(\Lambda_{k}^{\top} x\right)$ (obviously, lower dimensional projections defined by $\Lambda_{k}$ improve statistical properties of the nonparametric regression estimator).

## Kernel density estimators

* KDEs are sometimes introduced as "average shifted histograms" (ASH).
* High-dimensional KDEs suffer from "curse of dimensionality" because the optimal MISE is of order $O\left(n^{-4 /(d+4)}\right)$, where $d$ denotes the dimension.
* The various implementations of KDEs in R are not mutually compatible (for example, the bandwidth parameter used by one $R$ function typically does not have exactly the same meaning in other $R$ function).
* One should consider dimension reduction techniques before calculating high-dimensional KDEs.
Week 14 Projection pursuit

Given $p$-dimensional random vector $X$ with zero mean (and typically with unit variance, i.e., $\operatorname{Var}(X)=\mathcal{I}_{p}$ ), we try to find $\alpha \in \mathbb{R}^{p}$ such that $\alpha^{\top} X$ is "interesting".

Interestingness of projections $\alpha^{\top} X$ is measured by index $I(\alpha)$.
Example: PCA: $I(\alpha)=\operatorname{var}\left(\alpha^{\top} X\right)$ works only if the data set is not sphered.

In practice, we have the data matrix $\mathcal{X}$ and we optimize the (sample) projection pursuit index numerically.

## Friedman and Tukey index

Let $\widehat{f}_{h, \alpha}(z)$ denote the kernel density estimator of the pdf of the projection $Z=\alpha^{\top} X$, where $h$ denotes the bandwidth.

Friedman and Tukey (1974) proposed the index:

$$
I_{F T, h}(\alpha)=n^{-1} \sum_{i=1}^{n} \widehat{f}_{h, \alpha}\left(\alpha^{\top} X_{i}\right)
$$

that can be rewritten as $I_{F T, h}(\alpha)=\int \widehat{f}_{h, \alpha}(z) d F_{N}(z)$ (i.e., it estimates $\left.\int f(z) d F(z)=\int f^{2}(z) d z\right)$ leading to the maximization of $\int f^{2}(z) d z$.

The Friedman-Tukey index is minimal for a parabolic density and, by its maximization, we search for a distribution that is as far from the parabolic density as possible.

## Entropy index

An alternative approach is based on the (minus) entropy measure $\int f(z) \log f(z) d z$ leading to the entropy index:

$$
I_{E, h}(\alpha)=n^{-1} \sum_{i=1}^{n} \log \left\{\widehat{f}_{h, \alpha}\left(\alpha^{\top} X_{i}\right)\right\}
$$

that can be interpreted as an estimator of minus entropy $\int f(z) \log f(z) d z$.
The index is minimal for normal distribution and maximization of $I_{E, h}(\alpha)$ leads to non-normal projections.


The least and the most informative from 10000 randomly chosen directions (FT index) for Swiss bank notes $\rightarrow$ SMSeppbank.

## Jones and Sibson index

Jones and Sibson (1987) suggested to approximate the entropy index by a moment-based index:

$$
I_{J S}(\alpha)=\left\{\kappa_{3}^{2}\left(\alpha^{\top} X\right)+\kappa_{4}^{2}\left(\alpha^{\top} X\right) / 4\right\} / 12
$$

where $\kappa_{3}\left(\alpha^{\top} X\right)=E\left\{\left(\alpha^{\top} X\right)^{3}\right\}$ and $\kappa_{4}\left(\alpha^{\top} X\right)=E\left\{\left(\alpha^{\top} X\right)^{4}\right\}-3$ are cumulants of $\alpha^{\top} X$ (skewness and kurtosis).

The maximization of $I_{J S}(\alpha)$ also leads to the least-normal view of the data set.

## Computational aspects

The optimal projection $\alpha \in \mathbb{R}^{p}$ can be found by standard (iterative) optimization routines

The optimization task is not very simple because the parameter $\alpha$ is $p$-dimensional and the function $I(\alpha)$ has many local maxima.

In practice, one is interested in finding optimal one- and two-dimensional projections.

It is recommended to use various starting points in order to verify the stability of the result. Often, the optimization of $\alpha$ is used to define a guided tour through the data set.

