## Wolfgang Karl Härdle Léopold Simar

## Applied Multivariate Statistical Analysis

## Fourth Edition

Springer

Applied Multivariate Statistical Analysis

# Wolfgang Karl Härdle • Léopold Simar 

## Applied Multivariate Statistical Analysis

Fourth Edition

Wolfgang Karl Härdle<br>C.A.S.E. Centre f. Appl. Stat. \& Econ.<br>School of Business and Economics<br>Humboldt-Universität zu Berlin<br>Berlin, Germany

Léopold Simar<br>Center of Operations Research \&<br>Econometrics (CORE)<br>Katholieke Univeristeit Leuven Inst. Statistics<br>Leuven, Belgium

The majority of chapters have quantlet codes in Matlab or R. These quantlets may be downloaded from http://extras.springer.com or via a link on http://springer.com/978-3-662-45170-0 and from www.quantlet.de

ISBN 978-3-662-45170-0
ISBN 978-3-662-45171-7 (eBook)
DOI 10.1007/978-3-662-45171-7
Library of Congress Control Number: 2015933294
Mathematics Subject Classification (2000): 62H10, 62H12, 62H15, 62H17, 62H20, 62H25, 62H30, 62F25

Springer Heidelberg New York Dordrecht London © Springer-Verlag Berlin Heidelberg 2003, 2007, 2012, 2015
This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.
The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.
The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made.

Printed on acid-free paper
Springer-Verlag GmbH Berlin Heidelberg is part of Springer Science+Business Media (www.springer. com)

## Preface to the Fourth Edition

The fourth edition of this book on Applied Multivariate Statistical Analysis offers a new sub-chapter on Variable Selection by using least absolute shrinkage and selection operator (LASSO) and its general form the so-called Elastic Net.

All pictures and numerical examples have been now calculated in the (almost) standard language R \& MATLAB. The code for each picture is indicated with a small Q sign near the picture, e.g. a MVAdenbank denotes the corresponding quantlet for reproduction of Fig. 1.9, where we display the densities of the diagonal of genuine and counterfeit bank notes. We believe that these publicly available quantlets (see also http://sfb649.wiwi.hu-berlin.de/quantnet/) create a valuable contribution to distribution of knowledge in the statistical science. The symbols and notations have also been standardised. In the preparation of the fourth edition, we received valuable input from Dedy Dwi Prastyo, Petra Burdejova, Sergey Nasekin and Awdesch Melzer. We would like to thank them.

Berlin, Germany
Louvain la Neuve, Belgium
January 2014

Wolfgang Karl Härdle
Léopold Simar

## Preface to the Third Edition

The third edition of this book on Applied Multivariate Statistical Analysis offers the following new features.

1. A new Chap. 8 on Regression Models has been added.
2. Almost all numerical examples have been reproduced in MATLAB or R.

The chapter on regression models focuses on a core business of multivariate statistical analysis. This contribution has not been subject of a prominent discussion in earlier editions of this book. We now take the opportunity to cover classical themes of ANOVA and ANCOVA analysis. Categorical responses are presented in Sect. 8.2. The spectrum of log linear models for contingency tables is presented in Sect. 8.2.2, and applications to count data, e.g. in the economic and medical science are presented there. Logit models are discussed in great detail, and the numerical implementation in terms of matrix manipulations is presented.

The majority of pictures and numerical examples has been now calculated in the (almost) standard language $\mathrm{R} \&$ MATLAB. The code for each picture is indicated with a small 0 sign near the picture, e.g. a MVAdenbank denotes the corresponding quantlet for reproduction of Fig. 1.9, where we display the densities of the diagonal of genuine and counterfeit bank notes. We believe that these publicly available quantlets (see also www.quantlet.com) create a valuable contribution to distribution of knowledge in the statistical science. The symbols and notations have also been standardised. In the preparation of the third edition, we received valuable input from Song Song, Weining Wang and Mengmeng Guo. We would like to thank them.

Berlin, Germany
Louvain la Neuve, Belgium
Wolfgang Karl Härdle
Léopold Simar
June 2011

## Contents

Part I Descriptive Techniques
1 Comparison of Batches ..... 3
1.1 Boxplots ..... 4
1.2 Histograms ..... 11
1.3 Kernel Densities ..... 15
1.4 Scatterplots ..... 19
1.5 Chernoff-Flury Faces ..... 22
1.6 Andrews’ Curves ..... 29
1.7 Parallel Coordinates Plots ..... 32
1.8 Hexagon Plots ..... 37
1.9 Boston Housing ..... 40
1.10 Exercises ..... 48
Part II Multivariate Random Variables
2 A Short Excursion into Matrix Algebra ..... 53
2.1 Elementary Operations ..... 53
2.2 Spectral Decompositions ..... 60
2.3 Quadratic Forms ..... 62
2.4 Derivatives ..... 65
2.5 Partitioned Matrices ..... 66
2.6 Geometrical Aspects ..... 68
2.7 Exercises ..... 76
3 Moving to Higher Dimensions ..... 79
3.1 Covariance ..... 80
3.2 Correlation ..... 84
3.3 Summary Statistics ..... 89
3.4 Linear Model for Two Variables ..... 93
3.5 Simple Analysis of Variance ..... 100
3.6 Multiple Linear Model ..... 105
3.7 Boston Housing ..... 110
3.8 Exercises ..... 113
4 Multivariate Distributions ..... 117
4.1 Distribution and Density Function ..... 118
4.2 Moments and Characteristic Functions ..... 123
4.3 Transformations ..... 135
4.4 The Multinormal Distribution ..... 137
4.5 Sampling Distributions and Limit Theorems ..... 142
4.6 Heavy-Tailed Distributions ..... 149
4.7 Copulae ..... 166
4.8 Bootstrap ..... 176
4.9 Exercises ..... 179
5 Theory of the Multinormal ..... 183
5.1 Elementary Properties of the Multinormal ..... 183
5.2 The Wishart Distribution ..... 191
5.3 Hotelling's $T^{2}$-Distribution ..... 193
5.4 Spherical and Elliptical Distributions ..... 195
5.5 Exercises ..... 197
6 Theory of Estimation ..... 201
6.1 The Likelihood Function ..... 202
6.2 The Cramer-Rao Lower Bound ..... 206
6.3 Exercises ..... 210
7 Hypothesis Testing ..... 213
7.1 Likelihood Ratio Test ..... 214
7.2 Linear Hypothesis ..... 224
7.3 Boston Housing ..... 242
7.4 Exercises ..... 246
Part III Multivariate Techniques
8 Regression Models ..... 253
8.1 General ANOVA and ANCOVA Models ..... 255
8.1.1 ANOVA Models ..... 255
8.1.2 ANCOVA Models ..... 260
8.1.3 Boston Housing ..... 262
8.2 Categorical Responses ..... 263
8.2.1 Multinomial Sampling and Contingency Tables ..... 263
8.2.2 Log-Linear Models for Contingency Tables ..... 264
8.2.3 Testing Issues with Count Data ..... 268
8.2.4 Logit Models ..... 271
8.3 Exercises ..... 279
9 Variable Selection ..... 281
9.1 Lasso ..... 282
9.1.1 Lasso in the Linear Regression Model ..... 282
9.1.2 Lasso in High Dimensions ..... 292
9.1.3 Lasso in Logit Model ..... 293
9.2 Elastic Net ..... 297
9.2.1 Elastic Net in Linear Regression Model ..... 298
9.2.2 Elastic Net in Logit Model ..... 299
9.3 Group Lasso ..... 300
9.4 Exercises ..... 304
10 Decomposition of Data Matrices by Factors ..... 305
10.1 The Geometric Point of View ..... 306
10.2 Fitting the $p$-Dimensional Point Cloud ..... 307
10.3 Fitting the $n$-Dimensional Point Cloud ..... 310
10.4 Relations Between Subspaces ..... 312
10.5 Practical Computation ..... 314
10.6 Exercises ..... 317
11 Principal Components Analysis ..... 319
11.1 Standardised Linear Combination ..... 320
11.2 Principal Components in Practice ..... 324
11.3 Interpretation of the PCs ..... 327
11.4 Asymptotic Properties of the PCs ..... 331
11.5 Normalised Principal Components Analysis ..... 335
11.6 Principal Components as a Factorial Method ..... 336
11.7 Common Principal Components ..... 342
11.8 Boston Housing ..... 346
11.9 More Examples ..... 348
11.10 Exercises ..... 357
12 Factor Analysis ..... 359
12.1 The Orthogonal Factor Model ..... 360
12.2 Estimation of the Factor Model ..... 367
12.3 Factor Scores and Strategies ..... 376
12.4 Boston Housing ..... 378
12.5 Exercises ..... 382
13 Cluster Analysis ..... 385
13.1 The Problem ..... 386
13.2 The Proximity Between Objects ..... 387
13.3 Cluster Algorithms ..... 392
13.4 Boston Housing ..... 400
13.5 Exercises ..... 404
14 Discriminant Analysis ..... 407
14.1 Allocation Rules for Known Distributions ..... 407
14.2 Discrimination Rules in Practice ..... 415
14.3 Boston Housing ..... 421
14.4 Exercises ..... 423
15 Correspondence Analysis ..... 425
15.1 Motivation ..... 426
15.2 Chi-Square Decomposition ..... 428
15.3 Correspondence Analysis in Practice ..... 432
15.4 Exercises ..... 441
16 Canonical Correlation Analysis ..... 443
16.1 Most Interesting Linear Combination ..... 443
16.2 Canonical Correlation in Practice ..... 448
16.3 Exercises ..... 454
17 Multidimensional Scaling ..... 455
17.1 The Problem ..... 455
17.2 Metric MDS ..... 460
17.3 Nonmetric MDS ..... 465
17.4 Exercises ..... 472
18 Conjoint Measurement Analysis ..... 473
18.1 Introduction ..... 473
18.2 Design of Data Generation ..... 475
18.3 Estimation of Preference Orderings ..... 478
18.4 Exercises ..... 485
19 Applications in Finance ..... 487
19.1 Portfolio Choice ..... 487
19.2 Efficient Portfolio ..... 488
19.3 Efficient Portfolios in Practice ..... 496
19.4 The Capital Asset Pricing Model ..... 497
19.5 Exercises ..... 499
20 Computationally Intensive Techniques. ..... 501
20.1 Simplicial Depth ..... 502
20.2 Projection Pursuit ..... 505
20.3 Sliced Inverse Regression ..... 511
20.4 Support Vector Machines ..... 519
20.5 Classification and Regression Trees ..... 534
20.6 Boston Housing ..... 552
20.7 Exercises ..... 554
Part IV Appendix
21 Symbols and Notations ..... 557
22 Data ..... 561
22.1 Boston Housing Data ..... 561
22.2 Swiss Bank Notes ..... 562
22.3 Car Data ..... 562
22.4 Classic Blue Pullovers Data ..... 563
22.5 US Companies Data ..... 563
22.6 French Food Data ..... 563
22.7 Car Marks ..... 564
22.8 French Baccalauréat Frequencies ..... 564
22.9 Journaux Data ..... 564
22.10 US Crime Data ..... 565
22.11 Plasma Data ..... 566
22.12 WAIS Data ..... 566
22.13 ANOVA Data ..... 567
22.14 Timebudget Data ..... 567
22.15 Geopol Data ..... 568
22.16 US Health Data ..... 569
22.17 Vocabulary Data ..... 570
22.18 Athletic Records Data ..... 570
22.19 Unemployment Data ..... 570
22.20 Annual Population Data ..... 570
22.21 Bankruptcy Data I ..... 571
22.22 Bankruptcy Data II ..... 571
References ..... 573
Index ..... 577

## Part I <br> Descriptive Techniques

## Chapter 1 <br> Comparison of Batches

Multivariate statistical analysis is concerned with analysing and understanding data in high dimensions. We suppose that we are given a set $\left\{x_{i}\right\}_{i=1}^{n}$ of $n$ observations of a variable vector $X$ in $\mathbb{R}^{p}$. That is, we suppose that each observation $x_{i}$ has $p$ dimensions:

$$
x_{i}=\left(x_{i 1}, x_{i 2}, \ldots, x_{i p}\right),
$$

and that it is an observed value of a variable vector $X \in \mathbb{R}^{p}$. Therefore, $X$ is composed of $p$ random variables:

$$
X=\left(X_{1}, X_{2}, \ldots, X_{p}\right)
$$

where $X_{j}$, for $j=1, \ldots, p$, is a one-dimensional random variable. How do we begin to analyse this kind of data? Before we investigate questions on what inferences we can reach from the data, we should think about how to look at the data. This involves descriptive techniques. Questions that we could answer by descriptive techniques are:

- Are there components of $X$ that are more spread out than others?
- Are there some elements of $X$ that indicate sub-groups of the data?
- Are there outliers in the components of $X$ ?
- How "normal" is the distribution of the data?
- Are there "low-dimensional" linear combinations of $X$ that show "non-normal" behaviour?

One difficulty of descriptive methods for high-dimensional data is the human perceptional system. Point clouds in two dimensions are easy to understand and to interpret. With modern interactive computing techniques we have the possibility to see real time 3D rotations and thus to perceive also three-dimensional data. A "sliding technique" as described in Härdle and Scott (1992) may give insight
into four-dimensional structures by presenting dynamic 3D density contours as the fourth variable is changed over its range.

A qualitative jump in presentation difficulties occurs for dimensions greater than or equal to 5 , unless the high-dimensional structure can be mapped into lower-dimensional components (Klinke \& Polzehl, 1995). Features like clustered sub-groups or outliers, however, can be detected using a purely graphical analysis.

In this chapter, we investigate the basic descriptive and graphical techniques allowing simple exploratory data analysis. We begin the exploration of a data set using boxplots. A boxplot is a simple univariate device that detects outliers component by component and that can compare distributions of the data among different groups. Next, several multivariate techniques are introduced (Flury faces, Andrews' curves and parallel coordinates plots (PCPs)) which provide graphical displays addressing the questions formulated above. The advantages and the disadvantages of each of these techniques are stressed.

Two basic techniques for estimating densities are also presented: histograms and kernel densities. A density estimate gives a quick insight into the shape of the distribution of the data. We show that kernel density estimates (KDEs) overcome some of the drawbacks of the histograms.

Finally, scatterplots are shown to be very useful for plotting bivariate or trivariate variables against each other: they help to understand the nature of the relationship among variables in a data set and allow for the detection of groups or clusters of points. Draftman plots or matrix plots are the visualisation of several bivariate scatterplots on the same display. They help detect structures in conditional dependencies by brushing across the plots. Outliers and observations that need special attention may be discovered with Andrews curves and PCPs. This chapter ends with an explanatory analysis of the Boston Housing data.

### 1.1 Boxplots

Example 1.1 The Swiss bank data (see Chap. 22, Sect. 22.2) consists of 200 measurements on Swiss bank notes. The first half of these measurements are from genuine bank notes, the other half are from counterfeit bank notes.

The authorities measured, as indicated in Fig. 1.1,
$X_{1}=$ 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.


Fig. 1.1 An old Swiss 1000-franc bank note

These data are taken from Flury and Riedwyl (1988). The aim is to study how these measurements may be used in determining whether a bill is genuine or counterfeit.

The boxplot is a graphical technique that displays the distribution of variables. It helps us see the location, skewness, spread, tail length and outlying points.

It is particularly useful in comparing different batches. The boxplot is a graphical representation of the Five Number Summary. To introduce the Five Number Summary, let us consider for a moment a smaller, one-dimensional data set: the population of the 15 largest world cities in 2006 (Table 1.1).

In the Five Number Summary, we calculate the upper quartile $F_{U}$, the lower quartile $F_{L}$, the median and the extremes. Recall that order statistics $\left\{x_{(1)}, x_{(2)}, \ldots, x_{(n)}\right\}$ are a set of ordered values $x_{1}, x_{2}, \ldots, x_{n}$ where $x_{(1)}$ denotes the minimum and $x_{(n)}$ the maximum. The median $M$ typically cuts the set of observations in two equal parts, and is defined as

$$
M=\left\{\begin{array}{cc}
x_{\left(\frac{n+1}{2}\right)} & n \text { odd }  \tag{1.1}\\
\frac{1}{2}\left\{x_{\left(\frac{n}{2}\right)}+x_{\left(\frac{n}{2}+1\right)}\right\} & n \text { even }
\end{array} .\right.
$$

Table 1.1 The 15 largest world cities in 2006

| City | Country | Pop. $(10,000)$ | Order statistics |
| :--- | :--- | :--- | :--- |
| Tokyo | Japan | 3,420 | $x_{(15)}$ |
| Mexico city | Mexico | 2,280 | $x_{(14)}$ |
| Seoul | South Korea | 2,230 | $x_{(13)}$ |
| New York | USA | 2,190 | $x_{(12)}$ |
| Sao Paulo | Brazil | 2,020 | $x_{(11)}$ |
| Bombay | India | 1,985 | $x_{(10)}$ |
| Delhi | India | 1,970 | $x_{(9)}$ |
| Shanghai | China | 1,815 | $x_{(8)}$ |
| Los Angeles | USA | 1,800 | $x_{(7)}$ |
| Osaka | Japan | 1,680 | $x_{(6)}$ |
| Jakarta | Indonesia | 1,655 | $x_{(5)}$ |
| Calcutta | India | 1,565 | $x_{(4)}$ |
| Cairo | Egypt | 1,560 | $x_{(3)}$ |
| Manila | Philippines | 1,495 | $x_{(2)}$ |
| Karachi | Pakistan | 1,430 | $x_{(1)}$ |

The quartiles cut the set into four equal parts, which are often called fourths (that is why we use the letter $F$ ). Using a definition that goes back to Hoaglin, Mosteller, and Tukey (1983) the definition of a median can be generalised to fourths, eights, etc. Considering the order statistics we can define the depth of a data value $x_{(i)}$ as $\min \{i, n-i+1\}$. If $n$ is odd, the depth of the median is $\frac{n+1}{2}$. If $n$ is even, $\frac{n+1}{2}$ is a fraction. Thus, the median is determined to be the average between the two data values belonging to the next larger and smaller order statistics, i.e. $M=\frac{1}{2}\left\{x_{\left(\frac{n}{2}\right)}+x_{\left(\frac{n}{2}+1\right)}\right\}$. In our example, we have $n=15$ hence the median $M=x_{(8)}=1,815$.

We proceed in the same way to get the fourths. Take the depth of the median and calculate

$$
\text { depth of fourth }=\frac{\text { [depth of median] }+1}{2}
$$

with $[z]$ denoting the largest integer smaller than or equal to $z$. In our example this gives 4.5 and thus leads to the two fourths

$$
\begin{aligned}
& F_{L}=\frac{1}{2}\left\{x_{(4)}+x_{(5)}\right\} \\
& F_{U}=\frac{1}{2}\left\{x_{(11)}+x_{(12)}\right\}
\end{aligned}
$$

(recalling that a depth which is a fraction corresponds to the average of the two nearest data values).

Table 1.2 Five number summary

| $\#$ | 15 | World cities |  |  |
| :--- | :--- | :--- | :--- | :--- |
| $M$ | 8 |  | 1,815 |  |
| $F$ | 4.5 | 1,610 |  | 2,105 |
|  | 1 | 1,430 |  | 3,420 |

The $F$-spread, $d_{F}$, is defined as $d_{F}=F_{U}-F_{L}$. The outside bars

$$
\begin{gather*}
F_{U}+1.5 d_{F}  \tag{1.2}\\
F_{L}-1.5 d_{F} \tag{1.3}
\end{gather*}
$$

are the borders beyond which a point is regarded as an outlier. For the number of points outside these bars see Exercise 1.3. For the $n=15$ data points the fourths are $1610=\frac{1}{2}\left\{x_{(4)}+x_{(5)}\right\}$ and $2105=\frac{1}{2}\left\{x_{(11)}+x_{(12)}\right\}$. Therefore the $F$-spread and the upper and lower outside bars in the above example are calculated as follows:

$$
\begin{align*}
d_{F} & =F_{U}-F_{L}=2105-1610=495  \tag{1.4}\\
F_{L}-1.5 d_{F} & =1610-1.5 \cdot 495=867.5  \tag{1.5}\\
F_{U}+1.5 d_{F} & =2105+1.5 \cdot 495=2847.5 \tag{1.6}
\end{align*}
$$

Since Tokyo is beyond the outside bars it is considered to be an outlier. The minimum and the maximum are called the extremes. The mean is defined as

$$
\bar{x}=n^{-1} \sum_{i=1}^{n} x_{i},
$$

which is $1,939.7$ in our example. The mean is a measure of location. The median (1815), the fourths $(1610 ; 2105)$ and the extremes $(1430 ; 3420)$ constitute basic information about the data. The combination of these five numbers leads to the Five Number Summary as shown in Table 1.2. The depths of each of the five numbers have been added as an additional column.

## Construction of the Boxplot

1. Draw a box with borders (edges) at $F_{L}$ and $F_{U}$ (i.e. $50 \%$ of the data are in this box).
2. Draw the median as a solid line $(\mid)$ and the mean as a dotted line $(\mid)$.
3. Draw "whiskers" from each end of the box to the most remote point that is NOT an outlier.
4. Show outliers as either " $\star$ " or " $\bullet$ " depending on whether they are outside of $F_{U L} \pm$ $1.5 d_{F}$ or $F_{U L} \pm 3 d_{F}$ respectively (this feather is not contained in some software). Label them if possible.


Fig. 1.2 Boxplot for world cities $Q_{\text {mVAboxcity }}$

In the world cities example, the cut-off points (outside bars) are at 867.5 and 2847.5, hence we can draw whiskers to Karachi and Mexico City. We can see from Fig. 1.2 that the data are very skew: The upper half of the data (above the median) is more spread out than the lower half (below the median), the data contains one outlier marked as a circle and the mean (as a non-robust measure of location) is pulled away from the median.

Boxplots are very useful tools in comparing batches. The relative location of the distribution of different batches tells us a lot about the batches themselves. Before we come back to the Swiss bank data, let us compare the fuel economy of vehicles from different countries, see Fig. 1.3 and Table 22.3.

Example 1.2 The data are from the second column of Table 22.3 and show the mileage (miles per gallon) of American, Japanese and European cars. The five-number summaries for these data sets are $\{12,16.8,18.8,22,30\}$, $\{18,22,25,30.5,35\}$ and $\{14,19,23,25,28\}$ for American, Japanese and European cars, respectively. This reflects the information shown in Fig. 1.3. The following conclusions can be made:

- Japanese cars achieve higher fuel efficiency than US and European cars.
- There is one outlier, a very fuel-efficient car (VW-Rabbit Golf Diesel).
- The main body of the US car data (the box) lies below the Japanese car data.
- The worst Japanese car is more fuel-efficient than almost $50 \%$ of the US cars.
- The spread of the Japanese and the US cars are almost equal.
- The median of the Japanese data is above that of the European data and the US data.

Fig. 1.3 Boxplot for the mileage of American, Japanese and European cars (from left to right) O . MVAboxcar

Fig. 1.4 The $X_{6}$ variable of Swiss bank data (diagonal of bank notes) Q MVAboxbank6


Table 1.3 Five number summary

Table 1.4 Five number summary

| $\#$ | 100 | Counterfeit bank notes |  |  |
| :--- | :--- | :--- | :--- | :--- |
| $M$ | 50.5 |  | 139.5 |  |
| $F$ | 25.75 | 139.2 |  | 139.8 |
|  | 1 | 138.3 |  | 140.65 |

Fig. 1.5 The $X_{1}$ variable of Swiss bank data (length of bank notes) Q MVAboxbank1


One sees that the diagonals of the genuine bank notes tend to be larger. It is harder to see a clear distinction when comparing the length of the bank notes $X_{1}$, see Fig. 1.5. There are a few outliers in both plots. Almost all the observations of the diagonal of the genuine notes are above the ones from the counterfeit notes. There is one observation in Fig. 1.4 of the genuine notes that is almost equal to the median of the counterfeit notes. Can the parallel boxplot technique help us distinguish between the two types of bank notes?

| Sum Summary |
| :---: |
| $\hookrightarrow$ The median and mean bars are measures of locations. |
| $\hookrightarrow$ The relative location of the median (and the mean) in the box is a measure of how skewed it is. |
| $\hookrightarrow$ The length of the box and whiskers are a measure of spread. |
| $\hookrightarrow$ The length of the whiskers indicate the tail length of the distribution. |
| $\hookrightarrow$ The outlying points are indicated with a " $\star$ " or " $\bullet$ " depending on if they are outside of $F_{U L} \pm 1.5 d_{F}$ or $F_{U L} \pm 3 d_{F}$ respectively. |


|  | Summary (continued) |
| :---: | :--- |
| $\hookrightarrow$ | The boxplots do not indicate multi-modality or clusters. |
| $\hookrightarrow$ | If we compare the relative size and location of the boxes, we are <br> comparing distributions. |

### 1.2 Histograms

Histograms are density estimates. A density estimate gives a good impression of the distribution of the data. In contrast to boxplots, density estimates show possible multimodality of the data. The idea is to locally represent the data density by counting the number of observations in a sequence of consecutive intervals (bins) with origin $x_{0}$. Let $B_{j}\left(x_{0}, h\right)$ denote the bin of length $h$ which is the element of a bin grid starting at $x_{0}$ :

$$
B_{j}\left(x_{0}, h\right)=\left[x_{0}+(j-1) h, x_{0}+j h\right), \quad j \in \mathbb{Z},
$$

where [.,.) denotes a left closed and right open interval. If $\left\{x_{i}\right\}_{i=1}^{n}$ is an i.i.d. sample with density $f$, the histogram is defined as follows:

$$
\begin{equation*}
\hat{f}_{h}(x)=n^{-1} h^{-1} \sum_{j \in \mathbb{Z}} \sum_{i=1}^{n} \boldsymbol{I}\left\{x_{i} \in B_{j}\left(x_{0}, h\right)\right\} \boldsymbol{I}\left\{x \in B_{j}\left(x_{0}, h\right)\right\} . \tag{1.7}
\end{equation*}
$$

In sum (1.7) the first indicator function $\boldsymbol{I}\left\{x_{i} \in B_{j}\left(x_{0}, h\right)\right\}$ (see Symbols and Notation in Chap. 21) counts the number of observations falling into bin $B_{j}\left(x_{0}, h\right)$. The second indicator function is responsible for "localising" the counts around $x$. The parameter $h$ is a smoothing or localising parameter and controls the width of the histogram bins. An $h$ that is too large leads to very big blocks and thus to a very unstructured histogram. On the other hand, an $h$ that is too small gives a very variable estimate with many unimportant peaks.

The effect of $h$ is given in detail in Fig. 1.6. It contains the histogram (upper left) for the diagonal of the counterfeit bank notes for $x_{0}=137.8$ (the minimum of these observations) and $h=0.1$. Increasing $h$ to $h=0.2$ and using the same origin, $x_{0}=137.8$, results in the histogram shown in the lower left of the figure. This density histogram is somewhat smoother due to the larger $h$. The binwidth is next set to $h=0.3$ (upper right). From this histogram, one has the impression that the distribution of the diagonal is bimodal with peaks at about 138.5 and 139.9.


Fig. 1.6 Diagonal of counterfeit 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) $Q_{\text {MVAhisbank1 }}$

The detection of modes requires fine tuning of the binwidth. Using methods from smoothing methodology (Härdle, Müller, Sperlich, \& Werwatz, 2004) one can find an "optimal" binwidth $h$ for $n$ observations:

$$
h_{\mathrm{opt}}=\left(\frac{24 \sqrt{\pi}}{n}\right)^{1 / 3} .
$$

Unfortunately, the binwidth $h$ is not the only parameter determining the shapes of $\hat{f}$.
In Fig. 1.7, we show histograms with $x_{0}=137.65$ (upper left), $x_{0}=137.75$ (lower left), with $x_{0}=137.85$ (upper right), and $x_{0}=137.95$ (lower right). All the graphs have been scaled equally on the $y$-axis to allow comparison. One sees that-despite the fixed binwidth $h$-the interpretation is not facilitated. The shift of the origin $x_{0}$ (to four different locations) created four different histograms. This


Fig. 1.7 Diagonal of counterfeit 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) $Q$ MVAhisbank2
property of histograms strongly contradicts the goal of presenting data features. Obviously, the same data are represented quite differently by the four histograms. A remedy has been proposed by Scott (1985): "Average the shifted histograms!". The result is presented in Fig. 1.8.

Here all bank note observations (genuine and counterfeit) have been used. The (so-called) averaged shifted histogram is no longer dependent on the origin and shows a clear bimodality of the diagonals of the Swiss bank notes.


Fig. 1.8 Averaged shifted histograms based on all (counterfeit and genuine) Swiss bank notes: there are 2 shifts (upper left), 4 shifts (lower left), 8 shifts (upper right) and 16 shifts (lower right) Q MVAashbank

|  | $\hookrightarrow$ Modes of the density are detected with a histogram. |
| :--- | :--- |
|  | $\hookrightarrow$ Modes correspond to strong peaks in the histogram. |
|  | Histograms with the same $h$ need not be identical. They also <br> depend on the origin $x_{0}$ of the grid. |
|  | $\hookrightarrow$ |
|  | The influence of the origin $x_{0}$ is drastic. Changing $x_{0}$ creates <br> different looking histograms. |
|  | $\hookrightarrow$ | | The consequence of an $h$ that is too large is an unstructured |
| :--- |
| histogram that is too flat. |


| Summary (continued) |  |
| :---: | :--- |
| $\hookrightarrow$ | There is an "optimal" $h=(24 \sqrt{\pi} / n)^{1 / 3}$. |
| $\hookrightarrow$ | It is recommended to use averaged histograms. They are kernel <br> densities. |

### 1.3 Kernel Densities

The major difficulties of histogram estimation may be summarised in four critiques:

- determination of the binwidth $h$, which controls the shape of the histogram,
- choice of the bin origin $x_{0}$, which also influences to some extent the shape,
- loss of information since observations are replaced by the central point of the interval in which they fall,
- the underlying density function is often assumed to be smooth, but the histogram is not smooth.

Rosenblatt (1956), Whittle (1958) and Parzen (1962) developed an approach which avoids the last three difficulties. First, a smooth kernel function rather than a box is used as the basic building block. Second, the smooth function is centred directly over each observation. Let us study this refinement by supposing that $x$ is the centre value of a bin. The histogram can in fact be rewritten as

$$
\begin{equation*}
\hat{f}_{h}(x)=n^{-1} h^{-1} \sum_{i=1}^{n} \boldsymbol{I}\left(\left|x-x_{i}\right| \leq \frac{h}{2}\right) \tag{1.8}
\end{equation*}
$$

If we define $K(u)=\boldsymbol{I}\left(|u| \leq \frac{1}{2}\right)$, then (1.8) changes to

$$
\begin{equation*}
\hat{f_{h}}(x)=n^{-1} h^{-1} \sum_{i=1}^{n} K\left(\frac{x-x_{i}}{h}\right) . \tag{1.9}
\end{equation*}
$$

This is the general form of the kernel estimator. Allowing smoother kernel functions like the quartic kernel,

$$
K(u)=\frac{15}{16}\left(1-u^{2}\right)^{2} \quad \boldsymbol{I}(|u| \leq 1),
$$

and computing $x$ not only at bin centers gives us the kernel density estimator. Kernel estimators can also be derived via weighted averaging of rounded points (WARPing) or by averaging histograms with different origins, see Scott (1985). Table 1.5 introduces some commonly used kernels.

Table 1.5 Kernel functions

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

Different kernels generate different shapes of the estimated density. The most important parameter is the so-called bandwidth $h$, and can be optimised, for example, by cross-validation; see Härdle (1991) for details. The cross-validation method minimises the integrated squared error. This measure of discrepancy is based on the squared differences $\left\{\hat{f}_{h}(x)-f(x)\right\}^{2}$. Averaging these squared deviations over a grid of points $\left\{x_{l}\right\}_{l=1}^{L}$ leads to

$$
L^{-1} \sum_{l=1}^{L}\left\{\hat{f}_{h}\left(x_{l}\right)-f\left(x_{l}\right)\right\}^{2} .
$$

Asymptotically, if this grid size tends to zero, we obtain the integrated squared error:

$$
\int\left\{\hat{f_{h}}(x)-f(x)\right\}^{2} d x
$$

In practice, it turns out that the method consists of selecting a bandwidth that minimises the cross-validation function

$$
\int \hat{f}_{h}^{2}-2 \sum_{i=1}^{n} \hat{f}_{h, i}\left(x_{i}\right)
$$

where $\hat{f}_{h, i}$ is the density estimate obtained by using all datapoints except for the $i$-th observation. Both terms in the above function involve double sums. Computation may therefore be slow. There are many other density bandwidth selection methods. Probably the fastest way to calculate this is to refer to some reasonable reference distribution. The idea of using the Normal distribution as a reference, for example, goes back to Silverman (1986). The resulting choice of $h$ is called the rule of thumb.

For the Gaussian kernel from Table 1.5 and a Normal reference distribution, the rule of thumb is to choose

$$
\begin{equation*}
h_{G}=1.06 \hat{\sigma} n^{-1 / 5} \tag{1.10}
\end{equation*}
$$

Fig. 1.9 Densities of the diagonals of genuine and counterfeit bank notes. Automatic density estimates $\mathbf{Q}_{\text {MVAdenbank }}$

where $\hat{\sigma}=\sqrt{n^{-1} \sum_{i=1}^{n}\left(x_{i}-\bar{x}\right)^{2}}$ denotes the sample standard deviation. This choice of $h_{G}$ optimises the integrated squared distance between the estimator and the true density. For the quartic kernel, we need to transform (1.10). The modified rule of thumb is:

$$
\begin{equation*}
h_{Q}=2.62 \cdot h_{G} . \tag{1.11}
\end{equation*}
$$

Figure 1.9 shows the automatic density estimates for the diagonals of the counterfeit and genuine bank notes. The density on the left is the density corresponding to the diagonal of the counterfeit data. The separation is clearly visible, but there is also an overlap. The problem of distinguishing between the counterfeit and genuine bank notes is not solved by just looking at the diagonals of the notes. The question arises whether a better separation could be achieved using not only the diagonals, but one or two more variables of the data set. The estimation of higher dimensional densities is analogous to that of one dimensional. We show a two-dimensional density estimate for $X_{4}$ and $X_{5}$ in Fig. 1.10. The contour lines indicate the height of the density. One sees two separate distributions in this higher dimensional space, but they still overlap to some extent.

We can add one more dimension and give a graphical representation of a threedimensional density estimate, or more precisely an estimate of the joint distribution of $X_{4}, X_{5}$ and $X_{6}$. Figure 1.11 shows the contour areas at three different levels of the density: 0.2 (green), 0.4 (red) and 0.6 (blue) of this three-dimensional density estimate. One can clearly recognise two "ellipsoids" (at each level), but as before, they overlap. In Chap. 14 we will learn how to separate the two ellipsoids and how to develop a discrimination rule to distinguish between these data points.

Fig. 1.10 Contours of the density of $X_{5}$ and $X_{6}$ of genuine and counterfeit bank notes $\mathbf{Q}$ MVAcontbank2

Fig. 1.11 Contours of the density of $X_{4}, X_{5}, X_{6}$ of genuine and counterfeit bank notes $\mathbf{Q}$ MVAcontbank3


| $\hookrightarrow$ | Kernel densities estimate distribution densities by the kernel <br> method. |
| :---: | :--- |
| $\hookrightarrow$ | The bandwidth $h$ determines the degree of smoothness of the <br> estimate $\hat{f}$. |
|  | Kernel densities are smooth functions and they can graphically <br> represent distributions (up to three dimensions). |
|  | A simple (but not necessarily correct) way to find a good bandwidth <br> is to compute the rule of thumb bandwidth $h_{G}=1.06 \hat{\sigma} n^{-1 / 5}$. |
|  | This bandwidth is to be used only in combination with a Gaussian <br> kernel $\varphi$. |
| $\hookrightarrow$ | Kernel density estimates are a good descriptive tool for seeing <br> modes, location, skewness, tails, asymmetry, etc. |

### 1.4 Scatterplots

Scatterplots are bivariate or trivariate plots of variables against each other. They help us understand relationships among the variables of a data set. A downward-sloping scatter indicates that as we increase the variable on the horizontal axis, the variable on the vertical axis decreases. An analogous statement can be made for upwardsloping scatters.

Figure 1.12 plots the 5th column (upper inner frame) of the bank data against the 6th column (diagonal). The scatter is downward-sloping. As we already know from the previous section on marginal comparison (e.g. Fig. 1.9) a good separation between genuine and counterfeit bank notes is visible for the diagonal variable. The sub-cloud in the upper half (circles) of Fig. 1.12 corresponds to the true bank notes. As noted before, this separation is not distinct, since the two groups overlap somewhat.

This can be verified in an interactive computing environment by showing the index and coordinates of certain points in this scatterplot. In Fig. 1.12, the 70th observation in the merged data set is given as a thick circle, and it is from a genuine bank note. This observation lies well embedded in the cloud of counterfeit bank notes. One straightforward approach that could be used to tell the counterfeit from the genuine bank notes is to draw a straight line and define notes above this value as genuine. We would of course misclassify the 70th observation, but can we do better?


Fig. 1.12 2D scatterplot for $X_{5}$ vs. $X_{6}$ of the bank notes. Genuine notes are circles, counterfeit notes are stars $\mathbf{Q}$ MVAscabank56

Swiss bank notes


Fig. 1.13 3D scatterplot of the bank notes for $\left(X_{4}, X_{5}, X_{6}\right)$. Genuine notes are circles, counterfeit are stars $\mathbf{Q}$ MVAscabank456

If we extend the two-dimensional scatterplot by adding a third variable, e.g. $X_{4}$ (lower distance to inner frame), we obtain the scatterplot in three dimensions as shown in Fig. 1.13. It becomes apparent from the location of the point clouds that a better separation is obtained. We have rotated the three-dimensional data until this satisfactory 3D view was obtained. Later, we will see that the rotation is the same as bundling a high-dimensional observation into one or more linear combinations


Fig. 1.14 Draftman's plot of the bank notes. The pictures in the left-hand column show $\left(X_{3}, X_{4}\right),\left(X_{3}, X_{5}\right)$ and ( $X_{3}, X_{6}$ ), in the middle we have $\left(X_{4}, X_{5}\right)$ and ( $X_{4}, X_{6}$ ), and in the lower right ( $X_{5}, X_{6}$ ). The upper right half contains the corresponding density contour plots $\mathbf{a}$ MVAdrafbank4
of the elements of the observation vector. In other words, the "separation line" parallel to the horizontal coordinate axis in Fig. 1.12 is, in Fig. 1.13, a plane and no longer parallel to one of the axes. The formula for such a separation plane is a linear combination of the elements of the observation vector:

$$
\begin{equation*}
a_{1} x_{1}+a_{2} x_{2}+\cdots+a_{6} x_{6}=\text { const. } \tag{1.12}
\end{equation*}
$$

The algorithm that automatically finds the weights $\left(a_{1}, \ldots, a_{6}\right)$ will be investigated later on in Chap. 14.

Let us study yet another technique: the scatterplot matrix. If we want to draw all possible two-dimensional scatterplots for the variables, we can create a so-called draftman's plot (named after a draftman who prepares drafts for parliamentary discussions). Similar to a draftman's plot the scatterplot matrix helps in creating new ideas and in building knowledge about dependencies and structure.

Figure 1.14 shows a draftman's plot applied to the last four columns of the full bank data set. For ease of interpretation we have distinguished between the group of counterfeit and genuine bank notes by a different colour. As discussed several times
earlier, the separability of the two types of notes is different for different scatterplots. Not only is it difficult to perform this separation on, say, scatterplot $X_{3}$ vs. $X_{4}$, in addition the "separation line" is no longer parallel to one of the axes. The most obvious separation happens in the scatterplot in the lower right-hand side where indicated, as in Fig. 1.12, $X_{5}$ vs. $X_{6}$. The separation line here would be upwardsloping with an intercept at about $X_{6}=139$. The upper right half of the draftman's plot shows the density contours that we introduced in Sect. 1.3.

The power of the draftman's plot lies in its ability to show the internal connections of the scatter diagrams. Define a brush as a re-scalable rectangle that we can move via keyboard or mouse over the screen. Inside the brush we can highlight or colour observations. Suppose the technique is installed in such a way that as we move the brush in one scatter, the corresponding observations in the other scatters are also highlighted. By moving the brush, we can study conditional dependence.

If we brush (i.e. highlight or colour the observation with the brush), the $X_{5}$ vs. $X_{6}$ plot and move through the upper point cloud, we see that in other plots (e.g. $X_{3}$ vs. $X_{4}$ ), the corresponding observations are more embedded in the other sub-cloud.

| $\hookrightarrow$ | Scatterplots in two and three dimensions helps in identifying <br> separated points, outliers or sub-clusters. |
| :--- | :--- |
| $\hookrightarrow$ | Scatterplots help us in judging positive or negative dependencies. |
| $\hookrightarrow$ | Draftman scatterplot matrices help detect structures conditioned on <br> values of other variables. |
|  | As the brush of a scatterplot matrix moves through a point cloud, <br> we can study conditional dependence. |

### 1.5 Chernoff-Flury Faces

If we are given data in numerical form, we tend to also display it numerically. This was done in the preceding sections: an observation $x_{1}=(1,2)$ was plotted as the point $(1,2)$ in a two-dimensional coordinate system. In multivariate analysis we want to understand data in low dimensions (e.g. on a 2 D computer screen) although the structures are hidden in high dimensions. The numerical display of data structures using coordinates therefore ends at dimensions greater than three.

If we are interested in condensing a structure into 2 D elements, we have to consider alternative graphical techniques. The Chernoff-Flury faces, for example, provide such a condensation of high-dimensional information into a simple "face". In fact faces are a simple way of graphically displaying high-dimensional data. The size of the face elements like pupils, eyes, upper and lower hair line, etc. are assigned to certain variables. The idea of using faces goes back to Chernoff (1973) and has been further developed by Bernhard Flury. We follow the design described in Flury and Riedwyl (1988) which uses the following characteristics.