## Exercise: Swiss bank notes

(1) library (SMSdata); data(bank2)
(2) sphering (Mahalanobis transformation),
(3) generate randomly $N$ directions $\alpha_{1}, \ldots, \alpha_{N}$
(9) calculate the value $I\left(\alpha_{i}\right)$ for $i=1, \ldots, N$
(5) plot kde of the directions that both maximize and minimize the chosen index,
(6) compare the result obtained for PCA index with standard PCA analysis (this will work only without sphering),
(1) compare least and most informative projections obtained by JS and FT index,
(8) try to find the optimal direction using numerical optimization in R (optim()), compare results obtained by different algorithms (Nelder-Mead, BFGS, ...).

White noise analysis

White noise projections that are most similar to white noise are identified and discarded while the remaining informative projections are used to look for interesting relationships [Hui and Lindsay, 2010, Projection pursuit via white noise matrices, Sankhya B 72(2), 123--153.]

The White Noise Analysis (WNA) is based on the eigen-analysis of the standardized Fisher information matrix for the square transformed density estimated by the kernel method.

WNA is computationally simpler than the classical Projection Pursuit searching for low-dimensional least-normal projections.


## Ggobi: 1D and 2D guided tour

Guided tour through a multivariate data set is a sequence of low-dimensional projections that improve the chosen index.

```
library(SMSdata)
data(bank2)
library(rggobi) # using ggobi is easy if this works
ggobi(bank2)
write.csv(bank2,file="bank2.dat")
# START GGOBI AND LOAD DATA SET FROM CSV FILE
```


## R: tourr

Guided and grand tours work similarly as in Ggobi (same authors) but R does not allow interaction.
library (SMSdata)
data(bank2)
library(tourr)
animate(bank2, guided_tour(index_f=holes), display_xy()
, sphere=FALSE)
animate(bank2, guided_tour(index_f=cmass), display_xy())

Indices in Ggobi and tourr

## Holes

$$
I_{\text {Holes }}(\alpha)=\frac{1-\frac{1}{n} \exp \left(-z_{i} z_{i}^{\top} / 2\right)}{1-\exp (-p / 2)}
$$

where $z_{i}$ is the $i$-th row of $\mathcal{Z}=\mathcal{X} \alpha$ (the index works also for more dimensional projections).

## Central mass

$$
I_{\mathrm{CM}}(\alpha)=\frac{\frac{1}{n} \exp \left(-z_{i} z_{i}^{\top} / 2\right)-\exp (-p / 2)}{1-\exp (-p / 2)}
$$

is basically the opposite of $I_{\text {Holes }}$
Both indices are based on $\frac{1}{n} \exp \left(-z_{i} z_{i}^{\top} / 2\right)=\int \exp \left(-z_{i} z_{i}^{\top} / 2\right) d F_{n}(z)$ estimating $E \exp \left(-Z^{\top} Z / 2\right)$.

## R: tourr

\# tourr SHOULD work also with other type of graphics
animate_dist(bank2[95:106,], guided_tour(index_f=holes)) animate_image(bank2[95:106,],guided_tour(index_f=holes)) animate_pcp(bank2[95:106,], guided_tour(index_f=holes)) animate_scatmat(bank2[95:106,],guided_tour(index_f=holes)) animate_faces(bank2[95:106,],guided_tour(index_f=holes)) animate_stars(bank2[95:106,],guided_tour(index_f=holes)) animate_stereo(bank2[95:106,],guided_tour(index_f=holes)) animate_trails(bank2[95:106,],guided_tour(index_f=holes))

## Central mass and Holes

Cook, Buja, Cabrera (1993) Projection Pursuit Indexes Based on Orthonormal Function Expansions. Journal of Computational and Graphical Statistics 2(3), 225-250.

The derivation of both indeces is based on Fourier expansion of density function:

$$
f(x)=\sum_{i=0}^{\infty} a_{i} p_{i}(x)
$$

where $p_{i}(x)$ are (standardized) orthonormal polynomials with weight function $w(x)$ and $a_{i}=\left\langle f, p_{i}\right\rangle=\int f(x) p_{i}(x) w(x) d x$ are Fourier coefficients.

Fourier approximation of normal density


Fourier approximation from 100 observations


## Distance from Normal distribution

Cook, Buja \& Cabrera (1993) investigate the sample natural Hermite index for $M=0$, i.e., $I_{N, 0}=\left(\widehat{a}_{0}-b_{0}\right)^{2}$.

Clearly, the quadratic function $\left(\hat{a}_{0}-b_{0}\right)^{2}$ achieves its minimum $\widehat{a}_{0}-b_{0}$ and is maximized by extreme values of $a_{0}$.

Cook, Buja \& Cabrera (1993) show that, in a family of distributions with mean zero and variance at most one, $a_{0}$ is minimized by the central hole distribution:

$$
P(X=1)=0.5, \quad P(X=-1)=0.5
$$

and $a_{0}$ is maximized by the central mass distribution:

$$
P(X=0)=1
$$

## Original data matrix $\mathcal{X}$.

Sphered data matrix $\mathcal{Y}=\left(\mathcal{X}-\mathbf{1}_{n} \bar{x}^{\top}\right) \mathcal{S}^{-1 / 2}$.
Interesting linear combinations are:

$$
\mathcal{Y} \alpha=\left(\mathcal{X}-\mathbf{1}_{n} \bar{X}^{\top}\right) \mathcal{S}^{-1 / 2} \alpha=\mathcal{X} \mathcal{S}^{-1 / 2} \alpha+\text { const }=\mathcal{X} \alpha_{X}+\text { const. }
$$

## Central mass and holes

The central mass index in Ggobi looks for the rotation $\alpha$ maximizing $a_{0}(\alpha)$.
The holes index in Ggobi looks for the rotation $\alpha$ maximizing $-a_{0}(\alpha)$.
Distributions with very small or very large $a_{0}$ should have large distance (natural Hermite index) from Normal distribution.

Switching repeatedly between maximization of these two indeces leads to informative displays of the data set.

Example: Swiss bank notes in Ggobi and R (tourr).
Visual inference
Chow 14 Projection pursuit
statistical inference to better understand random class separations in high
dimensions, low sample size data, Computational Statistics 30: 293-316.
[The paper can be found using scholar.google.com.]
The problem: lower dimensional projections (especially based on LDA) can
be misleading (see Figure 1).
Example:
d=data.frame(matrix (rnorm(150), ncol=10))
animate (d,guided_tour (index_f=holes), display_xy(), sphere=TRUE)
Z. Hlávka (KPMS)

Chowdhury, Cook, Hofmann, Majumder, Lee \& Toth (2015) Using visual statistical inference to better understand random class separations in high [The paper can be found using scholar.google.com.]

The problem: lower dimensional projections (especially based on LDA) can be misleading (see Figure 1).

Example:

```
d=data.frame(matrix(rnorm(150),ncol=10))
animate(d,guided_tour(index_f=holes),display_xy(),sphere=TRUE)
```


## Visual inference

Proposed solution: use visual statistical inference via Amazon's
Mechanical Turk (the original "lived" from 1770-1854).


## Týden 14

## Gentle introduction

- kernel regression estimators,
- additive models,
- projection pursuit regression.

Sliced inverse regression:

- kernel regression estimators,
- additive models and projection pursuit regression,
- inverse regression curve,
- SIR,
- SIR II.
Z. Hlávka (KPMS)
generate three independent samples with the same $p$-variate distribution (using, e.g., rmvnorm(mvtnorm)),
(2) use function animate (tourr) to find interesting projections (preferably using the lda_pp index),
(3) plot the resulting projections (take care about scaling) denoting the three groups by different symbol-can you see some differences?repeat the simulation both for small and high dimension.

Exercise: simulated data set

## Kernel regression estimators

Suppose that we have independent observations $Y_{1}, \ldots, Y_{n}$ and the explanatory variable $X_{1}, \ldots, X_{n}$. The Nadaraya-Watson kernel regression estimator is defined as:

$$
\widehat{m}_{h}(x)=\frac{\sum_{i=1}^{n} K\left(\frac{x_{i}-x}{h}\right) Y_{i}}{\sum_{i=1}^{n} K\left(\frac{x_{i}-x}{h}\right)}=\frac{1}{n} \sum_{i=1}^{n} W_{h i}(x) Y_{i}
$$

It can be shown that the asymptotic MSE is:

$$
\operatorname{AMSE}(n, h)=\frac{1}{n h} C_{1}+h^{4} C_{2}
$$

where $C_{1}$ and $C_{2}$ are constants depending on the kernel function, the (derivatives of) the regression function and the density of $X$. Using the optimal bandwidth $h=C_{3} n^{-1 / 5}$, AMSE is of order $O\left(n^{-4 / 5}\right)$.