1. right eye size
2. right pupil size
3. position of right pupil
4. right eye slant
5. horizontal position of right eye
6. vertical position of right eye
7. curvature of right eyebrow
8. density of right eyebrow
9. horizontal position of right eyebrow
10. vertical position of right eyebrow
11. right upper hair line
12. right lower hair line
13. right face line
14. darkness of right hair
15. right hair slant
16. right nose line
17. right size of mouth
18. right curvature of mouth

19-36. like 1-18, only for the left side.
First, every variable that is to be coded into a characteristic face element is transformed into a $(0,1)$ scale, i.e. the minimum of the variable corresponds to 0 and the maximum to 1 . The extreme positions of the face elements therefore correspond to a certain "grin" or "happy" face element. Dark hair might be coded as 1, and blond hair as 0 and so on.

As an example, consider the observations 91-110 of the bank data. Recall that the bank data set consists of 200 observations of dimension 6 where, for example, $X_{6}$ is the diagonal of the note. If we assign the six variables to the following face elements

$$
\begin{aligned}
& X_{1}=1,19 \text { (eye sizes) } \\
& X_{2}=2,20 \text { (pupil sizes) } \\
& X_{3}=4,22 \text { (eye slants) }
\end{aligned}
$$



Fig. 1.15 Chernoff-Flury faces for observations $91-110$ of the bank notes $Q$ MVAfacebank10

$$
\begin{aligned}
& X_{4}=11,29 \text { (upper hair lines) } \\
& X_{5}=12,30 \text { (lower hair lines) } \\
& X_{6}=13,14,31,32 \text { (face lines and darkness of hair), }
\end{aligned}
$$

we obtain Fig. 1.15. Also recall that observations $1-100$ correspond to the genuine notes, and that observations $101-200$ correspond to the counterfeit notes. The counterfeit bank notes then correspond to the upper half of Fig. 1.15. In fact the faces for these observations look more grim and less happy. The variable $X_{6}$ (diagonal) already worked well in the boxplot in Fig. 1.4 in distinguishing between the counterfeit and genuine notes. Here, this variable is assigned to the face line and the darkness of the hair. That is why we clearly see a good separation within these 20 observations.

## Observations 1 to 50

年

Fig. 1.16 Chernoff-Flury faces for observations $1-50$ of the bank notes $\mathbb{Q}$ MVAfacebank 50

What happens if we include all 100 genuine and all 100 counterfeit bank notes in the Chernoff-Flury face technique? Figures 1.16 and 1.17 show the faces of the genuine bank notes with the same assignments as used before, and Figs. 1.18 and 1.19 show the faces of the counterfeit bank notes. Comparing Figs. 1.16 and 1.18 one clearly sees that the diagonal (face line) is longer for genuine bank notes. Equivalently coded is the hair darkness (diagonal) which is lighter (shorter) for the counterfeit bank notes. One sees that the faces of the genuine bank notes have a much darker appearance and have broader face lines. The faces in Figs. 1.16 and 1.17 are obviously different from the ones in Figs. 1.18 and 1.19.
Observations 51 to 100

Fig. 1.17 Chernoff-Flury faces for observations 51-100 of the bank notes $Q$ MVAfacebank50
101

Fig. 1.18 Chernoff-Flury faces for observations $101-150$ of the bank notes $Q$ MVAfacebank 50

## Observations 151 to 200

| 151 | 152 | 153 | 154 | 155 | 156 | 157 | 158 | 159 | 160 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  |  |  |  |  |  | $9$ | y |
| 161 | 162 | 163 | 164 | 165 | 166 | 167 | 168 | 169 | 170 |
| $\begin{aligned} & y \\ & y_{0} \\ & \theta \end{aligned}$ |  |  |  |  |  |  |  | 0 | 6 |
| 171 | 172 | 173 | 174 | 175 | 176 | 177 | 178 | 179 | 180 |
|  |  |  |  |  |  |  |  |  | y |
| 181 | 182 | 183 | 184 | 185 | 186 | 187 | 188 | 189 | 190 |
|  |  |  |  |  |  | 4 9 $b$ |  | $y$ 0 0 |  |
| 191 | 192 | 193 | 194 | 195 | 196 | 197 | 198 | 199 | 200 |
|  |  |  |  |  |  |  |  |  | 80 |

Fig. 1.19 Chernoff-Flury faces for observations $151-200$ of the bank notes $Q$ MVAfacebank50

| $\hookrightarrow$ | Faces can be used to detect sub-groups in multivariate data. |
| :--- | :--- |
| $\hookrightarrow$ | Sub-groups are characterised by similar looking faces. |
| $\hookrightarrow$ | Outliers are identified by extreme faces, e.g. dark hair, smile or a <br> happy face. |
|  | If one element of $X$ is unusual, the corresponding face element <br> significantly changes in shape. |

### 1.6 Andrews' Curves

The basic problem of graphical displays of multivariate data is the dimensionality. Scatterplots work well up to three dimensions (if we use interactive displays). More than three dimensions have to be coded into displayable 2D or 3D structures (e.g. faces). The idea of coding and representing multivariate data by curves was suggested by Andrews (1972). Each multivariate observation $X_{i}=\left(X_{i, 1}, \ldots, X_{i, p}\right)$ is transformed into a curve as follows:

$$
f_{i}(t)=\left\{\begin{array}{rlrl}
\frac{X_{i, 1}}{\sqrt{2}}+X_{i, 2} \sin (t)+X_{i, 3} \cos (t)+\cdots & &  \tag{1.13}\\
& +X_{i, p-1} \sin \left(\frac{p-1}{2} t\right)+X_{i, p} \cos \left(\frac{p-1}{2} t\right) & & \text { for } p \text { odd } \\
\frac{X_{i, 1}}{\sqrt{2}}+X_{i, 2} \sin (t)+X_{i, 3} \cos (t)+\cdots+X_{i, p} \sin \left(\frac{p}{2} t\right) & & \text { for } p \text { even }
\end{array}\right.
$$

the observation represents the coefficients of a so-called Fourier series $(t \in[-\pi, \pi])$.
Suppose that we have three-dimensional observations: $X_{1}=(0,0,1), X_{2}=$ $(1,0,0)$ and $X_{3}=(0,1,0)$. Here $p=3$ and the following representations correspond to the Andrews' curves:

$$
\begin{aligned}
f_{1}(t) & =\cos (t) \\
f_{2}(t) & =\frac{1}{\sqrt{2}} \text { and } \\
f_{3}(t) & =\sin (t)
\end{aligned}
$$

These curves are indeed quite distinct, since the observations $X_{1}, X_{2}$, and $X_{3}$ are the 3D unit vectors: each observation has mass only in one of the three dimensions. The order of the variables plays an important role.


Fig. 1.20 Andrews' curves of the observations 96-105 from the Swiss bank note data. The order of the variables is $1,2,3,4,5,6 \mathrm{O}$ MVAandcur

Example 1.3 Let us take the 96th observation of the Swiss bank note data set,

$$
X_{96}=(215.6,129.9,129.9,9.0,9.5,141.7)
$$

The Andrews' curve is by (1.13):

$$
\begin{aligned}
f_{96}(t)= & \frac{215.6}{\sqrt{2}}+129.9 \sin (t)+129.9 \cos (t)+9.0 \sin (2 t) \\
& +9.5 \cos (2 t)+141.7 \sin (3 t)
\end{aligned}
$$

Figure 1.20 shows the Andrews' curves for observations 96-105 of the Swiss bank note data set. We already know that the observations $96-100$ represent genuine bank notes, and that the observations $101-105$ represent counterfeit bank notes. We see that at least four curves differ from the others, but it is hard to tell which curve belongs to which group.

We know from Fig. 1.4 that the sixth variable is an important one. Therefore, the Andrews' curves are calculated again using a reversed order of the variables.

Example 1.4 Let us consider again the 96th observation of the Swiss bank note data set,

$$
X_{96}=(215.6,129.9,129.9,9.0,9.5,141.7)
$$

Fig. 1.21 Andrews' curves of the observations 96-105 from the Swiss bank note data. The order of the variables is $6,5,4,3,2,1 \mathrm{O}$ MVAandcur2


The Andrews' curve is computed using the reversed order of variables:

$$
\begin{aligned}
f_{96}(t)= & \frac{141.7}{\sqrt{2}}+9.5 \sin (t)+9.0 \cos (t)+129.9 \sin (2 t) \\
& +129.9 \cos (2 t)+215.6 \sin (3 t)
\end{aligned}
$$

In Fig. 1.21 the curves $f_{96}-f_{105}$ for observations $96-105$ are plotted. Instead of a difference in high frequency, now we have a difference in the intercept, which makes it more difficult for us to see the differences in observations.

This shows that the order of the variables plays an important role in the interpretation. If $X$ is high-dimensional, then the last variables will only have a small visible contribution to the curve: they fall into the high frequency part of the curve. To overcome this problem Andrews suggested using an order which is suggested by Principal Component Analysis. This technique will be treated in detail in Chap. 11. In fact, the sixth variable will appear there as the most important variable for discriminating between the two groups. If the number of observations is more than 20, there may be too many curves in one graph. This will result in an over plotting of curves or a bad "signal-to-ink-ratio", see Tufte (1983). It is therefore advisable to present multivariate observations via Andrews' curves only for a limited number of observations.

| Sum Summary |
| :---: |
| $\hookrightarrow$ Outliers appear as single Andrews’ curves that look different from the rest. |
| $\hookrightarrow$ A sub-group of data is characterised by a set of similar curves. |
| $\hookrightarrow$ The order of the variables plays an important role for interpretation. |
| $\hookrightarrow$ The order of variables may be optimised by Principal Component Analysis. |
| $\hookrightarrow$ For more than 20 observations we may obtain a bad "signal-to-ink ratio", i.e. too many curves are overlaid in one picture. |

### 1.7 Parallel Coordinates Plots

PCP is a method for representing high-dimensional data, see Inselberg (1985). Instead of plotting observations in an orthogonal coordinate system, PCP draws coordinates in parallel axes and connects them with straight lines. This method helps in representing data with more than four dimensions.

One first scales all variables to $\max =1$ and $\min =0$. The coordinate index $j$ is drawn onto the horizontal axis, and the scaled value of variable $x_{i j}$ is mapped onto the vertical axis. This way of representation is very useful for high-dimensional data. It is however also sensitive to the order of the variables, since certain trends in the data can be shown more clearly in one ordering than in another.

Example 1.5 Take, once again, the observations 96-105 of the Swiss bank notes. These observations are six dimensional, so we can't show them in a six-dimensional Cartesian coordinate system. Using the PCP technique, however, they can be plotted on parallel axes. This is shown in Fig. 1.22.

PCP can also be used for detecting linear dependencies between variables: if all the lines are of almost parallel dimensions $(p=2)$, there is a positive linear dependence between them. In Fig. 1.23 we display the two variables weight and displacement for the car data set in Sect. 22.3. The correlation coefficient $\rho$ introduced in Sect. 3.2 is 0.9. If all lines intersect visibly in the middle, there is evidence of a negative linear dependence between these two variables, see Fig. 1.24. In fact the correlation is $\rho=-0.82$ between two variables mileage and weight: The more the weight, the less the mileage.


Fig. 1.22 Parallel coordinates plot of observations 96-105 Q MVAparcoo1


Fig. 1.23 Parallel coordinates plot indicating strong positive dependence with $\rho=0.9, X_{1}=$ weight, $X_{2}=$ displacement $Q_{\text {MVApcp2 }}$

Another use of PCP is sub-groups detection. Lines converging to different discrete points indicate sub-groups. Figure 1.25 shows the last three variablesdisplacement, gear ratio for high gear and company's headquarters of the car data; we see convergence to the last variable. This last variable is the company's headquarters with three discrete values: USA, Japan and Europe. PCP can also be used for outlier detection. Figure 1.26 shows the variables headroom, rear seat

Fig. 1.24 Parallel coordinates plot showing strong negative dependence with $\rho=-0.82, X_{1}=$ mileage, $X_{2}=$ weight MVApcp3

Fig. 1.25 Parallel coordinates plot with sub-groups Q MVApcp4 $^{\text {a }}$

clearance and trunk (boot) space in the car data set. There are two outliers visible. The boxplot Fig. 1.27 confirms this.

PCPs have also possible shortcomings: We cannot distinguish observations when two lines cross at one point unless we distinguish them clearly (e.g. by different line style). In Fig. 1.28, observation A and B both have the same value at $j=2$. Two lines cross at one point here. At the 3rd and 4th dimension we cannot tell which line belongs to which observation. A dotted line for A and solid line for B could have helped there.

To solve this problem one uses an interpolation curve instead of straight lines, e.g. cubic curves as in Graham and Kennedy (2003). Figure 1.29 is a variant of Fig. 1.28. In Fig. 1.29, with a natural cubic spline, it is evident how to follow the curves and distinguish the observations. The real power of PCP comes though through colouring sub-groups.

Fig. 1.26 PCP for
$X_{1}=$ headroom, $X_{2}=$ rear seat clearance and $X_{3}=$ trunk space Q MVApcp5


Fig. 1.27 Boxplots for headroom, rear seat clearance and trunk space a MVApcp6


Example 1.6 Data in Fig. 1.30 are coloured according to $X_{13}$-car company's headquarters. Red stands for European car, green for Japan and black for US. This PCP with colouring can provide some information for us:

1. US cars (black) tend to have large value in $X_{7}, X_{8}, X_{9}, X_{10}, X_{11}$ (trunk (boot) space, weight, length, turning diameter, displacement), which means US cars are generally larger.
2. Japanese cars (green) have large value in $X_{3}, X_{4}$ (both for repair record), which means Japanese cars tend to be repaired less.

Fig. 1.28 PCP with intersection for given data points $\mathrm{A}=[0,2,3,2]$ and $B=[3,2,2,1]$ Q MVAрср7


Fig. 1.29 PCP with cubic spline interpolation Q MVApcp8

Parallel Coordinate Plot with Cubic Spline



Fig. 1.30 Parallel coordinates plot for car data $Q$ MVApcp1

|  | $\hookrightarrow$ |
| :--- | :--- |
|  | Parallel coordinates plots overcome the visualisation problem of |
| the Cartesian coordinate system for dimensions greater than 4. |  |

### 1.8 Hexagon Plots

This section closely follows the presentation of Lewin-Koh (2006). In geometry, a hexagon is a polygon with six edges and six vertices. Hexagon binning is a type of bivariate histogram with hexagon borders. It is useful for visualising the structure
of data sets entailing a large number of observations $n$. The concept of hexagon binning is as follows:

1. The $x y$ plane over the set $(\operatorname{range}(x)$, range $(y))$ is tessellated by a regular grid of hexagons.
2. The number of points falling in each hexagon is counted.
3. The hexagons with count $>0$ are plotted by using a colour ramp or varying the radius of the hexagon in proportion to the counts.

This algorithm is extremely fast and effective for displaying the structure of data sets even for $n \geq 10^{6}$. If the size of the grid and the cuts in the colour ramp are chosen in a clever fashion, then the structure inherent in the data should emerge in the binned plot. The same caveats apply to hexagon binning as histograms. Variance and bias vary in opposite directions with bin width, so we have to settle for finding the value of the bin width that yields the optimal compromise between variance and bias reduction. Clearly, if we increase the size of the grid, the hexagon plot appears to be smoother, but without some reasonable criterion on hand it remains difficult to say which bin width provides the "optimal" degree of smoothness. The default number of bins suggested by standard software is 30 .

Applications to some data sets are shown as follows. The data is taken from ALLBUS (2006)[ZA No.3762]. The number of respondents is 2,946. The following nine variables have been selected to analyse the relation between each pair of variables.

| $X_{1}:$ | Age |
| :--- | :--- |
| $X_{2}:$ | Net income |
| $X_{3}:$ | Time for television per day in minutes |
| $X_{4}:$ | Time for work per week in hours |
| $X_{5}:$ | Time for computer per week in hours |
| $X_{6}:$ | Days for illness yearly |
| $X_{7}:$ | Living space (square metres) |
| $X_{8}:$ | Size |
| $X_{9}:$ | Weight |

Firstly, we consider two variables $X_{1}=$ Age and $X_{2}=$ Net income in Fig. 1.31. The top left picture is a scatter plot. The second one is a hexagon plot with borders making it easier to see the separation between hexagons. Looking at these plots one can see that almost all individuals have a net monthly income of less than 2,000 EUR. Only two individuals earn more than 10,000 EUR per month.

Figure 1.32 shows the relation between $X_{1}$ and $X_{5}$. About $40 \%$ of respondents from 20 to 80 years old do not use a computer at least once per week. The respondent who deals with a computer 105 h each week was actually not in fulltime employment.


Fig. 1.31 Hexagon plots between $X_{1}$ and $X_{2}$ Q MVAageIncome


Fig. 1.32 Hexagon plot between $X_{1}$ and $X_{5} \mathrm{Q}_{\text {mVAageCom }}$

Clearly, people who earn modest incomes live in smaller flats. The trend here is relatively clear in Fig. 1.33. The larger the net income, the larger the flat. A few people do however earn high incomes but live in small flats.

Fig. 1.33 Hexagon plot between $X_{2}$ and $X_{7}$ a MVAincomeLi


|  | $\hookrightarrow$ |
| :--- | :--- |
|  | Hexagon binning is a type of bivariate histogram, used for visual- <br> ising large data. |
| $\hookrightarrow$ | Variance and bias vary in opposite directions with bin width. |
| $\hookrightarrow$ | Hexagons have the property of "symmetry of the nearest neigh- <br> bours" which lacks in square bins. |
| $\hookrightarrow$ | Hexagons are visually less biased for displaying densities than <br> other regular tessellations. |

### 1.9 Boston Housing

## Aim of the Analysis

The Boston Housing data set was analysed by Harrison and Rubinfeld (1978) who wanted to find out whether "clean air" had an influence on house prices. We will use this data set in this chapter and in most of the following chapters to illustrate the presented methodology. The data are described in Sect. 22.1.

Fig. 1.34 Parallel coordinates plot for Boston housing data O MVApcphousing


## What Can Be Seen from the PCPs

In order to highlight the relations of $X_{14}$ to the remaining 13 variables, we colour all of the observations with $X_{14}>\operatorname{median}\left(X_{14}\right)$ as red lines in Fig. 1.34. Some of the variables seem to be strongly related. The most obvious relation is the negative dependence between $X_{13}$ and $X_{14}$. It can also be argued that a strong dependence exists between $X_{12}$ and $X_{14}$ since no red lines are drawn in the lower part of $X_{12}$. The opposite can be said about $X_{11}$ : there are only red lines plotted in the lower part of this variable. Low values of $X_{11}$ induce high values of $X_{14}$.

For the PCP, the variables have been rescaled over the interval $[0,1]$ for better graphical representations. The PCP shows that the variables are not distributed in a symmetric manner. It can be clearly seen that the values of $X_{1}$ and $X_{9}$ are much more concentrated around 0 . Therefore it makes sense to consider transformations of the original data.

## The Scatterplot Matrix

One characteristic of PCPs is that many lines are drawn on top of each other. This problem is reduced by depicting the variables in pairs of scatterplots. Including all 14 variables in one large scatterplot matrix is possible, but makes it hard to see anything from the plots. Therefore, for illustratory purposes we will analyse only one such matrix from a subset of the variables in Fig. 1.35. On the basis of the PCP and the scatterplot matrix we would like to interpret each of the 13 variables and their eventual relation to the 14th variable. Included in the figure are images for

Fig. 1.35 Scatterplot matrix for variables $X_{1}, \ldots, X_{5}$ and $X_{14}$ of the Boston housing data Q MVAdrafthousing

Fig. 1.36 Scatterplot matrix for variables $\tilde{X}_{1}, \ldots, \tilde{X}_{5}$ and $\tilde{X}_{14}$ of the Boston housing data Q
MVAdrafthousingt

$X_{1}-X_{5}$ and $X_{14}$, although each variable is discussed in detail below. All references made to scatterplots in the following refer to Fig. 1.35.

## Per-Capita Crime Rate $\boldsymbol{X}_{1}$

Taking the logarithm makes the variable's distribution more symmetric. This can be seen in the boxplot of $\tilde{X}_{1}$ in Fig. 1.37 which shows that the median and the mean have moved closer to each other than they were for the original $X_{1}$. Plotting the KDE of $\tilde{X}_{1}=\log \left(X_{1}\right)$ would reveal that two sub-groups might exist with different mean values. However, taking a look at the scatterplots in Fig. 1.36 of the logarithms
which include $X_{1}$ does not clearly reveal such groups. Given that the scatterplot of $\log \left(X_{1}\right)$ vs. $\log \left(X_{14}\right)$ shows a relatively strong negative relation, it might be the case that the two sub-groups of $X_{1}$ correspond to houses with two different price levels. This is confirmed by the two boxplots shown to the right of the $X_{1}$ vs. $X_{2}$ scatterplot (in Fig. 1.35): the right boxplot's shape differs a lot from the black one's, having a much higher median and mean.

## Proportion of Residential Area Zoned for Large Lots $\boldsymbol{X}_{\mathbf{2}}$

It strikes the eye in Fig. 1.35 that there is a large cluster of observations for which $X_{2}$ is equal to 0 . It also strikes the eye that-as the scatterplot of $X_{1}$ vs. $X_{2}$ showsthere is a strong, though non-linear, negative relation between $X_{1}$ and $X_{2}$; almost all observations for which $X_{2}$ is high have an $X_{1}$-value close to zero, and vice versa, many observations for which $X_{2}$ is zero have quite a high per-capita crime rate $X_{1}$. This could be due to the location of the areas, e.g. urban districts might have a higher crime rate and at the same time it is unlikely that any residential land would be zoned in a generous manner.

As far as the house prices are concerned, it can be said that there seems to be no clear (linear) relation between $X_{2}$ and $X_{14}$, but it is obvious that the more expensive houses are situated in areas where $X_{2}$ is large (this can be seen from the two boxplots on the second position of the diagonal, where the red one has a clearly higher mean/median than the black one).

## Proportion of Non-retail Business Acres $\boldsymbol{X}_{3}$

The PCP (in Fig. 1.34) as well as the scatterplot of $X_{3}$ vs. $X_{14}$ shows an obvious negative relation between $X_{3}$ and $X_{14}$. The relationship between the logarithms of both variables seems to be almost linear. This negative relation might be explained by the fact that non-retail business sometimes causes annoying sounds and other pollution. Therefore, it seems reasonable to use $X_{3}$ as an explanatory variable for the prediction of $X_{14}$ in a linear-regression analysis.

As far as the distribution of $X_{3}$ is concerned, it can be said that the KDE of $X_{3}$ clearly has two peaks, which indicates that there are two sub-groups. According to the negative relation between $X_{3}$ and $X_{14}$ it could be the case that one sub-group corresponds to the more expensive houses and the other one to the cheaper houses.

## Charles River Dummy Variable $\boldsymbol{X}_{\mathbf{4}}$

The observation made from the PCP that there are more expensive houses than cheap houses situated on the banks of the Charles River is confirmed by inspecting the scatterplot matrix. Still, we might have some doubt that proximity to the river influences house prices. Looking at the original data set, it becomes clear that the
observations for which $X_{4}$ equals one are districts that are close to each other. Apparently, the Charles River does not flow through very many different districts. Thus, it may be pure coincidence that the more expensive districts are close to the Charles River-their high values might be caused by many other factors such as the pupil/teacher ratio or the proportion of non-retail business acres.

## Nitric Oxides Concentration $X_{5}$

The scatterplot of $X_{5}$ vs. $X_{14}$ and the separate boxplots of $X_{5}$ for more and less expensive houses reveal a clear negative relation between the two variables. As it was the main aim of the authors of the original study to determine whether pollution had an influence on housing prices, it should be considered very carefully whether $X_{5}$ can serve as an explanatory variable for price $X_{14}$. A possible reason against it being an explanatory variable is that people might not like to live in areas where the emissions of nitric oxides are high. Nitric oxides are emitted mainly by automobiles, by factories and from heating private homes. However, as one can imagine there are many good reasons besides nitric oxides not to live in urban or industrial areas. Noise pollution, for example, might be a much better explanatory variable for the price of housing units. As the emission of nitric oxides is usually accompanied by noise pollution, using $X_{5}$ as an explanatory variable for $X_{14}$ might lead to the false conclusion that people run away from nitric oxides, whereas in reality it is noise pollution that they are trying to escape.

## Average Number of Rooms per Dwelling $X_{6}$

The number of rooms per dwelling is a possible measure of the size of the houses. Thus we expect $X_{6}$ to be strongly correlated with $X_{14}$ (the houses' median price). Indeed-apart from some outliers-the scatterplot of $X_{6}$ vs. $X_{14}$ shows a point cloud which is clearly upward-sloping and which seems to be a realisation of a linear dependence of $X_{14}$ on $X_{6}$. The two boxplots of $X_{6}$ confirm this notion by showing that the quartiles, the mean and the median are all much higher for the red than for the black boxplot.

## Proportion of Owner-Occupied Units Built Prior to $1940 X_{7}$

There is no clear connection visible between $X_{7}$ and $X_{14}$. There could be a weak negative correlation between the two variables, since the (red) boxplot of $X_{7}$ for the districts whose price is above the median price indicates a lower mean and median than the (black) boxplot for the district whose price is below the median price. The fact that the correlation is not so clear could be explained by two opposing effects. On the one hand, house prices should decrease if the older houses are not in a good shape. On the other hand, prices could increase, because people often like older
houses better than newer houses, preferring their atmosphere of space and tradition. Nevertheless, it seems reasonable that the age of the houses has an influence on their price $X_{14}$.

Raising $X_{7}$ to the power of 2.5 reveals again that the data set might consist of two sub-groups. But in this case it is not obvious that the sub-groups correspond to more expensive or cheaper houses. One can furthermore observe a negative relation between $X_{7}$ and $X_{8}$. This could reflect the way the Boston metropolitan area developed over time; the districts with the newer buildings are further away from employment centers and industrial facilities.

## Weighted Distance to Five Boston Employment Centers $X_{8}$

Since most people like to live close to their place of work, we expect a negative relation between the distances to the employment centers and house prices. The scatterplot hardly reveals any dependence, but the boxplots of $X_{8}$ indicate that there might be a slightly positive relation as the red boxplot's median and mean are higher than the black ones. Again, there might be two effects in opposite directions at work here. The first is that living too close to an employment centre might not provide enough shelter from the pollution created there. The second, as mentioned above, is that people do not travel very far to their workplace.

## Index of Accessibility to Radial Highways $X_{9}$

The first obvious thing one can observe from the scatterplots, as well in the histograms and the KDEs, is that there are two sub-groups of districts containing $X_{9}$ values which are close to the respective group's mean. The scatterplots deliver no hint as to what might explain the occurrence of these two sub-groups. The boxplots indicate that for the cheaper and for the more expensive houses the average of $X_{9}$ is almost the same.

## Full-Value Property Tax $X_{10}$

$X_{10}$ shows behaviour similar to that of $X_{9}$ : two sub-groups exist. A downwardsloping curve seems to underlie the relation of $X_{10}$ and $X_{14}$. This is confirmed by the two boxplots drawn for $X_{10}$ : the red one has a lower mean and median than the black one.

## Pupil/Teacher Ratio $X_{11}$

The red and black boxplots of $X_{11}$ indicate a negative relation between $X_{11}$ and $X_{14}$. This is confirmed by inspection of the scatterplot of $X_{11}$ vs. $X_{14}$ : The point cloud is
downward sloping, i.e. the less teachers there are per pupil, the less people pay on median for their dwellings.

## Proportion of African-American $B, X_{12}=1000(B-0.63)^{2} I(B<0.63)$

Interestingly, $X_{12}$ is negatively-though not linearly-correlated with $X_{3}, X_{7}$ and $X_{11}$, whereas it is positively related with $X_{14}$. Looking at the data set reveals that for almost all districts $X_{12}$ takes on a value around 390 . Since $B$ cannot be larger than 0.63 , such values can only be caused by $B$ close to zero. Therefore, the higher $X_{12}$ is, the lower the actual proportion of African-Americans is. Among observations 405-470 there are quite a few that have a $X_{12}$ that is much lower than 390. This means that in these districts the proportion of African-Americans is above zero. We can observe two clusters of points in the scatterplots of $\log \left(X_{12}\right)$ : one cluster for which $X_{12}$ is close to 390 and a second one for which $X_{12}$ is between 3 and 100. When $X_{12}$ is positively related with another variable, the actual proportion of African-Americans is negatively correlated with this variable and vice versa. This means that African-Americans live in areas where there is a high proportion of nonretail business land, where there are older houses and where there is a high (i.e. bad) pupil/teacher ratio. It can be observed that districts with housing prices above the median can only be found where the proportion of African-Americans is virtually zero.

## Proportion of Lower Status of the Population $X_{13}$

Of all the variables $X_{13}$ exhibits the clearest negative relation with $X_{14}$-hardly any outliers show up. Taking the square root of $X_{13}$ and the logarithm of $X_{14}$ transforms the relation into a linear one.

## Transformations

Since most of the variables exhibit an asymmetry with a higher density on the lefthand side, the following transformations are proposed:

$$
\begin{aligned}
& \widetilde{X_{1}}=\log \left(X_{1}\right) \\
& \widetilde{X_{2}}=X_{2} / 10 \\
& \widetilde{X_{3}}=\log \left(X_{3}\right) \\
& \widetilde{X_{4}} \quad \text { none, since } X_{4} \text { is binary } \\
& \widetilde{X_{5}}=\log \left(X_{5}\right)
\end{aligned}
$$

$$
\begin{aligned}
& \widetilde{X_{6}}=\log \left(X_{6}\right) \\
& \widetilde{X_{7}}=X_{7}^{2.5} / 10000 \\
& \widetilde{X_{8}}=\log \left(X_{8}\right) \\
& \widetilde{X_{9}}=\log \left(X_{9}\right) \\
& \widetilde{X_{10}}=\log \left(X_{10}\right) \\
& \widetilde{X_{11}}=\exp \left(0.4 \times X_{11}\right) / 1000 \\
& \widetilde{X_{12}}=X_{12} / 100 \\
& \widetilde{X_{13}}=\sqrt{X_{13}} \\
& \widetilde{X_{14}}=\log \left(X_{14}\right)
\end{aligned}
$$

Taking the logarithm or raising the variables to the power of something smaller than one helps to reduce the asymmetry. This is due to the fact that lower values move further away from each other, whereas the distance between greater values is reduced by these transformations.

Figure 1.37 displays boxplots for the original mean variance scaled variables as well as for the proposed transformed variables. The transformed variables' boxplots are more symmetric and have less outliers than the original variables' boxplots.

Fig. 1.37 Boxplots for all of the variables from the Boston housing data before and after the proposed transformations Q MVAboxbhd


### 1.10 Exercises

Exercise 1.1 Is the upper extreme always an outlier?
Exercise 1.2 Is it possible for the mean or the median to lie outside of the fourths or even outside of the outside bars?

Exercise 1.3 Assume that the data are normally distributed $N(0,1)$. What percentage of the data do you expect to lie outside the outside bars?

Exercise 1.4 What percentage of the data do you expect to lie outside the outside bars if we assume that the data are normally distributed $N\left(0, \sigma^{2}\right)$ with unknown variance $\sigma^{2}$ ?

Exercise 1.5 How would the five-number summary of the 15 largest US cities differ from that of the 50 largest US cities? How would the five-number summary of 15 observations of $N(0,1)$-distributed data differ from that of 50 observations from the same distribution?

Exercise 1.6 Is it possible that all five numbers of the five-number summary could be equal? If so, under what conditions?

Exercise 1.7 Suppose we have 50 observations of $X \sim N(0,1)$ and another 50 observations of $Y \sim N(2,1)$. What would the 100 Flury faces look like if you had defined as face elements the face line and the darkness of hair? Do you expect any similar faces? How many faces do you think should look like observations of $Y$ even though they are $X$ observations?

Exercise 1.8 Draw a histogram for the mileage variable of the car data (Sect. 22.3). Do the same for the three groups (USA, Japan, and Europe). Do you obtain a similar conclusion as in the parallel boxplot in Fig. 1.3 for these data?

Exercise 1.9 Use some bandwidth selection criterion to calculate the optimally chosen bandwidth $h$ for the diagonal variable of the bank notes. Would it be better to have one bandwidth for the two groups?

Exercise 1.10 In Fig. 1.9 the densities overlap in the region of diagonal $\approx 140.4$. We partially observed this in the boxplot of Fig. 1.4. Our aim is to separate the two groups. Will we be able to do this effectively on the basis of this diagonal variable alone?

Exercise 1.11 Draw a parallel coordinates plot for the car data.
Exercise 1.12 How would you identify discrete variables (variables with only a limited number of possible outcomes) on a parallel coordinates plot?

Exercise 1.13 True or false: the height of the bars of a histogram are equal to the relative frequency with which observations fall into the respective bins.

Exercise 1.14 True or false: kernel density estimates must always take on a value between 0 and 1. (Hint: Which quantity connected with the density function has to
be equal to 1? Does this property imply that the density function has to always be less than 1?)

Exercise 1.15 Let the following data set represent the heights of 13 students taking the Applied Multivariate Statistical Analysis course:

$$
1.72,1.83,1.74,1.79,1.94,1.81,1.66,1.60,1.78,1.77,1.85,1.70,1.76
$$

1. Find the corresponding five-number summary.
2. Construct the boxplot.
3. Draw a histogram for this data set.

Exercise 1.16 Describe the unemployment data (see Table 22.19) that contain unemployment rates of all German Federal States using various descriptive techniques.

Exercise 1.17 Using yearly population data (see Sect. 22.20), generate

1. a boxplot (choose one of variables)
2. an Andrew's Curve (choose ten data points)
3. a scatterplot
4. a histogram (choose one of the variables)

What do these graphs tell you about the data and their structure?
Exercise 1.18 Make a draftman plot for the car data with the variables

$$
\begin{aligned}
& X_{1}=\text { price }, \\
& X_{2}=\text { mileage }, \\
& X_{8}=\text { weight }, \\
& X_{9}=\text { length } .
\end{aligned}
$$

Move the brush into the region of heavy cars. What can you say about price, mileage and length? Move the brush onto high fuel economy. Mark the Japanese, European and American cars. You should find the same condition as in boxplot Fig. 1.3.

Exercise 1.19 What is the form of a scatterplot of two independent random variables $X_{1}$ and $X_{2}$ with standard normal distribution?

Exercise 1.20 Rotate a three-dimensional standard normal point cloud in 3D space. Does it "almost look the same from all sides"? Can you explain why or why not?

Exercise 1.21 There are many reasons for using hexagons to visualise the structure of data.

1. Hexagons have the property of "symmetry of nearest neighbours" which lacks in square bins.


Fig. 1.38 Hexagon binning algorithm $Q_{\text {MVAhexaAl }}$
2. Hexagons have the maximum number of sides that a polygon can have for a regular tessellation of the plane.
3. Hexagons are visually less biased for displaying densities than other regular tessellations.

The hexagon binning algorithm is as follows:

1. Decrease $y$-axis variable by a factor of $\sqrt{3}$ (making the calculation more quickly)
2. Create a dual lattice (circle and star lines in Fig. 1.38)
3. Bin each point into a pair of near neighbour rectangles
4. Choose the closest of the rectangle centers (adjusting for $\sqrt{3}$ )

The rectangles created from dual lattice have length $h_{x}$ (bin width of hexagons) and height $h_{y}=\sqrt{3} h_{x}$. From these rectangles we can get hexagons with bin width $h_{x}$. The first point of the star lattice has coordinates $x_{0}$ and $y_{0}$. The other star points will have coordinates $x_{0}+k_{1} h_{x}$ and $y_{0}+l_{1} h_{y}$, where $k_{1}, l_{1}=1,2, \ldots$ The first point of the circle lattice has coordinates $x_{0}+\frac{h_{x}}{2}$ and $y_{0}+\frac{\sqrt{3} h_{x}}{2}$. Other circle points are calculated like star points. Suppose an arbitrary point with coordinates $x, y$ lies in the intersection of two near neighbour rectangles. What's the distance from this point to one of two corners?

Part II
Multivariate Random Variables

## Chapter 2 <br> A Short Excursion into Matrix Algebra

This chapter serves as a reminder of basic concepts of matrix algebra, which are particularly useful in multivariate analysis. It also introduces the notations used in this book for vectors and matrices. Eigenvalues and eigenvectors play an important role in multivariate techniques. In Sects. 2.2 and 2.3, we present the spectral decomposition of matrices and consider the maximisation (minimisation) of quadratic forms given some constraints.

In analysing the multivariate normal distribution, partitioned matrices appear naturally. Some of the basic algebraic properties are given in Sect. 2.5. These properties will be heavily used in Chaps. 4 and 5.

The geometry of the multinormal and the geometric interpretation of the multivariate techniques (Part III) intensively uses the notion of angles between two vectors, the projection of a point on a vector and the distances between two points. These ideas are introduced in Sect. 2.6.

### 2.1 Elementary Operations

A matrix $\mathcal{A}$ is a system of numbers with $n$ rows and $p$ columns:

$$
\mathcal{A}=\left(\begin{array}{ccccc}
a_{11} & a_{12} & \ldots & \ldots & \ldots \\
\vdots & & a_{1 p} \\
\vdots & \vdots & \ddots & & \\
\vdots & \vdots & & \ddots & \\
\vdots & \vdots & & & \vdots \\
\vdots \\
a_{n 1} & a_{n 2} & \ldots & \ldots & \ldots
\end{array}\right)
$$

Table 2.1 Special matrices and vectors

| Name | Definition | Notation | Example |
| :---: | :---: | :---: | :---: |
| Scalar | $p=n=1$ | $a$ | 3 |
| Column vector | $p=1$ | $a$ | $\binom{1}{3}$ |
| Row vector | $n=1$ | $a^{\top}$ | $\left(\begin{array}{ll}1 & 3\end{array}\right)$ |
| 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 | $n=p$ | $\mathcal{A}(p \times p)$ | $\left(\begin{array}{ll}2 & 0 \\ 0 & 2\end{array}\right)$ |
| 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 | $a_{i j}=1, n=p$ | $1_{n} 1_{n}^{\top}$ | $\left(\begin{array}{ll}1 & 1 \\ 1 & 1\end{array}\right)$ |
| Symmetric matrix | $a_{i j}=a_{j i}$ |  | $\left(\begin{array}{ll}1 & 2 \\ 2 & 3\end{array}\right)$ |
| Null matrix | $a_{i j}=0$ | 0 | $\left(\begin{array}{ll}0 & 0 \\ 0 & 0\end{array}\right)$ |
| Upper triangular matrix | $a_{i j}=0, i<j$ |  | $\left(\begin{array}{lll}1 & 2 & 4 \\ 0 & 1 & 3 \\ 0 & 0 & 1\end{array}\right)$ |
| Idempotent matrix | $\mathcal{A} \mathcal{A}=\mathcal{A}$ |  | $\left(\begin{array}{lll}1 & 0 & 0 \\ 0 & \frac{1}{2} & \frac{1}{2} \\ 0 & \frac{1}{2} & \frac{1}{2}\end{array}\right)$ |
| Orthogonal matrix | $\mathcal{A}^{\top} \mathcal{A}=\mathcal{I}=\mathcal{A} \mathcal{A}^{\top}$ |  | $\left(\begin{array}{rr}\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}}\end{array}\right)$ |

We also write $\left(a_{i j}\right)$ for $\mathcal{A}$ and $\mathcal{A}(n \times p)$ to indicate the numbers of rows and columns. Vectors are matrices with one column and are denoted as $x$ or $x(p \times 1)$. Special matrices and vectors are defined in Table 2.1. Note that we use small letters for scalars as well as for vectors.

## Matrix Operations

Elementary operations are summarised below:

$$
\begin{aligned}
\mathcal{A}^{\top} & =\left(a_{j i}\right) \\
\mathcal{A}+\mathcal{B} & =\left(a_{i j}+b_{i j}\right) \\
\mathcal{A}-\mathcal{B} & =\left(a_{i j}-b_{i j}\right) \\
c \cdot \mathcal{A} & =\left(c \cdot a_{i j}\right) \\
\mathcal{A} \cdot \mathcal{B} & =\mathcal{A}(n \times p) \mathcal{B}(p \times m)=\mathcal{C}(n \times m)=\left(c_{i j}\right)=\left(\sum_{j=1}^{p} a_{i j} b_{j k}\right) .
\end{aligned}
$$

## Properties of Matrix Operations

$$
\begin{aligned}
\mathcal{A}+\mathcal{B} & =\mathcal{B}+\mathcal{A} \\
\mathcal{A}(\mathcal{B}+\mathcal{C}) & =\mathcal{A B}+\mathcal{A C} \\
\mathcal{A}(\mathcal{B C}) & =(\mathcal{A B}) \mathcal{C} \\
\left(\mathcal{A}^{\top}\right)^{\top} & =\mathcal{A} \\
(\mathcal{A B})^{\top} & =\mathcal{B}^{\top} \mathcal{A}^{\top}
\end{aligned}
$$

## Matrix Characteristics

## Rank

The $\operatorname{rank}$, $\operatorname{rank}(\mathcal{A})$, of a matrix $\mathcal{A}(n \times p)$ is defined as the maximum number of linearly independent rows (columns). A set of $k$ rows $a_{j}$ of $\mathcal{A}(n \times p)$ are said to be linearly independent if $\sum_{j=1}^{k} c_{j} a_{j}=0_{p}$ implies $c_{j}=0, \forall j$, where $c_{1}, \ldots, c_{k}$ are scalars. In other words no rows in this set can be expressed as a nontrivial linear combination of the $(k-1)$ remaining rows.