## Kernel regression estimates in R :

Typically 1 response and 1 or 2 explanatory variables
1D ksmooth(), locpoly(KernSmooth)
2D sm.regression(sm)

```
library(sm); library(MSES); data(athletic)
# most simple univariate kernel regression estimates
# (better functions exist in other libraries)
plot(athletic[,"Marathon"],athletic[,"100m"])
lines(ksmooth(athletic[,"Marathon"],athletic[,"100m"],
    kernel="normal", bandwidth=20),col="red",lwd=2)
```

library(KernSmooth)
plot(athletic[,"Marathon"],athletic[,"100m"])
lines (locpoly (athletic [, "Marathon"], athletic [, "100m"], bandwidth=10),
col="red", lwd=2)

## Week 14 Kernel regression

\#\# bivariate kernel regression
library(sm)
sm.regression(athletic[,c("Marathon", "400m")],athletic[,"100m"])
sm.regression(athletic[,c("Marathon","400m")], athletic[,"100m"],
display="image")

The asymptotic properties of the kernel regression estimator are bad for high-dimensional explanatory variable (curse of dimensionality). Moreover it is difficult to plot the resulting estimator for $p>2$.
sm.regression(athletic[,"Marathon"], athletic[,"100m"])
\#\# bivariate kernel density estimator
library(MASS)
plot(athletic[,"Marathon"], athletic[,"100m"])
d1=kde2d(athletic[,"Marathon"], athletic[,"100m"])
image (d1, zlim $=c(0,0.13))$
persp(d1, phi $=30$, theta $=20, d=5$ )
contour (d1)
\# add original points
points(athletic[,"Marathon"], athletic[,"100m"])
\# add kernel regression line
lines(ksmooth(athletic[,"Marathon"], athletic[,"100m"],
kernel="normal", bandwidth=20), col="red", lwd=2)

## Week 14 Kernel regression <br> Curse of dimensionality (from Wikipedia)

One way to illustrate the "vastness" of high-dimensional Euclidean space is to compare the proportion of an inscribed hypersphere with radius $r$ and dimension $d$, to that of a hypercube with edges of length $2 r$. The volume of such a sphere is: $\frac{2 r^{d} \pi^{d / 2}}{d \Gamma(d / 2)}$. The volume of the cube would be: $(2 r)^{d}$. As the dimension $d$ of the space increases, the hypersphere becomes an insignificant volume relative to that of the hypercube. This can clearly be seen by comparing the proportions as the dimension $d$ goes to infinity: $\frac{\pi^{d / 2}}{d 2^{d-1} \Gamma(d / 2)} \rightarrow 0$ as $d \rightarrow \infty$. Furthermore, the distance between the center and the corners is $r \sqrt{d}$, which increases without bound for fixed $r$.

In this sense, nearly all of the high-dimensional space is "far away" from the centre. To put it another way, the high-dimensional unit hypercube can be said to consist almost entirely of the "corners" of the hypercube with almost no "middle".

## Additive model

In order to avoid the curse of dimensionality, it can be useful to consider additive model (AM) with response $p$-dimensional explanatory variable $X$ :

$$
E(Y \mid X=x)=\sum_{j=1}^{p} f_{j}\left(x_{j}\right)+c
$$

where $c=E(Y)$ and the (univariate) additive components are centered, i.e., $E\left\{f_{j}\left(X_{j}\right)\right\}=0$ for $1 \leq j \leq p$.

The components of the additive model (and its various generalizations) are usually estimated by iterative algorithms (backfitting).

## Week 14 <br> Sliced inverse regression

## Sliced inverse regression

Sliced Inverse Regression (SIR) is a dimension reduction technique that can be described as a generalization of projection pursuit regression.