## Trace

The trace of a matrix $\mathcal{A}(p \times p)$ is the sum of its diagonal elements

$$
\operatorname{tr}(\mathcal{A})=\sum_{i=1}^{p} a_{i i}
$$

## Determinant

The determinant is an important concept of matrix algebra. For a square matrix $\mathcal{A}$, it is defined as:

$$
\operatorname{det}(\mathcal{A})=|\mathcal{A}|=\sum(-1)^{|\tau|} a_{1 \tau(1)} \ldots a_{p \tau(p)}
$$

the summation is over all permutations $\tau$ of $\{1,2, \ldots, p\}$, and $|\tau|=0$ if the permutation can be written as a product of an even number of transpositions and $|\tau|=1$ otherwise. Some properties of determinant of a matrix are:

$$
\begin{aligned}
\left|\mathcal{A}^{\top}\right| & =|\mathcal{A}| \\
|\mathcal{A B}| & =|\mathcal{A}| \cdot|\mathcal{A}| \\
|c \mathcal{A}| & =c^{n}|\mathcal{A}| .
\end{aligned}
$$

Example 2.1 In the case of $p=2, \mathcal{A}=\left(\begin{array}{ll}a_{11} & a_{12} \\ a_{21} & a_{22}\end{array}\right)$ and we can permute the digits " 1 " and " 2 " once or not at all. So,

$$
|\mathcal{A}|=a_{11} a_{22}-a_{12} a_{21}
$$

## Transpose

For $\mathcal{A}(n \times p)$ and $\mathcal{B}(p \times n)$

$$
\left(\mathcal{A}^{\top}\right)^{\top}=\mathcal{A}, \text { and }(\mathcal{A B})^{\top}=\mathcal{B}^{\top} \mathcal{A}^{\top} .
$$

## Inverse

If $|\mathcal{A}| \neq 0$ and $\mathcal{A}(p \times p)$, then the inverse $\mathcal{A}^{-1}$ exists:

$$
\mathcal{A} \mathcal{A}^{-1}=\mathcal{A}^{-1} \mathcal{A}=\mathcal{I}_{p}
$$

For small matrices, the inverse of $\mathcal{A}=\left(a_{i j}\right)$ can be calculated as

$$
\mathcal{A}^{-1}=\frac{\mathcal{C}}{|\mathcal{A}|},
$$

where $\mathcal{C}=\left(c_{i j}\right)$ is the adjoint matrix of $\mathcal{A}$. The elements $c_{j i}$ of $\mathcal{C}^{\top}$ are the co-factors of $\mathcal{A}$ :

$$
c_{j i}=(-1)^{i+j}\left|\begin{array}{cccccc}
a_{11} & \ldots & a_{1(j-1)} & a_{1(j+1)} & \ldots & a_{1 p} \\
\vdots & & & & & \\
a_{(i-1) 1} & \ldots & a_{(i-1)(j-1)} & a_{(i-1)(j+1)} & \ldots & a_{(i-1) p} \\
a_{(i+1) 1} & \ldots & a_{(i+1)(j-1)} & a_{(i+1)(j+1)} & \ldots & a_{(i+1) p} \\
\vdots & & & & & \\
a_{p 1} & \ldots & a_{p(j-1)} & a_{p(j+1)} & \ldots & a_{p p}
\end{array}\right| .
$$

The relationship between determinant and inverse of matrix $\mathcal{A}$ is $\left|\mathcal{A}^{-1}\right|=|\mathcal{A}|^{-1}$.

## G-Inverse

A more general concept is the $G$-inverse (Generalised Inverse) $\mathcal{A}^{-}$which satisfies the following:

$$
\mathcal{A} \mathcal{A}^{-} \mathcal{A}=\mathcal{A} .
$$

Later we will see that there may be more than one $G$-inverse.
Example 2.2 The generalised inverse can also be calculated for singular matrices. We have:

$$
\left(\begin{array}{ll}
1 & 0 \\
0 & 0
\end{array}\right)\left(\begin{array}{ll}
1 & 0 \\
0 & 0
\end{array}\right)\left(\begin{array}{ll}
1 & 0 \\
0 & 0
\end{array}\right)=\left(\begin{array}{ll}
1 & 0 \\
0 & 0
\end{array}\right),
$$

which means that the generalised inverse of $\mathcal{A}=\left(\begin{array}{ll}1 & 0 \\ 0 & 0\end{array}\right)$ is $\mathcal{A}^{-}=\left(\begin{array}{ll}1 & 0 \\ 0 & 0\end{array}\right)$ even though the inverse matrix of $\mathcal{A}$ does not exist in this case.

## Eigenvalues, Eigenvectors

Consider a $(p \times p)$ matrix $\mathcal{A}$. If there a scalar $\lambda$ and a vector $\gamma$ exists such as

$$
\begin{equation*}
\mathcal{A} \gamma=\lambda \gamma, \tag{2.1}
\end{equation*}
$$

then we call

$$
\begin{array}{ll}
\lambda & \text { an eigenvalue } \\
\gamma & \text { an eigenvector. }
\end{array}
$$

It can be proven that an eigenvalue $\lambda$ is a root of the $p$-th order polynomial $\mid \mathcal{A}-$ $\lambda I_{p} \mid=0$. Therefore, there are up to $p$ eigenvalues $\lambda_{1}, \lambda_{2}, \ldots, \lambda_{p}$ of $\mathcal{A}$. For each eigenvalue $\lambda_{j}$, a corresponding eigenvector $\gamma_{j}$ exists given by Eq. (2.1). Suppose the matrix $\mathcal{A}$ has the eigenvalues $\lambda_{1}, \ldots, \lambda_{p}$. Let $\Lambda=\operatorname{diag}\left(\lambda_{1}, \ldots, \lambda_{p}\right)$.

The determinant $|\mathcal{A}|$ and the trace $\operatorname{tr}(\mathcal{A})$ can be rewritten in terms of the eigenvalues:

$$
\begin{align*}
|\mathcal{A}| & =|\Lambda|=\prod_{j=1}^{p} \lambda_{j}  \tag{2.2}\\
\operatorname{tr}(\mathcal{A}) & =\operatorname{tr}(\Lambda)=\sum_{j=1}^{p} \lambda_{j} . \tag{2.3}
\end{align*}
$$

An idempotent matrix $\mathcal{A}$ (see the definition in Table 2.1) can only have eigenvalues in $\{0,1\}$ therefore $\operatorname{tr}(\mathcal{A})=\operatorname{rank}(\mathcal{A})=$ number of eigenvalues $\neq 0$.
Example 2.3 Let us consider the matrix $\mathcal{A}=\left(\begin{array}{ccc}1 & 0 & 0 \\ 0 & \frac{1}{2} & \frac{1}{2} \\ 0 & \frac{1}{2} & \frac{1}{2}\end{array}\right)$. It is easy to verify that $\mathcal{A} \mathcal{A}=\mathcal{A}$ which implies that the matrix $\mathcal{A}$ is idempotent.

We know that the eigenvalues of an idempotent matrix are equal to 0 or 1 . In this case, the eigenvalues of $\mathcal{A}$ are $\lambda_{1}=1, \lambda_{2}=1$, and $\lambda_{3}=0$ since $\left(\begin{array}{lll}1 & 0 & 0 \\ 0 & \frac{1}{2} & \frac{1}{2} \\ 0 & \frac{1}{2} & \frac{1}{2}\end{array}\right)\left(\begin{array}{l}1 \\ 0 \\ 0\end{array}\right)=$ $1\left(\begin{array}{l}1 \\ 0 \\ 0\end{array}\right),\left(\begin{array}{lll}1 & 0 & 0 \\ 0 & \frac{1}{2} & \frac{1}{2} \\ 0 & \frac{1}{2} & \frac{1}{2}\end{array}\right)\left(\begin{array}{r}0 \\ \frac{\sqrt{2}}{2} \\ \frac{\sqrt{2}}{2}\end{array}\right)=1\left(\begin{array}{r}0 \\ \frac{\sqrt{2}}{2} \\ \frac{\sqrt{2}}{2}\end{array}\right)$, and $\left(\begin{array}{lll}1 & 0 & 0 \\ 0 & \frac{1}{2} & \frac{1}{2} \\ 0 & \frac{1}{2} & \frac{1}{2}\end{array}\right)\left(\begin{array}{r}0 \\ \frac{\sqrt{2}}{2} \\ -\frac{\sqrt{2}}{2}\end{array}\right)=0\left(\begin{array}{r}0 \\ \frac{\sqrt{2}}{2} \\ -\frac{\sqrt{2}}{2}\end{array}\right)$.

Using formulas (2.2) and (2.3), we can calculate the trace and the determinant of $\mathcal{A}$ from the eigenvalues: $\operatorname{tr}(\mathcal{A})=\lambda_{1}+\lambda_{2}+\lambda_{3}=2,|\mathcal{A}|=\lambda_{1} \lambda_{2} \lambda_{3}=0$, and $\operatorname{rank}(\mathcal{A})=2$.

## Properties of Matrix Characteristics

$\mathcal{A}(n \times n), \mathcal{B}(n \times n), c \in \mathbb{R}$

$$
\begin{align*}
\operatorname{tr}(\mathcal{A}+\mathcal{B}) & =\operatorname{tr} \mathcal{A}+\operatorname{tr} \mathcal{B}  \tag{2.4}\\
\operatorname{tr}(c \mathcal{A}) & =c \operatorname{tr} \mathcal{A} \tag{2.5}
\end{align*}
$$

$$
\begin{align*}
|c \mathcal{A}| & =c^{n}|\mathcal{A}|  \tag{2.6}\\
|\mathcal{A B}| & =|\mathcal{B A}|=|\mathcal{A}||\mathcal{B}| \tag{2.7}
\end{align*}
$$

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

$$
\begin{align*}
\operatorname{tr}(\mathcal{A} \cdot \mathcal{B}) & =\operatorname{tr}(\mathcal{B} \cdot \mathcal{A})  \tag{2.8}\\
\operatorname{rank}(\mathcal{A}) & \leq \min (n, p) \\
\operatorname{rank}(\mathcal{A}) & \geq 0  \tag{2.9}\\
\operatorname{rank}(\mathcal{A}) & =\operatorname{rank}\left(\mathcal{A}^{\top}\right)  \tag{2.10}\\
\operatorname{rank}\left(\mathcal{A}^{\top} A\right) & =\operatorname{rank}(\mathcal{A})  \tag{2.11}\\
\operatorname{rank}(\mathcal{A}+\mathcal{B}) & \leq \operatorname{rank}(\mathcal{A})+\operatorname{rank}(\mathcal{B})  \tag{2.12}\\
\operatorname{rank}(\mathcal{A B}) & \leq \min \{\operatorname{rank}(\mathcal{A}), \operatorname{rank}(\mathcal{B})\} \tag{2.13}
\end{align*}
$$

$\mathcal{A}(n \times p), \mathcal{B}(p \times q), \mathcal{C}(q \times n)$

$$
\begin{align*}
\operatorname{tr}(\mathcal{A B C}) & =\operatorname{tr}(\mathcal{B C \mathcal { A } )} \\
& =\operatorname{tr}(\mathcal{C A B})  \tag{2.14}\\
\operatorname{rank}(\mathcal{A B C}) & =\operatorname{rank}(\mathcal{B}) \quad \text { for nonsingular } \mathcal{A}, \mathcal{C} \tag{2.15}
\end{align*}
$$

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

$$
\begin{align*}
\left|\mathcal{A}^{-1}\right| & =|\mathcal{A}|^{-1}  \tag{2.16}\\
\operatorname{rank}(\mathcal{A}) & =p \quad \text { if and only if } \mathcal{A} \text { is nonsingular. } \tag{2.17}
\end{align*}
$$

\(\left.$$
\begin{array}{|l|l|}\hline \hookrightarrow & \text { The determinant }|\mathcal{A}| \text { is the product of the eigenvalues of } \mathcal{A} . \\
\hline & \hookrightarrow \text { The inverse of a matrix } \mathcal{A} \text { exists if }|\mathcal{A}| \neq 0 . \\
\hline & \hookrightarrow \text { The trace } \operatorname{tr}(\mathcal{A}) \text { is the sum of the eigenvalues of } \mathcal{A} . \\
\hline & \hookrightarrow\end{array}
$$ \begin{array}{l}The sum of the traces of two matrices equals the trace of the sum <br>

of the two matrices.\end{array}\right]\)|  | The trace $\operatorname{tr}(\mathcal{A B})$ equals $\operatorname{tr}(\mathcal{B A})$. |
| :--- | :--- |
|  | $\hookrightarrow$ |
|  | The $\operatorname{rank}(\mathcal{A})$ is the maximal number of linearly independent rows <br> (columns) of $\mathcal{A}$. |

### 2.2 Spectral Decompositions

The computation of eigenvalues and eigenvectors is an important issue in the analysis of matrices. The spectral decomposition or Jordan decomposition links the structure of a matrix to the eigenvalues and the eigenvectors.

Theorem 2.1 (Jordan Decomposition) Each symmetric matrix $\mathcal{A}(p \times p)$ can be written as

$$
\begin{equation*}
\mathcal{A}=\Gamma \Lambda \Gamma^{\top}=\sum_{j=1}^{p} \lambda_{j} \gamma_{j} \gamma_{j}^{\top} \tag{2.18}
\end{equation*}
$$

where

$$
\Lambda=\operatorname{diag}\left(\lambda_{1}, \ldots, \lambda_{p}\right)
$$

and where

$$
\Gamma=\left(\gamma_{1}, \gamma_{2}, \ldots, \gamma_{p}\right)
$$

is an orthogonal matrix consisting of the eigenvectors $\gamma_{j}$ of $\mathcal{A}$.
Example 2.4 Suppose that $\mathcal{A}=\left(\begin{array}{ll}1 & 2 \\ 2 & 3\end{array}\right)$. The eigenvalues are found by solving $\mid \mathcal{A}-$ $\lambda \mathcal{I} \mid=0$. This is equivalent to

$$
\left|\begin{array}{cc}
1-\lambda & 2 \\
2 & 3-\lambda
\end{array}\right|=(1-\lambda)(3-\lambda)-4=0 .
$$

Hence, the eigenvalues are $\lambda_{1}=2+\sqrt{5}$ and $\lambda_{2}=2-\sqrt{5}$. The eigenvectors are $\gamma_{1}=(0.5257,0.8506)^{\top}$ and $\gamma_{2}=(0.8506,-0.5257)^{\top}$. They are orthogonal since $\gamma_{1}^{\top} \gamma_{2}=0$.

Using spectral decomposition, we can define powers of a matrix $\mathcal{A}(p \times p)$. Suppose $\mathcal{A}$ is a symmetric matrix with positive eigenvalues. Then by Theorem 2.1

$$
\mathcal{A}=\Gamma \Lambda \Gamma^{\top},
$$

and we define for some $\alpha \in \mathbb{R}$

$$
\begin{equation*}
\mathcal{A}^{\alpha}=\Gamma \Lambda^{\alpha} \Gamma^{\top}, \tag{2.19}
\end{equation*}
$$

where $\Lambda^{\alpha}=\operatorname{diag}\left(\lambda_{1}^{\alpha}, \ldots, \lambda_{p}^{\alpha}\right)$. In particular, we can easily calculate the inverse of the matrix $\mathcal{A}$. Suppose that the eigenvalues of $\mathcal{A}$ are positive. Then with $\alpha=-1$, we obtain the inverse of $\mathcal{A}$ from

$$
\begin{equation*}
\mathcal{A}^{-1}=\Gamma \Lambda^{-1} \Gamma^{\top} . \tag{2.20}
\end{equation*}
$$

Another interesting decomposition which is later used is given in the following theorem.

Theorem 2.2 (Singular Value Decomposition) Each matrix $\mathcal{A}(n \times p)$ with rank $r$ can be decomposed as

$$
\mathcal{A}=\Gamma \Lambda \Delta^{\top},
$$

where $\Gamma(n \times r)$ and $\Delta(p \times r)$. Both $\Gamma$ and $\Delta$ are column orthonormal, i.e. $\Gamma^{\top} \Gamma=$ $\Delta^{\top} \Delta=\mathcal{I}_{r}$ and $\Lambda=\operatorname{diag}\left(\lambda_{1}^{1 / 2}, \ldots, \lambda_{r}^{1 / 2}\right), \lambda_{j}>0$. The values $\lambda_{1}, \ldots, \lambda_{r}$ are the nonzero eigenvalues of the matrices $\mathcal{A} \mathcal{A}^{\top}$ and $\mathcal{A}^{\top} \mathcal{A} . \Gamma$ and $\Delta$ consist of the corresponding $r$ eigenvectors of these matrices.

This is obviously a generalisation of Theorem 2.1 (Jordan decomposition). With Theorem 2.2, we can find a $G$-inverse $\mathcal{A}^{-}$of $\mathcal{A}$. Indeed, define $\mathcal{A}^{-}=\Delta \Lambda^{-1} \Gamma^{\top}$. Then $\mathcal{A} \mathcal{A}^{-} \mathcal{A}=\Gamma \Lambda \Delta^{\top}=\mathcal{A}$. Note that the $G$-inverse is not unique.

Example 2.5 In Example 2.2, we showed that the generalised inverse of $\mathcal{A}=$ $\left(\begin{array}{ll}1 & 0 \\ 0 & 0\end{array}\right)$ is $\mathcal{A}^{-}\left(\begin{array}{ll}1 & 0 \\ 0 & 0\end{array}\right)$. The following also holds

$$
\left(\begin{array}{ll}
1 & 0 \\
0 & 0
\end{array}\right)\left(\begin{array}{ll}
1 & 0 \\
0 & 8
\end{array}\right)\left(\begin{array}{ll}
1 & 0 \\
0 & 0
\end{array}\right)=\left(\begin{array}{ll}
1 & 0 \\
0 & 0
\end{array}\right)
$$

which means that the matrix $\left(\begin{array}{ll}1 & 0 \\ 0 & 8\end{array}\right)$ is also a generalised inverse of $\mathcal{A}$.


## Summary

$\hookrightarrow$ The Jordan decomposition gives a representation of a symmetric matrix in terms of eigenvalues and eigenvectors.
$\hookrightarrow$ The eigenvectors belonging to the largest eigenvalues indicate the "main direction" of the data.

| Summary (continued) |  |
| :---: | :--- |
| $\hookrightarrow$ | The Jordan decomposition allows one to easily compute the power <br> of a symmetric matrix $\mathcal{A}: \mathcal{A}^{\alpha}=\Gamma \Lambda^{\alpha} \Gamma^{\top}$. |
| $\hookrightarrow$ | The singular value decomposition (SVD) is a generalisation of the <br>  <br>  <br> Jordan decomposition to non-quadratic matrices. |

### 2.3 Quadratic Forms

A quadratic form $Q(x)$ is built from a symmetric matrix $\mathcal{A}(p \times p)$ and a vector $x \in \mathbb{R}^{p}$ :

$$
\begin{equation*}
Q(x)=x^{\top} \mathcal{A} x=\sum_{i=1}^{p} \sum_{j=1}^{p} a_{i j} x_{i} x_{j} \tag{2.21}
\end{equation*}
$$

## Definiteness of Quadratic Forms and Matrices

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

A matrix $\mathcal{A}$ is called positive definite (semidefinite) if the corresponding quadratic form $Q($.$) is positive definite (semidefinite). We write \mathcal{A}>0(\geq 0)$.

Quadratic forms can always be diagonalised, as the following result shows.
Theorem 2.3 If $\mathcal{A}$ is symmetric and $Q(x)=x^{\top} \mathcal{A} x$ is the corresponding quadratic form, then there exists a transformation $x \mapsto \Gamma^{\top} x=y$ such that

$$
x^{\top} \mathcal{A} x=\sum_{i=1}^{p} \lambda_{i} y_{i}^{2},
$$

where $\lambda_{i}$ are the eigenvalues of $\mathcal{A}$.
Proof $\mathcal{A}=\Gamma \Lambda \Gamma^{\top}$. By Theorem 2.1 and $y=\Gamma^{\top} \alpha$ we have that $x^{\top} \mathcal{A} x=$ $x^{\top} \Gamma \Lambda \Gamma^{\top} x=y^{\top} \Lambda y=\sum_{i=1}^{p} \lambda_{i} y_{i}^{2}$.

Positive definiteness of quadratic forms can be deduced from positive eigenvalues.

Theorem 2.4 $\mathcal{A}>0$ if and only if all $\lambda_{i}>0, i=1, \ldots, p$.
Proof $0<\lambda_{1} y_{1}^{2}+\cdots+\lambda_{p} y_{p}^{2}=x^{\top} \mathcal{A} x$ for all $x \neq 0$ by Theorem 2.3.

Corollary 2.1 If $\mathcal{A}>0$, then $\mathcal{A}^{-1}$ exists and $|\mathcal{A}|>0$.
Example 2.6 The quadratic form $Q(x)=x_{1}^{2}+x_{2}^{2}$ corresponds to the matrix $\mathcal{A}=$ $\left(\begin{array}{ll}1 & 0 \\ 0 & 1\end{array}\right)$ with eigenvalues $\lambda_{1}=\lambda_{2}=1$ and is thus positive definite. The quadratic form $Q(x)=\left(x_{1}-x_{2}\right)^{2}$ corresponds to the matrix $\mathcal{A}=\left(\begin{array}{rr}1 & -1 \\ -1 & 1\end{array}\right)$ with eigenvalues $\lambda_{1}=2, \lambda_{2}=0$ and is positive semidefinite. The quadratic form $Q(x)=x_{1}^{2}-x_{2}^{2}$ with eigenvalues $\lambda_{1}=1, \lambda_{2}=-1$ is indefinite.

In the statistical analysis of multivariate data, we are interested in maximising quadratic forms given some constraints.

Theorem 2.5 If $\mathcal{A}$ and $\mathcal{B}$ are symmetric and $\mathcal{B}>0$, then the maximum of $\frac{x^{\top} \mathcal{A} x}{x^{\top} \mathcal{B} x}$ is given by the largest eigenvalue of $\mathcal{B}^{-1} \mathcal{A}$. More generally,

$$
\max _{x} \frac{x^{\top} \mathcal{A} x}{x^{\top} \mathcal{B} x}=\lambda_{1} \geq \lambda_{2} \geq \cdots \geq \lambda_{p}=\min _{x} \frac{x^{\top} \mathcal{A} x}{x^{\top} \mathcal{B} x}
$$

where $\lambda_{1}, \ldots, \lambda_{p}$ denote the eigenvalues of $\mathcal{B}^{-1} \mathcal{A}$. The vector which maximises (minimises) $\frac{x^{\top} \mathcal{A} x}{x^{\top} \mathcal{B} x}$ is the eigenvector of $\mathcal{B}^{-1} \mathcal{A}$ which corresponds to the largest (smallest) eigenvalue of $\mathcal{B}^{-1} \mathcal{A}$. If $x^{\top} \mathcal{B} x=1$, we get

$$
\max _{x} x^{\top} \mathcal{A} x=\lambda_{1} \geq \lambda_{2} \geq \cdots \geq \lambda_{p}=\min _{x} x^{\top} \mathcal{A} x
$$

Proof Denote norm of vector $x$ as $\|x\|=\sqrt{x^{\top} x}$. By definition, $\mathcal{B}^{1 / 2}=$ $\Gamma_{\mathcal{B}} \Lambda_{\mathcal{B}}^{1 / 2} \Gamma_{\mathcal{B}}^{\top}$ is symmetric. Then $x^{\top} \mathcal{B} x=\left\|x^{\top} \mathcal{B}^{1 / 2}\right\|^{2}=\left\|\mathcal{B}^{1 / 2} x\right\|^{2}$. Set $y=$ $\frac{\mathcal{B}^{1 / 2} x}{\left\|\mathcal{B}^{1 / 2} x\right\|}$, then

$$
\begin{equation*}
\max _{x} \frac{x^{\top} \mathcal{A} x}{x^{\top} \mathcal{B} x}=\max _{\left\{y: y^{\top} y=1\right\}} y^{\top} \mathcal{B}^{-1 / 2} \mathcal{A} \mathcal{B}^{-1 / 2} y . \tag{2.22}
\end{equation*}
$$

From Theorem 2.1, let

$$
\mathcal{B}^{-1 / 2} \mathcal{A} \mathcal{B}^{-1 / 2}=\Gamma \Lambda \Gamma^{\top}
$$

be the spectral decomposition of $\mathcal{B}^{-1 / 2} \mathcal{A} \mathcal{B}^{-1 / 2}$. Set

$$
z=\Gamma^{\top} y, \text { then } z^{\top} z=y^{\top} \Gamma \Gamma^{\top} y=y^{\top} y .
$$

Thus (2.22) is equivalent to

$$
\max _{\left\{z: z^{\top} z=1\right\}} z^{\top} \Lambda z=\max _{\left\{z: z^{\top} z=1\right\}} \sum_{i=1}^{p} \lambda_{i} z_{i}^{2} .
$$

But

$$
\max _{z} \sum \lambda_{i} z_{i}^{2} \leq \lambda_{1} \underbrace{\max _{z} \sum z_{i}^{2}}_{=1}=\lambda_{1}
$$

The maximum is thus obtained by $z=(1,0, \ldots, 0)^{\top}$, i.e.

$$
y=\gamma_{1}, \text { hence } x=\mathcal{B}^{-1 / 2} \gamma_{1} .
$$

Since $\mathcal{B}^{-1} \mathcal{A}$ and $\mathcal{B}^{-1 / 2} \mathcal{A} \mathcal{B}^{-1 / 2}$ have the same eigenvalues, the proof is complete.
To maximise (minimise) $x^{\top} \mathcal{A} x$ under $x^{\top} \mathcal{B} x=1$, below is another proof using the Lagrange method.

$$
\max _{x} x^{\top} \mathcal{A} x=\max _{x}\left[x^{\top} \mathcal{A} x-\lambda\left(x^{\top} \mathcal{B} x-1\right)\right] .
$$

The first derivative of it in respect to $x$ is equal to 0 :

$$
2 \mathcal{A} x-2 \lambda \mathcal{B} x=0
$$

so

$$
\mathcal{B}^{-1} \mathcal{A} x=\lambda x
$$

By the definition of eigenvector and eigenvalue, our maximiser $x^{*}$ is $\mathcal{B}^{-1} \mathcal{A}$ 's eigenvector corresponding to eigenvalue $\lambda$. So

$$
\max _{\left\{x: x^{\top} \mathcal{B} x=1\right\}} x^{\top} \mathcal{A} x=\max _{\left\{x: x^{\top} \mathcal{B} x=1\right\}} x^{\top} \mathcal{B} \mathcal{B}^{-1} \mathcal{A} x=\max _{\left\{x: x^{\top} \mathcal{B} x=1\right\}} x^{\top} \mathcal{B} \lambda x=\max \lambda
$$

which is just the maximum eigenvalue of $\mathcal{B}^{-1} \mathcal{A}$, and we choose the corresponding eigenvector as our maximiser $x^{*}$.
Example 2.7 Consider the matrices $\mathcal{A}=\left(\begin{array}{ll}1 & 2 \\ 2 & 3\end{array}\right) \quad$ and $\quad \mathcal{B}=\left(\begin{array}{ll}1 & 0 \\ 0 & 1\end{array}\right)$, we calculate $\mathcal{B}^{-1} \mathcal{A}=\left(\begin{array}{ll}1 & 2 \\ 2 & 3\end{array}\right)$. The biggest eigenvalue of the matrix $\mathcal{B}^{-1} \mathcal{A}$ is $2+$ $\sqrt{5}$. This means that the maximum of $x^{\top} \mathcal{A} x$ under the constraint $x^{\top} \mathcal{B} x=1$ is $2+\sqrt{5}$. Notice that the constraint $x^{\top} \mathcal{B} x=1$ corresponds to our choice of $\mathcal{B}$, to the points which lie on the unit circle $x_{1}^{2}+x_{2}^{2}=1$.

| $\hookrightarrow$ | A quadratic form can be described by a symmetric matrix $\mathcal{A}$. |
| :--- | :--- |
|  | Quadratic forms can always be diagonalised. |
|  | $\hookrightarrow$ | | Positive definiteness of a quadratic form is equivalent to positive- |
| :--- |
| ness of the eigenvalues of the matrix $\mathcal{A}$. |

### 2.4 Derivatives

For later sections of this book, it will be useful to introduce matrix notation for derivatives of a scalar function of a vector $x$, i.e. $f(x)$, with respect to $x$. Consider $f: \mathbb{R}^{p} \rightarrow \mathbb{R}$ and a $(p \times 1)$ vector $x$, then $\frac{\partial f(x)}{\partial x}$ is the column vector of partial derivatives $\left\{\frac{\partial f(x)}{\partial x_{j}}\right\}, j=1, \ldots, p$ and $\frac{\partial f(x)}{\partial x^{\top}}$ is the row vector of the same derivative $\left(\frac{\partial f(x)}{\partial x}\right.$ is called the gradient of $\left.f\right)$.

We can also introduce second order derivatives: $\frac{\partial^{2} f(x)}{\partial x \partial x^{\top}}$ is the $(p \times p)$ matrix of elements $\frac{\partial^{2} f(x)}{\partial x_{i} \partial x_{j}}, i=1, \ldots, p$ and $j=1, \ldots, p\left(\frac{\partial^{2} f(x)}{\partial x \partial x^{\top}}\right.$ is called the Hessian of $\left.f\right)$.

Suppose that $a$ is a $(p \times 1)$ vector and that $\mathcal{A}=\mathcal{A}^{\top}$ is a $(p \times p)$ matrix. Then

$$
\begin{gather*}
\frac{\partial a^{\top} x}{\partial x}=\frac{\partial x^{\top} a}{\partial x}=a,  \tag{2.23}\\
\frac{\partial x^{\top} \mathcal{A} x}{\partial x}=2 \mathcal{A} x . \tag{2.24}
\end{gather*}
$$

The Hessian of the quadratic form $Q(x)=x^{\top} \mathcal{A} x$ is:

$$
\begin{equation*}
\frac{\partial^{2} x^{\top} \mathcal{A} x}{\partial x \partial x^{\top}}=2 \mathcal{A} . \tag{2.25}
\end{equation*}
$$

Example 2.8 Consider the matrix

$$
\mathcal{A}=\left(\begin{array}{ll}
1 & 2 \\
2 & 3
\end{array}\right) .
$$

From formulas (2.24) and (2.25) it immediately follows that the gradient of $Q(x)=$ $x^{\top} \mathcal{A} x$ is

$$
\frac{\partial x^{\top} \mathcal{A} x}{\partial x}=2 \mathcal{A} x=2\left(\begin{array}{ll}
1 & 2 \\
2 & 3
\end{array}\right) x=\left(\begin{array}{ll}
2 x & 4 x \\
4 x & 6 x
\end{array}\right)
$$

and the Hessian is

$$
\frac{\partial^{2} x^{\top} \mathcal{A} x}{\partial x \partial x^{\top}}=2 \mathcal{A}=2\left(\begin{array}{ll}
1 & 2 \\
2 & 3
\end{array}\right)=\left(\begin{array}{ll}
2 & 4 \\
4 & 6
\end{array}\right) .
$$

### 2.5 Partitioned Matrices

Very often we will have to consider certain groups of rows and columns of a matrix $\mathcal{A}(n \times p)$. In the case of two groups, we have

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

where $\mathcal{A}_{i j}\left(n_{i} \times p_{j}\right), i, j=1,2, n_{1}+n_{2}=n$ and $p_{1}+p_{2}=p$.
If $\mathcal{B}(n \times p)$ is partitioned accordingly, we have:

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

An important particular case is the square matrix $\mathcal{A}(p \times p)$, partitioned in such a way that $\mathcal{A}_{11}$ and $\mathcal{A}_{22}$ are both square matrices (i.e. $n_{j}=p_{j}, j=1,2$ ). It can be verified that when $\mathcal{A}$ is non-singular $\left(\mathcal{A} \mathcal{A}^{-1}=\mathcal{I}_{p}\right)$ :

$$
\mathcal{A}^{-1}=\left(\begin{array}{ll}
\mathcal{A}^{11} & \mathcal{A}^{12}  \tag{2.26}\\
\mathcal{A}^{21} & \mathcal{A}^{22}
\end{array}\right)
$$

where

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

An alternative expression can be obtained by reversing the positions of $\mathcal{A}_{11}$ and $\mathcal{A}_{22}$ in the original matrix.

The following results will be useful if $\mathcal{A}_{11}$ is non-singular:

$$
\begin{equation*}
|\mathcal{A}|=\left|\mathcal{A}_{11}\right|\left|\mathcal{A}_{22}-\mathcal{A}_{21} \mathcal{A}_{11}^{-1} \mathcal{A}_{12}\right|=\left|\mathcal{A}_{11}\right|\left|\mathcal{A}_{22 \cdot 1}\right| . \tag{2.27}
\end{equation*}
$$

If $\mathcal{A}_{22}$ is non-singular, we have that:

$$
\begin{equation*}
|\mathcal{A}|=\left|\mathcal{A}_{22}\right|\left|\mathcal{A}_{11}-\mathcal{A}_{12} \mathcal{A}_{22}^{-1} \mathcal{A}_{21}\right|=\left|\mathcal{A}_{22}\right|\left|\mathcal{A}_{11 \cdot 2}\right| \tag{2.28}
\end{equation*}
$$

A useful formula is derived from the alternative expressions for the inverse and the determinant. For instance let

$$
\mathcal{B}=\left(\begin{array}{cc}
1 & b^{\top} \\
a & \mathcal{A}
\end{array}\right)
$$

where $a$ and $b$ are $(p \times 1)$ vectors and $\mathcal{A}$ is non-singular. We then have:

$$
\begin{equation*}
|\mathcal{B}|=\left|\mathcal{A}-a b^{\top}\right|=|\mathcal{A}|\left|1-b^{\top} \mathcal{A}^{-1} a\right| \tag{2.29}
\end{equation*}
$$

and equating the two expressions for $\mathcal{B}^{22}$, we obtain the following:

$$
\begin{equation*}
\left(\mathcal{A}-a b^{\top}\right)^{-1}=\mathcal{A}^{-1}+\frac{\mathcal{A}^{-1} a b^{\top} \mathcal{A}^{-1}}{1-b^{\top} \mathcal{A}^{-1} a} \tag{2.30}
\end{equation*}
$$

Example 2.9 Let's consider the matrix

$$
\mathcal{A}=\left(\begin{array}{ll}
1 & 2 \\
2 & 2
\end{array}\right)
$$

We can use formula (2.26) to calculate the inverse of a partitioned matrix, i.e. $\mathcal{A}^{11}=$ $-1, \mathcal{A}^{12}=\mathcal{A}^{21}=1, \mathcal{A}^{22}=-1 / 2$. The inverse of $\mathcal{A}$ is

$$
\mathcal{A}^{-1}=\left(\begin{array}{cc}
-1 & 1 \\
1 & -0.5
\end{array}\right) .
$$

It is also easy to calculate the determinant of $\mathcal{A}$ :

$$
|\mathcal{A}|=|1||2-4|=-2 .
$$

Let $\mathcal{A}(n \times p)$ and $\mathcal{B}(p \times n)$ be any two matrices and suppose that $n \geq p$. From (2.27) and (2.28) we can conclude that

$$
\left|\begin{array}{cc}
-\lambda \mathcal{I}_{n}-\mathcal{A}  \tag{2.31}\\
\mathcal{B} & \mathcal{I}_{p}
\end{array}\right|=(-\lambda)^{n-p}\left|\mathcal{B A}-\lambda \mathcal{I}_{p}\right|=\left|\mathcal{A B}-\lambda \mathcal{I}_{n}\right| .
$$

Since both determinants on the right-hand side of (2.31) are polynomials in $\lambda$, we find that the $n$ eigenvalues of $\mathcal{A B}$ yield the $p$ eigenvalues of $\mathcal{B A}$ plus the eigenvalue $0, n-p$ times.

The relationship between the eigenvectors is described in the next theorem.
Theorem 2.6 For $\mathcal{A}(n \times p)$ and $\mathcal{B}(p \times n)$, the nonzero eigenvalues of $\mathcal{A B}$ and $\mathcal{B A}$ are the same and have the same multiplicity. If $x$ is an eigenvector of $\mathcal{A B}$ for an eigenvalue $\lambda \neq 0$, then $y=\mathcal{B} x$ is an eigenvector of $\mathcal{B} \mathcal{A}$.

Corollary 2.2 For $\mathcal{A}(n \times p), \mathcal{B}(q \times n), a(p \times 1)$, and $b(q \times 1)$ we have

$$
\operatorname{rank}\left(\mathcal{A} a b^{\top} \mathcal{B}\right) \leq 1
$$

The nonzero eigenvalue, if it exists, equals $b^{\top} \mathcal{B} \mathcal{A} a$ (with eigenvector $\mathcal{A} a$ ).
Proof Theorem 2.6 asserts that the eigenvalues of $\mathcal{A} a b^{\top} \mathcal{B}$ are the same as those of $b^{\top} \mathcal{B} \mathcal{A} a$. Note that the matrix $b^{\top} \mathcal{B} \mathcal{A} a$ is a scalar and hence it is its own eigenvalue $\lambda_{1}$.

Applying $\mathcal{A} a b^{\top} \mathcal{B}$ to $\mathcal{A} a$ yields

$$
\left(\mathcal{A} a b^{\top} \mathcal{B}\right)(\mathcal{A} a)=(\mathcal{A} a)\left(b^{\top} \mathcal{B} \mathcal{A} a\right)=\lambda_{1} \mathcal{A} a
$$

### 2.6 Geometrical Aspects

## Distance

Let $x, y \in \mathbb{R}^{p}$. A distance $d$ is defined as a function
$d: \mathbb{R}^{2 p} \rightarrow \mathbb{R}_{+} \quad$ which fulfills $\left\{\begin{array}{ll}d(x, y)>0 & \forall x \neq y \\ d(x, y)=0 & \text { if and only if } x=y . \\ d(x, y) \leq d(x, z)+d(z, y) & \forall x, y, z\end{array}\right.$.
A Euclidean distance $d$ between two points $x$ and $y$ is defined as

$$
\begin{equation*}
d^{2}(x, y)=(x-y)^{T} \mathcal{A}(x-y) \tag{2.32}
\end{equation*}
$$

where $\mathcal{A}$ is a positive definite matrix $(\mathcal{A}>0) . \mathcal{A}$ is called a metric.

Fig. 2.1 Distance $d$


Fig. 2.2 Iso-distance sphere


Fig. 2.3 Iso-distance ellipsoid


Example 2.10 A particular case is when $\mathcal{A}=\mathcal{I}_{p}$, i.e.

$$
\begin{equation*}
d^{2}(x, y)=\sum_{i=1}^{p}\left(x_{i}-y_{i}\right)^{2} . \tag{2.33}
\end{equation*}
$$

Figure 2.1 illustrates this definition for $p=2$.
Note that the sets $E_{d}=\left\{x \in \mathbb{R}^{p} \mid\left(x-x_{0}\right)^{\top}\left(x-x_{0}\right)=d^{2}\right\}$, i.e. the spheres with radius $d$ and centre $x_{0}$, are the Euclidean $\mathcal{I}_{p}$ iso-distance curves from the point $x_{0}$ (see Fig. 2.2).

The more general distance (2.32) with a positive definite matrix $\mathcal{A}(\mathcal{A}>0)$ leads to the iso-distance curves

$$
\begin{equation*}
E_{d}=\left\{x \in \mathbb{R}^{p} \mid\left(x-x_{0}\right)^{\top} \mathcal{A}\left(x-x_{0}\right)=d^{2}\right\}, \tag{2.34}
\end{equation*}
$$

i.e. ellipsoids with centre $x_{0}$, matrix $\mathcal{A}$ and constant $d$ (see Fig. 2.3).

Let $\gamma_{1}, \gamma_{2}, \ldots, \gamma_{p}$ be the orthonormal eigenvectors of $\mathcal{A}$ corresponding to the eigenvalues $\lambda_{1} \geq \lambda_{2} \geq \cdots \geq \lambda_{p}$. The resulting observations are given in the next theorem.

Theorem 2.7 (i) The principal axes of $E_{d}$ are in the direction of $\gamma_{i} ; i=$ $1, \ldots, p$.
(ii) The half-lengths of the axes are $\sqrt{\frac{d^{2}}{\lambda_{i}}} ; i=1, \ldots, p$.
(iii) The rectangle surrounding the ellipsoid $E_{d}$ is defined by the following inequalities:

$$
x_{0 i}-\sqrt{d^{2} a^{i i}} \leq x_{i} \leq x_{0 i}+\sqrt{d^{2} a^{i i}}, \quad i=1, \ldots, p,
$$

where $a^{i i}$ is the $(i, i)$ element of $\mathcal{A}^{-1}$. By the rectangle surrounding the ellipsoid $E_{d}$ we mean the rectangle whose sides are parallel to the coordinate axis.

It is easy to find the coordinates of the tangency points between the ellipsoid and its surrounding rectangle parallel to the coordinate axes. Let us find the coordinates of the tangency point that are in the direction of the $j$-th coordinate axis (positive direction).

For ease of notation, we suppose the ellipsoid is centred around the origin ( $x_{0}=$ $0)$. If not, the rectangle will be shifted by the value of $x_{0}$.

The coordinate of the tangency point is given by the solution to the following problem:

$$
\begin{equation*}
x=\arg \max _{x^{\top} \mathcal{A} x=d^{2}} e_{j}^{\top} x \tag{2.35}
\end{equation*}
$$

where $e_{j}^{\top}$ is the $j$-th column of the identity matrix $\mathcal{I}_{p}$. The coordinate of the tangency point in the negative direction would correspond to the solution of the min problem: by symmetry, it is the opposite value of the former.

The solution is computed via the Lagrangian $L=e_{j}^{\top} x-\lambda\left(x^{\top} \mathcal{A} x-d^{2}\right)$ which by (2.23) leads to the following system of equations:

$$
\begin{align*}
& \frac{\partial L}{\partial x}=e_{j}-2 \lambda \mathcal{A} x=0  \tag{2.36}\\
& \frac{\partial L}{\partial \lambda}=x^{\top} \mathcal{A} x-d^{2}=0 \tag{2.37}
\end{align*}
$$

This gives $x=\frac{1}{2 \lambda} \mathcal{A}^{-1} e_{j}$, or componentwise

$$
\begin{equation*}
x_{i}=\frac{1}{2 \lambda} a^{i j}, i=1, \ldots, p \tag{2.38}
\end{equation*}
$$

where $a^{i j}$ denotes the $(i, j)$-th element of $\mathcal{A}^{-1}$.

Premultiplying (2.36) by $x^{\top}$, we have from (2.37):

$$
x_{j}=2 \lambda d^{2}
$$

Comparing this to the value obtained by (2.38), for $i=j$ we obtain $2 \lambda=\sqrt{\frac{a^{i j}}{d^{2}}}$. We choose the positive value of the square root because we are maximising $e_{j}^{\top} x$. A minimum would correspond to the negative value. Finally, we have the coordinates of the tangency point between the ellipsoid and its surrounding rectangle in the positive direction of the $j$-th axis:

$$
\begin{equation*}
x_{i}=\sqrt{\frac{d^{2}}{a^{i j}}} a^{i j}, i=1, \ldots, p \tag{2.39}
\end{equation*}
$$

The particular case where $i=j$ provides statement (iii) in Theorem 2.7.

## Remark: Usefulness of Theorem 2.7

Theorem 2.7 will prove to be particularly useful in many subsequent chapters. First, it provides a helpful tool for graphing an ellipse in two dimensions. Indeed, knowing the slope of the principal axes of the ellipse, their half-lengths and drawing the rectangle inscribing the ellipse, allows one to quickly draw a rough picture of the shape of the ellipse.

In Chap. 7, it is shown that the confidence region for the vector $\mu$ of a multivariate normal population is given by a particular ellipsoid whose parameters depend on sample characteristics. The rectangle inscribing the ellipsoid (which is much easier to obtain) will provide the simultaneous confidence intervals for all of the components in $\mu$.

In addition it will be shown that the contour surfaces of the multivariate normal density are provided by ellipsoids whose parameters depend on the mean vector and on the covariance matrix. We will see that the tangency points between the contour ellipsoids and the surrounding rectangle are determined by regressing one component on the $(p-1)$ other components. For instance, in the direction of the $j$-th axis, the tangency points are given by the intersections of the ellipsoid contours with the regression line of the vector of $(p-1)$ variables (all components except the $j$-th) on the $j$-th component.

## Norm of a Vector

Consider a vector $x \in \mathbb{R}^{p}$. The norm or length of $x$ (with respect to the metric $\mathcal{I}_{p}$ ) is defined as

$$
\|x\|=d\left(0_{p}, x\right)=\sqrt{x^{\top} x}
$$

If $\|x\|=1, x$ is called a unit vector. A more general norm can be defined with respect to the metric $\mathcal{A}$ :

$$
\|x\|_{\mathcal{A}}=\sqrt{x^{\top} \mathcal{A} x}
$$

## Angle Between Two Vectors

Consider two vectors $x$ and $y \in \mathbb{R}^{p}$. The angle $\theta$ between $x$ and $y$ is defined by the cosine of $\theta$ :

$$
\begin{equation*}
\cos \theta=\frac{x^{\top} y}{\|x\|\|y\|} \tag{2.40}
\end{equation*}
$$

see Fig. 2.4. Indeed for $p=2, x=\binom{x_{1}}{x_{2}}$ and $y=\binom{y_{1}}{y_{2}}$, we have

$$
\begin{align*}
& \|x\| \cos \theta_{1}=x_{1} ;\|y\| \cos \theta_{2}=y_{1}  \tag{2.41}\\
& \|x\| \sin \theta_{1}=x_{2} ;\|y\| \sin \theta_{2}=y_{2},
\end{align*}
$$

therefore,

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

Fig. 2.4 Angle between vectors


Fig. 2.5 Projection


Remark 2.1 If $x^{\top} y=0$, then the angle $\theta$ is equal to $\frac{\pi}{2}$. From trigonometry, we know that the cosine of $\theta$ equals the length of the base of a triangle $\left(\left\|p_{x}\right\|\right)$ divided by the length of the hypotenuse $(\|x\|)$. Hence, we have

$$
\begin{equation*}
\left\|p_{x}\right\|=\left\|x|\| \cos \theta|=\frac{\left|x^{\top} y\right|}{\|y\|}\right. \tag{2.42}
\end{equation*}
$$

where $p_{x}$ is the projection of $x$ on $y$ (which is defined below). It is the coordinate of $x$ on the $y$ vector, see Fig. 2.5.

The angle can also be defined with respect to a general metric $\mathcal{A}$

$$
\begin{equation*}
\cos \theta=\frac{x^{\top} \mathcal{A} y}{\|x\|_{\mathcal{A}}\|y\|_{\mathcal{A}}} \tag{2.43}
\end{equation*}
$$

If $\cos \theta=0$ then $x$ is orthogonal to $y$ with respect to the metric $\mathcal{A}$.
Example 2.11 Assume that there are two centred (i.e. zero mean) data vectors. The cosine of the angle between them is equal to their correlation (defined in (3.8)). Indeed for $x$ and $y$ with $\bar{x}=\bar{y}=0$ we have

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

according to formula (2.40).

## Rotations

When we consider a point $x \in \mathbb{R}^{p}$, we generally use a $p$-coordinate system to obtain its geometric representation, like in Fig. 2.1 for instance. There will be situations in multivariate techniques where we will want to rotate this system of coordinates by the angle $\theta$.

Consider for example the point $P$ with coordinates $x=\left(x_{1}, x_{2}\right)^{\top}$ in $\mathbb{R}^{2}$ with respect to a given set of orthogonal axes. Let $\Gamma$ be a $(2 \times 2)$ orthogonal matrix where

$$
\begin{equation*}
\Gamma=\binom{\cos \theta \sin \theta}{-\sin \theta \cos \theta} . \tag{2.44}
\end{equation*}
$$

If the axes are rotated about the origin through an angle $\theta$ in a clockwise direction, the new coordinates of $P$ will be given by the vector $y$

$$
\begin{equation*}
y=\Gamma x \tag{2.45}
\end{equation*}
$$

and a rotation through the same angle in a anti-clockwise direction gives the new coordinates as

$$
\begin{equation*}
y=\Gamma^{\top} x . \tag{2.46}
\end{equation*}
$$

More generally, premultiplying a vector $x$ by an orthogonal matrix $\Gamma$ geometrically corresponds to a rotation of the system of axes, so that the first new axis is determined by the first row of $\Gamma$. This geometric point of view will be exploited in Chaps. 11 and 12.

## Column Space and Null Space of a Matrix

Define for $\mathcal{X}(n \times p)$

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

the space generated by the columns of $\mathcal{X}$ or the column space of $\mathcal{X}$. Note that $C(\mathcal{X}) \subseteq \mathbb{R}^{n}$ and $\operatorname{dim}\{C(\mathcal{X})\}=\operatorname{rank}(\mathcal{X})=r \leq \min (n, p)$.

$$
\operatorname{Ker}(\mathcal{X}) \stackrel{\text { def }}{=} N(\mathcal{X})=\left\{y \in \mathbb{R}^{p} \mid \mathcal{X} y=0\right\}
$$

is the null space of $\mathcal{X}$. Note that $N(\mathcal{X}) \subseteq \mathbb{R}^{p}$ and that $\operatorname{dim}\{N(\mathcal{X})\}=p-r$.
Remark $2.2 N\left(\mathcal{X}^{\top}\right)$ is the orthogonal complement of $C(\mathcal{X})$ in $\mathbb{R}^{n}$, i.e. given a vector $b \in \mathbb{R}^{n}$ it will hold that $x^{\top} b=0$ for all $x \in C(\mathcal{X})$, if and only if $b \in N\left(\mathcal{X}^{\top}\right)$.

Example 2.12 Let $\mathcal{X}=\left(\begin{array}{lll}2 & 3 & 5 \\ 4 & 6 & 7 \\ 6 & 8 & 6 \\ 8 & 2 & 4\end{array}\right)$. It is easy to show (e.g. by calculating the determinant of $\mathcal{X}$ ) that $\operatorname{rank}(\mathcal{X})=3$. Hence, the column space of $\mathcal{X}$ is $C(\mathcal{X})=\mathbb{R}^{3}$. The null space of $\mathcal{X}$ contains only the zero vector $(0,0,0)^{\top}$ and its dimension is equal to $\operatorname{rank}(\mathcal{X})-3=0$.

For $\mathcal{X}=\left(\begin{array}{lll}2 & 3 & 1 \\ 4 & 6 & 2 \\ 6 & 8 & 3 \\ 8 & 2 & 4\end{array}\right)$, the third column is a multiple of the first one and the matrix $\mathcal{X}$ cannot be of full rank. Noticing that the first two columns of $\mathcal{X}$ are independent, we see that $\operatorname{rank}(\mathcal{X})=2$. In this case, the dimension of the columns space is 2 and the dimension of the null space is 1 .

## Projection Matrix

A matrix $\mathcal{P}(n \times n)$ is called an (orthogonal) projection matrix in $\mathbb{R}^{n}$ if and only if $\mathcal{P}=\mathcal{P}^{\top}=\mathcal{P}^{2}(\mathcal{P}$ is idempotent $)$. Let $b \in \mathbb{R}^{n}$. Then $a=\mathcal{P} b$ is the projection of $b$ on $C(\mathcal{P})$.

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

Consider $\mathcal{X}(n \times p)$ and let

$$
\begin{equation*}
\mathcal{P}=\mathcal{X}\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1} \mathcal{X}^{\top} \tag{2.47}
\end{equation*}
$$

and $\mathcal{Q}=\mathcal{I}_{n}-\mathcal{P}$. It's easy to check that $\mathcal{P}$ and $\mathcal{Q}$ are idempotent and that

$$
\begin{equation*}
\mathcal{P} \mathcal{X}=\mathcal{X} \text { and } \mathcal{Q X}=0 \tag{2.48}
\end{equation*}
$$

Since the columns of $\mathcal{X}$ are projected onto themselves, the projection matrix $\mathcal{P}$ projects any vector $b \in \mathbb{R}^{n}$ onto $C(\mathcal{X})$. Similarly, the projection matrix $\mathcal{Q}$ projects any vector $b \in \mathbb{R}^{n}$ onto the orthogonal complement of $C(\mathcal{X})$.

Theorem 2.8 Let $\mathcal{P}$ be the projection (2.47) and $\mathcal{Q}$ its orthogonal complement. Then:
(i) $x=\mathcal{P b}$ entails $x \in C(\mathcal{X})$,
(ii) $y=\mathcal{Q} b$ means that $y^{\top} x=0 \forall x \in C(\mathcal{X})$.

Proof (i) holds, since $x=\mathcal{X}\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1} \mathcal{X}^{\top} b=\mathcal{X} a$, where $a=\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1} \mathcal{X}^{\top} b \in$ $\mathbb{R}^{p}$.
(ii) follows from $y=b-\mathcal{P} b$ and $x=\mathcal{X} a$. Hence $y^{\top} x=b^{\top} \mathcal{X} a-$ $b^{\top} \mathcal{X}\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1} \mathcal{X}^{\top} \mathcal{X} a=0$.

Remark 2.3 Let $x, y \in \mathbb{R}^{n}$ and consider $p_{x} \in \mathbb{R}^{n}$, the projection of $x$ on $y$ (see Fig. 2.5). With $\mathcal{X}=y$ we have from (2.47)

$$
\begin{equation*}
p_{x}=y\left(y^{\top} y\right)^{-1} y^{\top} x=\frac{y^{\top} x}{\|y\|^{2}} y \tag{2.49}
\end{equation*}
$$

and we can easily verify that

$$
\left\|p_{x}\right\|=\sqrt{p_{x}^{\top} p_{x}}=\frac{\left|y^{\top} x\right|}{\|y\|} .
$$

See again Remark 2.1.

| Sum Sumary |
| :---: |
| $\hookrightarrow$ A distance between two $p$-dimensional points $x$ and $y$ is a quadratic form $(x-y)^{\top} \mathcal{A}(x-y)$ in the vectors of differences $(x-y)$. A distance defines the norm of a vector. |
| $\hookrightarrow$ Iso-distance curves of a point $x_{0}$ are all those points that have the same distance from $x_{0}$. Iso-distance curves are ellipsoids whose principal axes are determined by the direction of the eigenvectors of $\mathcal{A}$. The half-length of principal axes is proportional to the inverse of the roots of the eigenvalues of $\mathcal{A}$. |
| $\hookrightarrow$ The angle between two vectors $x$ and $y$ is given by $\cos \theta=$ $\frac{x^{\top} \mathcal{A} y}{\\|x\\|_{\mathcal{A}}\left\\|_{y}\right\\|_{\mathcal{A}}}$ w.r.t. the metric $\mathcal{A}$. |
| $\hookrightarrow$ For the Euclidean distance with $\mathcal{A}=\mathcal{I}$ the correlation between two centred data vectors $x$ and $y$ is given by the cosine of the angle between them, i.e. $\cos \theta=r_{X Y}$. |
| $\hookrightarrow$ The projection $\mathcal{P}=\mathcal{X}\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1} \mathcal{X}^{\top}$ is the projection onto the column space $C(\mathcal{X})$ of $\mathcal{X}$. |
| $\hookrightarrow$ The projection of $x \in \mathbb{R}^{n}$ on $y \in \mathbb{R}^{n}$ is given by $p_{x}=\frac{y^{\top} x}{\\|y\\|^{2}} y$. |

### 2.7 Exercises

Exercise 2.1 Compute the determinant for a $(3 \times 3)$ matrix.
Exercise 2.2 Suppose that $|\mathcal{A}|=0$. Is it possible that all eigenvalues of $\mathcal{A}$ are positive?

Exercise 2.3 Suppose that all eigenvalues of some (square) matrix $\mathcal{A}$ are different from zero. Does the inverse $\mathcal{A}^{-1}$ of $\mathcal{A}$ exist?

Exercise 2.4 Write a program that calculates the Jordan decomposition of the matrix

$$
\mathcal{A}=\left(\begin{array}{lll}
1 & 2 & 3 \\
2 & 1 & 2 \\
3 & 2 & 1
\end{array}\right)
$$

Check Theorem 2.1 numerically.
Exercise 2.5 Prove (2.23), (2.24) and (2.25).
Exercise 2.6 Show that a projection matrix only has eigenvalues in $\{0,1\}$.
Exercise 2.7 Draw some iso-distance ellipsoids for the metric $\mathcal{A}=\Sigma^{-1}$ of Example 3.13.

Exercise 2.8 Find a formula for $\left|\mathcal{A}+a a^{\top}\right|$ and for $\left(\mathcal{A}+a a^{\top}\right)^{-1}$. (Hint: use the inverse partitioned matrix with $\mathcal{B}=\left(\begin{array}{cc}1 & -a^{\top} \\ a & \mathcal{A}\end{array}\right)$.)

Exercise 2.9 Prove the Binomial inverse theorem for two non-singular matrices $\mathcal{A}(p \times p)$ and $\mathcal{B}(p \times p):(\mathcal{A}+\mathcal{B})^{-1}=\mathcal{A}^{-1}-\mathcal{A}^{-1}\left(\mathcal{A}^{-1}+\mathcal{B}^{-1}\right)^{-1} \mathcal{A}^{-1}$. (Hint: use (2.26) with $\mathcal{C}=\left(\begin{array}{ll}\mathcal{A} & I_{p} \\ -I_{p} & \mathcal{B}^{-1}\end{array}\right)$.)

## Chapter 3 <br> Moving to Higher Dimensions

We have seen in the previous chapters how very simple graphical devices can help in understanding the structure and dependency of data. The graphical tools were based on either univariate (bivariate) data representations or on "slick" transformations of multivariate information perceivable by the human eye. Most of the tools are extremely useful in a modelling step, but unfortunately, do not give the full picture of the data set. One reason for this is that the graphical tools presented capture only certain dimensions of the data and do not necessarily concentrate on those dimensions or sub-parts of the data under analysis that carry the maximum structural information. In Part III of this book, powerful tools for reducing the dimension of a data set will be presented. In this chapter, as a starting point, simple and basic tools are used to describe dependency. They are constructed from elementary facts of probability theory and introductory statistics (e.g. the covariance and correlation between two variables).

Sections 3.1 and 3.2 show how to handle these concepts in a multivariate setup and how a simple test on correlation between two variables can be derived. Since linear relationships are involved in these measures, Sect. 3.4 presents the simple linear model for two variables and recalls the basic $t$-test for the slope. In Sect. 3.5, a simple example of one-factorial analysis of variance introduces the notations for the well-known $F$-test.

Due to the power of matrix notation, all of this can easily be extended to a more general multivariate setup. Section 3.3 shows how matrix operations can be used to define summary statistics of a data set and for obtaining the empirical moments of linear transformations of the data. These results will prove to be very useful in most of the chapters in Part III.

Finally, matrix notation allows us to introduce the flexible multiple linear model, where more general relationships among variables can be analysed. In Sect. 3.6, the least squares adjustment of the model and the usual test statistics are presented with their geometric interpretation. Using these notations, the ANOVA model is just a particular case of the multiple linear model.

### 3.1 Covariance

Covariance is a measure of dependency between random variables. Given two (random) variables $X$ and $Y$ the (theoretical) covariance is defined by:

$$
\begin{equation*}
\sigma_{X Y}=\operatorname{Cov}(X, Y)=\mathrm{E}(X Y)-(\mathrm{E} X)(\mathrm{E} Y) \tag{3.1}
\end{equation*}
$$

The precise definition of expected values is given in Chap.4. If $X$ and $Y$ are independent of each other, the covariance $\operatorname{Cov}(X, Y)$ is necessarily equal to zero, see Theorem 3.1. The converse is not true. The covariance of $X$ with itself is the variance:

$$
\sigma_{X X}=\operatorname{Var}(X)=\operatorname{Cov}(X, X)
$$

If the variable $X$ is $p$-dimensional multivariate, e.g. $X=\left(\begin{array}{c}X_{1} \\ \vdots \\ X_{p}\end{array}\right)$, then the theoretical covariances among all the elements are put into matrix form, i.e. the covariance matrix:

$$
\Sigma=\left(\begin{array}{ccc}
\sigma_{X_{1} X_{1}} & \ldots & \sigma_{X_{1} X_{p}} \\
\vdots & \ddots & \vdots \\
\sigma_{X_{p} X_{1}} & \ldots & \sigma_{X_{p} X_{p}}
\end{array}\right) .
$$

Properties of covariance matrices will be detailed in Chap. 4. Empirical versions of these quantities are:

$$
\begin{align*}
& s_{X Y}=\frac{1}{n} \sum_{i=1}^{n}\left(x_{i}-\bar{x}\right)\left(y_{i}-\bar{y}\right)  \tag{3.2}\\
& s_{X X}=\frac{1}{n} \sum_{i=1}^{n}\left(x_{i}-\bar{x}\right)^{2} . \tag{3.3}
\end{align*}
$$

For small $n$, say $n \leq 20$, we should replace the factor $\frac{1}{n}$ in (3.2) and (3.3) by $\frac{1}{n-1}$ in order to correct for a small bias. For a $p$-dimensional random variable, one obtains the empirical covariance matrix (see Sect. 3.3 for properties and details)

$$
\mathcal{S}=\left(\begin{array}{ccc}
s_{X_{1} X_{1}} & \ldots & s_{X_{1} X_{p}} \\
\vdots & \ddots & \vdots \\
s_{X_{p} X_{1}} & \ldots & s_{X_{p} X_{p}}
\end{array}\right) .
$$

For a scatterplot of two variables the covariances measure "how close the scatter is to a line". Mathematical details follow but it should already be understood here that in this sense covariance measures only "linear dependence".

Example 3.1 If $\mathcal{X}$ is the entire bank data set, one obtains the covariance matrix $\mathcal{S}$ as indicated below:

$$
\mathcal{S}=\left(\begin{array}{rrrrrr}
0.14 & 0.03 & 0.02 & -0.10 & -0.01 & 0.08  \tag{3.4}\\
0.03 & 0.12 & 0.10 & 0.21 & 0.10 & -0.21 \\
0.02 & 0.10 & 0.16 & 0.28 & 0.12 & -0.24 \\
-0.10 & 0.21 & 0.28 & 2.07 & 0.16 & -1.03 \\
-0.01 & 0.10 & 0.12 & 0.16 & 0.64 & -0.54 \\
0.08 & -0.21 & -0.24 & -1.03 & -0.54 & 1.32
\end{array}\right) .
$$

The empirical covariance between $X_{4}$ and $X_{5}$, i.e. $s_{X_{4} X_{5}}$, is found in row 4 and column 5. The value is $s_{X_{4} X_{5}}=0.16$. Is it obvious that this value is positive? In Exercise 3.1 we will discuss this question further.

If $\mathcal{X}_{f}$ denotes the counterfeit bank notes, we obtain:

$$
\mathcal{S}_{f}=\left(\begin{array}{rrrrrr}
0.123 & 0.031 & 0.023 & -0.099 & 0.019 & 0.011  \tag{3.5}\\
0.031 & 0.064 & 0.046 & -0.024 & -0.012 & -0.005 \\
0.024 & 0.046 & 0.088 & -0.018 & 0.000 & 0.034 \\
-0.099 & -0.024 & -0.018 & 1.268 & -0.485 & 0.236 \\
0.019 & -0.012 & 0.000 & -0.485 & 0.400 & -0.022 \\
0.011 & -0.005 & 0.034 & 0.236 & -0.022 & 0.308
\end{array}\right) .
$$

For the genuine $\mathcal{X}_{g}$, we have:

$$
\mathcal{S}_{g}=\left(\begin{array}{rrrrrr}
0.149 & 0.057 & 0.057 & 0.056 & 0.014 & 0.005  \tag{3.6}\\
0.057 & 0.131 & 0.085 & 0.056 & 0.048 & -0.043 \\
0.057 & 0.085 & 0.125 & 0.058 & 0.030 & -0.024 \\
0.056 & 0.056 & 0.058 & 0.409 & -0.261 & -0.000 \\
0.014 & 0.049 & 0.030 & -0.261 & 0.417 & -0.074 \\
0.005 & -0.043 & -0.024 & -0.000 & -0.074 & 0.198
\end{array}\right) .
$$

Note that the covariance between $X_{4}$ (distance of the frame to the lower border) and $X_{5}$ (distance of the frame to the upper border) is negative in both (3.5) and (3.6). Why would this happen? In Exercise 3.2 we will discuss this question in more detail.

At first sight, the matrices $\mathcal{S}_{f}$ and $\mathcal{S}_{g}$ look different, but they create almost the same scatterplots (see the discussion in Sect. 1.4). Similarly, the common principal component analysis in Chap. 11 suggests a joint analysis of the covariance structure as in Flury and Riedwyl (1988).

Fig. 3.1 Scatterplot of variables $X_{4}$ vs. $X_{5}$ of the entire bank data set $\mathbf{Q}$ MVAscabank45


Scatterplots with point clouds that are "upward-sloping", like the one in the upper left of Fig. 1.14, show variables with positive covariance. Scatterplots with "downward-sloping" structure have negative covariance. In Fig. 3.1 we show the scatterplot of $X_{4}$ vs. $X_{5}$ of the entire bank data set. The point cloud is upwardsloping. However, the two sub-clouds of counterfeit and genuine bank notes are downward-sloping.

Example 3.2 A textile shop manager is studying the sales of "classic blue" pullovers over ten different periods. He observes the number of pullovers sold ( $X_{1}$ ), variation in price ( $X_{2}$, in EUR), the advertisement costs in local newspapers ( $X_{3}$, in EUR) and the presence of a sales assistant ( $X_{4}$, in hours per period). Over the periods, he observes the following data matrix:

$$
\mathcal{X}=\left(\begin{array}{rrrr}
230 & 125 & 200 & 109 \\
181 & 99 & 55 & 107 \\
165 & 97 & 105 & 98 \\
150 & 115 & 85 & 71 \\
97 & 120 & 0 & 82 \\
192 & 100 & 150 & 103 \\
181 & 80 & 85 & 111 \\
189 & 90 & 120 & 93 \\
172 & 95 & 110 & 86 \\
170 & 125 & 130 & 78
\end{array}\right) .
$$

He is convinced that the price must have a large influence on the number of pullovers sold. So he makes a scatterplot of $X_{2}$ vs. $X_{1}$, see Fig.3.2. A rough impression

Fig. 3.2 Scatterplot of variables $X_{2}$ vs. $X_{1}$ of the pullovers data set $\mathbf{Q}$ MVAscapullı

Pullovers Data

is that the cloud is somewhat downward-sloping. A computation of the empirical covariance yields

$$
s_{X_{1} X_{2}}=\frac{1}{9} \sum_{i=1}^{10}\left(X_{1 i}-\bar{X}_{1}\right)\left(X_{2 i}-\bar{X}_{2}\right)=-80.02
$$

a negative value as expected.
Note: The covariance function is scale dependent. Thus, if the prices in this example were in Japanese Yen (JPY), we would obtain a different answer (see Exercise 3.16). A measure of (linear) dependence independent of the scale is the correlation, which we introduce in the next section.

|  | $\hookrightarrow$ The covariance is a measure of dependence. |
| :--- | :--- |
|  | $\hookrightarrow$ Covariance measures only linear dependence. |
|  | $\hookrightarrow$ Covariance is scale dependent. |
|  | $\hookrightarrow$ There are non-linear dependencies that have zero covariance. |
|  | Zero covariance does not imply independence. |


| Summary (continued) |  |
| :--- | :--- |
| $\hookrightarrow$ | Independence implies zero covariance. |
| $\hookrightarrow$ | Negative covariance corresponds to downward-sloping scatter- <br> plots. |
| $\hookrightarrow$ | Positive covariance corresponds to upward-sloping scatterplots. |
| $\hookrightarrow$ | The covariance of a variable with itself is its variance |
|  | $\operatorname{Cov}(X, X)=\sigma_{X X}=\sigma_{X}^{2}$. |

### 3.2 Correlation

The correlation between two variables $X$ and $Y$ is defined from the covariance as the following:

$$
\begin{equation*}
\rho_{X Y}=\frac{\operatorname{Cov}(X, Y)}{\sqrt{\operatorname{Var}(X) \operatorname{Var}(Y)}} \tag{3.7}
\end{equation*}
$$

The advantage of the correlation is that it is independent of the scale, i.e. changing the variables' scale of measurement does not change the value of the correlation. Therefore, the correlation is more useful as a measure of association between two random variables than the covariance. The empirical version of $\rho_{X Y}$ is as follows:

$$
\begin{equation*}
r_{X Y}=\frac{s_{X Y}}{\sqrt{s_{X X} s_{Y Y}}} \tag{3.8}
\end{equation*}
$$

The correlation is in absolute value always less than 1. It is zero if the covariance is zero and vice versa. For $p$-dimensional vectors $\left(X_{1}, \ldots, X_{p}\right)^{\top}$ we have the theoretical correlation matrix

$$
\mathcal{P}=\left(\begin{array}{ccc}
\rho_{X_{1} X_{1}} & \ldots & \rho_{X_{1} X_{p}} \\
\vdots & \ddots & \vdots \\
\rho_{X_{p} X_{1}} & \ldots & \rho_{X_{p} X_{p}}
\end{array}\right)
$$

and its empirical version, the empirical correlation matrix which can be calculated from the observations,

$$
\mathcal{R}=\left(\begin{array}{ccc}
r_{X_{1} X_{1}} & \ldots & r_{X_{1} X_{p}} \\
\vdots & \ddots & \vdots \\
r_{X_{p} X_{1}} & \ldots & r_{X_{p} X_{p}}
\end{array}\right)
$$

Example 3.3 We obtain the following correlation matrix for the genuine bank notes:

$$
\mathcal{R}_{g}=\left(\begin{array}{rrrrrr}
1.00 & 0.41 & 0.41 & 0.22 & 0.05 & 0.03  \tag{3.9}\\
0.41 & 1.00 & 0.66 & 0.24 & 0.20 & -0.25 \\
0.41 & 0.66 & 1.00 & 0.25 & 0.13 & -0.14 \\
0.22 & 0.24 & 0.25 & 1.00 & -0.63 & -0.00 \\
0.05 & 0.20 & 0.13 & -0.63 & 1.00 & -0.25 \\
0.03 & -0.25 & -0.14 & -0.00 & -0.25 & 1.00
\end{array}\right)
$$

and for the counterfeit bank notes:

$$
\mathcal{R}_{f}=\left(\begin{array}{rrrrrr}
1.00 & 0.35 & 0.24 & -0.25 & 0.08 & 0.06  \tag{3.10}\\
0.35 & 1.00 & 0.61 & -0.08 & -0.07 & -0.03 \\
0.24 & 0.61 & 1.00 & -0.05 & 0.00 & 0.20 \\
-0.25 & -0.08 & -0.05 & 1.00 & -0.68 & 0.37 \\
0.08 & -0.07 & 0.00 & -0.68 & 1.00 & -0.06 \\
0.06 & -0.03 & 0.20 & 0.37 & -0.06 & 1.00
\end{array}\right) .
$$

As noted before for $\operatorname{Cov}\left(X_{4}, X_{5}\right)$, the correlation between $X_{4}$ (distance of the frame to the lower border) and $X_{5}$ (distance of the frame to the upper border) is negative. This is natural, since the covariance and correlation always have the same sign (see also Exercise 3.17).

Why is the correlation an interesting statistic to study? It is related to independence of random variables, which we shall define more formally later on. For the moment we may think of independence as the fact that one variable has no influence on another.

Theorem 3.1 If $X$ and $Y$ are independent, then $\rho(X, Y)=\operatorname{Cov}(X, Y)=0$.


In general, the converse is not true, as the following example shows.
Example 3.4 Consider a standard normally-distributed random variable $X$ and a random variable $Y=X^{2}$, which is surely not independent of $X$. Here we have

$$
\operatorname{Cov}(X, Y)=\mathrm{E}(X Y)-\mathrm{E}(X) \mathrm{E}(Y)=\mathrm{E}\left(X^{3}\right)=0
$$

(because $\mathrm{E}(X)=0$ and $\mathrm{E}\left(X^{2}\right)=1$ ). Therefore $\rho(X, Y)=0$, as well. This example also shows that correlations and covariances measure only linear dependence. The quadratic dependence of $Y=X^{2}$ on $X$ is not reflected by these measures of dependence.

Remark 3.1 For two normal random variables, the converse of Theorem 3.1 is true: zero covariance for two normally distributed random variables implies independence. This will be shown later in Corollary 5.2.

Theorem 3.1 enables us to check for independence between the components of a bivariate normal random variable. That is, we can use the correlation and test whether it is zero. The distribution of $r_{X Y}$ for an arbitrary $(X, Y)$ is unfortunately complicated. The distribution of $r_{X Y}$ will be more accessible if $(X, Y)$ are jointly normal (see Chap. 5). If we transform the correlation by Fisher's $Z$-transformation,

$$
\begin{equation*}
W=\frac{1}{2} \log \left(\frac{1+r_{X Y}}{1-r_{X Y}}\right), \tag{3.11}
\end{equation*}
$$

we obtain a variable that has a more accessible distribution. Under the hypothesis that $\rho=0, W$ has an asymptotic normal distribution. Approximations of the expectation and variance of $W$ are given by the following:

$$
\begin{align*}
\mathrm{E}(W) & \approx \frac{1}{2} \log \left(\frac{1+\rho_{X Y}}{1-\rho_{X Y}}\right)  \tag{3.12}\\
\operatorname{Var}(W) & \approx \frac{1}{(n-3)} .
\end{align*}
$$

The distribution is given in Theorem 3.2.

## Theorem 3.2

$$
\begin{equation*}
Z=\frac{W-E(W)}{\sqrt{\operatorname{Var}(W)}} \xrightarrow{\mathcal{L}} N(0,1) . \tag{3.13}
\end{equation*}
$$

The symbol " $\xrightarrow{\mathcal{L}}$ " denotes convergence in distribution, which will be explained in more detail in Chap. 4.

Theorem 3.2 allows us to test different hypotheses on correlation. We can fix the level of significance $\alpha$ (the probability of rejecting a true hypothesis) and reject the hypothesis if the difference between the hypothetical value and the calculated value of $Z$ is greater than the corresponding critical value of the normal distribution. The following example illustrates the procedure.

Example 3.5 Let's study the correlation between mileage ( $X_{2}$ ) and weight $\left(X_{8}\right)$ for the car data set (22.3) where $n=74$. We have $r_{X_{2} X_{8}}=-0.823$. Our conclusions from the boxplot in Fig. 1.3 ("Japanese cars generally have better mileage than the others") needs to be revised. From Fig. 3.3 and $r_{X_{2} X_{8}}$, we can see that mileage is highly correlated with weight, and that the Japanese cars in the sample are in fact all lighter than the others.

If we want to know whether $\rho_{X_{2} X_{8}}$ is significantly different from $\rho_{0}=0$, we apply Fisher's $Z$-transform (3.11). This gives us

$$
w=\frac{1}{2} \log \left(\frac{1+r_{X_{2} X_{8}}}{1-r_{X_{2} X_{8}}}\right)=-1.166 \quad \text { and } \quad z=\frac{-1.166-0}{\sqrt{\frac{1}{71}}}=-9.825,
$$

Fig. 3.3 Mileage $\left(X_{2}\right)$ vs. weight ( $X_{8}$ ) of US (star), European (plus signs) and Japanese (circle) cars 0 MVAscacar

i.e. a highly significant value to reject the hypothesis that $\rho=0$ (the $2.5 \%$ and $97.5 \%$ quantiles of the normal distribution are -1.96 and 1.96 , respectively). If we want to test the hypothesis that, say, $\rho_{0}=-0.75$, we obtain:

$$
z=\frac{-1.166-(-0.973)}{\sqrt{\frac{1}{71}}}=-1.627 .
$$

This is a non-significant value at the $\alpha=0.05$ level for $z$ since it is between the critical values at the $5 \%$ significance level (i.e. $-1.96<z<1.96$ ).

Example 3.6 Let us consider again the pullovers data set from Example 3.2. Consider the correlation between the presence of the sales assistants $\left(X_{4}\right)$ vs. the number of sold pullovers ( $X_{1}$ ) (see Fig. 3.4). Here we compute the correlation as

$$
r_{X_{1} X_{4}}=0.633
$$

The $Z$-transform of this value is

$$
\begin{equation*}
w=\frac{1}{2} \log \left(\frac{1+r_{X_{1} X_{4}}}{1-r_{X_{1} X_{4}}}\right)=0.746 \tag{3.14}
\end{equation*}
$$

The sample size is $n=10$, so for the hypothesis $\rho_{X_{1} X_{4}}=0$, the statistic to consider is:

$$
\begin{equation*}
z=\sqrt{7}(0.746-0)=1.974 \tag{3.15}
\end{equation*}
$$

which is just statistically significant at the $5 \%$ level (i.e. 1.974 is just a little larger than 1.96).

Fig. 3.4 Hours of sales assistants ( $X_{4}$ ) vs. sales ( $X_{1}$ ) of pullovers $\mathbf{Q}$ MVAscapull2


Remark 3.2 The normalising and variance stabilising properties of $W$ are asymptotic. In addition the use of $W$ in small samples (for $n \leq 25$ ) is improved by Hotelling's transform (Hotelling, 1953):

$$
W^{*}=W-\frac{3 W+\tanh (W)}{4(n-1)} \quad \text { with } \quad \operatorname{Var}\left(W^{*}\right)=\frac{1}{n-1}
$$

The transformed variable $W^{*}$ is asymptotically distributed as a normal distribution.
Example 3.7 From the preceding remark, we obtain $w^{*}=0.6663$ and $\sqrt{10-1} w^{*}=1.9989$ for the preceding Example 3.6. This value is significant at the $5 \%$ level.
Remark 3.3 Note that the Fisher's $Z$-transform is the inverse of the hyperbolic tangent function: $W=\tanh ^{-1}\left(r_{X Y}\right)$; equivalently $r_{X Y}=\tanh (W)=\frac{e^{2 W}-1}{e^{2 W}+1}$.
Remark 3.4 Under the assumptions of normality of $X$ and $Y$, we may test their independence $\left(\rho_{X Y}=0\right)$ using the exact $t$-distribution of the statistic

$$
T=r_{X Y} \sqrt{\frac{n-2}{1-r_{X Y}^{2}}} \stackrel{\rho_{X Y}=0}{\sim} t_{n-2} .
$$

Setting the probability of the first error type to $\alpha$, we reject the null hypothesis $\rho_{X Y}=0$ if $|T| \geq t_{1-\alpha / 2 ; n-2}$.

|  | $\hookrightarrow$ |
| :--- | :--- |
|  | The correlation is a standardised measure of dependence. |
|  | $\hookrightarrow$ |
|  | The absolute value of the correlation is always less or equal to one. |
|  | $\hookrightarrow$ | There are non-linear dependencies that have zero correlation.

### 3.3 Summary Statistics

This section focuses on the representation of basic summary statistics (means, covariances and correlations) in matrix notation, since we often apply linear transformations to data. The matrix notation allows us to derive instantaneously the corresponding characteristics of the transformed variables. The Mahalanobis transformation is a prominent example of such linear transformations.

Assume that we have observed $n$ realisations of a $p$-dimensional random variable; we have a data matrix $\mathcal{X}(n \times p)$ :

$$
\mathcal{X}=\left(\begin{array}{ccc}
x_{11} & \cdots & x_{1 p}  \tag{3.16}\\
\vdots & & \vdots \\
\vdots & & \vdots \\
x_{n 1} & \cdots & x_{n p}
\end{array}\right)
$$

The rows $x_{i}=\left(x_{i 1}, \ldots, x_{i p}\right) \in \mathbb{R}^{p}$ denote the $i$ th observation of a $p$-dimensional random variable $X \in \mathbb{R}^{p}$.

The statistics that were briefly introduced in Sects. 3.1 and 3.2 can be rewritten in matrix form as follows. The "centre of gravity" of the $n$ observations in $\mathbb{R}^{p}$ is given by the vector $\bar{x}$ of the means $\bar{x}_{j}$ of the $p$ variables:

$$
\bar{x}=\left(\begin{array}{c}
\bar{x}_{1}  \tag{3.17}\\
\vdots \\
\bar{x}_{p}
\end{array}\right)=n^{-1} \mathcal{X}^{\top} 1_{n}
$$

The dispersion of the $n$ observations can be characterised by the covariance matrix of the $p$ variables. The empirical covariances defined in (3.2) and (3.3) are the elements of the following matrix:

$$
\begin{equation*}
\mathcal{S}=n^{-1} \mathcal{X}^{\top} \mathcal{X}-\bar{x} \bar{x}^{\top}=n^{-1}\left(\mathcal{X}^{\top} \mathcal{X}-n^{-1} \mathcal{X}^{\top} 1_{n} 1_{n}^{\top} \mathcal{X}\right) \tag{3.18}
\end{equation*}
$$

Note that this matrix is equivalently defined by

$$
\mathcal{S}=\frac{1}{n} \sum_{i=1}^{n}\left(x_{i}-\bar{x}\right)\left(x_{i}-\bar{x}\right)^{\top} .
$$

The covariance formula (3.18) can be rewritten as $\mathcal{S}=n^{-1} \mathcal{X}^{\top} \mathcal{H} \mathcal{X}$ with the centering matrix

$$
\begin{equation*}
\mathcal{H}=\mathcal{I}_{n}-n^{-1} 1_{n} 1_{n}^{\top} . \tag{3.19}
\end{equation*}
$$

Note that the centering matrix is symmetric and idempotent. Indeed,

$$
\begin{aligned}
\mathcal{H}^{2} & =\left(\mathcal{I}_{n}-n^{-1} 1_{n} 1_{n}^{\top}\right)\left(\mathcal{I}_{n}-n^{-1} 1_{n} 1_{n}^{\top}\right) \\
& =\mathcal{I}_{n}-n^{-1} 1_{n} 1_{n}^{\top}-n^{-1} 1_{n} 1_{n}^{\top}+\left(n^{-1} 1_{n} 1_{n}^{\top}\right)\left(n^{-1} 1_{n} 1_{n}^{\top}\right) \\
& =\mathcal{I}_{n}-n^{-1} 1_{n} 1_{n}^{\top}=\mathcal{H} .
\end{aligned}
$$

As a consequence $\mathcal{S}$ is positive semidefinite, i.e.

$$
\begin{equation*}
\mathcal{S} \geq 0 \tag{3.20}
\end{equation*}
$$

Indeed for all $a \in \mathbb{R}^{p}$,

$$
\begin{aligned}
a^{\top} \mathcal{S} a & =n^{-1} a^{\top} \mathcal{X}^{\top} \mathcal{H} \mathcal{X} a \\
& =n^{-1}\left(a^{\top} \mathcal{X}^{\top} \mathcal{H}^{\top}\right)(\mathcal{H X} a) \quad \text { since } \mathcal{H}^{\top} \mathcal{H}=\mathcal{H}, \\
& =n^{-1} y^{\top} y=n^{-1} \sum_{j=1}^{p} y_{j}^{2} \geq 0
\end{aligned}
$$

for $y=\mathcal{H X} a$. It is well known from the one-dimensional case that $n^{-1} \sum_{i=1}^{n}\left(x_{i}-\right.$ $\bar{x})^{2}$ as an estimate of the variance exhibits a bias of the order $n^{-1}$ (Breiman, 1973). In the multi-dimensional case, $\mathcal{S}_{u}=\frac{n}{n-1} \mathcal{S}$ is an unbiased estimate of the true covariance. (This will be shown in Example 4.15.)

The sample correlation coefficient between the $i$ th and $j$ th variables is $r_{X_{i} X_{j}}$, see (3.8). If $\mathcal{D}=\operatorname{diag}\left(s_{X_{i} X_{i}}\right)$, then the correlation matrix is

$$
\begin{equation*}
\mathcal{R}=\mathcal{D}^{-1 / 2} \mathcal{S D}^{-1 / 2} \tag{3.21}
\end{equation*}
$$

where $\mathcal{D}^{-1 / 2}$ is a diagonal matrix with elements $\left(s_{X_{i} X_{i}}\right)^{-1 / 2}$ on its main diagonal.
Example 3.8 The empirical covariances are calculated for the pullover data set.
The vector of the means of the four variables in the dataset is $\bar{x}=$ $(172.7,104.6,104.0,93.8)^{\top}$.

The sample covariance matrix is $\mathcal{S}=\left(\begin{array}{rrrr}1037.2 & -80.2 & 1430.7 & 271.4 \\ -80.2 & 219.8 & 92.1 & -91.6 \\ 1430.7 & 92.1 & 2624 & 210.3 \\ 271.4 & -91.6 & 210.3 & 177.4\end{array}\right)$.

The unbiased estimate of the variance $(n=10)$ is equal to

$$
\mathcal{S}_{u}=\frac{10}{9} \mathcal{S}=\left(\begin{array}{rrrr}
1152.5 & -88.9 & 1589.7 & 301.6 \\
-88.9 & 244.3 & 102.3 & -101.8 \\
1589.7 & 102.3 & 2915.6 & 233.7 \\
301.6 & -101.8 & 233.7 & 197.1
\end{array}\right)
$$

The sample correlation matrix is $\mathcal{R}=\left(\begin{array}{rrll}1 & -0.17 & 0.87 & 0.63 \\ -0.17 & 1 & 0.12 & -0.46 \\ 0.87 & 0.12 & 1 & 0.31 \\ 0.63 & -0.46 & 0.31 & 1\end{array}\right)$.