The idea is to find EDR-directions (i.e., projections of explanatory variables) suitable for nonparametric regression estimator for the response.

Given a response variable $Y$ and a (random) vector $X \in \mathbb{R}^{p}$ of explanatory variables, SIR is based on the model:

$$
Y=m\left(\beta_{1}^{\top} X, \ldots, \beta_{k}^{\top} X, \varepsilon\right)
$$

where $\beta_{1}, \ldots, \beta_{k}$ are unknown projection vectors

Projection pursuit regression

Projection pursuit regression [Friedman, J.H. and Stuetzle, W. (1981) Projection Pursuit Regression. Journal of the American Statistical Association, 76, 817-823]:

$$
E(Y \mid X=x)=\sum_{j=1}^{r} f_{j}\left(\beta_{j}^{\top} x\right)+c,
$$

applies the additive model on projections of explanatory variables, i.e., it reduces the dimensionality of the space of explanatory variables (keeping in mind that we model the conditional expectation of $Y$ ).

Implementation in R: function $\operatorname{ppr}()$ in library stats.

## Centered inverse regression curve

Recall that $Y=m\left(\beta_{1}^{\top} X, \ldots, \beta_{k}^{\top} X, \varepsilon\right)$.
According to Theorem 20.1 in [Härdle and Simar, Applied Multivariate Statistical Analysis, 4th edition] we have that: "Under some assumptions, the ( $p$-dimensional) centered inverse regression curve $E(X \mid Y=y)-E X$ lies in the linear subspace spanned by $\Sigma \beta_{i}, i=1, \ldots, k, \Sigma=\operatorname{Var} X$."

It follows that for $Z=\Sigma^{-1 / 2}(X-E X)$, the standardized inverse regression curve $m_{1}(y)=E(Z \mid Y=y)$ lies in a linear subspace spanned by $\eta_{i}=\Sigma^{1 / 2} \beta_{i}$.

The idea of SIR algorithm is to generate points lying on the inverse regression curve and then estimate the linear subspace...

## SIR Algorithm (part 1)

The algorithm to estimate the EDR-directions via SIR is as follows:

- Standardize $x$ :

$$
z_{i}=\hat{\Sigma}^{-1 / 2}\left(x_{i}-\bar{x}\right)
$$

- Divide the range of $y_{i}$ in $S$ non-overlapping intervals (slices) Sttk , $s=1, \ldots, S . n_{s}$ denotes the number of observations within slice Sttk and $I_{S t t k_{s}}$ is the indicator function for this slice $\left(n_{s}=\sum_{i=1}^{n} I_{S t t k_{s}}\left(y_{i}\right)\right)$ :
- Compute the mean of $z_{i}$ over all slices. This is a crude estimate $\widehat{m}_{1}$ for the inverse regression curve $m_{1}$ :

$$
\bar{z}_{s}=\frac{1}{n_{s}} \sum_{i=1}^{n} z_{i} I_{S t t k_{s}}\left(y_{i}\right)
$$

- Calculate the estimate for $\operatorname{Var}\left\{m_{1}(y)\right\}$ :

$$
\widehat{V}=n^{-1} \sum_{s=1}^{S} n_{s} \bar{z}_{s} \bar{z}_{s}^{\top}
$$

- Identify the eigenvalues $\hat{\lambda}_{i}$ and eigenvectors $\hat{\eta}_{i}$ of $\widehat{V}$.
- Transform the standardized EDR-directions $\hat{\eta}_{i}$ back to the original scale. Now the estimates for the EDR-directions are given by

$$
\hat{\beta}_{i}=\hat{\Sigma}^{-1 / 2} \hat{\eta}_{i}
$$




Plot of the true response versus the true indices. The monotonic and the convex shapes can be clearly seen $\rightarrow$ MVAsirdata


SIR algorithm works quite well (although the IR curve may not span the entire EDR space).
Z. Hlávka (KPMS)

SIR II

In some situations SIR does not find EDR directions because the inverse regression curve does not have to span the entire EDR space.

Example: Suppose that $\left(X_{1}, X_{2}\right)^{\top} \sim N\left(0, \mathcal{I}_{2}\right)$ and $Y=X_{2}^{2}$.
Notice that the EDR space is spanned by $\beta_{1}=(0,1)^{\top}$ and the IR curve is $E\left(X_{1} \mid y\right)=E\left(X_{2} \mid y\right)=0$.