## Linear Transformation

In many practical applications we need to study linear transformations of the original data. This motivates the question of how to calculate summary statistics after such linear transformations.

Let $\mathcal{A}$ be a $(q \times p)$ matrix and consider the transformed data matrix

$$
\begin{equation*}
\mathcal{Y}=\mathcal{X} \mathcal{A}^{\top}=\left(y_{1}, \ldots, y_{n}\right)^{\top} \tag{3.22}
\end{equation*}
$$

The row $y_{i}=\left(y_{i 1}, \ldots, y_{i q}\right) \in \mathbb{R}^{q}$ can be viewed as the $i$ th observation of a $q$-dimensional random variable $Y=\mathcal{A} X$. In fact we have $y_{i}=x_{i} \mathcal{A}^{\top}$. We immediately obtain the mean and the empirical covariance of the variables (columns) forming the data matrix $\mathcal{Y}$ :

$$
\begin{align*}
\bar{y} & =\frac{1}{n} \mathcal{Y}^{\top} 1_{n}=\frac{1}{n} \mathcal{A} \mathcal{X}^{\top} 1_{n}=\mathcal{A} \bar{x}  \tag{3.23}\\
\mathcal{S}_{\mathcal{Y}} & =\frac{1}{n} \mathcal{Y}^{\top} \mathcal{H Y}=\frac{1}{n} \mathcal{A} \mathcal{X}^{\top} \mathcal{H X} \mathcal{A}^{\top}=\mathcal{A} \mathcal{S}_{\mathcal{X}} \mathcal{A}^{\top} \tag{3.24}
\end{align*}
$$

Note that if the linear transformation is non-homogeneous, i.e.

$$
y_{i}=\mathcal{A} x_{i}+b \quad \text { where } \quad b(q \times 1)
$$

only (3.23) changes: $\bar{y}=\mathcal{A} \bar{x}+b$. The formulas (3.23) and (3.24) are useful in the particular case of $q=1$, i.e. $y=\mathcal{X} a$, i.e. $y_{i}=a^{\top} x_{i} ; i=1, \ldots, n$ :

$$
\begin{aligned}
\bar{y} & =a^{\top} \bar{x} \\
\mathcal{S}_{y} & =a^{\top} \mathcal{S}_{\mathcal{X}} a .
\end{aligned}
$$

Example 3.9 Suppose that $\mathcal{X}$ is the pullover data set. The manager wants to compute his mean expenses for advertisement ( $X_{3}$ ) and sales assistant $\left(X_{4}\right)$.

Suppose that the sales assistant charges an hourly wage of 10 EUR. Then the shop manager calculates the expenses $Y$ as $Y=X_{3}+10 X_{4}$. Formula (3.22) says that this is equivalent to defining the matrix $\mathcal{A}(4 \times 1)$ as:

$$
\mathcal{A}=(0,0,1,10) .
$$

Using formulas (3.23) and (3.24), it is now computationally very easy to obtain the sample mean $\bar{y}$ and the sample variance $\mathcal{S}_{y}$ of the overall expenses:

$$
\begin{gathered}
\bar{y}=\mathcal{A} \bar{x}=(0,0,1,10)\left(\begin{array}{r}
172.7 \\
104.6 \\
104.0 \\
93.8
\end{array}\right)=1042.0 \\
\mathcal{S}_{\mathcal{Y}}=\mathcal{A S}_{\mathcal{X}} \mathcal{A}^{\top}=(0,0,1,10)\left(\begin{array}{rrrr}
1152.5 & -88.9 & 1589.7 & 301.6 \\
-88.9 & 244.3 & 102.3 & -101.8 \\
1589.7 & 102.3 & 2915.6 & 233.7 \\
301.6-101.8 & 233.7 & 197.1
\end{array}\right)\left(\begin{array}{r}
0 \\
0 \\
1 \\
10
\end{array}\right) \\
\\
=2915.6+4674+19710=27299.6
\end{gathered}
$$

## Mahalanobis Transformation

A special case of this linear transformation is

$$
\begin{equation*}
z_{i}=\mathcal{S}^{-1 / 2}\left(x_{i}-\bar{x}\right), \quad i=1, \ldots, n \tag{3.25}
\end{equation*}
$$

Note that for the transformed data matrix $\mathcal{Z}=\left(z_{1}, \ldots, z_{n}\right)^{\top}$,

$$
\begin{equation*}
\mathcal{S}_{\mathcal{Z}}=n^{-1} \mathcal{Z}^{\top} \mathcal{H} \mathcal{Z}=\mathcal{I}_{p} \tag{3.26}
\end{equation*}
$$

So the Mahalanobis transformation eliminates the correlation between the variables and standardises the variance of each variable. If we apply (3.24) using $\mathcal{A}=\mathcal{S}^{-1 / 2}$, we obtain the identity covariance matrix as indicated in (3.26).

|  | Summary |
| :--- | :--- |$\quad$| The centre of gravity of a data matrix is given by its mean vector |
| :--- |
| $\bar{x}=n^{-1} \mathcal{X}^{\top} 1_{n}$. |

### 3.4 Linear Model for Two Variables

We have looked several times now at downward and upward-sloping scatterplots. What does the eye define here as a slope? Suppose that we can construct a line corresponding to the general direction of the cloud. The sign of the slope of this line would correspond to the upward and downward directions. Call the variable on the vertical axis $Y$ and the one on the horizontal axis $X$. A slope line is a linear relationship between $X$ and $Y$ :

$$
\begin{equation*}
y_{i}=\alpha+\beta x_{i}+\varepsilon_{i}, i=1, \ldots, n \tag{3.27}
\end{equation*}
$$

Here, $\alpha$ is the intercept and $\beta$ is the slope of the line. The errors (or deviations from the line) are denoted as $\varepsilon_{i}$ and are assumed to have zero mean and finite variance $\sigma^{2}$. The task of finding $(\alpha, \beta)$ in (3.27) is referred to as a linear adjustment.

In Sect. 3.6 we shall derive estimators for $\alpha$ and $\beta$ more formally, as well as accurately describe what a "good" estimator is. For now, one may try to find a "good" estimator $(\hat{\alpha}, \hat{\beta})$ via graphical techniques. A very common numerical and statistical technique is to use those $\hat{\alpha}$ and $\hat{\beta}$ that minimise:

$$
\begin{equation*}
(\hat{\alpha}, \hat{\beta})=\arg \min _{(\alpha, \beta)} \sum_{i=1}^{n}\left(y_{i}-\alpha-\beta x_{i}\right)^{2} . \tag{3.28}
\end{equation*}
$$

The solution to this task are the estimators:

$$
\begin{align*}
& \hat{\beta}=\frac{s_{X Y}}{s_{X X}}  \tag{3.29}\\
& \hat{\alpha}=\bar{y}-\hat{\beta} \bar{x} . \tag{3.30}
\end{align*}
$$

The variance of $\hat{\beta}$ is:

$$
\begin{equation*}
\operatorname{Var}(\hat{\beta})=\frac{\sigma^{2}}{n \cdot s_{X X}} \tag{3.31}
\end{equation*}
$$

The standard error (SE) of the estimator is the square root of (3.31),

$$
\begin{equation*}
\operatorname{SE}(\hat{\beta})=\{\operatorname{Var}(\hat{\beta})\}^{1 / 2}=\frac{\sigma}{\left(n \cdot s_{X X}\right)^{1 / 2}} \tag{3.32}
\end{equation*}
$$

We can use this formula to test the hypothesis that $\beta=0$. In an application the variance $\sigma^{2}$ has to be estimated by an estimator $\hat{\sigma}^{2}$ that will be given below. Under a normality assumption of the errors, the $t$-test for the hypothesis $\beta=0$ works as follows.

One computes the statistic

$$
\begin{equation*}
t=\frac{\hat{\beta}}{\operatorname{SE}(\hat{\beta})} \tag{3.33}
\end{equation*}
$$

and rejects the hypothesis at a $5 \%$ significance level if $|t| \geq t_{0.975 ; n-2}$, where the $97.5 \%$ quantile of the Student's $t_{n-2}$ distribution is clearly the $95 \%$ critical value for the two-sided test. For $n \geq 30$, this can be replaced by 1.96 , the $97.5 \%$ quantile of the normal distribution. An estimator $\hat{\sigma}^{2}$ of $\sigma^{2}$ will be given in the following.
Example 3.10 Let us apply the linear regression model (3.27) to the "classic blue" pullovers. The sales manager believes that there is a strong dependence on the


Fig. 3.5 Regression of sales $\left(X_{1}\right)$ on price $\left(X_{2}\right)$ of pullovers Q MVAregpull
number of sales as a function of price. He computes the regression line as shown in Fig. 3.5.

How good is this fit? This can be judged via goodness-of-fit measures. Define

$$
\begin{equation*}
\widehat{y}_{i}=\hat{\alpha}+\hat{\beta} x_{i} \tag{3.34}
\end{equation*}
$$

as the predicted value of $y$ as a function of $x$. With $\hat{y}$ the textile shop manager in the above example can predict sales as a function of prices $x$. The variation in the response variable is:

$$
\begin{equation*}
n s_{Y Y}=\sum_{i=1}^{n}\left(y_{i}-\bar{y}\right)^{2} \tag{3.35}
\end{equation*}
$$

The variation explained by the linear regression (3.27) with the predicted values (3.34) is:

$$
\begin{equation*}
\sum_{i=1}^{n}\left(\widehat{y}_{i}-\bar{y}\right)^{2} \tag{3.36}
\end{equation*}
$$

The residual sum of squares, the minimum in (3.28), is given by:

$$
\begin{equation*}
\mathrm{RSS}=\sum_{i=1}^{n}\left(y_{i}-\widehat{y}_{i}\right)^{2} \tag{3.37}
\end{equation*}
$$

An unbiased estimator $\widehat{\sigma}^{2}$ of $\sigma^{2}$ is given by $\operatorname{RSS} /(n-2)$.
The following relation holds between (3.35) and (3.37):

$$
\begin{align*}
\sum_{i=1}^{n}\left(y_{i}-\bar{y}\right)^{2} & =\sum_{i=1}^{n}\left(\widehat{y}_{i}-\bar{y}\right)^{2}+\sum_{i=1}^{n}\left(y_{i}-\widehat{y}_{i}\right)^{2}  \tag{3.38}\\
\text { Total variation } & =\text { Explained variation }+ \text { Unexplained variation. }
\end{align*}
$$

The coefficient of determination is $r^{2}$ :

$$
\begin{equation*}
r^{2}=\frac{\sum_{i=1}^{n}\left(\widehat{y}_{i}-\bar{y}\right)^{2}}{\sum_{i=1}^{n}\left(y_{i}-\bar{y}\right)^{2}}=\frac{\text { explained variation }}{\text { total variation }} \tag{3.39}
\end{equation*}
$$

The coefficient of determination increases with the proportion of explained variation by the linear relation (3.27). In the extreme cases where $r^{2}=1$, all of the variation is explained by the linear regression (3.27). The other extreme, $r^{2}=0$, is where the empirical covariance is $s_{X Y}=0$. The coefficient of determination can be rewritten as

$$
\begin{equation*}
r^{2}=1-\frac{\sum_{i=1}^{n}\left(y_{i}-\hat{y}_{i}\right)^{2}}{\sum_{i=1}^{n}\left(y_{i}-\bar{y}\right)^{2}} . \tag{3.40}
\end{equation*}
$$

From (3.39), it can be seen that in the linear regression (3.27), $r^{2}=r_{X Y}^{2}$ is the square of the correlation between $X$ and $Y$.

Example 3.11 For the above pullover example, we estimate

$$
\hat{\alpha}=210.774 \quad \text { and } \quad \hat{\beta}=-0.364
$$

The coefficient of determination is

$$
r^{2}=0.028
$$

The textile shop manager concludes that sales are not influenced very much by the price (in a linear way).

The geometrical representation of formula (3.38) can be graphically evaluated using Fig. 3.6. This plot shows a section of the linear regression of the "sales" on "price" for the pullovers data. The distance between any point and the overall mean is given by the distance between the point and the regression line and the distance between the regression line and the mean. The sums of these two distances represent the total variance (solid blue lines from the observations to the overall

Fig. 3.6 Regression of sales $\left(X_{1}\right)$ on price $\left(X_{2}\right)$ of pullovers. The overall mean is given by the dashed line $\mathbf{0}$ MVAregzoom


Fig. 3.7 Regression of $X_{5}$ (upper inner frame) on $X_{4}$ (lower inner frame) for genuine bank notes $\mathbf{Q}$ MVAregbank

mean), i.e. the explained variance (distance from the regression curve to the mean) and the unexplained variance (distance from the observation to the regression line), respectively.

In general the regression of $Y$ on $X$ is different from that of $X$ on $Y$. We will demonstrate this, once again, using the Swiss bank notes data.

Example 3.12 The least squares fit of the variables $X_{4}(X)$ and $X_{5}(Y)$ from the genuine bank notes are calculated. Figure 3.7 shows the fitted line if $X_{5}$ is approximated by a linear function of $X_{4}$. In this case the parameters are

$$
\hat{\alpha}=15.464 \quad \text { and } \quad \hat{\beta}=-0.638
$$

If we predict $X_{4}$ by a function of $X_{5}$ instead, we would arrive at a different intercept and slope

$$
\hat{\alpha}=14.666 \quad \text { and } \quad \hat{\beta}=-0.626
$$

The linear regression of $Y$ on $X$ is given by minimising (3.28), i.e. the vertical errors $\varepsilon_{i}$. The linear regression of $X$ on $Y$ does the same, but here the errors to be minimised in the least squares sense are measured horizontally. As seen in Example 3.12, the two least squares lines are different although both measure (in a certain sense) the slope of the cloud of points.

As shown in the next example, there is still one other way to measure the main direction of a cloud of points: it is related to the spectral decomposition of covariance matrices.

Example 3.13 Suppose that we have the following covariance matrix:

$$
\Sigma=\left(\begin{array}{ll}
1 & \rho \\
\rho & 1
\end{array}\right)
$$

Figure 3.8 shows a scatterplot of a sample of two normal random variables with such a covariance matrix (with $\rho=0.8$ ).

The eigenvalues of $\Sigma$ are, as was shown in Example 2.4, solutions to:

$$
\left|\begin{array}{cc}
1-\lambda & \rho \\
\rho & 1-\lambda
\end{array}\right|=0 .
$$

Fig. 3.8 Scatterplot for a sample of two correlated normal random variables (sample size $n=150$, $\rho=0.8$ ) Q MVAcorrnorm


Hence, $\lambda_{1}=1+\rho$ and $\lambda_{2}=1-\rho$. Therefore $\Lambda=\operatorname{diag}(1+\rho, 1-\rho)$. The eigenvector corresponding to $\lambda_{1}=1+\rho$ can be computed from the system of linear equations:

$$
\left(\begin{array}{ll}
1 & \rho \\
\rho & 1
\end{array}\right)\binom{x_{1}}{x_{2}}=(1+\rho)\binom{x_{1}}{x_{2}}
$$

or

$$
\begin{aligned}
& x_{1}+\rho x_{2}=x_{1}+\rho x_{1} \\
& \rho x_{1}+x_{2}=x_{2}+\rho x_{2}
\end{aligned}
$$

and thus

$$
x_{1}=x_{2} .
$$

The first (standardised) eigenvector is

$$
\gamma_{1}=\binom{1 / \sqrt{2}}{1 / \sqrt{2}}
$$

The direction of this eigenvector is the diagonal in Fig. 3.8 and captures the main variation in this direction. We shall come back to this interpretation in Chap. 11. The second eigenvector (orthogonal to $\gamma_{1}$ ) is

$$
\gamma_{2}=\binom{1 / \sqrt{2}}{-1 / \sqrt{2}} .
$$

So finally

$$
\Gamma=\left(\gamma_{1}, \gamma_{2}\right)=\left(\begin{array}{rr}
1 / \sqrt{2} & 1 / \sqrt{2} \\
1 / \sqrt{2} & -1 / \sqrt{2}
\end{array}\right)
$$

and we can check our calculation by

$$
\Sigma=\Gamma \Lambda \Gamma^{\top}
$$

The first eigenvector captures the main direction of a point cloud. The linear regression of $Y$ on $X$ and $X$ on $Y$ accomplished, in a sense, the same thing. In general the direction of the eigenvector and the least squares slope are different. The reason is that the least squares estimator minimises either vertical or horizontal errors (in (3.28)), whereas the first eigenvector corresponds to a minimisation that is orthogonal to the eigenvector (see Chap. 11).

| SLed Summary |
| :---: |
| $\hookrightarrow$ The linear regression $y=\alpha+\beta x+\varepsilon$ models a linear relation between two one-dimensional variables. |
| $\hookrightarrow$ The sign of the slope $\hat{\beta}$ is the same as that of the covariance and the correlation of $x$ and $y$. |
| $\hookrightarrow$ A linear regression predicts values of $Y$ given a possible observation $x$ of $X$. |
| $\hookrightarrow$ The coefficient of determination $r^{2}$ measures the amount of variation in $Y$ which is explained by a linear regression on $X$. |
| $\hookrightarrow$ If the coefficient of determination is $r^{2}=1$, then all points lie on one line. |
| $\hookrightarrow$ The regression line of $X$ on $Y$ and the regression line of $Y$ on $X$ are in general different. |
| $\hookrightarrow$ The $t$-test for the hypothesis $\beta=0$ is $t=\frac{\hat{\beta}}{\operatorname{SE}(\hat{\beta})}$, where $\operatorname{SE}(\hat{\beta})=$ $\frac{\hat{\sigma}}{\left(n \cdot s_{X X}\right)^{1 / 2}}$. |
| $\hookrightarrow$ The $t$-test rejects the null hypothesis $\beta=0$ at the level of significance $\alpha$ if $\|t\| \geq t_{1-\alpha / 2 ; n-2}$ where $t_{1-\alpha ; n-2}$ is the $1-\alpha / 2$ quantile of the Student's $t$-distribution with $(n-2)$ degrees of freedom. |
| $\hookrightarrow$ The standard error $\operatorname{SE}(\hat{\beta})$ increases/decreases with less/more spread in the $X$ variables. |
| $\hookrightarrow$ The direction of the first eigenvector of the covariance matrix of a two-dimensional point cloud is different from the least squares regression line. |

### 3.5 Simple Analysis of Variance

In a simple (i.e. one-factorial) analysis of variance (ANOVA), it is assumed that the average values of the response variable $y$ are induced by one simple factor. Suppose that this factor takes on $p$ values and that for each factor level, we have $m=n / p$ observations. The sample is of the form given in Table 3.1, where all of the observations are independent.

The goal of a simple ANOVA is to analyse the observation structure

$$
\begin{equation*}
y_{k l}=\mu_{l}+\varepsilon_{k l} \text { for } k=1, \ldots, m, \text { and } l=1, \ldots, p . \tag{3.41}
\end{equation*}
$$

Each factor has a mean value $\mu_{l}$. Each observation $y_{k l}$ is assumed to be a sum of the corresponding factor mean value $\mu_{l}$ and a zero mean random error $\varepsilon_{k l}$. The linear

Table 3.1 Observation structure of a simple ANOVA

| Sample element | Factor levels $l$ |  |  |  |  |
| :--- | :---: | :---: | :---: | :---: | :---: |
| 1 | $y_{11}$ | $\cdots$ | $y_{1 l}$ | $\cdots$ | $y_{1 p}$ |
| 2 | $\vdots$ |  | $\vdots$ |  | $\vdots$ |
| $\vdots$ | $\vdots$ |  | $\vdots$ |  | $\vdots$ |
| $k$ | $y_{k 1}$ | $\cdots$ | $y_{k l}$ | $\cdots$ | $y_{k p}$ |
| $\vdots$ | $\vdots$ |  | $\vdots$ |  | $\vdots$ |
| $m=n / p$ | $y_{m 1}$ | $\cdots$ | $y_{m l}$ | $\cdots$ | $y_{m p}$ |

Table 3.2 Pullover sales as function of marketing strategy

| Shop | Marketing strategy |  |  |
| :--- | ---: | ---: | ---: |
| $k$ | Factor $l$ |  |  |
|  | 1 | 2 | 3 |
| 1 | 9 | 10 | 18 |
| 2 | 11 | 15 | 14 |
| 3 | 10 | 11 | 17 |
| 4 | 12 | 15 | 9 |
| 5 | 7 | 15 | 14 |
| 6 | 11 | 13 | 17 |
| 7 | 12 | 7 | 16 |
| 8 | 10 | 15 | 14 |
| 9 | 11 | 13 | 17 |
| 10 | 13 | 10 | 15 |

regression model falls into this scheme with $m=1, p=n$ and $\mu_{i}=\alpha+\beta x_{i}$, where $x_{i}$ is the $i$ th level value of the factor.

Example 3.14 The "classic blue" pullover company analyses the effect of three marketing strategies

1. advertisement in local newspaper,
2. presence of sales assistant,
3. luxury presentation in shop windows.

All of these strategies are tried in ten different shops. The resulting sale observations are given in Table 3.2.

There are $p=3$ factors and $n=m p=30$ observations in the data. The "classic blue" pullover company wants to know whether all three marketing strategies have the same mean effect or whether there are differences. Having the same effect means that all $\mu_{l}$ in (3.41) equal one value, $\mu$. The hypothesis to be tested is therefore

$$
H_{0}: \mu_{l}=\mu \text { for } l=1, \ldots, p
$$

The alternative hypothesis, that the marketing strategies have different effects, can be formulated as

$$
H_{1}: \mu_{l} \neq \mu_{l^{\prime}} \text { for some } l \text { and } l^{\prime}
$$

This means that one marketing strategy is better than the others.
The method used to test this problem is to compute as in (3.38) the total variation and to decompose it into the sources of variation. This gives:

$$
\begin{equation*}
\sum_{l=1}^{p} \sum_{k=1}^{m}\left(y_{k l}-\bar{y}\right)^{2}=m \sum_{l=1}^{p}\left(\bar{y}_{l}-\bar{y}\right)^{2}+\sum_{l=1}^{p} \sum_{k=1}^{m}\left(y_{k l}-\bar{y}_{l}\right)^{2} \tag{3.42}
\end{equation*}
$$

The total variation (sum of squares $=\mathrm{SS}$ ) is:

$$
\begin{equation*}
\mathrm{SS}(\text { reduced })=\sum_{l=1}^{p} \sum_{k=1}^{m}\left(y_{k l}-\bar{y}\right)^{2} \tag{3.43}
\end{equation*}
$$

where $\bar{y}=n^{-1} \sum_{l=1}^{p} \sum_{k=1}^{m} y_{k l}$ is the overall mean. Here the total variation is denoted as SS (reduced), since in comparison with the model under the alternative $H_{1}$, we have a reduced set of parameters. In fact there is 1 parameter $\mu=\mu_{l}$ under $H_{0}$. Under $H_{1}$, the "full" model, we have three parameters, namely the three different means $\mu_{l}$.

The variation under $H_{1}$ is therefore:

$$
\begin{equation*}
\mathrm{SS}(\mathrm{full})=\sum_{l=1}^{p} \sum_{k=1}^{m}\left(y_{k l}-\bar{y}_{l}\right)^{2} \tag{3.44}
\end{equation*}
$$

where $\bar{y}_{l}=m^{-1} \sum_{k=1}^{m} y_{k l}$ is the mean of each factor $l$. The hypothetical model $H_{0}$ is called reduced, since it has (relative to $H_{1}$ ) fewer parameters.

The $F$-test of the linear hypothesis is used to compare the difference in the variations under the reduced model $H_{0}$ (3.43) and the full model $H_{1}$ (3.44) to the variation under the full model $H_{1}$ :

$$
\begin{equation*}
F=\frac{\{\operatorname{SS}(\text { reduced })-\operatorname{SS}(\text { full })\} /\{d f(r)-d f(f)\}}{\operatorname{SS}(\text { full }) / d f(f)} \tag{3.45}
\end{equation*}
$$

Here $d f(f)$ and $d f(r)$ denote the degrees of freedom under the full model and the reduced model, respectively. The degrees of freedom are essential in specifying the shape of the $F$-distribution. They have a simple interpretation: $d f(\cdot)$ is equal to the number of observations minus the number of parameters in the model.

From Example 3.14, $p=3$ parameters are estimated under the full model, i.e. $d f(f)=n-p=30-3=27$. Under the reduced model, there is one parameter to estimate, namely the overall mean, i.e. $d f(r)=n-1=29$. We can compute

$$
\mathrm{SS}(\text { reduced })=260.3
$$

and

$$
S S(\text { full })=157.7
$$

The $F$-statistic (3.45) is therefore

$$
F=\frac{(260.3-157.7) / 2}{157.7 / 27}=8.78
$$

This value needs to be compared to the quantiles of the $F_{2,27}$ distribution. Looking up the critical values in a $F$-distribution shows that the test statistic above is highly significant. We conclude that the marketing strategies have different effects.

## The F-Test in a Linear Regression Model

The $t$-test of a linear regression model can be put into this framework. For a linear regression model (3.27), the reduced model is the one with $\beta=0$ :

$$
y_{i}=\alpha+0 \cdot x_{i}+\varepsilon_{i} .
$$

The reduced model has $n-1$ degrees of freedom and one parameter, the intercept $\alpha$.
The full model is given by $\beta \neq 0$,

$$
y_{i}=\alpha+\beta \cdot x_{i}+\varepsilon_{i}
$$

and has $n-2$ degrees of freedom, since there are two parameters $(\alpha, \beta)$.
The SS(reduced) equals

$$
\mathrm{SS}(\text { reduced })=\sum_{i=1}^{n}\left(y_{i}-\bar{y}\right)^{2}=\text { total variation } .
$$

The SS(full) equals

$$
\mathrm{SS}(\mathrm{full})=\sum_{i=1}^{n}\left(y_{i}-\hat{y_{i}}\right)^{2}=\mathrm{RSS}=\text { unexplained variation. }
$$

The $F$-test is therefore, from (3.45),

$$
\begin{align*}
F & =\frac{(\text { total variation }- \text { unexplained variation }) / 1}{(\text { unexplained variation }) /(n-2)}  \tag{3.46}\\
& =\frac{\text { explained variation }}{(\text { unexplained variation }) /(n-2)} . \tag{3.47}
\end{align*}
$$

Using the estimators $\hat{\alpha}$ and $\hat{\beta}$ the explained variation is:

$$
\begin{aligned}
\sum_{i=1}^{n}\left(\hat{y}_{i}-\bar{y}\right)^{2} & =\sum_{i=1}^{n}\left(\hat{\alpha}+\hat{\beta} x_{i}-\bar{y}\right)^{2} \\
& =\sum_{i=1}^{n}\left\{(\bar{y}-\hat{\beta} \bar{x})+\hat{\beta} x_{i}-\bar{y}\right\}^{2} \\
& =\sum_{i=1}^{n} \hat{\beta}^{2}\left(x_{i}-\bar{x}\right)^{2} \\
& =\hat{\beta}^{2} n s_{X X}
\end{aligned}
$$

From (3.32) the $F$-ratio (3.46) is therefore:

$$
\begin{align*}
F & =\frac{\hat{\beta}^{2} n s_{X X}}{\operatorname{RSS} /(n-2)}  \tag{3.48}\\
& =\left(\frac{\hat{\beta}}{\operatorname{SE}(\hat{\beta})}\right)^{2} \tag{3.49}
\end{align*}
$$

The $t$-test statistic (3.33) is just the square root of the $F$-statistic (3.49).
Note, using (3.39) the $F$-statistic can be rewritten as

$$
F=\frac{r^{2} / 1}{\left(1-r^{2}\right) /(n-2)}
$$

In the pullover Example 3.11, we obtain $F=\frac{0.028}{0.972} \frac{8}{1}=0.2305$, so that the null hypothesis $\beta=0$ cannot be rejected. We conclude therefore that there is only a minor influence of prices on sales.

| $\hookrightarrow$ | Simple ANOVA models an output $Y$ as a function of one factor. |
| :--- | :--- |
|  | $\hookrightarrow$ | The reduced model is the hypothesis of equal means.

### 3.6 Multiple Linear Model

The simple linear model and the analysis of variance model can be viewed as a particular case of a more general linear model where the variations of one variable $y$ are explained by $p$ explanatory variables $x$ respectively. Let $y(n \times 1)$ and $\mathcal{X}(n \times p)$ be a vector of observations on the response variable and a data matrix on the $p$ explanatory variables. An important application of the developed theory is the least squares fitting. The idea is to approximate $y$ by a linear combination $\hat{y}$ of columns of $\mathcal{X}$, i.e. $\hat{y} \in C(\mathcal{X})$. The problem is to find $\hat{\beta} \in \mathbb{R}^{p}$ such that $\hat{y}=\mathcal{X} \hat{\beta}$ is the best fit of $y$ in the least-squares sense. The linear model can be written as

$$
\begin{equation*}
y=\mathcal{X} \beta+\varepsilon, \tag{3.50}
\end{equation*}
$$

where $\varepsilon$ are the errors. The least squares solution is given by $\hat{\beta}$ :

$$
\begin{equation*}
\hat{\beta}=\arg \min _{\beta}(y-\mathcal{X} \beta)^{\top}(y-\mathcal{X} \beta)=\arg \min _{\beta} \varepsilon^{\top} \varepsilon \tag{3.51}
\end{equation*}
$$

Suppose that $\left(\mathcal{X}^{\top} \mathcal{X}\right)$ is of full rank and thus invertible. Minimising the expression (3.51) with respect to $\beta$ yields:

$$
\begin{equation*}
\hat{\beta}=\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1} \mathcal{X}^{\top} y \tag{3.52}
\end{equation*}
$$

The fitted value $\hat{y}=\mathcal{X} \hat{\beta}=\mathcal{X}\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1} \mathcal{X}^{\top} y=\mathcal{P} y$ is the projection of $y$ onto $C(\mathcal{X})$ as computed in (2.47).

The least squares residuals are

$$
e=y-\hat{y}=y-\mathcal{X} \hat{\beta}=\mathcal{Q} y=\left(\mathcal{I}_{n}-\mathcal{P}\right) y .
$$

The vector $e$ is the projection of $y$ onto the orthogonal complement of $C(\mathcal{X})$.
Remark 3.5 A linear model with an intercept $\alpha$ can also be written in this framework. The approximating equation is:

$$
y_{i}=\alpha+\beta_{1} x_{i 1}+\cdots+\beta_{p} x_{i p}+\varepsilon_{i} ; i=1, \ldots, n .
$$

This can be written as:

$$
y=\mathcal{X}^{*} \beta^{*}+\varepsilon
$$

where $\mathcal{X}^{*}=\left(1_{n} \mathcal{X}\right)$ (we add a column of ones to the data). We have by (3.52):

$$
\hat{\beta}^{*}=\binom{\hat{\alpha}}{\hat{\beta}}=\left(\mathcal{X}^{* \top} \mathcal{X}^{*}\right)^{-1} \mathcal{X}^{* \top} y .
$$

Example 3.15 Let us come back to the "classic blue" pullovers example. In Example 3.11, we considered the regression fit of the sales $X_{1}$ on the price $X_{2}$ and concluded that there was only a small influence of sales by changing the prices. A linear model incorporating all three variables allows us to approximate sales as a linear function of price $\left(X_{2}\right)$, advertisement ( $X_{3}$ ) and presence of sales assistants $\left(X_{4}\right)$ simultaneously. Adding a column of ones to the data (in order to estimate the intercept $\alpha$ ) leads to

$$
\hat{\alpha}=65.670 \text { and } \widehat{\beta}_{1}=-0.216, \widehat{\beta}_{2}=0.485, \widehat{\beta}_{3}=0.844
$$

The coefficient of determination is computed as before in (3.40) and is:

$$
r^{2}=1-\frac{e^{\top} e}{\sum\left(y_{i}-\bar{y}\right)^{2}}=0.907
$$

We conclude that the variation of $X_{1}$ is well approximated by the linear relation.

Remark 3.6 The coefficient of determination is influenced by the number of regressors. For a given sample size $n$, the $r^{2}$ value will increase by adding more regressors into the linear model. The value of $r^{2}$ may therefore be high even if possibly irrelevant regressors are included. An adjusted coefficient of determination for $p$ regressors and a constant intercept ( $p+1$ parameters) is

$$
\begin{equation*}
r_{\mathrm{adj}}^{2}=r^{2}-\frac{p\left(1-r^{2}\right)}{n-(p+1)} \tag{3.53}
\end{equation*}
$$

Example 3.16 The corrected coefficient of determination for Example 3.15 is

$$
\begin{aligned}
r_{\mathrm{adj}}^{2} & =0.907-\frac{3\left(1-0.907^{2}\right)}{10-3-1} \\
& =0.818
\end{aligned}
$$

This means that $81.8 \%$ of the variation of the response variable is explained by the explanatory variables.

Note that the linear model (3.50) is very flexible and can model non-linear relationships between the response $y$ and the explanatory variables $x$. For example, a quadratic relation in one variable $x$ could be included. Then $y_{i}=\alpha+\beta_{1} x_{i}+$ $\beta_{2} x_{i}^{2}+\varepsilon_{i}$ could be written in matrix notation as in (3.50), $y=\mathcal{X} \beta+\varepsilon$ where

$$
\mathcal{X}=\left(\begin{array}{ccc}
1 & x_{1} & x_{1}^{2} \\
1 & x_{2} & x_{2}^{2} \\
\vdots & \vdots & \vdots \\
1 & x_{n} & x_{n}^{2}
\end{array}\right)
$$

## Properties of $\hat{\boldsymbol{\beta}}$

When $y_{i}$ is the $i$ th observation of a random variable $Y$, the errors are also random. Under standard assumptions (independence, zero mean and constant variance $\sigma^{2}$ ), inference can be conducted on $\beta$. Using the properties of Chap. 4 , it is easy to prove:

$$
\begin{aligned}
& \mathrm{E}(\hat{\beta})=\beta \\
& \operatorname{Var}(\hat{\beta})=\sigma^{2}\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1}
\end{aligned}
$$

The analogue of the $t$-test for the multivariate linear regression situation is

$$
t=\frac{\widehat{\beta}_{j}}{\operatorname{SE}\left(\widehat{\beta}_{j}\right)}
$$

The standard error of each coefficient $\widehat{\beta}_{j}$ is given by the square root of the diagonal elements of the matrix $\operatorname{Var}(\hat{\beta})$. In standard situations, the variance $\sigma^{2}$ of the error $\varepsilon$ is not known. For linear model with intercept, one may estimate it by

$$
\hat{\sigma}^{2}=\frac{1}{n-(p+1)}(y-\hat{y})^{\top}(y-\hat{y}),
$$

where $p$ is the dimension of $\beta$. In testing $\beta_{j}=0$ we reject the hypothesis at the significance level $\alpha$ if $|t| \geq t_{1-\alpha / 2 ; n-(p+1)}$. More general issues on testing linear models are addressed in Chap. 7.

## The ANOVA Model in Matrix Notation

The simple ANOVA problem (Sect.3.5) may also be rewritten in matrix terms. Recall the definition of a vector of ones from (2.1) and define a vector of zeros as $0_{n}$. Then construct the following $(n \times p)$ matrix (here $p=3$ ),

$$
\mathcal{X}=\left(\begin{array}{ccc}
1_{m} & 0_{m} & 0_{m}  \tag{3.54}\\
0_{m} & 1_{m} & 0_{m} \\
0_{m} & 0_{m} & 1_{m}
\end{array}\right),
$$

where $m=10$. Equation (3.41) then reads as follows.
The parameter vector is $\beta=\left(\mu_{1}, \mu_{2}, \mu_{3}\right)^{\top}$. The data set from Example 3.14 can therefore be written as a linear model $y=\mathcal{X} \beta+\varepsilon$ where $y \in \mathbb{R}^{n}$ with $n=m \cdot p$ is the stacked vector of the columns of Table 3.1. The projection into the column space $C(\mathcal{X})$ of (3.54) yields the least-squares estimator $\hat{\beta}=\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1} \mathcal{X}^{\top} y$. Note that $\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1}=(1 / 10) \mathcal{I}_{3}$ and that $\mathcal{X}^{\top} y=(106,124,151)^{\top}$ is the sum $\sum_{k=1}^{m} y_{k j}$ for each factor, i.e. the three column sums of Table 3.1. The least squares estimator is therefore the vector $\hat{\beta}_{H_{1}}=\left(\hat{\mu}_{1}, \hat{\mu}_{2}, \hat{\mu}_{3}\right)=(10.6,12.4,15.1)^{\top}$ of sample means for each factor level $j=1,2,3$. Under the null hypothesis of equal mean values $\mu_{1}=\mu_{2}=\mu_{3}=\mu$, we estimate the parameters under the same constraints. This can be put into the form of a linear constraint:

$$
\begin{aligned}
& -\mu_{1}+\mu_{2}=0 \\
& -\mu_{1}+\mu_{3}=0 .
\end{aligned}
$$

This can be written as $\mathcal{A} \beta=a$, where

$$
a=\binom{0}{0}
$$

and

$$
\mathcal{A}=\left(\begin{array}{lll}
-1 & 1 & 0 \\
-1 & 0 & 1
\end{array}\right) .
$$

The constrained least-squares solution can be shown (Exercise 3.24) to be given by:

$$
\begin{equation*}
\hat{\beta}_{H_{0}}=\hat{\beta}_{H_{1}}-\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1} \mathcal{A}^{\top}\left\{\mathcal{A}\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1} \mathcal{A}^{\top}\right\}^{-1}\left(\mathcal{A} \hat{\beta}_{H_{1}}-a\right) . \tag{3.55}
\end{equation*}
$$

It turns out that (3.55) amounts to simply calculating the overall mean $\bar{y}=12.7$ of the response variable $y: \hat{\beta}_{H_{0}}=(12.7,12.7,12.7)^{\top}$.

The $F$-test that has already been applied in Example 3.14 can be written as

$$
\begin{equation*}
F=\frac{\left\{\left\|y-\mathcal{X} \hat{\beta}_{H_{0}}\right\|^{2}-\left\|y-\mathcal{X} \hat{\beta}_{H_{1}}\right\|^{2}\right\} / 2}{\left\|y-\mathcal{X} \hat{\beta}_{H_{1}}\right\|^{2} / 27} \tag{3.56}
\end{equation*}
$$

which gives the same significant value 8.78. Note that again we compare the $\operatorname{RSS}_{H_{0}}$ of the reduced model to the $\mathrm{RSS}_{H_{1}}$ of the full model. It corresponds to comparing the lengths of projections into different column spaces. This general approach in testing linear models is described in detail in Chap. 7.

| Summary |
| :---: |
| $\hookrightarrow$ The relation $y=\mathcal{X} \beta+e$ models a linear relation between a onedimensional variable $Y$ and a $p$-dimensional variable $X$. $\mathcal{P} y$ gives the best linear regression fit of the vector $y$ onto $C(\mathcal{X})$. The least squares parameter estimator is $\hat{\beta}=\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1} \mathcal{X}^{\top} y$. |
| $\hookrightarrow$ The simple ANOVA model can be written as a linear model. |
| $\hookrightarrow$ The ANOVA model can be tested by comparing the length of the projection vectors. |
| $\hookrightarrow$ The test statistic of the $F$-test can be written as $\frac{\left\{\left\\|y-\mathcal{X} \hat{\beta}_{H_{0}}\right\\|^{2}-\left\\|y-\mathcal{X} \hat{\beta}_{H_{1}}\right\\|^{2}\right\} /\{d f(r)-d f(f)\}}{\left\\|y-\mathcal{X} \hat{\beta}_{H_{1}}\right\\|^{2} / d f(f)}$ |
| $\hookrightarrow$ The adjusted coefficient of determination is $r_{\mathrm{adj}}^{2}=r^{2}-\frac{p\left(1-r^{2}\right)}{n-(p+1)}$ |

### 3.7 Boston Housing

The main statistics presented so far can be computed for the data matrix $\mathcal{X}(506 \times 14)$ from our Boston Housing data set. The sample means and the sample medians of each variable are displayed in Table 3.3. The table also provides the unbiased estimates of the variance of each variable and the corresponding standard deviations. The comparison of the means and the medians confirms the asymmetry of the components of $\mathcal{X}$ that was pointed out in Sect. 1.9.

The (unbiased) sample covariance matrix is given by the following ( $14 \times 14$ ) matrix $\mathcal{S}_{n}$ :
$\left(\begin{array}{rrrrrrrrrrrrrr}73.99 & -40.22 & 23.99 & -0.12 & 0.42 & -1.33 & 85.41 & -6.88 & 46.85 & 844.82 & 5.40 & -302.38 & 27.99 & -30.72 \\ -40.22 & 543.94 & -85.41 & -0.25 & -1.40 & 5.11 & -373.90 & 32.63 & -63.35 & -1236.45 & -19.78 & 373.72 & -68.78 & 77.32 \\ 23.99 & -85.41 & 47.06 & 0.11 & 0.61 & -1.89 & 124.51 & -10.23 & 35.55 & 833.36 & 5.69 & -223.58 & 29.58 & -30.52 \\ -0.12 & -0.25 & 0.11 & 0.06 & 0.00 & 0.02 & 0.62 & -0.05 & -0.02 & -1.52 & -0.07 & 1.13 & -0.10 & 0.41 \\ 0.42 & -1.40 & 0.61 & 0.00 & 0.01 & -0.02 & 2.39 & -0.19 & 0.62 & 13.05 & 0.05 & -4.02 & 0.49 & -0.46 \\ -1.33 & 5.11 & -1.89 & 0.02 & -0.02 & 0.49 & -4.75 & 0.30 & -1.28 & -34.58 & -0.54 & 8.22 & -3.08 & 4.49 \\ 85.41 & -373.90 & 124.51 & 0.62 & 2.39 & -4.75 & 792.36 & -44.33 & 111.77 & 2402.69 & 15.94 & -702.94 & 121.08 & -97.59 \\ -6.88 & 32.63 & -10.23 & -0.05 & -0.19 & 0.30 & -44.33 & 4.43 & -9.07 & -189.66 & -1.06 & 56.04 & -7.47 & 4.84 \\ 46.85 & -63.35 & 35.55 & -0.02 & 0.62 & -1.28 & 111.77 & -9.07 & 75.82 & 1335.76 & 8.76 & -353.28 & 30.39 & -30.56 \\ 844.82 & -1236.45 & 833.36 & -1.52 & 13.05 & -34.58 & 2402.69 & -189.66 & 1335.76 & 28404.76 & 168.15 & -6797.91 & 654.71 & -726.26 \\ 5.40 & -19.78 & 5.69-0.07 & 0.05 & -0.54 & 15.94 & -1.06 & 8.76 & 168.15 & 4.69 & -35.06 & 5.78 & -10.11 \\ -302.38 & 373.72 & -223.58 & 1.13 & -4.02 & 8.22 & -702.94 & 56.04 & -353.28 & -6797.91 & -35.06 & 8334.75 & -238.67 & 279.99 \\ 27.99 & -68.78 & 29.58 & -0.10 & 0.49 & -3.08 & 121.08 & -7.47 & 30.39 & 654.71 & 5.78 & -238.67 & 50.99 & -48.45 \\ -30.72 & 77.32 & -30.52 & 0.41 & -0.46 & 4.49 & -97.59 & 4.84 & -30.56 & -726.26 & -10.11 & 279.99 & -48.45 & 84.59\end{array}\right)$,
and the corresponding correlation matrix $\mathcal{R}(14 \times 14)$ is:

Table 3.3 Descriptive statistics for the Boston Housing data set Q MVAdescbh

| $X$ | $\bar{x}$ | $\operatorname{Median}(X)$ | $\operatorname{Var}(X)$ | $\operatorname{Std}(X)$ |
| :--- | ---: | :---: | ---: | ---: |
| $X_{1}$ | 3.61 | 0.26 | 73.99 | 8.60 |
| $X_{2}$ | 11.36 | 0.00 | 543.94 | 23.32 |
| $X_{3}$ | 11.14 | 9.69 | 47.06 | 6.86 |
| $X_{4}$ | 0.07 | 0.00 | 0.06 | 0.25 |
| $X_{5}$ | 0.55 | 0.54 | 0.01 | 0.12 |
| $X_{6}$ | 6.28 | 6.21 | 0.49 | 0.70 |
| $X_{7}$ | 68.57 | 77.50 | 792.36 | 28.15 |
| $X_{8}$ | 3.79 | 3.21 | 4.43 | 2.11 |
| $X_{9}$ | 9.55 | 5.00 | 75.82 | 8.71 |
| $X_{10}$ | 408.24 | 330.00 | $28,405.00$ | 168.54 |
| $X_{11}$ | 18.46 | 19.05 | 4.69 | 2.16 |
| $X_{12}$ | 356.67 | 391.44 | $8,334.80$ | 91.29 |
| $X_{13}$ | 12.65 | 11.36 | 50.99 | 7.14 |
| $X_{14}$ | 22.53 | 21.20 | 84.59 | 9.20 |

$$
\left(\begin{array}{rrrrrrrrrrrrrr}
1.00-0.20 & 0.41 & -0.06 & 0.42 & -0.22 & 0.35 & -0.38 & 0.63 & 0.58 & 0.29 & -0.39 & 0.46 & -0.39 \\
-0.20 & 1.00 & -0.53 & -0.04 & -0.52 & 0.31 & -0.57 & 0.66 & -0.31 & -0.31 & -0.39 & 0.18 & -0.41 & 0.36 \\
0.41-0.53 & 1.00 & 0.06 & 0.76 & -0.39 & 0.64 & -0.71 & 0.60 & 0.72 & 0.38 & -0.36 & 0.60 & -0.48 \\
-0.06-0.04 & 0.06 & 1.00 & 0.09 & 0.09 & 0.09 & -0.10 & -0.01 & -0.04 & -0.12 & 0.05 & -0.05 & 0.18 \\
0.42-0.52 & 0.76 & 0.09 & 1.00 & -0.30 & 0.73 & -0.77 & 0.61 & 0.67 & 0.19 & -0.38 & 0.59 & -0.43 \\
-0.22 & 0.31 & -0.39 & 0.09 & -0.30 & 1.00 & -0.24 & 0.21 & -0.21 & -0.29 & -0.36 & 0.13 & -0.61 & 0.70 \\
0.35-0.57 & 0.64 & 0.09 & 0.73-0.24 & 1.00-0.75 & 0.46 & 0.51 & 0.26 & -0.27 & 0.60 & -0.38 \\
-0.38 & 0.66 & -0.71 & -0.10 & -0.77 & 0.21 & -0.75 & 1.00 & -0.49 & -0.53 & -0.23 & 0.29 & -0.50 & 0.25 \\
0.63-0.31 & 0.60-0.01 & 0.61 & -0.21 & 0.46 & -0.49 & 1.00 & 0.91 & 0.46 & -0.44 & 0.49 & -0.38 \\
0.58-0.31 & 0.72-0.04 & 0.67 & -0.29 & 0.51 & -0.53 & 0.91 & 1.00 & 0.46 & -0.44 & 0.54 & -0.47 \\
0.29-0.39 & 0.38 & -0.12 & 0.19-0.36 & 0.26 & -0.23 & 0.46 & 0.46 & 1.00 & -0.18 & 0.37 & -0.51 \\
-0.39 & 0.18 & -0.36 & 0.05 & -0.38 & 0.13 & -0.27 & 0.29 & -0.44 & -0.44 & -0.18 & 1.00 & -0.37 & 0.33 \\
0.46-0.41 & 0.60-0.05 & 0.59-0.61 & 0.60 & -0.50 & 0.49 & 0.54 & 0.37 & -0.37 & 1.00 & -0.74 \\
-0.39 & 0.36 & -0.48 & 0.18 & -0.43 & 0.70-0.38 & 0.25 & -0.38 & -0.47 & -0.51 & 0.33 & -0.74 & 1.00
\end{array}\right)
$$

Analysing $\mathcal{R}$ confirms most of the comments made from examining the scatterplot matrix in Chap. 1. In particular, the correlation between $X_{14}$ (the value of the house) and all the other variables is given by the last row (or column) of $\mathcal{R}$. The highest correlations (in absolute values) are in decreasing order $X_{13}, X_{6}, X_{11}, X_{10}$, etc.

Using the Fisher's $Z$-transform on each of the correlations between $X_{14}$ and the other variables would confirm that all are significantly different from zero, except the correlation between $X_{14}$ and $X_{4}$ (the indicator variable for the Charles River). We know, however, that the correlation and Fisher's $Z$-transform are not appropriate for binary variable.

The same descriptive statistics can be calculated for the transformed variables (transformations were motivated in Sect. 1.9). The results are given in Table 3.4 and as can be seen, most of the variables are now more symmetric. Note that the

Table 3.4 Descriptive statistics for the Boston Housing data set after the transformation $\mathbf{Q}$ MVAdescbh

| $\tilde{X}$ | $\overline{\tilde{x}}$ | Median( $\tilde{X})$ | $\operatorname{Var}(\tilde{X})$ | $\operatorname{Std}(\tilde{X})$ |
| :---: | :---: | :---: | :---: | :---: |
| $\widetilde{X}_{1}$ | -0.78 | -1.36 | 4.67 | 2.16 |
| $\widetilde{X}_{2}$ | 1.14 | 0.00 | 5.44 | 2.33 |
| $\widetilde{X}_{3}$ | 2.16 | 2.27 | 0.60 | 0.78 |
| $\widetilde{X}_{4}$ | 0.07 | 0.00 | 0.06 | 0.25 |
| $\widetilde{X}_{5}$ | -0.61 | -0.62 | 0.04 | 0.20 |
| $\widetilde{X}_{6}$ | 1.83 | 1.83 | 0.01 | 0.11 |
| $\widetilde{X}_{7}$ | 5.06 | 5.29 | 12.72 | 3.57 |
| ${ }_{\sim}^{X_{8}}$ | 1.19 | 1.17 | 0.29 | 0.54 |
| $\widetilde{X}^{\widetilde{X}}$ | 1.87 | 1.61 | 0.77 | 0.87 |
| $\widetilde{X}_{10}$ | 5.93 | 5.80 | 0.16 | 0.40 |
| $\widetilde{X}_{11}$ | 2.15 | 2.04 | 1.86 | 1.36 |
| $\widetilde{X}_{12}$ | 3.57 | 3.91 | 0.83 | 0.91 |
| $\widetilde{X}_{13}$ | 3.42 | 3.37 | 0.97 | 0.99 |
| $\widetilde{X}_{14}$ | 3.03 | 3.05 | 0.17 | 0.41 |

covariances and the correlations are sensitive to these non-linear transformations. For example, the correlation matrix is now

Notice that some of the correlations between $\widetilde{X}_{14}$ and the other variables have increased.

If we want to explain the variations of the price $\widetilde{X}_{14}$ by the variation of all the other variables $\widetilde{X}_{1}, \ldots, \widetilde{X}_{13}$ we could estimate the linear model

$$
\begin{equation*}
\widetilde{X}_{14}=\beta_{0}+\sum_{j=1}^{13} \beta_{j} \widetilde{X}_{j}+\varepsilon \tag{3.57}
\end{equation*}
$$

The result is given in Table 3.5.

Table 3.5 Linear regression results for all variables of Boston Housing data set $\mathbf{Q}$ MVAlinregbh

| Variable | $\hat{\beta}_{j}$ | $\operatorname{SE}\left(\hat{\beta}_{j}\right)$ | $t$ | $p$-Value |
| :--- | ---: | ---: | ---: | :--- |
| Constant | 4.1769 | 0.3790 | 11.020 | 0.0000 |
| $\widetilde{X}_{1}$ | -0.0146 | 0.0117 | -1.254 | 0.2105 |
| $\widetilde{X}_{2}$ | 0.0014 | 0.0056 | 0.247 | 0.8051 |
| $\widetilde{X}_{3}$ | -0.0127 | 0.0223 | -0.570 | 0.5692 |
| $\widetilde{X}_{4}$ | 0.1100 | 0.0366 | 3.002 | 0.0028 |
| $\widetilde{X}_{5}$ | -0.2831 | 0.1053 | -2.688 | 0.0074 |
| $\widetilde{X}_{6}$ | 0.4211 | 0.1102 | 3.822 | 0.0001 |
| $\widetilde{X}_{7}$ | 0.0064 | 0.0049 | 1.317 | 0.1885 |
| $\widetilde{X}_{8}$ | -0.1832 | 0.0368 | -4.977 | 0.0000 |
| $\widetilde{\widetilde{X}}_{9}$ | 0.0684 | 0.0225 | 3.042 | 0.0025 |
| $\widetilde{X}_{10}$ | -0.2018 | 0.0484 | -4.167 | 0.0000 |
| $\widetilde{X}_{11}$ | -0.0400 | 0.0081 | -4.946 | 0.0000 |
| $\widetilde{X}_{12}$ | 0.0445 | 0.0115 | 3.882 | 0.0001 |
| $\widetilde{X}_{13}$ | -0.2626 | 0.0161 | -16.320 | 0.0000 |

The value of $r^{2}(0.765)$ and $r_{a d j}^{2}(0.759)$ show that most of the variance of $X_{14}$ is explained by the linear model (3.57).

Again we see that the variations of $\widetilde{X}_{14}$ are mostly explained by (in decreasing order of the absolute value of the $t$-statistic) $\widetilde{X}_{13}, \widetilde{X}_{8}, \widetilde{X}_{11}, \widetilde{X}_{10}, \widetilde{X}_{12}, \widetilde{X}_{6}, \widetilde{X}_{9}, \widetilde{X}_{4}$ and $\widetilde{X}_{5}$. The other variables $\widetilde{X}_{1}, \widetilde{X}_{2}, \widetilde{X}_{3}$ and $\widetilde{X}_{7}$ seem to have little influence on the variations of $\widetilde{X}_{14}$. This will be confirmed by the testing procedures that will be developed in Chap. 7.

### 3.8 Exercises

Exercise 3.1 The covariance $s_{X_{4} X_{5}}$ between $X_{4}$ and $X_{5}$ for the entire bank data set is positive. Given the definitions of $X_{4}$ and $X_{5}$, we would expect a negative covariance. Using Fig. 3.1 can you explain why $s_{X_{4} X_{5}}$ is positive?

Exercise 3.2 Consider the two sub-clouds of counterfeit and genuine bank notes in Fig. 3.1 separately. Do you still expect $s_{X_{4} X_{5}}$ (now calculated separately for each cloud) to be positive?

Exercise 3.3 We remarked that for two normal random variables, zero covariance implies independence. Why does this remark not apply to Example 3.4?

Exercise 3.4 Compute the covariance between the variables

$$
\begin{aligned}
& X_{2}=\text { miles per gallon }, \\
& X_{8}=\text { weight }
\end{aligned}
$$

from the car data set (Table 22.3). What sign do you expect the covariance to have?
Exercise 3.5 Compute the correlation matrix of the variables in Example 3.2. Comment on the sign of the correlations and test the hypothesis

$$
\rho_{X_{1} X_{2}}=0 .
$$

Exercise 3.6 Suppose you have observed a set of observations $\left\{x_{i}\right\}_{i=1}^{n}$ with $\bar{x}=0$, $s_{X X}=1$ and $n^{-1} \sum_{i=1}^{n}\left(x_{i}-\bar{x}\right)^{3}=0$. Define the variable $y_{i}=x_{i}^{2}$. Can you immediately tell whether $r_{X Y} \neq 0$ ?

Exercise 3.7 Find formulas (3.29) and (3.30) for $\hat{\alpha}$ and $\hat{\beta}$ by differentiating the objective function in (3.28) w.r.t. $\alpha$ and $\beta$.

Exercise 3.8 How many sales does the textile manager expect with a "classic blue" pullover price of $x=105$ ?

Exercise 3.9 What does a scatterplot of two random variables look like for $r^{2}=1$ and $r^{2}=0$ ?

Exercise 3.10 Prove the variance decomposition (3.38) and show that the coefficient of determination is the square of the simple correlation between $X$ and $Y$.

Exercise 3.11 Make a boxplot for the residuals $\varepsilon_{i}=y_{i}-\hat{\alpha}-\hat{\beta} x_{i}$ for the "classic blue" pullovers data. If there are outliers, identify them and run the linear regression again without them. Do you obtain a stronger influence of price on sales?

Exercise 3.12 Under what circumstances would you obtain the same coefficients from the linear regression lines of $Y$ on $X$ and of $X$ on $Y$ ?

Exercise 3.13 Treat the design of Example 3.14 as if there were thirty shops and not ten. Define $x_{i}$ as the index of the shop, i.e. $x_{i}=i, i=1,2, \ldots, 30$. The null hypothesis is a constant regression line, $E Y=\mu$. What does the alternative regression curve look like?

Exercise 3.14 Perform the test in Exercise 3.13 for the shop example with a 0.99 significance level. Do you still reject the hypothesis of equal marketing strategies?

Exercise 3.15 Compute an approximate confidence interval for $\rho_{X_{1} X_{4}}$ in Example 3.2. Hint: start from a confidence interval for $\tanh ^{-1}\left(\rho_{X_{1} X_{4}}\right)$ and then apply the inverse transformation.

Exercise 3.16 In Example 3.2, using the exchange rate of 1 EUR $=106 \mathrm{JPY}$, compute the same empirical covariance using prices in Japanese Yen rather than in Euros. Is there a significant difference? Why?

Exercise 3.17 Why does the correlation have the same sign as the covariance?
Exercise 3.18 Show that $\operatorname{rank}(\mathcal{H})=\operatorname{tr}(\mathcal{H})=n-1$.
Exercise 3.19 Show that $\mathcal{X}_{*}=\mathcal{H X} \mathcal{D}^{-1 / 2}$ is the standardised data matrix, i.e. $\bar{x}_{*}=0$ and $\mathcal{S}_{\mathcal{X}_{*}}=\mathcal{R}_{\mathcal{X}}$.

Exercise 3.20 Compute for the pullovers data the regression of $X_{1}$ on $X_{2}, X_{3}$ and of $X_{1}$ on $X_{2}, X_{4}$. Which one has the better coefficient of determination?

Exercise 3.21 Compare for the pullovers data the coefficient of determination for the regression of $X_{1}$ on $X_{2}$ (Example 3.11), of $X_{1}$ on $X_{2}, X_{3}$ (Exercise 3.20) and of $X_{1}$ on $X_{2}, X_{3}, X_{4}$ (Example 3.15). Observe that this coefficient is increasing with the number of predictor variables. Is this always the case?

Exercise 3.22 Consider the ANOVA problem (Sect. 3.5) again. Establish the constraint Matrix $\mathcal{A}$ for testing $\mu_{1}=\mu_{2}$. Test this hypothesis via an analog of (3.55) and (3.56).
 $\frac{\partial f(\beta)}{\partial \beta}=0$.)
Exercise 3.24 Consider the linear model $Y=\mathcal{X} \beta+\varepsilon$ where $\hat{\beta}=\arg \min _{\beta} \varepsilon^{\top} \varepsilon$ is subject to the linear constraints $\mathcal{A} \hat{\beta}=a$ where $\mathcal{A}(q \times p),(q \leq p)$ is of rank $q$ and
a is of dimension $(q \times 1)$. Show that $\hat{\beta}=\hat{\beta}_{\text {OLS }}-\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1} \mathcal{A}^{\top}\left(\mathcal{A}\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1} \mathcal{A}^{\top}\right)^{-1}$ $\left(\mathcal{A} \hat{\beta}_{\text {OLS }}-a\right)$ where $\hat{\beta}_{\text {OLS }}=\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1} \mathcal{X}^{\top} y$. (Hint, let $f(\beta, \lambda)=(y-x \beta)^{\top}(y-$ $x \beta)-\lambda^{\top}(\mathcal{A} \beta-a)$ where $\lambda \in \mathbb{R}^{q}$ and solve $\frac{\partial f(\beta, \lambda)}{\partial \beta}=0$ and $\frac{\partial f(\beta, \lambda)}{\partial \lambda}=0$.)
Exercise 3.25 Compute the covariance matrix $\mathcal{S}=\operatorname{Cov}(\mathcal{X})$ where $\mathcal{X}$ denotes the matrix of observations on the counterfeit bank notes. Make a Jordan decomposition of $\mathcal{S}$. Why are all of the eigenvalues positive?
Exercise 3.26 Compute the covariance of the counterfeit notes after they are linearly transformed by the vector $a=(1,1,1,1,1,1)^{\top}$.

## Chapter 4 <br> Multivariate Distributions

The preceding chapter showed that by using the two first moments of a multivariate distribution (the mean and the covariance matrix), a lot of information on the relationship between the variables can be made available. Only basic statistical theory was used to derive tests of independence or of linear relationships. In this chapter we give an introduction to the basic probability tools useful in statistical multivariate analysis.

Means and covariances share many interesting and useful properties, but they represent only part of the information on a multivariate distribution. Section 4.1 presents the basic probability tools used to describe a multivariate random variable, including marginal and conditional distributions and the concept of independence. In Sect. 4.2, basic properties on means and covariances (marginal and conditional ones) are derived.

Since many statistical procedures rely on transformations of a multivariate random variable, Sect. 4.3 proposes the basic techniques needed to derive the distribution of transformations with a special emphasis on linear transforms. As an important example of a multivariate random variable, Sect. 4.4 defines the multinormal distribution. It will be analysed in more detail in Chap. 5 along with most of its "companion" distributions that are useful in making multivariate statistical inferences.

The normal distribution plays a central role in statistics because it can be viewed as an approximation and limit of many other distributions. The basic justification relies on the central limit theorem presented in Sect.4.5. We present this central theorem in the framework of sampling theory. A useful extension of this theorem is also given: it is an approximate distribution to transformations of asymptotically normal variables. The increasing power of computers today makes it possible to consider alternative approximate sampling distributions. These are based on resampling techniques and are suitable for many general situations. Section 4.8 gives an introduction to the ideas behind bootstrap approximations.

### 4.1 Distribution and Density Function

Let $X=\left(X_{1}, X_{2}, \ldots, X_{p}\right)^{\top}$ be a random vector. The cumulative distribution function (cdf) of $X$ is defined by

$$
F(x)=\mathrm{P}(X \leq x)=\mathrm{P}\left(X_{1} \leq x_{1}, X_{2} \leq x_{2}, \ldots, X_{p} \leq x_{p}\right)
$$

For continuous $X$, a nonnegative probability density function (pdf) $f$ exists that

$$
\begin{equation*}
F(x)=\int_{-\infty}^{x} f(u) d u \tag{4.1}
\end{equation*}
$$

Note that

$$
\int_{-\infty}^{\infty} f(u) d u=1
$$

Most of the integrals appearing below are multidimensional. For instance, $\int_{-\infty}^{x} f(u) d u$ means $\int_{-\infty}^{x_{p}} \ldots \int_{-\infty}^{x_{1}} f\left(u_{1}, \ldots, u_{p}\right) d u_{1} \ldots d u_{p}$. Note also that the cdf $F$ is differentiable with

$$
f(x)=\frac{\partial^{p} F(x)}{\partial x_{1} \cdots \partial x_{p}} .
$$

For discrete $X$, the values of this random variable are concentrated on a countable or finite set of points $\left\{c_{j}\right\}_{j \in J}$, the probability of events of the form $\{X \in D\}$ can then be computed as

$$
\mathrm{P}(X \in D)=\sum_{\left\{j: c_{j} \in D\right\}} \mathrm{P}\left(X=c_{j}\right) .
$$

If we partition $X$ as $X=\left(X_{1}, X_{2}\right)^{\top}$ with $X_{1} \in \mathbb{R}^{k}$ and $X_{2} \in \mathbb{R}^{p-k}$, then the function

$$
\begin{equation*}
F_{X_{1}}\left(x_{1}\right)=\mathrm{P}\left(X_{1} \leq x_{1}\right)=F\left(x_{11}, \ldots, x_{1 k}, \infty, \ldots, \infty\right) \tag{4.2}
\end{equation*}
$$

is called the marginal $c d f . F=F(x)$ is called the joint cdf. For continuous $X$ the marginal pdf can be computed from the joint density by "integrating out" the variable not of interest.

$$
\begin{equation*}
f_{X_{1}}\left(x_{1}\right)=\int_{-\infty}^{\infty} f\left(x_{1}, x_{2}\right) d x_{2} \tag{4.3}
\end{equation*}
$$

The conditional pdf of $X_{2}$ given $X_{1}=x_{1}$ is given as

$$
\begin{equation*}
f\left(x_{2} \mid x_{1}\right)=\frac{f\left(x_{1}, x_{2}\right)}{f_{X_{1}}\left(x_{1}\right)} \tag{4.4}
\end{equation*}
$$

Example 4.1 Consider the pdf

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

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

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

The marginal densities are

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

The conditional densities are therefore

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

Note that these conditional pdf's are nonlinear in $x_{1}$ and $x_{2}$ although the joint pdf has a simple (linear) structure.

Independence of two random variables is defined as follows.
Definition 4.1 $X_{1}$ and $X_{2}$ are independent iff $f(x)=f\left(x_{1}, x_{2}\right)=$ $f_{X_{1}}\left(x_{1}\right) f_{X_{2}}\left(x_{2}\right)$.

That is, $X_{1}$ and $X_{2}$ are independent if the conditional pdf's are equal to the marginal densities, ie. $f\left(x_{1} \mid x_{2}\right)=f_{X_{1}}\left(x_{1}\right)$ and $f\left(x_{2} \mid x_{1}\right)=f_{X_{2}}\left(x_{2}\right)$. Independence can be interpreted as follows: knowing $X_{2}=x_{2}$ does not change the probability assessments on $X_{1}$, and conversely.


Different joint pdf's may have the same marginal pdf's.

Example 4.2 Consider the pdf's

$$
f\left(x_{1}, x_{2}\right)=1, \quad 0<x_{1}, x_{2}<1,
$$

and

$$
f\left(x_{1}, x_{2}\right)=1+\alpha\left(2 x_{1}-1\right)\left(2 x_{2}-1\right), \quad 0<x_{1}, x_{2}<1, \quad-1 \leq \alpha \leq 1 .
$$

We compute in both cases the marginal pdf's as

$$
f_{X_{1}}\left(x_{1}\right)=1, \quad f_{X_{2}}\left(x_{2}\right)=1
$$

Indeed

$$
\int_{0}^{1} 1+\alpha\left(2 x_{1}-1\right)\left(2 x_{2}-1\right) d x_{2}=1+\alpha\left(2 x_{1}-1\right)\left[x_{2}^{2}-x_{2}\right]_{0}^{1}=1 .
$$

Hence we obtain identical marginals from different joint distributions.
Let us study the concept of independence using the bank notes example. Consider the variables $X_{4}$ (lower inner frame) and $X_{5}$ (upper inner frame). From Chap. 3, we already know that they have significant correlation, so they are almost surely not independent. Kernel estimates of the marginal densities, $\hat{f}_{X_{4}}$ and $\hat{f}_{X_{5}}$, are given in Fig. 4.1. In Fig. 4.2 (left) we show the product of these two densities. The kernel density technique was presented in Sect. 1.3. If $X_{4}$ and $X_{5}$ are independent, this product $\hat{f}_{X_{4}} \cdot \hat{f}_{X_{5}}$ should be roughly equal to $\hat{f}\left(x_{4}, x_{5}\right)$, the estimate of the joint density of ( $X_{4}, X_{5}$ ). Comparing the two graphs in Fig. 4.2 reveals that the two densities are different. The two variables $X_{4}$ and $X_{5}$ are therefore not independent.

An elegant concept of connecting marginals with joint cdfs is given by copulae. Copulae are important in Value-at-Risk calculations and are an essential tool in quantitative finance (Härdle, Hautsch, \& Overbeck, 2009).

For simplicity of presentation we concentrate on the $p=2$ dimensional case. A two-dimensional copula is a function $C:[0,1]^{2} \rightarrow[0,1]$ with the following properties:

- For every $u \in[0,1]: C(0, u)=C(u, 0)=0$.
- For every $u \in[0,1]: C(u, 1)=u$ and $C(1, u)=u$.
- 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 $u_{2} \leq v_{2}$ :

$$
C\left(v_{1}, v_{2}\right)-C\left(v_{1}, u_{2}\right)-C\left(u_{1}, v_{2}\right)+C\left(u_{1}, u_{2}\right) \geq 0 .
$$

The usage of the name "copula" for the function $C$ is explained by the following theorem.


Fig. 4.1 Univariate estimates of the density of $X_{4}$ (left) and $X_{5}$ (right) of the bank notes 0 MVAdenbank2


Fig. 4.2 The product of univariate density estimates (left) and the joint density estimate (right) for $X_{4}$ (left) and $X_{5}$ of the bank notes $\mathbf{Q}$ MVAdenbank3

Theorem 4.1 (Sklar's Theorem) Let $F$ be a joint distribution function with marginal distribution functions $F_{X_{1}}$ and $F_{X_{2}}$. Then a copula $C$ exists with

$$
\begin{equation*}
F\left(x_{1}, x_{2}\right)=C\left\{F_{X_{1}}\left(x_{1}\right), F_{X_{2}}\left(x_{2}\right)\right\} \tag{4.5}
\end{equation*}
$$

for every $x_{1}, x_{2} \in \mathbb{R}$. If $F_{X_{1}}$ and $F_{X_{2}}$ are continuous, then $C$ is unique. On the other hand, if $C$ is a copula and $F_{X_{1}}$ and $F_{X_{2}}$ are distribution functions, then the function $F$ defined by (4.5) is a joint distribution function with marginals $F_{X_{1}}$ and $F_{X_{2}}$.

With Sklar's Theorem, the use of the name "copula" becomes obvious. It was chosen to describe "a function that links a multidimensional distribution to its onedimensional margins" and appeared in the mathematical literature for the first time in Sklar (1959).

Example 4.3 The structure of independence implies that the product of the distribution functions $F_{X_{1}}$ and $F_{X_{2}}$ equals their joint distribution function $F$,

$$
\begin{equation*}
F\left(x_{1}, x_{2}\right)=F_{X_{1}}\left(x_{1}\right) \cdot F_{X_{2}}\left(x_{2}\right) \tag{4.6}
\end{equation*}
$$

Thus, we obtain the independence copula $C=\Pi$ from

$$
\Pi\left(u_{1}, \ldots, u_{n}\right)=\prod_{i=1}^{n} u_{i}
$$

Theorem 4.2 Let $X_{1}$ and $X_{2}$ be random variables with continuous distribution functions $F_{X_{1}}$ and $F_{X_{2}}$ and the joint distribution function $F$. Then $X_{1}$ and $X_{2}$ are independent if and only if $C_{X_{1}, X_{2}}=\Pi$.

Proof From Sklar's Theorem we know that there exists an unique copula $C$ with

$$
\begin{equation*}
\mathrm{P}\left(X_{1} \leq x_{1}, X_{2} \leq x_{2}\right)=F\left(x_{1}, x_{2}\right)=C\left\{F_{X_{1}}\left(x_{1}\right), F_{X_{2}}\left(x_{2}\right)\right\} \tag{4.7}
\end{equation*}
$$

Independence can be seen using (4.5) for the joint distribution function $F$ and the definition of $\Pi$,

$$
\begin{equation*}
F\left(x_{1}, x_{2}\right)=C\left\{F_{X_{1}}\left(x_{1}\right), F_{X_{2}}\left(x_{2}\right)\right\}=F_{X_{1}}\left(x_{1}\right) F_{X_{2}}\left(x_{2}\right) . \tag{4.8}
\end{equation*}
$$

Example 4.4 The Gumbel-Hougaard family of copulae (Nelsen, 1999) is given by the function

$$
\begin{equation*}
C_{\theta}(u, v)=\exp \left[-\left\{(-\log u)^{\theta}+(-\log v)^{\theta}\right\}^{1 / \theta}\right] . \tag{4.9}
\end{equation*}
$$

The parameter $\theta$ may take all values in the interval $[1, \infty)$. The Gumbel-Hougaard copulae are suited to describe bivariate extreme value distributions.

For $\theta=1$, the expression (4.9) reduces to the product copula, i.e. $C_{1}(u, v)=$ $\Pi(u, v)=u v$. For $\theta \rightarrow \infty$ one finds for the Gumbel-Hougaard copula:

$$
C_{\theta}(u, v) \longrightarrow \min (u, v)=M(u, v)
$$

where the function $M$ is also a copula such that $C(u, v) \leq M(u, v)$ for arbitrary copula $C$. The copula $M$ is called the Fréchet-Hoeffding upper bound.

Similarly, we obtain the Fréchet-Hoeffding lower bound $W(u, v)=\max (u+$ $v-1,0)$ which satisfies $W(u, v) \leq C(u, v)$ for any other copula $C$.


## Summary

$\hookrightarrow$ The cumulative distribution function (cdf) is defined as $F(x)=$ $\mathrm{P}(X<x)$.
$\hookrightarrow$ If a probability density function (pdf) $f$ exists then $F(x)=$ $\int_{-\infty}^{x} f(u) d u$.
$\hookrightarrow$ The pdf integrates to one, i.e. $\int_{-\infty}^{\infty} f(x) d x=1$.
$\hookrightarrow$ Let $X=\left(X_{1}, X_{2}\right)^{\top}$ be partitioned into sub-vectors $X_{1}$ and $X_{2}$ with joint cdf $F$. Then $F_{X_{1}}\left(x_{1}\right)=\mathrm{P}\left(X_{1} \leq x_{1}\right)$ is the marginal cdf of $X_{1}$. The marginal pdf of $X_{1}$ is obtained by $f_{X_{1}}\left(x_{1}\right)=\int_{-\infty}^{\infty} f\left(x_{1}, x_{2}\right) d x_{2}$. Different joint pdf's may have the same marginal pdf's.
$\hookrightarrow$ The conditional pdf of $X_{2}$ given $X_{1}=x_{1}$ is defined as $f\left(x_{2} \mid\right.$ $\left.x_{1}\right)=\frac{f\left(x_{1}, x_{2}\right)}{f_{X_{1}}\left(x_{1}\right)}$.
$\hookrightarrow$ Two random variables $X_{1}$ and $X_{2}$ are called independent iff $f\left(x_{1}, x_{2}\right)=f_{X_{1}}\left(x_{1}\right) f_{X_{2}}\left(x_{2}\right)$. This is equivalent to $f\left(x_{2} \mid x_{1}\right)=$ $f_{X_{2}}\left(x_{2}\right)$.
$\hookrightarrow$ Different joint pdf's may have identical marginal pdf's.
$\hookrightarrow$ Copula is a function which connects marginals to form joint cdfs.

### 4.2 Moments and Characteristic Functions

## Moments: Expectation and Covariance Matrix

If $X$ is a random vector with density $f(x)$ then the expectation of $X$ is

$$
\mathrm{E} X=\left(\begin{array}{c}
\mathrm{E} X_{1}  \tag{4.10}\\
\vdots \\
\mathrm{E} X_{p}
\end{array}\right)=\int x f(x) d x=\left(\begin{array}{c}
\int x_{1} f(x) d x \\
\vdots \\
\int x_{p} f(x) d x
\end{array}\right)=\mu
$$

Accordingly, the expectation of a matrix of random elements has to be understood component by component. The operation of forming expectations is linear:

$$
\begin{equation*}
\mathrm{E}(\alpha X+\beta Y)=\alpha \mathrm{E} X+\beta \mathrm{E} Y \tag{4.11}
\end{equation*}
$$

If $\mathcal{A}(q \times p)$ is a matrix of real numbers, we have:

$$
\begin{equation*}
\mathrm{E}(\mathcal{A} X)=\mathcal{A} \mathrm{E} X \tag{4.12}
\end{equation*}
$$

When $X$ and $Y$ are independent,

$$
\begin{equation*}
\mathrm{E}\left(X Y^{\top}\right)=\mathrm{E} X \mathrm{E} Y^{\top} \tag{4.13}
\end{equation*}
$$

The matrix

$$
\begin{equation*}
\operatorname{Var}(X)=\Sigma=\mathrm{E}(X-\mu)(X-\mu)^{\top} \tag{4.14}
\end{equation*}
$$

is the (theoretical) covariance matrix. We write for a vector $X$ with mean vector $\mu$ and covariance matrix $\Sigma$,

$$
\begin{equation*}
X \sim(\mu, \Sigma) \tag{4.15}
\end{equation*}
$$

The $(p \times q)$ matrix

$$
\begin{equation*}
\Sigma_{X Y}=\operatorname{Cov}(X, Y)=\mathrm{E}(X-\mu)(Y-v)^{\top} \tag{4.16}
\end{equation*}
$$

is the covariance matrix of $X \sim\left(\mu, \Sigma_{X X}\right)$ and $Y \sim\left(v, \Sigma_{Y Y}\right)$. Note that $\Sigma_{X Y}=\Sigma_{Y X}^{\top}$ and that $Z=\binom{X}{Y}$ has covariance $\Sigma_{Z Z}=\binom{\Sigma_{X X} \Sigma_{X Y}}{\Sigma_{Y X} \Sigma_{Y Y}}$. From

$$
\begin{equation*}
\operatorname{Cov}(X, Y)=\mathrm{E}\left(X Y^{\top}\right)-\mu \nu^{\top}=\mathrm{E}\left(X Y^{\top}\right)-\mathrm{E} X \mathrm{E} Y^{\top} \tag{4.17}
\end{equation*}
$$

it follows that $\operatorname{Cov}(X, Y)=0$ in the case where $X$ and $Y$ are independent. We often say that $\mu=\mathrm{E}(X)$ is the first order moment of $X$ and that $\mathrm{E}\left(X X^{\top}\right)$ provides the second order moments of $X$ :

$$
\begin{equation*}
\mathrm{E}\left(X X^{\top}\right)=\left\{\mathrm{E}\left(X_{i} X_{j}\right)\right\}, \text { for } i=1, \ldots, p \text { and } j=1, \ldots, p \tag{4.18}
\end{equation*}
$$

## Properties of the Covariance Matrix $\Sigma=\operatorname{Var}(X)$

$$
\begin{align*}
& \Sigma=\left(\sigma_{X_{i} X_{j}}\right), \quad \sigma_{X_{i} X_{j}}=\operatorname{Cov}\left(X_{i}, X_{j}\right), \quad \sigma_{X_{i} X_{i}}=\operatorname{Var}\left(X_{i}\right)  \tag{4.19}\\
& \Sigma=\mathrm{E}\left(X X^{\top}\right)-\mu \mu^{\top}  \tag{4.20}\\
& \Sigma \geq 0 \tag{4.21}
\end{align*}
$$

## Properties of Variances and Covariances

$$
\begin{align*}
& \operatorname{Var}\left(a^{\top} X\right)=a^{\top} \operatorname{Var}(X) a=\sum_{i, j} a_{i} a_{j} \sigma_{X_{i} X_{j}}  \tag{4.22}\\
& \operatorname{Var}(\mathcal{A} X+b)=\mathcal{A} \operatorname{Var}(X) \mathcal{A}^{\top}  \tag{4.23}\\
& \operatorname{Cov}(X+Y, Z)=\operatorname{Cov}(X, Z)+\operatorname{Cov}(Y, Z)  \tag{4.24}\\
& \operatorname{Var}(X+Y)=\operatorname{Var}(X)+\operatorname{Cov}(X, Y)+\operatorname{Cov}(Y, X)+\operatorname{Var}(Y)  \tag{4.25}\\
& \operatorname{Cov}(\mathcal{A} X, \mathcal{B} Y)=\mathcal{A} \operatorname{Cov}(X, Y) \mathcal{B}^{\top} . \tag{4.26}
\end{align*}
$$

Let us compute these quantities for a specific joint density.
Example 4.5 Consider the pdf of Example 4.1. The mean vector $\mu=\binom{\mu_{1}}{\mu_{2}}$ is

$$
\begin{aligned}
\mu_{1} & =\iint x_{1} f\left(x_{1}, x_{2}\right) d x_{1} d x_{2}=\int_{0}^{1} \int_{0}^{1} x_{1}\left(\frac{1}{2} x_{1}+\frac{3}{2} x_{2}\right) d x_{1} d x_{2} \\
& =\int_{0}^{1} x_{1}\left(\frac{1}{2} x_{1}+\frac{3}{4}\right) d x_{1}=\frac{1}{2}\left[\frac{x_{1}^{3}}{3}\right]_{0}^{1}+\frac{3}{4}\left[\frac{x_{1}^{2}}{2}\right]_{0}^{1} \\
& =\frac{1}{6}+\frac{3}{8}=\frac{4+9}{24}=\frac{13}{24}, \\
\mu_{2} & =\iint x_{2} f\left(x_{1}, x_{2}\right) d x_{1} d x_{2}=\int_{0}^{1} \int_{0}^{1} x_{2}\left(\frac{1}{2} x_{1}+\frac{3}{2} x_{2}\right) d x_{1} d x_{2} \\
& =\int_{0}^{1} x_{2}\left(\frac{1}{4}+\frac{3}{2} x_{2}\right) d x_{2}=\frac{1}{4}\left[\frac{x_{2}^{2}}{2}\right]_{0}^{1}+\frac{3}{2}\left[\frac{x_{2}^{3}}{3}\right]_{0}^{1} \\
& =\frac{1}{8}+\frac{1}{2}=\frac{1+4}{8}=\frac{5}{8} .
\end{aligned}
$$

The elements of the covariance matrix are

$$
\begin{aligned}
& \sigma_{X_{1} X_{1}}=\mathrm{E} X_{1}^{2}-\mu_{1}^{2} \text { with } \\
& \mathrm{E} X_{1}^{2}=\int_{0}^{1} \int_{0}^{1} x_{1}^{2}\left(\frac{1}{2} x_{1}+\frac{3}{2} x_{2}\right) d x_{1} d x_{2}=\frac{1}{2}\left[\frac{x_{1}^{4}}{4}\right]_{0}^{1}+\frac{3}{4}\left[\frac{x_{1}^{3}}{3}\right]_{0}^{1}=\frac{3}{8} \\
& \sigma_{X_{2} X_{2}}=\mathrm{E} X_{2}^{2}-\mu_{2}^{2} \text { with } \\
& \mathrm{E} X_{2}^{2}=\int_{0}^{1} \int_{0}^{1} x_{2}^{2}\left(\frac{1}{2} x_{1}+\frac{3}{2} x_{2}\right) d x_{1} d x_{2}=\frac{1}{4}\left[\frac{x_{2}^{3}}{3}\right]_{0}^{1}+\frac{3}{2}\left[\frac{x_{2}^{4}}{4}\right]_{0}^{1}=\frac{11}{24} \\
& \sigma_{X_{1} X_{2}}=\mathrm{E}\left(X_{1} X_{2}\right)-\mu_{1} \mu_{2} \quad \text { with }
\end{aligned}
$$

$$
\begin{aligned}
\mathrm{E}\left(X_{1} X_{2}\right) & =\int_{0}^{1} \int_{0}^{1} x_{1} x_{2}\left(\frac{1}{2} x_{1}+\frac{3}{2} x_{2}\right) d x_{1} d x_{2}=\int_{0}^{1}\left(\frac{1}{6} x_{2}+\frac{3}{4} x_{2}^{2}\right) d x_{2} \\
& =\frac{1}{6}\left[\frac{x_{2}^{2}}{2}\right]_{0}^{1}+\frac{3}{4}\left[\frac{x_{2}^{3}}{3}\right]_{0}^{1}=\frac{1}{3}
\end{aligned}
$$

Hence the covariance matrix is

$$
\Sigma=\left(\begin{array}{ll}
0.0815 & 0.0052 \\
0.0052 & 0.0677
\end{array}\right)
$$

## Conditional Expectations

The conditional expectations are

$$
\begin{equation*}
\mathrm{E}\left(X_{2} \mid x_{1}\right)=\int x_{2} f\left(x_{2} \mid x_{1}\right) d x_{2} \quad \text { and } \quad \mathrm{E}\left(X_{1} \mid x_{2}\right)=\int x_{1} f\left(x_{1} \mid x_{2}\right) d x_{1} \tag{4.27}
\end{equation*}
$$

$\mathrm{E}\left(X_{2} \mid x_{1}\right)$ represents the location parameter of the conditional pdf of $X_{2}$ given that $X_{1}=x_{1}$. In the same way, we can define $\operatorname{Var}\left(X_{2} \mid X_{1}=x_{1}\right)$ as a measure of the dispersion of $X_{2}$ given that $X_{1}=x_{1}$. We have from (4.20) that

$$
\operatorname{Var}\left(X_{2} \mid X_{1}=x_{1}\right)=\mathrm{E}\left(X_{2} X_{2}^{\top} \mid X_{1}=x_{1}\right)-\mathrm{E}\left(X_{2} \mid X_{1}=x_{1}\right) \mathrm{E}\left(X_{2}^{\top} \mid X_{1}=x_{1}\right)
$$

Using the conditional covariance matrix, the conditional correlations may be defined as:

$$
\rho_{X_{2} X_{3} \mid X_{1}=x_{1}}=\frac{\operatorname{Cov}\left(X_{2}, X_{3} \mid X_{1}=x_{1}\right)}{\sqrt{\operatorname{Var}\left(X_{2} \mid X_{1}=x_{1}\right) \operatorname{Var}\left(X_{3} \mid X_{1}=x_{1}\right)}}
$$

These conditional correlations are known as partial correlations between $X_{2}$ and $X_{3}$, conditioned on $X_{1}$ being equal to $x_{1}$.
Example 4.6 Consider the following pdf

$$
f\left(x_{1}, x_{2}, x_{3}\right)=\frac{2}{3}\left(x_{1}+x_{2}+x_{3}\right) \text { where } 0<x_{1}, x_{2}, x_{3}<1 .
$$

Note that the pdf is symmetric in $x_{1}, x_{2}$ and $x_{3}$ which facilitates the computations. For instance,

$$
\begin{array}{ll}
f\left(x_{1}, x_{2}\right)=\frac{2}{3}\left(x_{1}+x_{2}+\frac{1}{2}\right) & 0<x_{1}, x_{2}<1 \\
f\left(x_{1}\right)=\frac{2}{3}\left(x_{1}+1\right) & 0<x_{1}<1
\end{array}
$$

and the other marginals are similar. We also have

$$
\begin{gathered}
f\left(x_{1}, x_{2} \mid x_{3}\right)=\frac{x_{1}+x_{2}+x_{3}}{x_{3}+1}, 0<x_{1}, x_{2}<1 \\
f\left(x_{1} \mid x_{3}\right)=\frac{x_{1}+x_{3}+\frac{1}{2}}{x_{3}+1}, 0<x_{1}<1 .
\end{gathered}
$$

It is easy to compute the following moments:
$\mathrm{E}\left(X_{i}\right)=\frac{5}{9} ; \mathrm{E}\left(X_{i}^{2}\right)=\frac{7}{18} ; \mathrm{E}\left(X_{i} X_{j}\right)=\frac{11}{36} \quad(i \neq j$ and $i, j=1,2,3)$
$\mathrm{E}\left(X_{1} \mid X_{3}=x_{3}\right)=\mathrm{E}\left(X_{2} \mid X_{3}=x_{3}\right)=\frac{1}{12}\left(\frac{6 x_{3}+7}{x_{3}+1}\right) ;$
$\mathrm{E}\left(X_{1}^{2} \mid X_{3}=x_{3}\right)=\mathrm{E}\left(X_{2}^{2} \mid X_{3}=x_{3}\right)=\frac{1}{12}\left(\frac{4 x_{3}+5}{x_{3}+1}\right)$
and
$\mathrm{E}\left(X_{1} X_{2} \mid X_{3}=x_{3}\right)=\frac{1}{12}\left(\frac{3 x_{3}+4}{x_{3}+1}\right)$.
Note that the conditional means of $X_{1}$ and of $X_{2}$, given $X_{3}=x_{3}$, are not linear in $x_{3}$. From these moments we obtain:

$$
\Sigma=\left(\begin{array}{rrr}
\frac{13}{162} & -\frac{1}{324} & -\frac{1}{324} \\
-\frac{1}{324} & \frac{13}{162} & -\frac{1}{324} \\
-\frac{1}{324} & -\frac{1}{324} & \frac{13}{162}
\end{array}\right) \text { in particular } \rho_{X_{1} X_{2}}=-\frac{1}{26} \approx-0.0385 .
$$

The conditional covariance matrix of $X_{1}$ and $X_{2}$, given $X_{3}=x_{3}$ is

$$
\operatorname{Var}\left(\left.\binom{X_{1}}{X_{2}} \right\rvert\, X_{3}=x_{3}\right)=\left(\begin{array}{ll}
\frac{12 x_{3}^{2}+24 x_{3}+11}{144\left(x_{3}+1\right)^{2}} & \frac{-1}{144\left(x_{3}+1\right)^{2}} \\
\frac{-1}{144\left(x_{3}+1\right)^{2}} & \frac{12 x_{3}^{2}+24 x_{3}+11}{144\left(x_{3}+1\right)^{2}}
\end{array}\right) .
$$

In particular, the partial correlation between $X_{1}$ and $X_{2}$, given that $X_{3}$ is fixed at $x_{3}$, is given by $\rho_{X_{1} X_{2} \mid X_{3}=x_{3}}=-\frac{1}{12 x_{3}^{2}+24 x_{3}+11}$ which ranges from -0.0909 to -0.0213 when $x_{3}$ goes from 0 to 1 . Therefore, in this example, the partial correlation may be larger or smaller than the simple correlation, depending on the value of the condition $X_{3}=x_{3}$.