SIR II algorithm uses the (inverse) conditional variance $\operatorname{Var}(X \mid y)$ instead of the inverse regression curve. In practice, it is recommended to use SIR and SIR II jointly.


EDR directions for US companies (for market values) $\qquad$
Z. Hlávka (KPMS) NMST539 406 / 413

## Simulated example

Example:Let us simulate a data set with
$X \sim N_{4}\left(0, \mathcal{I}_{4}\right), Y=\left(X_{1}+3 X_{2}\right)^{2}+\left(X_{3}-X_{4}\right)^{4}+\varepsilon$ and $\varepsilon \sim N(0,1)$ and use the SIR and SIR II technique to find the EDR directions.

The true response variable depends on the explanatory variables nonlinearly through the linear combinations $X \beta_{1}=X_{1}+3 X_{2}$ and $X \beta_{2}=X_{3}-X_{4}$, where $\beta_{1}=(1,3,0,0)^{\top}$ and $\beta_{2}=(0,0,3,-4)^{\top}$.

We simulate altogether 200 observations.
$\rightarrow$ SMSsir2simu


SIR and SIR II applied on the simulated data set. Screeplot and scatterplots of first three indices against the response. $\rightarrow$ SMSsir2simu

## $m$ <br> SIR

* SIR serves as dimension reduction tool for regression problems.
* Inverse regression helps to avoid the curse of dimensionality.
* The dimension reduction can be conducted without estimation of the regression function $y=m(x)$.
* SIR searches for the eefective dimension reduction (EDR) by computing the inverse regression IR.
* SIR II bases the EDR on computing the inverse conditional variance.
* In certain circumstances, SIR might miss EDR directions that are found by SIR II.

Modifications

SAVE The algorithm sliced average variance estimates is based on the conditional variance matrix (similarly as SIR II).
pHd The method of principal Hessian directions is based on the Hessian matrix $E\left\{(Z-E Z)(X-E X)(X-E X)^{\top}\right\}$, where the vector $Z$ is given either by the response $Y$ or by the linear model residuals.

R library dr:
method: This character string specifies the method of fitting. The options include "sir", "save", "phdy", "phdres" and "ire".

## Závěr

Opakování a shrnutí:

- shrnutí,
- informace o zkoušce.


## Summary

## Multivariate distributions:

- random vector and its characteristics,
- multinormal, spherical and elliptical distributions, copulas.

Estimation and testing: maximum likelihood techniques.

## Analysis of multivariate data:

- summary statistics, principal components,
- factor analysis, canonical correlations,
- discriminant analysis, cluster analysis,
- correspondence analysis, projection pursuit,
- projection pursuit regression, SIR.

NMST539


[^0]:    bodyx $=\mathrm{m}[1]+\mathrm{c}(-1,1) *$ sqrt $(\mathrm{S}[1,1] * 2 * q \mathrm{f}(0.95,2, \mathrm{n}-2) /(\mathrm{n}-2))$
    bodyy $=m[2]+c(-1,1) *$ sqrt $(S[2,2] * 2 * q f(0.95,2, n-2) /(n-2))$
    polygon(x=bodyx[c(1, $1,2,2,1)], y=$ bodyy $[c(1,2,2,1,1)]$,border="blue")

[^1]:    library(SMDdata);library(MASS); data(carmean2)

    X=cmdscale(dist(carmean2))
    plot(X,type="n")
    text(X,labels=row.names(carmean2))

[^2]:    Week 10-11 Discriminant analysis
    Example:

    ```
    library(MASS);library(MSES); data(bank2)
    1da.b2=lda(bank2,pf<-rep(c("Prave","Fales"),each=100))
    lda.b2
    table(predict(lda.b2,bank2)$class,pf)
    qda.b2=qda(bank2,pf)
    qda.b2
    table(predict(lda.b2,bank2)$class,pf)
    ?lda
    ?qda
    ```

    Note: applying lda() with $x_{i}, x_{i}^{2}$, and $x_{i} x_{j}$ is similar (but not equivalent)
    to qda().