Example 4.7 Consider the following joint pdf

$$
f\left(x_{1}, x_{2}, x_{3}\right)=2 x_{2}\left(x_{1}+x_{3}\right) ; 0<x_{1}, x_{2}, x_{3}<1 .
$$

Note the symmetry of $x_{1}$ and $x_{3}$ in the pdf and that $X_{2}$ is independent of $\left(X_{1}, X_{3}\right)$. It immediately follows that

$$
\begin{gathered}
f\left(x_{1}, x_{3}\right)=\left(x_{1}+x_{3}\right) \quad 0<x_{1}, x_{3}<1 \\
f\left(x_{1}\right)=x_{1}+\frac{1}{2} \\
f\left(x_{2}\right)=2 x_{2} ; \\
f\left(x_{3}\right)=x_{3}+\frac{1}{2}
\end{gathered}
$$

Simple computations lead to

$$
\mathrm{E}(X)=\left(\begin{array}{c}
\frac{7}{12} \\
\frac{2}{3} \\
\frac{7}{12}
\end{array}\right) \text { and } \Sigma=\left(\begin{array}{rrr}
\frac{11}{144} & 0 & -\frac{1}{144} \\
0 & \frac{1}{18} & 0 \\
-\frac{1}{144} & 0 & \frac{11}{144}
\end{array}\right) .
$$

Let us analyse the conditional distribution of $\left(X_{1}, X_{2}\right)$ given $X_{3}=x_{3}$. We have

$$
\begin{array}{rl}
f\left(x_{1}, x_{2} \mid x_{3}\right) & =\frac{4\left(x_{1}+x_{3}\right) x_{2}}{2 x_{3}+1} \\
0<x_{1}, x_{2}<1 \\
f\left(x_{1} \mid x_{3}\right) & =2\left(\frac{x_{1}+x_{3}}{2 x_{3}+1}\right) \\
f\left(x_{2} \mid x_{3}\right) & 0<x_{1}<1 \\
f\left(x_{2}\right)=2 x_{2} & 0<x_{2}<1
\end{array}
$$

so that again $X_{1}$ and $X_{2}$ are independent conditional on $X_{3}=x_{3}$. In this case

$$
\begin{aligned}
\mathrm{E}\left(\left.\binom{X_{1}}{X_{2}} \right\rvert\, X_{3}=x_{3}\right) & =\binom{\frac{1}{3}\left(\frac{2+3 x_{3}}{1+2 x_{3}}\right)}{\frac{2}{3}} \\
\operatorname{Var}\left(\left.\binom{X_{1}}{X_{2}} \right\rvert\, X_{3}=x_{3}\right) & =\left(\begin{array}{ll}
\frac{1}{18}\left(\frac{6 x_{3}^{2}+6 x_{3}+1}{\left(2 x_{3}+1\right)^{2}}\right) & 0 \\
0 & \frac{1}{18}
\end{array}\right) .
\end{aligned}
$$

## Properties of Conditional Expectations

Since $\mathrm{E}\left(X_{2} \mid X_{1}=x_{1}\right)$ is a function of $x_{1}$, say $h\left(x_{1}\right)$, we can define the random variable $h\left(X_{1}\right)=\mathrm{E}\left(X_{2} \mid X_{1}\right)$. The same can be done when defining the random variable $\operatorname{Var}\left(X_{2} \mid X_{1}\right)$. These two random variables share some interesting properties:

$$
\begin{align*}
\mathrm{E}\left(X_{2}\right) & =\mathrm{E}\left\{\mathrm{E}\left(X_{2} \mid X_{1}\right)\right\}  \tag{4.28}\\
\operatorname{Var}\left(X_{2}\right) & =\mathrm{E}\left\{\operatorname{Var}\left(X_{2} \mid X_{1}\right)\right\}+\operatorname{Var}\left\{\mathrm{E}\left(X_{2} \mid X_{1}\right)\right\} \tag{4.29}
\end{align*}
$$

Example 4.8 Consider the following pdf

$$
f\left(x_{1}, x_{2}\right)=2 e^{-\frac{x_{2}}{x_{1}}} ; 0<x_{1}<1, x_{2}>0 .
$$

It is easy to show that

$$
\begin{gathered}
f\left(x_{1}\right)=2 x_{1} \text { for } 0<x_{1}<1 ; \quad \mathrm{E}\left(X_{1}\right)=\frac{2}{3} \text { and } \operatorname{Var}\left(X_{1}\right)=\frac{1}{18} \\
f\left(x_{2} \mid x_{1}\right)=\frac{1}{x_{1}} e^{-\frac{x_{2}}{x_{1}}} \text { for } x_{2}>0 ; \quad \mathrm{E}\left(X_{2} \mid X_{1}\right)=X_{1} \text { and } \operatorname{Var}\left(X_{2} \mid X_{1}\right)=X_{1}^{2}
\end{gathered}
$$

Without explicitly computing $f\left(x_{2}\right)$, we can obtain:

$$
\begin{aligned}
\mathrm{E}\left(X_{2}\right) & =\mathrm{E}\left\{\mathrm{E}\left(X_{2} \mid X_{1}\right)\right\}=\mathrm{E}\left(X_{1}\right)=\frac{2}{3} \\
\operatorname{Var}\left(X_{2}\right) & =\mathrm{E}\left\{\operatorname{Var}\left(X_{2} \mid X_{1}\right)\right\}+\operatorname{Var}\left\{\mathrm{E}\left(X_{2} \mid X_{1}\right)\right\} \\
& =\mathrm{E}\left(X_{1}^{2}\right)+\operatorname{Var}\left(X_{1}\right)=\frac{2}{4}+\frac{1}{18}=\frac{10}{18} .
\end{aligned}
$$

The conditional expectation $\mathrm{E}\left(X_{2} \mid X_{1}\right)$ viewed as a function $h\left(X_{1}\right)$ of $X_{1}$ (known as the regression function of $X_{2}$ on $X_{1}$ ), can be interpreted as a conditional approximation of $X_{2}$ by a function of $X_{1}$. The error term of the approximation is then given by:

$$
U=X_{2}-\mathrm{E}\left(X_{2} \mid X_{1}\right)
$$

Theorem 4.3 Let $X_{1} \in \mathbb{R}^{k}$ and $X_{2} \in \mathbb{R}^{p-k}$ and $U=X_{2}-E\left(X_{2} \mid X_{1}\right)$. Then we have:

1. $E(U)=0$
2. $E\left(X_{2} \mid X_{1}\right)$ is the best approximation of $X_{2}$ by a function $h\left(X_{1}\right)$ of $X_{1}$ where $h$ : $\mathbb{R}^{k} \longrightarrow \mathbb{R}^{p-k}$. "Best" is the minimum mean squared error (MSE) sense, where

$$
\operatorname{MSE}(h)=E\left[\left\{X_{2}-h\left(X_{1}\right)\right\}^{\top}\left\{X_{2}-h\left(X_{1}\right)\right\}\right] .
$$

## Characteristic Functions

The characteristic function (cf) of a random vector $X \in \mathbb{R}^{p}$ (respectively its density $f(x))$ is defined as

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

where $\mathbf{i}$ is the complex unit: $\mathbf{i}^{2}=-1$. The cf has the following properties:

$$
\begin{equation*}
\varphi_{X}(0)=1 \text { and }\left|\varphi_{X}(t)\right| \leq 1 \tag{4.30}
\end{equation*}
$$

If $\varphi$ is absolutely integrable, i.e. the integral $\int_{-\infty}^{\infty}|\varphi(x)| d x$ exists and is finite, then

$$
\begin{equation*}
f(x)=\frac{1}{(2 \pi)^{p}} \int_{-\infty}^{\infty} e^{-\mathrm{i} t^{\top} x} \varphi_{X}(t) d t \tag{4.31}
\end{equation*}
$$

If $X=\left(X_{1}, X_{2}, \ldots, X_{p}\right)^{\top}$, then for $t=\left(t_{1}, t_{2}, \ldots, t_{p}\right)^{\top}$

$$
\begin{equation*}
\varphi_{X_{1}}\left(t_{1}\right)=\varphi_{X}\left(t_{1}, 0, \ldots, 0\right), \ldots, \varphi_{X_{p}}\left(t_{p}\right)=\varphi_{X}\left(0, \ldots, 0, t_{p}\right) \tag{4.32}
\end{equation*}
$$

If $X_{1}, \ldots, X_{p}$ are independent random variables, then for $t=\left(t_{1}, t_{2}, \ldots, t_{p}\right)^{\top}$

$$
\begin{equation*}
\varphi_{X}(t)=\varphi_{X_{1}}\left(t_{1}\right) \cdot \ldots \cdot \varphi_{X_{p}}\left(t_{p}\right) \tag{4.33}
\end{equation*}
$$

If $X_{1}, \ldots, X_{p}$ are independent random variables, then for $t \in \mathbb{R}$

$$
\begin{equation*}
\varphi_{X_{1}+\cdots+X_{p}}(t)=\varphi_{X_{1}}(t) \cdot \ldots \cdot \varphi_{X_{p}}(t) . \tag{4.34}
\end{equation*}
$$

The characteristic function can recover all the cross-product moments of any order: $\forall j_{k} \geq 0, k=1, \ldots, p$ and for $t=\left(t_{1}, \ldots, t_{p}\right)^{\top}$ we have

$$
\begin{equation*}
\mathrm{E}\left(X_{1}^{j_{1}} \cdot \ldots \cdot X_{p}^{j_{p}}\right)=\frac{1}{\mathbf{i}^{j_{1}+\cdots+j_{p}}}\left[\frac{\partial \varphi_{X}(t)}{\partial t_{1}^{j_{1}} \ldots \partial t_{p}^{j_{p}}}\right]_{t=0} \tag{4.35}
\end{equation*}
$$

Example 4.9 The cf of the density in Example 4.5 is given by

$$
\begin{aligned}
\varphi_{X}(t) & =\int_{0}^{1} \int_{0}^{1} e^{\mathbf{i t}^{\top} x} f(x) d x \\
& =\int_{0}^{1} \int_{0}^{1}\left\{\cos \left(t_{1} x_{1}+t_{2} x_{2}\right)+\mathbf{i} \sin \left(t_{1} x_{1}+t_{2} x_{2}\right)\right\}\left(\frac{1}{2} x_{1}+\frac{3}{2} x_{2}\right) d x_{1} d x_{2},
\end{aligned}
$$

$$
\begin{aligned}
= & \frac{0.5 e^{\mathbf{i} t_{1}}\left(3 \mathbf{i} t_{1}-3 \mathbf{i} e^{\mathbf{i} t_{2}} t_{1}+\mathbf{i} t_{2}-\mathbf{i} e^{\mathbf{i} t_{2}} t_{2}+t_{1} t_{2}-4 e^{\mathbf{i} t_{2}} t_{1} t_{2}\right)}{t_{1}^{2} t_{2}^{2}} \\
& -\frac{0.5\left(3 \mathbf{i} t_{1}-3 \mathbf{i} e^{\mathbf{i} t_{2}} t_{1}+\mathbf{i} t_{2}-\mathbf{i} e^{\mathbf{i} t_{2}} t_{2}-3 e^{\mathbf{i} t_{2}} t_{1} t_{2}\right)}{t_{1}{ }^{2} t_{2}{ }^{2}} .
\end{aligned}
$$

Example 4.10 Suppose $X \in \mathbb{R}^{1}$ follows the density of the standard normal distribution

$$
f_{X}(x)=\frac{1}{\sqrt{2 \pi}} \exp \left(-\frac{x^{2}}{2}\right)
$$

(see Sect. 4.4) then the cf can be computed via

$$
\begin{aligned}
\varphi_{X}(t) & =\frac{1}{\sqrt{2 \pi}} \int_{-\infty}^{\infty} e^{\mathbf{i} t x} \exp \left(-\frac{x^{2}}{2}\right) d x \\
& =\frac{1}{\sqrt{2 \pi}} \int_{-\infty}^{\infty} \exp \left\{-\frac{1}{2}\left(x^{2}-2 \mathbf{i} t x+\mathbf{i}^{2} t^{2}\right)\right\} \exp \left\{\frac{1}{2} \mathbf{i}^{2} t^{2}\right\} d x \\
& =\exp \left(-\frac{t^{2}}{2}\right) \int_{-\infty}^{\infty} \frac{1}{\sqrt{2 \pi}} \exp \left\{-\frac{(x-\mathbf{i} t)^{2}}{2}\right\} d x \\
& =\exp \left(-\frac{t^{2}}{2}\right)
\end{aligned}
$$

since $\mathbf{i}^{2}=-1$ and $\int \frac{1}{\sqrt{2 \pi}} \exp \left\{-\frac{(x-\mathbf{i} t)^{2}}{2}\right\} d x=1$.
A variety of distributional characteristics can be computed from $\varphi_{X}(t)$. The standard normal distribution has a very simple cf, as was seen in Example 4.10. Deviations from normal covariance structures can be measured by the deviations from the cf (or characteristics of it). In Table 4.1 we give an overview of the cf's for a variety of distributions.

Theorem 4.4 (Cramer-Wold) The distribution of $X \in \mathbb{R}^{p}$ is completely determined by the set of all (one-dimensional) distributions of $t^{\top} X$ where $t \in \mathbb{R}^{p}$.

Table 4.1 Characteristic functions for some common distributions

|  | pdf | cf |
| :--- | :--- | :--- |
| Uniform | $f(x)=\boldsymbol{I}(x \in[a, b]) /(b-a)$ | $\varphi_{X}(t)=\left(e^{\mathbf{i} b t}-e^{\mathbf{i} a t}\right) /(b-a) \mathbf{i} t$ |
| $N_{1}\left(\mu, \sigma^{2}\right)$ | $f(x)=\left(2 \pi \sigma^{2}\right)^{-1 / 2} \exp \left\{-(x-\mu)^{2} / 2 \sigma^{2}\right\}$ | $\varphi_{X}(t)=e^{\mathbf{i} \mu t-\sigma^{2} t^{2} / 2}$ |
| $\chi^{2}(n)$ | $f(x)=\boldsymbol{I}(x>0) x^{n / 2-1} e^{-x / 2} /\left\{\Gamma(n / 2) 2^{n / 2}\right\}$ | $\varphi_{X}(t)=(1-2 \mathbf{i} t)^{-n / 2}$ |
| $N_{p}(\mu, \Sigma)$ | $f(x)=\|2 \pi \Sigma\|^{-1 / 2} \exp \left\{-(x-\mu)^{\top} \Sigma(x-\mu) / 2\right\}$ | $\varphi_{X}(t)=e^{\mathbf{i} t^{\top} \mu-t^{\top} \Sigma t / 2}$ |

This theorem says that we can determine the distribution of $X$ in $\mathbb{R}^{p}$ by specifying all of the one-dimensional distributions of the linear combinations

$$
\sum_{j=1}^{p} t_{j} X_{j}=t^{\top} X, \quad t=\left(t_{1}, t_{2}, \ldots, t_{p}\right)^{\top}
$$

## Cumulant Functions

Moments $m_{k}=\int x^{k} f(x) d x$ often help in describing distributional characteristics. The normal distribution in $d=1$ dimension is completely characterised by its standard normal density $f=\varphi$ and the moment parameters are $\mu=m_{1}$ and $\sigma^{2}=m_{2}-m_{1}^{2}$. Another helpful class of parameters are the cumulants or semiinvariants of a distribution. In order to simplify notation we concentrate here on the one-dimensional ( $d=1$ ) case.

For a given one-dimensional random variable $X$ with density $f$ and finite moments of order $k$ the characteristic function $\varphi_{X}(t)=\mathrm{E}\left(e^{\mathbf{i} t X}\right)$ has the derivative

$$
\frac{1}{\mathbf{i}^{j}}\left[\frac{\partial^{j} \log \left\{\varphi_{X}(t)\right\}}{\partial t^{j}}\right]_{t=0}=\kappa_{j}, \quad j=1, \ldots, k
$$

The values $\kappa_{j}$ are called cumulants or semi-invariants since $\kappa_{j}$ does not change (for $j>1$ ) under a shift transformation $X \mapsto X+a$. The cumulants are natural parameters for dimension reduction methods, in particular the Projection Pursuit method (see Sect. 20.2).

The relationship between the first $k$ moments $m_{1}, \ldots, m_{k}$ and the cumulants is given by

$$
\kappa_{k}=(-1)^{k-1}\left|\begin{array}{cccc}
m_{1} & 1 & \ldots & 0  \tag{4.36}\\
m_{2} & \binom{1}{0} m_{1} & \ldots & \\
\vdots & \vdots & \ddots & \vdots \\
m_{k}\binom{k-1}{0} & m_{k-1} & \cdots\binom{k-1}{k-2} m_{1}
\end{array}\right|
$$

Example 4.11 Suppose that $k=1$, then formula (4.36) above yields

$$
\kappa_{1}=m_{1}
$$

For $k=2$ we obtain

$$
\kappa_{2}=-\left|\begin{array}{cc}
m_{1} & 1 \\
m_{2}\binom{1}{0} m_{1}
\end{array}\right|=m_{2}-m_{1}^{2} .
$$

For $k=3$ we have to calculate

$$
\kappa_{3}=\left|\begin{array}{ccc}
m_{1} & 1 & 0 \\
m_{2} & m_{1} & 1 \\
m_{3} & m_{2} & 2 m_{1}
\end{array}\right|
$$

Calculating the determinant we have:

$$
\begin{align*}
\kappa_{3} & =m_{1}\left|\begin{array}{ll}
m_{1} & 1 \\
m_{2} & 2 m_{1}
\end{array}\right|-m_{2}\left|\begin{array}{cc}
1 & 0 \\
m_{2} & 2 m_{1}
\end{array}\right|+m_{3}\left|\begin{array}{cc}
1 & 0 \\
m_{1} & 1
\end{array}\right| \\
& =m_{1}\left(2 m_{1}^{2}-m_{2}\right)-m_{2}\left(2 m_{1}\right)+m_{3} \\
& =m_{3}-3 m_{1} m_{2}+2 m_{1}^{3} . \tag{4.37}
\end{align*}
$$

Similarly one calculates

$$
\begin{equation*}
\kappa_{4}=m_{4}-4 m_{3} m_{1}-3 m_{2}^{2}+12 m_{2} m_{1}^{2}-6 m_{1}^{4} . \tag{4.38}
\end{equation*}
$$

The same type of process is used to find the moments from the cumulants:

$$
\begin{align*}
& m_{1}=\kappa_{1} \\
& m_{2}=\kappa_{2}+\kappa_{1}^{2} \\
& m_{3}=\kappa_{3}+3 \kappa_{2} \kappa_{1}+\kappa_{1}^{3} \\
& m_{4}=\kappa_{4}+4 \kappa_{3} \kappa_{1}+3 \kappa_{2}^{2}+6 \kappa_{2} \kappa_{1}^{2}+\kappa_{1}^{4} \tag{4.39}
\end{align*}
$$

A very simple relationship can be observed between the semi-invariants and the central moments $\mu_{k}=\mathrm{E}(X-\mu)^{k}$, where $\mu=m_{1}$ as defined before. In fact, $\kappa_{2}=\mu_{2}, \kappa_{3}=\mu_{3}$ and $\kappa_{4}=\mu_{4}-3 \mu_{2}^{2}$.

Skewness $\gamma_{3}$ and kurtosis $\gamma_{4}$ are defined as:

$$
\begin{align*}
& \gamma_{3}=\mathrm{E}(X-\mu)^{3} / \sigma^{3} \\
& \gamma_{4}=\mathrm{E}(X-\mu)^{4} / \sigma^{4} \tag{4.40}
\end{align*}
$$

The skewness and kurtosis determine the shape of one-dimensional distributions. The skewness of a normal distribution is 0 and the kurtosis equals 3 . The relation of these parameters to the cumulants is given by:

$$
\begin{equation*}
\gamma_{3}=\frac{\kappa_{3}}{\kappa_{2}^{3 / 2}} \tag{4.41}
\end{equation*}
$$

From (4.39) and Example 4.11

$$
\begin{equation*}
\gamma_{4}=\frac{\kappa_{4}+3 \kappa_{2}^{2}+\kappa_{1}^{4}-m_{1}^{4}}{\sigma^{4}}=\frac{\kappa_{4}+3 \kappa_{2}^{2}}{\kappa_{2}^{2}}=\frac{\kappa_{4}}{\kappa_{2}^{2}}+3 \tag{4.42}
\end{equation*}
$$

These relations will be used later in Sect. 20.2 on Projection Pursuit to determine deviations from normality.
$\left.\begin{array}{|l|l|}\hline \hookrightarrow & \begin{array}{l}\text { The expectation of a random vector } X \text { is } \mu=\int x f(x) d x \text {, the } \\ \text { covariance matrix } \Sigma=\operatorname{Var}(X)=\mathrm{E}(X-\mu)(X-\mu)^{\top} .\end{array} \\ & X \sim(\mu, \Sigma) \text {. We denote }\end{array}\right\}$

### 4.3 Transformations

Suppose that $X$ has pdf $f_{X}(x)$. What is the pdf of $Y=3 X$ ? Or if $X=$ $\left(X_{1}, X_{2}, X_{3}\right)^{\top}$, what is the pdf of

$$
Y=\left(\begin{array}{c}
3 X_{1} \\
X_{1}-4 X_{2} \\
X_{3}
\end{array}\right) ?
$$

This is a special case of asking for the pdf of $Y$ when

$$
\begin{equation*}
X=u(Y) \tag{4.43}
\end{equation*}
$$

for a one-to-one transformation $u: \mathbb{R}^{p} \rightarrow \mathbb{R}^{p}$. Define the Jacobian of $u$ as

$$
\mathcal{J}=\left(\frac{\partial x_{i}}{\partial y_{j}}\right)=\left(\frac{\partial u_{i}(y)}{\partial y_{j}}\right)
$$

and let $\operatorname{abs}(|\mathcal{J}|)$ be the absolute value of the determinant of this Jacobian. The pdf of $Y$ is given by

$$
\begin{equation*}
f_{Y}(y)=\operatorname{abs}(|\mathcal{J}|) \cdot f_{X}\{u(y)\} \tag{4.44}
\end{equation*}
$$

Using this we can answer the introductory questions, namely

$$
\left(x_{1}, \ldots, x_{p}\right)^{\top}=u\left(y_{1}, \ldots, y_{p}\right)=\frac{1}{3}\left(y_{1}, \ldots, y_{p}\right)^{\top}
$$

with

$$
\mathcal{J}=\left(\begin{array}{ccc}
\frac{1}{3} & & 0 \\
& \ddots & \\
& & \\
0 & & \frac{1}{3}
\end{array}\right)
$$

and hence abs $(|\mathcal{J}|)=\left(\frac{1}{3}\right)^{p}$. So the pdf of $Y$ is $\frac{1}{3^{p}} f_{X}\left(\frac{y}{3}\right)$.
This introductory example is a special case of

$$
Y=\mathcal{A} X+b, \text { where } \mathcal{A} \text { is nonsingular. }
$$

The inverse transformation is

$$
X=\mathcal{A}^{-1}(Y-b) .
$$

Therefore

$$
\mathcal{J}=\mathcal{A}^{-1}
$$

and hence

$$
\begin{equation*}
f_{Y}(y)=\operatorname{abs}\left(|\mathcal{A}|^{-1}\right) f_{X}\left\{\mathcal{A}^{-1}(y-b)\right\} . \tag{4.45}
\end{equation*}
$$

Example 4.12 Consider $X=\left(X_{1}, X_{2}\right) \in \mathbb{R}^{2}$ with density $f_{X}(x)=f_{X}\left(x_{1}, x_{2}\right)$,

$$
\mathcal{A}=\left(\begin{array}{rr}
1 & 1 \\
1 & -1
\end{array}\right), \quad b=\binom{0}{0} .
$$

Then

$$
Y=\mathcal{A} X+b=\binom{X_{1}+X_{2}}{X_{1}-X_{2}}
$$

and

$$
|\mathcal{A}|=-2, \quad \operatorname{abs}\left(|\mathcal{A}|^{-1}\right)=\frac{1}{2}, \quad \mathcal{A}^{-1}=-\frac{1}{2}\left(\begin{array}{rr}
-1 & -1 \\
-1 & 1
\end{array}\right) .
$$

Hence

$$
\begin{align*}
f_{Y}(y) & =\operatorname{abs}\left(|\mathcal{A}|^{-1}\right) \cdot f_{X}\left(\mathcal{A}^{-1} y\right) \\
& =\frac{1}{2} f_{X}\left\{\frac{1}{2}\left(\begin{array}{rr}
1 & 1 \\
1 & -1
\end{array}\right)\binom{y_{1}}{y_{2}}\right\} \\
& =\frac{1}{2} f_{X}\left\{\frac{1}{2}\left(y_{1}+y_{2}\right), \frac{1}{2}\left(y_{1}-y_{2}\right)\right\} . \tag{4.46}
\end{align*}
$$

Example 4.13 Consider $X \in \mathbb{R}^{1}$ with density $f_{X}(x)$ and $Y=\exp (X)$. According to (4.43) $x=u(y)=\log (y)$ and hence the Jacobian is

$$
\mathcal{J}=\frac{d x}{d y}=\frac{1}{y}
$$

The pdf of $Y$ is therefore:

$$
f_{Y}(y)=\frac{1}{y} f_{X}\{\log (y)\}
$$



## Summary

$\hookrightarrow$ If $X$ has pdf $f_{X}(x)$, then a transformed random vector $Y$, i.e. $X=$ $u(Y)$, has pdf $f_{Y}(y)=\operatorname{abs}(|\mathcal{J}|) \cdot f_{X}\{u(y)\}$, where $\mathcal{J}$ denotes the Jacobian $\mathcal{J}=\left(\frac{\partial u\left(y_{i}\right)}{\partial y_{j}}\right)$.
$\hookrightarrow$ In the case of a linear relation $Y=\mathcal{A} X+b$ the pdf's of $X$ and $Y$ are related via $f_{Y}(y)=\operatorname{abs}\left(|\mathcal{A}|^{-1}\right) f_{X}\left\{\mathcal{A}^{-1}(y-b)\right\}$.

### 4.4 The Multinormal Distribution

The multinormal distribution with mean $\mu$ and covariance $\Sigma>0$ has the density

$$
\begin{equation*}
f(x)=|2 \pi \Sigma|^{-1 / 2} \exp \left\{-\frac{1}{2}(x-\mu)^{\top} \Sigma^{-1}(x-\mu)\right\} \tag{4.47}
\end{equation*}
$$

We write $X \sim N_{p}(\mu, \Sigma)$.
How is this multinormal distribution with mean $\mu$ and covariance $\Sigma$ related to the multivariate standard normal $N_{p}\left(0, \mathcal{I}_{p}\right)$ ? Through a linear transformation using the results of Sect. 4.3, as shown in the next theorem.

Theorem 4.5 Let $X \sim N_{p}(\mu, \Sigma)$ and $Y=\Sigma^{-1 / 2}(X-\mu)$ (Mahalanobis transformation). Then

$$
Y \sim N_{p}\left(0, \mathcal{I}_{p}\right),
$$

i.e. the elements $Y_{j} \in \mathbb{R}$ are independent, one-dimensional $N(0,1)$ variables.

Proof Note that $(X-\mu)^{\top} \Sigma^{-1}(X-\mu)=Y^{\top} Y$. Application of (4.45) gives $\mathcal{J}=$ $\Sigma^{1 / 2}$, hence

$$
\begin{equation*}
f_{Y}(y)=(2 \pi)^{-p / 2} \exp \left(-\frac{1}{2} y^{\top} y\right) \tag{4.48}
\end{equation*}
$$

which is by (4.47) the pdf of a $N_{p}\left(0, \mathcal{I}_{p}\right)$.

Note that the above Mahalanobis transformation yields in fact a random variable $Y=\left(Y_{1}, \ldots, Y_{p}\right)^{\top}$ composed of independent one-dimensional $Y_{j} \sim N_{1}(0,1)$ since

$$
\begin{aligned}
f_{Y}(y) & =\frac{1}{(2 \pi)^{p / 2}} \exp \left(-\frac{1}{2} y^{\top} y\right) \\
& =\prod_{j=1}^{p} \frac{1}{\sqrt{2 \pi}} \exp \left(-\frac{1}{2} y_{j}^{2}\right) \\
& =\prod_{j=1}^{p} f_{Y_{j}}\left(y_{j}\right) .
\end{aligned}
$$

Here each $f_{Y_{j}}(y)$ is a standard normal density $\frac{1}{\sqrt{2 \pi}} \exp \left(-\frac{y^{2}}{2}\right)$. From this it is clear that $\mathrm{E}(Y)=0$ and $\operatorname{Var}(Y)=\mathcal{I}_{p}$.

How can we create $N_{p}(\mu, \Sigma)$ variables on the basis of $N_{p}\left(0, \mathcal{I}_{p}\right)$ variables? We use the inverse linear transformation

$$
\begin{equation*}
X=\Sigma^{1 / 2} Y+\mu \tag{4.49}
\end{equation*}
$$

Using (4.11) and (4.23) we can also check that $\mathrm{E}(X)=\mu$ and $\operatorname{Var}(X)=\Sigma$. The following theorem is useful because it presents the distribution of a variable after it has been linearly transformed. The proof is left as an exercise.

Theorem 4.6 Let $X \sim N_{p}(\mu, \Sigma)$ and $\mathcal{A}(p \times p)$, $c \in \mathbb{R}^{p}$, where $\mathcal{A}$ is nonsingular.
Then $Y=\mathcal{A} X+c$ is again a $p$-variate Normal, i.e.

$$
\begin{equation*}
Y \sim N_{p}\left(\mathcal{A} \mu+c, \mathcal{A} \Sigma \mathcal{A}^{\top}\right) \tag{4.50}
\end{equation*}
$$

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

From (4.47) we see that the density of the $N_{p}(\mu, \Sigma)$ distribution is constant on ellipsoids of the form

$$
\begin{equation*}
(x-\mu)^{\top} \Sigma^{-1}(x-\mu)=d^{2} \tag{4.51}
\end{equation*}
$$

Example 4.14 Figure 4.3 shows the contour ellipses of a two-dimensional normal distribution. Note that these contour ellipses are the iso-distance curves (2.34) from the mean of this normal distribution corresponding to the metric $\Sigma^{-1}$.


Fig. 4.3 Scatterplot of a normal sample and contour ellipses for $\mu=\binom{3}{2}$ and $\Sigma=\left(\begin{array}{cc}1 & -1.5 \\ -1.5 & 4\end{array}\right)$ Q MVAcontnorm

According to Theorem 2.7 in Sect. 2.6 the half-lengths of the axes in the contour ellipsoid are $\sqrt{d^{2} \lambda_{i}}$ where $\lambda_{i}$ are the eigenvalues of $\Sigma$. If $\Sigma$ is a diagonal matrix, the rectangle circumscribing the contour ellipse has sides with length $2 d \sigma_{i}$ and is thus naturally proportional to the standard deviations of $X_{i}(i=1,2)$.

The distribution of the quadratic form in (4.51) is given in the next theorem.
Theorem 4.7 If $X \sim N_{p}(\mu, \Sigma)$, then the variable $U=(X-\mu)^{\top} \Sigma^{-1}(X-\mu)$ has a $\chi_{p}^{2}$ distribution.
Theorem 4.8 The characteristic function (cf) of a multinormal $N_{p}(\mu, \Sigma)$ is given by

$$
\begin{equation*}
\varphi_{X}(t)=\exp \left(i t^{\top} \mu-\frac{1}{2} t^{\top} \Sigma t\right) \tag{4.52}
\end{equation*}
$$

We can check Theorem 4.8 by transforming the cf back:

$$
\begin{aligned}
f(x) & =\frac{1}{(2 \pi)^{p}} \int \exp \left(-\mathbf{i} t^{\top} x+\mathbf{i} t^{\top} \mu-\frac{1}{2} t^{\top} \Sigma t\right) d t \\
& =\frac{1}{\left|2 \pi \Sigma^{-1}\right|^{1 / 2}|2 \pi \Sigma|^{1 / 2}} \int \exp \left[-\frac{1}{2}\left\{t^{\top} \Sigma t+2 \mathbf{i} t^{\top}(x-\mu)-(x-\mu)^{\top} \Sigma^{-1}(x-\mu)\right\}\right]
\end{aligned}
$$

$$
\begin{aligned}
& \times \exp \left[-\frac{1}{2}\left\{(x-\mu)^{\top} \Sigma^{-1}(x-\mu)\right\}\right] d t \\
= & \frac{1}{|2 \pi \Sigma|^{1 / 2}} \exp \left[-\frac{1}{2}\left\{(x-\mu)^{\top} \Sigma(x-\mu)\right\}\right]
\end{aligned}
$$

since

$$
\begin{aligned}
& \int \frac{1}{\left|2 \pi \Sigma^{-1}\right|^{1 / 2}} \exp \left[-\frac{1}{2}\left\{t^{\top} \Sigma t+2 \mathbf{i} t^{\top}(x-\mu)-(x-\mu)^{\top} \Sigma^{-1}(x-\mu)\right\}\right] d t \\
& \quad=\int \frac{1}{\left|2 \pi \Sigma^{-1}\right|^{1 / 2}} \exp \left[-\frac{1}{2}\left\{\left(t+\mathbf{i} \Sigma^{-1}(x-\mu)\right)^{\top} \Sigma\left(t+\mathbf{i} \Sigma^{-1}(x-\mu)\right)\right\}\right] d t \\
& \quad=1
\end{aligned}
$$

Note that if $Y \sim N_{p}\left(0, \mathcal{I}_{p}\right)$, then

$$
\begin{aligned}
\varphi_{Y}(t) & =\exp \left(-\frac{1}{2} t^{\top} \mathcal{I}_{p} t\right)=\exp \left(-\frac{1}{2} \sum_{i=1}^{p} t_{i}^{2}\right) \\
& =\varphi_{Y_{1}}\left(t_{1}\right) \cdot \ldots \cdot \varphi_{Y_{p}}\left(t_{p}\right)
\end{aligned}
$$

which is consistent with (4.33).

## Singular Normal Distribution

Suppose that we have $\operatorname{rank}(\Sigma)=k<p$, where $p$ is the dimension of $X$. We define the (singular) density of $X$ with the aid of the $G$-Inverse $\Sigma^{-}$of $\Sigma$,

$$
\begin{equation*}
f(x)=\frac{(2 \pi)^{-k / 2}}{\left(\lambda_{1} \cdots \lambda_{k}\right)^{1 / 2}} \exp \left\{-\frac{1}{2}(x-\mu)^{\top} \Sigma^{-}(x-\mu)\right\} \tag{4.53}
\end{equation*}
$$

where

1. $x$ lies on the hyperplane $\mathcal{N}^{\top}(x-\mu)=0$ with $\mathcal{N}(p \times(p-k)): \mathcal{N}^{\top} \Sigma=0$ and $\mathcal{N}^{\top} \mathcal{N}=\mathcal{I}_{k}$.
2. $\Sigma^{-}$is the $G$-Inverse of $\Sigma$, and $\lambda_{1}, \ldots, \lambda_{k}$ are the nonzero eigenvalues of $\Sigma$.

What is the connection to a multinormal with $k$-dimensions? If

$$
\begin{equation*}
Y \sim N_{k}\left(0, \Lambda_{1}\right) \text { and } \Lambda_{1}=\operatorname{diag}\left(\lambda_{1}, \ldots, \lambda_{k}\right) \tag{4.54}
\end{equation*}
$$

then an orthogonal matrix $\mathcal{B}(p \times k)$ with $\mathcal{B}^{\top} \mathcal{B}=\mathcal{I}_{k}$ exists that means $X=\mathcal{B} Y+\mu$ where $X$ has a singular pdf of the form (4.53).

## Gaussian Copula

In Examples 4.3 and 4.4 we have introduced copulae. Another important copula is the Gaussian or normal copula,

$$
\begin{equation*}
C_{\rho}(u, v)=\int_{-\infty}^{\Phi_{1}^{-1}(u)} \int_{-\infty}^{\Phi_{2}^{-1}(v)} f_{\rho}\left(x_{1}, x_{2}\right) d x_{2} d x_{1} \tag{4.55}
\end{equation*}
$$

see Embrechts, McNeil, and Straumann (1999). In (4.55), $f_{\rho}$ denotes the bivariate normal density function with correlation $\rho$ for $n=2$. The functions $\Phi_{1}$ and $\Phi_{2}$ in (4.55) refer to the corresponding one-dimensional standard normal cdfs of the margins.

In the case of vanishing correlation, $\rho=0$, the Gaussian copula becomes

$$
\begin{aligned}
C_{0}(u, v) & =\int_{-\infty}^{\Phi_{1}^{-1}(u)} f_{X_{1}}\left(x_{1}\right) d x_{1} \int_{-\infty}^{\Phi_{2}^{-1}(v)} f_{X_{2}}\left(x_{2}\right) d x_{2} \\
& =u v \\
& =\Pi(u, v)
\end{aligned}
$$

## Summary

$\hookrightarrow \quad$ The pdf of a $p$-dimensional multinormal $X \sim N_{p}(\mu, \Sigma)$ is

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

The contour curves of a multinormal are ellipsoids with halflengths proportional to $\sqrt{\lambda_{i}}$, where $\lambda_{i}$ denotes the eigenvalues of $\Sigma(i=1, \ldots, p)$.
$\hookrightarrow$ The Mahalanobis transformation transforms $X \sim N_{p}(\mu, \Sigma)$ to $Y=\Sigma^{-1 / 2}(X-\mu) \sim N_{p}\left(0, \mathcal{I}_{p}\right)$. Going in the other direction, one can create a $X \sim N_{p}(\mu, \Sigma)$ from $Y \sim N_{p}\left(0, \mathcal{I}_{p}\right)$ via $X=$ $\Sigma^{1 / 2} Y+\mu$.
$\hookrightarrow$ If the covariance matrix $\Sigma$ is singular (i.e. $\operatorname{rank}(\Sigma)<p$ ), then it defines a singular normal distribution.

| Summary (continued) |
| :---: |
| $\qquad \hookrightarrow$ The Gaussian copula is given by |
| $C_{\rho}(u, v)=\int_{-\infty}^{\Phi_{1}^{-1}(u)} \int_{-\infty}^{\Phi_{2}^{-1}(v)} f_{\rho}\left(x_{1}, x_{2}\right) d x_{2} d x_{1}$. |
| $\hookrightarrow$ The density of a singular normal distribution is given by |
| $\frac{(2 \pi)^{-k / 2}}{\left(\lambda_{1} \cdots \lambda_{k}\right)^{1 / 2}} \exp \left\{-\frac{1}{2}(x-\mu)^{\top} \Sigma^{-}(x-\mu)\right\}$. |

### 4.5 Sampling Distributions and Limit Theorems

In multivariate statistics, we observe the values of a multivariate random variable $X$ and obtain a sample $\left\{x_{i}\right\}_{i=1}^{n}$, as described in Chap. 3. Under random sampling, these observations are considered to be realisations of a sequence of i.i.d. random variables $X_{1}, \ldots, X_{n}$, where each $X_{i}$ is a $p$-variate random variable which replicates the parent or population random variable $X$. Some notational confusion is hard to avoid: $X_{i}$ is not the $i$ th component of $X$, but rather the $i$ th replicate of the $p$-variate random variable $X$ which provides the $i$ th observation $x_{i}$ of our sample.

For a given random sample $X_{1}, \ldots, X_{n}$, the idea of statistical inference is to analyse the properties of the population variable $X$. This is typically done by analysing some characteristic $\theta$ of its distribution, like the mean, covariance matrix, etc. Statistical inference in a multivariate setup is considered in more detail in Chaps. 6 and 7.

Inference can often be performed using some observable function of the sample $X_{1}, \ldots, X_{n}$, i.e. a statistics. Examples of such statistics were given in Chap. 3: the sample mean $\bar{x}$, the sample covariance matrix $\mathcal{S}$. To get an idea of the relationship between a statistics and the corresponding population characteristic, one has to derive the sampling distribution of the statistic. The next example gives some insight into the relation of $(\bar{x}, S)$ to $(\mu, \Sigma)$.

Example 4.15 Consider an iid sample of $n$ random vectors $X_{i} \in \mathbb{R}^{p}$ where $\mathrm{E}\left(X_{i}\right)=\mu$ and $\operatorname{Var}\left(X_{i}\right)=\Sigma$. The sample mean $\bar{x}$ and the covariance matrix $\mathcal{S}$ have already been defined in Sect. 3.3. It is easy to prove the following results:

$$
\begin{aligned}
& \mathrm{E}(\bar{x})=n^{-1} \sum_{i=1}^{n} \mathrm{E}\left(X_{i}\right)=\mu \\
& \operatorname{Var}(\bar{x})=n^{-2} \sum_{i=1}^{n} \operatorname{Var}\left(X_{i}\right)=n^{-1} \Sigma=\mathrm{E}\left(\bar{x} \bar{x}^{\top}\right)-\mu \mu^{\top}
\end{aligned}
$$

$$
\begin{aligned}
\mathrm{E}(\mathcal{S}) & =n^{-1} \mathrm{E}\left\{\sum_{i=1}^{n}\left(X_{i}-\bar{x}\right)\left(X_{i}-\bar{x}\right)^{\top}\right\} \\
& =n^{-1} \mathrm{E}\left\{\sum_{i=1}^{n} X_{i} X_{i}^{\top}-n \bar{x} \bar{x}^{\top}\right\} \\
& =n^{-1}\left\{n\left(\Sigma+\mu \mu^{\top}\right)-n\left(n^{-1} \Sigma+\mu \mu^{\top}\right)\right\} \\
& =\frac{n-1}{n} \Sigma .
\end{aligned}
$$

This shows in particular that $\mathcal{S}$ is a biased estimator of $\Sigma$. By contrast, $\mathcal{S}_{u}=\frac{n}{n-1} \mathcal{S}$ is an unbiased estimator of $\Sigma$.

Statistical inference often requires more than just the mean and/or the variance of a statistic. We need the sampling distribution of the statistics to derive confidence intervals or to define rejection regions in hypothesis testing for a given significance level. Theorem 4.9 gives the distribution of the sample mean for a multinormal population.

Theorem 4.9 Let $X_{1}, \ldots, X_{n}$ be i.i.d. with $X_{i} \sim N_{p}(\mu, \Sigma)$. Then $\bar{x} \sim$ $N_{p}\left(\mu, n^{-1} \Sigma\right)$.

Proof $\bar{x}=n^{-1} \sum_{i=1}^{n} X_{i}$ is a linear combination of independent normal variables, so it has a normal distribution (see Chap.5). The mean and the covariance matrix were given in the preceding example.

With multivariate statistics, the sampling distributions of the statistics are often more difficult to derive than in the preceding Theorem. In addition they might be so complicated that approximations have to be used. These approximations are provided by limit theorems. Since they are based on asymptotic limits, the approximations are only valid when the sample size is large enough. In spite of this restriction, they make complicated situations rather simple. The following central limit theorem shows that even if the parent distribution is not normal, when the sample size $n$ is large, the sample mean $\bar{x}$ has an approximate normal distribution.

Theorem 4.10 (Central Limit Theorem (CLT)) Let $X_{1}, X_{2}, \ldots, X_{n}$ be i.i.d. with $X_{i} \sim(\mu, \Sigma)$. Then the distribution of $\sqrt{n}(\bar{x}-\mu)$ is asymptotically $N_{p}(0, \Sigma)$, i.e.

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

The symbol " $\xrightarrow{\mathcal{L}}$ " denotes convergence in distribution which means that the distribution function of the random vector $\sqrt{n}(\bar{x}-\mu)$ converges to the distribution function of $N_{p}(0, \Sigma)$.

Example 4.16 Assume that $X_{1}, \ldots, X_{n}$ are i.i.d. and that they have Bernoulli distributions where $p=\frac{1}{2}$ (this means that $\left.P\left(X_{i}=1\right)=\frac{1}{2}, P\left(X_{i}=0\right)=\frac{1}{2}\right)$. Then $\mu=p=\frac{1}{2}$ and $\Sigma=p(1-p)=\frac{1}{4}$. Hence,

$$
\sqrt{n}\left(\bar{x}-\frac{1}{2}\right) \xrightarrow{\mathcal{L}} N_{1}\left(0, \frac{1}{4}\right) \quad \text { as } \quad n \longrightarrow \infty
$$

The results are shown in Fig. 4.4 for varying sample sizes.
Example 4.17 Now consider a two-dimensional random sample $X_{1}, \ldots, X_{n}$ that is i.i.d. and created from two independent Bernoulli distributions with $p=0.5$. The joint distribution is given by $P\left(X_{i}=(0,0)^{\top}\right)=\frac{1}{4}, P\left(X_{i}=(0,1)^{\top}\right)=\frac{1}{4}, P\left(X_{i}=\right.$ $\left.(1,0)^{\top}\right)=\frac{1}{4}, P\left(X_{i}=(1,1)^{\top}\right)=\frac{1}{4}$. Here we have

$$
\sqrt{n}\left\{\bar{x}-\binom{\frac{1}{2}}{\frac{1}{2}}\right\}=N_{2}\left(\binom{0}{0},\left(\begin{array}{cc}
\frac{1}{4} & 0 \\
0 & \frac{1}{4}
\end{array}\right)\right) \quad \text { as } \quad n \longrightarrow \infty
$$

Figure 4.5 displays the estimated two-dimensional density for different sample sizes.

The asymptotic normal distribution is often used to construct confidence intervals for the unknown parameters. A confidence interval at the level $1-\alpha, \alpha \in(0,1)$, is an interval that covers the true parameter with probability $1-\alpha$ :

$$
P\left(\theta \in\left[\widehat{\theta}_{l}, \widehat{\theta}_{u}\right]\right)=1-\alpha
$$

where $\theta$ denotes the (unknown) parameter and $\widehat{\theta}_{l}$ and $\widehat{\theta}_{u}$ are the lower and upper confidence bounds, respectively.
Example 4.18 Consider the i.i.d. random variables $X_{1}, \ldots, X_{n}$ with $X_{i} \sim\left(\mu, \sigma^{2}\right)$ and $\sigma^{2}$ known. Since we have $\sqrt{n}(\bar{x}-\mu) \xrightarrow{\mathcal{L}} N\left(0, \sigma^{2}\right)$ from the CLT, it follows that

$$
P\left(-u_{1-\alpha / 2} \leq \sqrt{n} \frac{(\bar{x}-\mu)}{\sigma} \leq u_{1-\alpha / 2}\right) \longrightarrow 1-\alpha, \quad \text { as } \quad n \longrightarrow \infty
$$

where $u_{1-\alpha / 2}$ denotes the $(1-\alpha / 2)$-quantile of the standard normal distribution. Hence the interval

$$
\left[\bar{x}-\frac{\sigma}{\sqrt{n}} u_{1-\alpha / 2}, \bar{x}+\frac{\sigma}{\sqrt{n}} u_{1-\alpha / 2}\right]
$$

is an approximate $(1-\alpha)$-confidence interval for $\mu$.
But what can we do if we do not know the variance $\sigma^{2}$ ? The following corollary gives the answer.

Fig. 4.4 The CLT for Bernoulli distributed random variables. Sample size $n=5$ (up) and $n=35$ (down) 0 MVAcltbern


Corollary 4.1 If $\hat{\Sigma}$ is a consistent estimate for $\Sigma$, then the CLT still holds, namely

$$
\sqrt{n} \widehat{\Sigma}^{-1 / 2}(\bar{x}-\mu) \xrightarrow{\mathcal{L}} N_{p}(0, \mathcal{I}) \quad \text { as } n \longrightarrow \infty .
$$



Fig. 4.5 The CLT in the two-dimensional case. Sample size $n=5$ (left) and $n=85$ (right) 0 . MVAcltbern2

Example 4.19 Consider the i.i.d. random variables $X_{1}, \ldots, X_{n}$ with $X_{i} \sim\left(\mu, \sigma^{2}\right)$, and now with an unknown variance $\sigma^{2}$. From Corollary 4.1 using $\widehat{\sigma}^{2}=\frac{1}{n} \sum_{i=1}^{n}\left(x_{i}-\right.$ $\bar{x})^{2}$ we obtain

$$
\sqrt{n}\left(\frac{\bar{x}-\mu}{\hat{\sigma}}\right) \xrightarrow{\mathcal{L}} N(0,1) \quad \text { as } \quad n \longrightarrow \infty .
$$

Hence we can construct an approximate ( $1-\alpha$ )-confidence interval for $\mu$ using the variance estimate $\hat{\sigma}^{2}$ :

$$
C_{1-\alpha}=\left[\bar{x}-\frac{\hat{\sigma}}{\sqrt{n}} u_{1-\alpha / 2}, \bar{x}+\frac{\hat{\sigma}}{\sqrt{n}} u_{1-\alpha / 2}\right] .
$$

Note that by the CLT

$$
P\left(\mu \in C_{1-\alpha}\right) \longrightarrow 1-\alpha \quad \text { as } \quad n \longrightarrow \infty
$$

Remark 4.1 One may wonder how large should $n$ be in practice to provide reasonable approximations. There is no definite answer to this question: it mainly depends on the problem at hand (the shape of the distribution of the $X_{i}$ and the dimension of $X_{i}$ ). If the $X_{i}$ are normally distributed, the normality of $\bar{x}$ is achieved from $n=1$. In most situations, however, the approximation is valid in onedimensional problems for $n$ larger than, say, 50 .

## Transformation of Statistics

Often in practical problems, one is interested in a function of parameters for which one has an asymptotically normal statistic. Suppose for instance that we are interested in a cost function depending on the mean $\mu$ of the process: $f(\mu)=$ $\mu^{\top} \mathcal{A} \mu$ where $\mathcal{A}>0$ is given. To estimate $\mu$ we use the asymptotically normal statistic $\bar{x}$. The question is: how does $f(\bar{x})$ behave? More generally, what happens to a statistic $t$ that is asymptotically normal when we transform it by a function $f(t)$ ? The answer is given by the following theorem.
Theorem 4.11 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.

$$
\begin{equation*}
\sqrt{n}\{f(t)-f(\mu)\} \xrightarrow{\mathcal{L}} N_{q}\left(0, \mathcal{D}^{\top} \Sigma \mathcal{D}\right) \quad \text { for } \quad n \longrightarrow \infty, \tag{4.56}
\end{equation*}
$$

where

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

is the $(p \times q)$ matrix of all partial derivatives.
Example 4.20 We are interested in seeing how $f(\bar{x})=\bar{x}^{\top} \mathcal{A} \bar{x}$ behaves asymptotically with respect to the quadratic cost function of $\mu, f(\mu)=\mu^{\top} \mathcal{A} \mu$, where $\mathcal{A}>0$.

$$
D=\left.\frac{\partial f(\bar{x})}{\partial \bar{x}}\right|_{\bar{x}=\mu}=2 \mathcal{A} \mu .
$$

By Theorem 4.11 we have

$$
\sqrt{n}\left(\bar{x}^{\top} \mathcal{A} \bar{x}-\mu^{\top} \mathcal{A} \mu\right) \xrightarrow{\mathcal{L}} N_{1}\left(0,4 \mu^{\top} \mathcal{A} \Sigma \mathcal{A} \mu\right) .
$$

Example 4.21 Suppose

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

We have by the CLT (Theorem 4.10) for $n \rightarrow \infty$ that

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

Suppose that we would like to compute the distribution of $\binom{\bar{x}_{1}^{2}-\bar{x}_{2}}{\bar{x}_{1}+3 \bar{x}_{2}}$. According to Theorem 4.11 we have to consider $f=\left(f_{1}, f_{2}\right)^{\top}$ with

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

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

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

Thus

$$
\mathcal{D}=\left(\begin{array}{rr}
0 & 1 \\
-1 & 3
\end{array}\right)
$$

The covariance is
which yields

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

Example 4.22 Let us continue the previous example by adding one more component to the function $f$. Since $q=3>p=2$, we might expect a singular normal distribution. Consider $f=\left(f_{1}, f_{2}, f_{3}\right)^{\top}$ with

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

From this we have that

$$
\mathcal{D}=\left(\begin{array}{rrr}
0 & 1 & 0 \\
-1 & 3 & 0
\end{array}\right) \quad \text { and thus } \quad \mathcal{D}^{\top} \Sigma \mathcal{D}=\left(\begin{array}{rrr}
1 & -\frac{7}{2} & 0 \\
-\frac{7}{2} & 13 & 0 \\
0 & 0 & 0
\end{array}\right)
$$

The limit is in fact a singular normal distribution!

| Su Summary |
| :---: |
| $\hookrightarrow$ If $X_{1}, \ldots, X_{n}$ are i.i.d. random vectors with $X_{i} \sim N_{p}(\mu, \Sigma)$, then $\bar{x} \sim N_{p}\left(\mu, \frac{1}{n} \Sigma\right)$. |
| $\hookrightarrow$ 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). |
| $\hookrightarrow$ 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: $\bar{x} \pm \frac{\hat{\sigma}}{\sqrt{n}} u_{1-\alpha / 2}$. |
| $\hookrightarrow$ 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. |

### 4.6 Heavy-Tailed Distributions

Heavy-tailed distributions were first introduced by the Italian-born Swiss economist Pareto and extensively studied by Paul Lévy. Although in the beginning these distributions were mainly studied theoretically, nowadays they have found many applications in areas as diverse as finance, medicine, seismology, structural engineering. More concretely, they have been used to model returns of assets in financial markets, stream flow in hydrology, precipitation and hurricane damage in meteorology, earthquake prediction in seismology, pollution, material strength, teletraffic and many others.

A distribution is called heavy-tailed if it has higher probability density in its tail area compared with a normal distribution with same mean $\mu$ and variance $\sigma^{2}$. Figure 4.6 demonstrates the differences of the pdf curves of a standard Gaussian distribution and a Cauchy distribution with location parameter $\mu=0$ and scale parameter $\sigma=1$. The graphic shows that the probability density of the Cauchy distribution is much higher than that of the Gaussian in the tail part, while in the area around the centre, the probability density of the Cauchy distribution is much lower.

In terms of kurtosis, a heavy-tailed distribution has kurtosis greater than 3 (see Chap. 4, formula (4.40)), which is called leptokurtic, in contrast to mesokurtic distribution (kurtosis $=3$ ) and platykurtic distribution (kurtosis $<3$ ). Since univariate heavy-tailed distributions serve as basics for their multivariate counterparts and their density properties have been proved useful even in multivariate cases, we will start from introducing some univariate heavy-tailed distributions. Then we will move on to analyse their multivariate counterparts and their tail behaviour.


Fig. 4.6 Comparison of the pdf of a standard Gaussian (blue) and a Cauchy distribution (red) with location parameter 0 and scale parameter 10 MVAgausscauchy


Fig. 4.7 pdf (left) and cdf (right) of GH $(\lambda=0.5$, black), HYP (red), and NIG (blue) with $\alpha=1, \beta=0, \delta=1, \mu=0$ O MVAghdis

## Generalised Hyperbolic Distribution

The generalised hyperbolic distribution was introduced by Barndorff-Nielsen and at first applied to model grain size distributions of wind blown sands. Today one of its most important uses is in stock price modelling and market risk measurement. The name of the distribution is derived from the fact that its log-density forms a hyperbola, while the log-density of the normal distribution is a parabola (Fig. 4.7).

The density of a one-dimensional generalised hyperbolic (GH) distribution for $x \in \mathbb{R}$ is

$$
\begin{align*}
& f_{\mathrm{GH}}(x ; \lambda, \alpha, \beta, \delta, \mu) \\
& \quad=\frac{\left(\sqrt{\alpha^{2}-\beta^{2}} / \delta\right)^{\lambda}}{\sqrt{2 \pi} K_{\lambda}\left(\delta \sqrt{\alpha^{2}-\beta^{2}}\right)} \frac{K_{\lambda-1 / 2}\left\{\alpha \sqrt{\delta^{2}+(x-\mu)^{2}}\right\}}{\left.\sqrt{\delta^{2}+(x-\mu)^{2}} / \alpha\right)^{1 / 2-\lambda}} e^{\beta(x-\mu)} \tag{4.57}
\end{align*}
$$

where $K_{\lambda}$ is a modified Bessel function of the third kind with index $\lambda$

$$
\begin{equation*}
K_{\lambda}(x)=\frac{1}{2} \int_{0}^{\infty} y^{\lambda-1} e^{-\frac{x}{2}\left(y+y^{-1}\right)} d y \tag{4.58}
\end{equation*}
$$

The domain of variation of the parameters is $\mu \in \mathbb{R}$ and

$$
\begin{array}{ll}
\delta \geq 0,|\beta|<\alpha, & \text { if } \quad \lambda>0 \\
\delta>0,|\beta|<\alpha, & \text { if } \quad \lambda=0 \\
\delta>0,|\beta| \leq \alpha, & \text { if } \quad \lambda<0
\end{array}
$$

The generalised hyperbolic distribution has the following mean and variance

$$
\begin{align*}
\mathrm{E}[X]= & \mu+\frac{\delta \beta}{\sqrt{\alpha^{2}-\beta^{2}}} \frac{K_{\lambda+1}\left(\delta \sqrt{\alpha^{2}-\beta^{2}}\right)}{K_{\lambda}\left(\delta \sqrt{\alpha^{2}-\beta^{2}}\right)}  \tag{4.59}\\
\operatorname{Var}[X]= & \delta^{2}\left[\frac{K_{\lambda+1}\left(\delta \sqrt{\alpha^{2}-\beta^{2}}\right)}{\delta \sqrt{\alpha^{2}-\beta^{2}} K_{\lambda}\left(\delta \sqrt{\alpha^{2}-\beta^{2}}\right)}+\frac{\beta^{2}}{\alpha^{2}-\beta^{2}}\left[\frac{K_{\lambda+2}\left(\delta \sqrt{\alpha^{2}-\beta^{2}}\right)}{K_{\lambda}\left(\delta \sqrt{\alpha^{2}-\beta^{2}}\right)}\right.\right. \\
& \left.\left.-\left\{\frac{K_{\lambda+1}\left(\delta \sqrt{\alpha^{2}-\beta^{2}}\right)}{K_{\lambda}\left(\delta \sqrt{\alpha^{2}-\beta^{2}}\right)}\right\}^{2}\right]\right] \tag{4.60}
\end{align*}
$$

Where $\mu$ and $\delta$ play important roles in the density's location and scale respectively. With specific values of $\lambda$, we obtain different sub-classes of GH such as hyperbolic (HYP) or normal-inverse Gaussian (NIG) distribution.

For $\lambda=1$ we obtain the hyperbolic distributions (HYP)

$$
\begin{equation*}
f_{\mathrm{HYP}}(x ; \alpha, \beta, \delta, \mu)=\frac{\sqrt{\alpha^{2}-\beta^{2}}}{2 \alpha \delta K_{1}\left(\delta \sqrt{\alpha^{2}-\beta^{2}}\right)} e^{\left\{-\alpha \sqrt{\delta^{2}+(x-\mu)^{2}}+\beta(x-\mu)\right\}} \tag{4.61}
\end{equation*}
$$

where $x, \mu \in \mathbb{R}, \delta \geq 0$ and $|\beta|<\alpha$. For $\lambda=-1 / 2$ we obtain the NIG distribution

$$
\begin{equation*}
f_{\mathrm{NIG}}(x ; \alpha, \beta, \delta, \mu)=\frac{\alpha \delta}{\pi} \frac{K_{1}\left(\alpha \sqrt{\left.\left(\delta^{2}+(x-\mu)^{2}\right)\right)}\right.}{\sqrt{\delta^{2}+(x-\mu)^{2}}} e^{\left\{\delta \sqrt{\alpha^{2}-\beta^{2}}+\beta(x-\mu)\right\}} \tag{4.62}
\end{equation*}
$$

## Student's $\boldsymbol{t}$-Distribution

The $t$-distribution was first analysed by Gosset (1908) who published it under pseudonym "Student" by request of his employer. Let $X$ be a normally distributed random variable with mean $\mu$ and variance $\sigma^{2}$, and $Y$ be the random variable such that $Y^{2} / \sigma^{2}$ has a chi-square distribution with $n$ degrees of freedom. Assume that $X$ and $Y$ are independent, then

$$
\begin{equation*}
t \stackrel{\text { def }}{=} \frac{X \sqrt{n}}{Y} \tag{4.63}
\end{equation*}
$$

is distributed as Student's $t$ with $n$ degrees of freedom. The $t$-distribution has the following density function

$$
\begin{equation*}
f_{t}(x ; n)=\frac{\Gamma\left(\frac{n+1}{2}\right)}{\sqrt{n \pi} \Gamma\left(\frac{n}{2}\right)}\left(1+\frac{x^{2}}{n}\right)^{-\frac{n+1}{2}} \tag{4.64}
\end{equation*}
$$

where $n$ is the number of degrees of freedom, $-\infty<x<\infty$, and $\Gamma$ is the gamma function:

$$
\begin{equation*}
\Gamma(\alpha)=\int_{0}^{\infty} x^{\alpha-1} e^{-x} d x \tag{4.65}
\end{equation*}
$$

The mean, variance, skewness and kurtosis of Student's $t$-distribution ( $n>4$ ) are:

$$
\begin{aligned}
\mu & =0 \\
\sigma^{2} & =\frac{n}{n-2} \\
\text { Skewness } & =0 \\
\text { Kurtosis } & =3+\frac{6}{n-4} .
\end{aligned}
$$

The $t$-distribution is symmetric around 0 , which is consistent with the fact that its mean is 0 and skewness is also 0 (Fig. 4.8).

Student's $t$-distribution approaches the normal distribution as $n$ increases, since

$$
\begin{equation*}
\lim _{n \rightarrow \infty} f_{t}(x ; n)=\frac{1}{\sqrt{2 \pi}} e^{-\frac{x^{2}}{2}} \tag{4.66}
\end{equation*}
$$

In practice the $t$-distribution is widely used, but its flexibility of modelling is restricted because of the integer-valued tail index.

In the tail area of the $t$-distribution, $x$ is proportional to $|x|^{-(n+1)}$. In Fig. 4.13 we compared the tail-behaviour of $t$-distribution with different degrees of freedom. With higher degree of freedom, the $t$-distribution decays faster.


Fig. 4.8 pdf (left) and cdf (right) of $t$-distribution with different degrees of freedom ( t 3 stands for $t$-distribution with degree of freedom 3) Q MVAtdis

## Laplace Distribution

The univariate Laplace distribution with mean zero was introduced by Laplace (1774). The Laplace distribution can be defined as the distribution of differences between two independent variates with identical exponential distributions. Therefore it is also called the double exponential distribution (Fig. 4.9).

The Laplace distribution with mean $\mu$ and scale parameter $\theta$ has the pdf

$$
\begin{equation*}
f_{\text {Laplace }}(x ; \mu, \theta)=\frac{1}{2 \theta} e^{-\frac{|x-\mu|}{\theta}} \tag{4.67}
\end{equation*}
$$

and the cdf

$$
\begin{equation*}
F_{\text {Laplace }}(x ; \mu, \theta)=\frac{1}{2}\left\{1+\operatorname{sign}(x-\mu)\left(1-e^{-\frac{|x-\mu|}{\theta}}\right)\right\}, \tag{4.68}
\end{equation*}
$$

where sign is sign function. The mean, variance, skewness and kurtosis of the Laplace distribution are

$$
\begin{aligned}
\mu & =\mu \\
\sigma^{2} & =2 \theta^{2} \\
\text { Skewness } & =0 \\
\text { Kurtosis } & =6
\end{aligned}
$$

With mean 0 and $\theta=1$, we obtain the standard Laplace distribution


Fig. 4.9 pdf (left) and cdf (right) of Laplace distribution with zero mean and different scale parameters (L1 stands for Laplace distribution with $\theta=1$ ) $Q$ MVAlaplacedis

$$
\begin{align*}
& f(x)=\frac{e^{-|x|}}{2}  \tag{4.69}\\
& F(x)=\left\{\begin{array}{lll}
\frac{e^{x}}{2} & \text { for } & x<0 \\
1-\frac{e^{-x}}{2} & \text { for } & x \geq 0
\end{array}\right. \tag{4.70}
\end{align*}
$$

## Cauchy Distribution

The Cauchy distribution is motivated by the following example.
Example 4.23 A gangster has just robbed a bank. As he runs to a point $s$ metres away from the wall of the bank, a policeman reaches the crime scene behind the wall of the bank. The robber turns back and starts to shoot but he is such a poor shooter that the angle of his fire (marked in Fig. 4.10 as $\alpha$ ) is uniformly distributed. The bullets hit the wall at distance $x$ (from the centre). Obviously the distribution of $x$, the random variable where the bullet hits the wall, is of vital knowledge to the policeman in order to identify the location of the gangster. (Should the policeman calculate the mean or the median of the observed bullet hits $\left\{x_{i}\right\}_{i=1}^{n}$ in order to identify the location of the robber?)

Since $\alpha$ is uniformly distributed:

$$
f(\alpha)=\frac{1}{\pi} I(\alpha \in[-\pi / 2, \pi / 2])
$$



Fig. 4.10 Introduction to Cauchy distribution—robber vs. policeman
and

$$
\begin{aligned}
\tan \alpha & =\frac{x}{s} \\
\alpha & =\arctan \left(\frac{x}{s}\right) \\
d \alpha & =\frac{1}{s} \frac{1}{1+\left(\frac{x}{s}\right)^{2}} d x
\end{aligned}
$$

For a small interval $d \alpha$, the probability is given by

$$
\begin{aligned}
f(\alpha) d \alpha & =\frac{1}{\pi} d \alpha \\
& =\frac{1}{s \pi} \frac{1}{1+\left(\frac{x}{s}\right)^{2}} d x
\end{aligned}
$$

with

$$
\int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \frac{1}{\pi} d \alpha=1
$$

$$
\begin{aligned}
\int_{-\infty}^{\infty} \frac{1}{s \pi} \frac{1}{1+\left(\frac{x}{s}\right)^{2}} d x & =\frac{1}{\pi}\left\{\arctan \left(\frac{x}{s}\right)\right\}_{-\infty}^{\infty} \\
& =\frac{1}{\pi}\left\{\frac{\pi}{2}-\left(-\frac{\pi}{2}\right)\right\} \\
& =1
\end{aligned}
$$

So the pdf of $x$ can be written as:

$$
f(x)=\frac{1}{s \pi} \frac{1}{1+\left(\frac{x}{s}\right)^{2}}
$$

The general formula for the pdf and cdf of the Cauchy distribution is

$$
\begin{align*}
& f_{\text {Cauchy }}(x ; m, s)=\frac{1}{s \pi} \frac{1}{1+\left(\frac{x-m}{s}\right)^{2}}  \tag{4.71}\\
& F_{\text {Cauchy }}(x ; m, s)=\frac{1}{2}+\frac{1}{\pi} \arctan \left(\frac{x-m}{s}\right) \tag{4.72}
\end{align*}
$$

where $m$ and $s$ are location and scale parameter respectively. The case in the above example where $m=0$ and $s=1$ is called the standard Cauchy distribution with pdf and cdf as following,

$$
\begin{align*}
f_{\text {Cauchy }}(x) & =\frac{1}{\pi\left(1+x^{2}\right)}  \tag{4.73}\\
F_{\text {Cauchy }}(x ; m, s) & =\frac{1}{2}+\frac{\arctan (x)}{\pi} \tag{4.74}
\end{align*}
$$

The mean, variance, skewness and kurtosis of Cauchy distribution are all undefined, since its moment generating function diverges. But it has mode and median, both equal to the location parameter $m$ (Fig. 4.11).

## Mixture Model

Mixture modelling concerns modelling a statistical distribution by a mixture (or weighted sum) of different distributions. For many choices of component density functions, the mixture model can approximate any continuous density to arbitrary accuracy, provided that the number of component density functions is sufficiently large and the parameters of the model are chosen correctly. The pdf of a mixture distribution consists of $n$ distributions and can be written as:


Fig. 4.11 pdf (left) and cdf (right) of Cauchy distribution with $m=0$ and different scale parameters ( C 1 stands for Cauchy distribution with $s=1$ ) Q MVAcauchy

$$
\begin{equation*}
f(x)=\sum_{l=1}^{L} w_{l} p_{l}(x) \tag{4.75}
\end{equation*}
$$

under the constraints:

$$
\begin{aligned}
0 \leq w_{l} & \leq 1 \\
\sum_{l=1}^{L} w_{l} & =1 \\
\int p_{l}(x) d x & =1
\end{aligned}
$$

where $p_{l}(x)$ is the pdf of the $l$ 'th component density and $w_{l}$ is a weight. The mean, variance, skewness and kurtosis of a mixture are

$$
\begin{align*}
\mu & =\sum_{l=1}^{L} w_{l} \mu_{l}  \tag{4.76}\\
\sigma^{2} & =\sum_{l=1}^{L} w_{l}\left\{\sigma_{l}^{2}+\left(\mu_{l}-\mu\right)^{2}\right\}  \tag{4.77}\\
\text { Skewness } & =\sum_{l=1}^{L} w_{l}\left\{\left(\frac{\sigma_{l}}{\sigma}\right)^{3} S K_{l}+\frac{3 \sigma_{l}^{2}\left(\mu_{l}-\mu\right)}{\sigma^{3}}+\left(\frac{\mu_{l}-\mu}{\sigma}\right)^{3}\right\} \tag{4.78}
\end{align*}
$$

$$
\begin{align*}
\text { Kurtosis }= & \sum_{l=1}^{L} w_{l}\left\{\left(\frac{\sigma_{l}}{\sigma}\right)^{4} K_{l}+\frac{6\left(\mu_{l}-\mu\right)^{2} \sigma_{l}^{2}}{\sigma^{4}}+\frac{4\left(\mu_{l}-\mu\right) \sigma_{l}^{3}}{\sigma^{4}} S K_{l}\right. \\
& \left.+\left(\frac{\mu_{l}-\mu}{\sigma}\right)^{4}\right\}, \tag{4.79}
\end{align*}
$$

where $\mu_{l}, \sigma_{l}, S K_{l}$ and $K_{l}$ are respectively mean, variance, skewness and kurtosis of $l$ 'th distribution.

Mixture models are ubiquitous in virtually every facet of statistical analysis, machine learning and data mining. For data sets comprising continuous variables, the most common approach involves mixture distributions having Gaussian components.

The pdf for a Gaussian mixture is:

$$
\begin{equation*}
f_{\mathrm{GM}}(x)=\sum_{l=1}^{L} \frac{w_{l}}{\sqrt{2 \pi} \sigma_{l}} e^{-\frac{\left(x-\mu_{l}\right)^{2}}{2 \sigma_{l}^{2}}} . \tag{4.80}
\end{equation*}
$$

For a Gaussian mixture consisting of Gaussian distributions with mean 0 , this can be simplified to:

$$
\begin{equation*}
f_{\mathrm{GM}}(x)=\sum_{l=1}^{L} \frac{w_{l}}{\sqrt{2 \pi} \sigma_{l}} e^{-\frac{x^{2}}{2 \sigma_{l}^{2}}}, \tag{4.81}
\end{equation*}
$$

with variance, skewness and kurtosis

$$
\begin{align*}
\sigma^{2} & =\sum_{l=1}^{L} w_{l} \sigma_{l}^{2}  \tag{4.82}\\
\text { Skewness } & =0  \tag{4.83}\\
\text { Kurtosis } & =\sum_{l=1}^{L} w_{l}\left(\frac{\sigma_{l}}{\sigma}\right)^{4} 3 \tag{4.84}
\end{align*}
$$

Example 4.24 Consider a Gaussian Mixture which is $80 \% N(0,1)$ and $20 \% N(0,9)$. The pdf of $N(0,1)$ and $N(0,9)$ are (Fig. 4.12):

$$
\begin{aligned}
& f_{N(0,1)}(x)=\frac{1}{\sqrt{2 \pi}} e^{-\frac{x^{2}}{2}} \\
& f_{N(0,9)}(x)=\frac{1}{3 \sqrt{2 \pi}} e^{-\frac{x^{2}}{18}}
\end{aligned}
$$

so the pdf of the Gaussian Mixture is


Fig. 4.12 pdf (left) and cdf (right) of a Gaussian mixture (Example 4.24) Q mVAmixture
Table 4.2 Basic statistics of $t$, Laplace and Cauchy distribution

|  | $t$ | Laplace | Cauchy |
| :--- | :--- | :--- | :--- |
| Mean | 0 | $\mu$ | Not defined |
| Variance | $\frac{n}{n-2}$ | $2 \theta^{2}$ | Not defined |
| Skewness | 0 | 0 | Not defined |
| Kurtosis | $3+\frac{6}{n-4}$ | 6 | Not defined |

$$
f_{\mathrm{GM}}(x)=\frac{1}{5 \sqrt{2 \pi}}\left(4 e^{-\frac{x^{2}}{2}}+\frac{1}{3} e^{-\frac{x^{2}}{18}}\right)
$$

Notice that the Gaussian Mixture is not a Gaussian distribution:

$$
\begin{aligned}
\mu & =0 \\
\sigma^{2} & =0.8 \times 1+0.2 \times 9=2.6 \\
\text { Skewness } & =0 \\
\text { Kurtosis } & =0.8 \times\left(\frac{1}{\sqrt{2.6}}\right)^{4} \times 3+0.2 \times\left(\frac{\sqrt{9}}{\sqrt{2.6}}\right)^{4} \times 3=7.54
\end{aligned}
$$

The kurtosis of this Gaussian mixture is higher than 3.
A summary of the basic statistics is given in Tables 4.2 and 4.3.

Table 4.3 Basic statistics of GH distribution and mixture model

| GH | $\mu+\frac{\delta \beta}{\sqrt{\alpha^{2}+\beta^{2}}} \frac{K_{\lambda+1}\left(\delta \sqrt{\alpha^{2}+\beta^{2}}\right)}{K_{\lambda}\left(\delta \sqrt{\alpha^{2}+\beta^{2}}\right)}$ |
| :--- | :--- |
| Mean | $\delta^{2}\left[\frac{K_{\lambda+1}\left(\delta \sqrt{\alpha^{2}+\beta^{2}}\right)}{\delta \sqrt{\alpha^{2}+\beta^{2}} K_{\lambda}\left(\delta \sqrt{\alpha^{2}+\beta^{2}}\right)}+\frac{\beta^{2}}{\alpha^{2}+\beta^{2}}\left[\frac{K_{\lambda+2}\left(\delta \sqrt{\alpha^{2}+\beta^{2}}\right)}{K_{\lambda}\left(\delta \sqrt{\alpha^{2}+\beta^{2}}\right)}-\left\{\frac{K_{\lambda+1}\left(\delta \sqrt{\alpha^{2}+\beta^{2}}\right)}{K_{\lambda}\left(\delta \sqrt{\alpha^{2}+\beta^{2}}\right)}\right\}^{2}\right]\right]$ |
| Variance | $\sum_{l=1}^{L} w_{l} \mu_{l}$ |
| Mixture | $\sum_{l=1}^{L} w_{l}\left\{\sigma_{l}^{2}+\left(\mu_{l}-\mu\right)^{2}\right\}$ |
| Mean | $\sum_{l=1}^{L} w_{l}\left\{\left(\frac{\sigma_{l}}{\sigma}\right)^{3} S K_{l}+\frac{3 \sigma_{l}^{2}\left(\mu_{l}-\mu\right)}{\sigma^{3}}+\left(\frac{\mu l-\mu}{\sigma}\right)^{3}\right\}$ |
| Variance $\}$ |  |
| Kurwness | $\sum_{l=1}^{L} w_{l}\left\{\left(\frac{\sigma_{l}}{\sigma}\right)^{4} K_{l}+\frac{6\left(\mu_{l}-\mu\right)^{2} \sigma_{l}^{2}}{\sigma^{4}}+\frac{4\left(\mu_{l}-\mu\right) \sigma_{l}^{3}}{\sigma^{4}} S K_{l}+\left(\frac{\mu_{l}-\mu}{\sigma}\right)^{4}\right\}$ |

## Multivariate Generalised Hyperbolic Distribution

The multivariate Generalised Hyperbolic Distribution $\left(\mathrm{GH}_{d}\right)$ has the following pdf

$$
\begin{align*}
& f_{\mathrm{GH}_{d}}(x ; \lambda, \alpha, \beta, \delta, \Delta, \mu)=a_{d} \frac{K_{\lambda-\frac{d}{2}}\left\{\alpha \sqrt{\delta^{2}+(x-\mu)^{\top} \Delta^{-1}(x-\mu)}\right\}}{\left\{\alpha^{-1} \sqrt{\delta^{2}+(x-\mu)^{\top} \Delta^{-1}(x-\mu)}\right\}^{\frac{d}{2}-\lambda}} e^{\beta^{\top}(x-\mu)}  \tag{4.85}\\
& a_{d}=a_{d}(\lambda, \alpha, \beta, \delta, \Delta)=\frac{\left(\sqrt{\alpha^{2}-\beta^{\top} \Delta \beta} / \delta\right)^{\lambda}}{(2 \pi)^{\frac{d}{2}} K_{\lambda}\left(\delta \sqrt{\alpha^{2}-\beta^{\top} \Delta \beta}\right.}, \tag{4.86}
\end{align*}
$$

and characteristic function

$$
\begin{align*}
\phi(t)= & \left(\frac{\alpha^{2}-\beta^{\top} \Delta \beta}{\alpha^{2}-\beta^{\top} \Delta \beta+\frac{1}{2} t^{\top} \Delta t-i \beta^{\top} \Delta t}\right)^{\frac{\lambda}{2}} \\
& \times \frac{K_{\lambda}\left(\delta \sqrt{\alpha^{2}-\beta^{\top} \Delta \beta^{\top}+\frac{1}{2} t^{\top} \Delta t-i \beta^{\top} \Delta t}\right)}{K_{\lambda}\left(\delta \sqrt{\alpha^{2}-\beta^{\top} \Delta \beta^{\top}}\right)} \tag{4.87}
\end{align*}
$$

These parameters have the following domain of variation:

$$
\begin{array}{ll}
\lambda \in \mathbb{R}, & \beta, \mu \in \mathbb{R}^{d} \\
\delta>0, & \alpha>\beta^{\top} \Delta \beta \\
\Delta \in \mathbb{R}^{d \times d} & \text { positive definite matrix } \\
|\Delta|=1
\end{array}
$$

For $\lambda=\frac{d+1}{2}$ we obtain the multivariate hyperbolic (HYP) distribution; for $\lambda=-\frac{1}{2}$ we get the multivariate normal inverse Gaussian (NIG) distribution.

Blæsild and Jensen (1981) introduced a second parameterisation ( $\zeta, \Pi, \Sigma$ ), where

$$
\begin{align*}
\zeta & =\delta \sqrt{\alpha^{2}-\beta^{\top} \Delta \beta}  \tag{4.88}\\
\Pi & =\beta \sqrt{\frac{\Delta}{\alpha^{2}-\beta^{\top} \Delta \beta}}  \tag{4.89}\\
\Sigma & =\delta^{2} \Delta \tag{4.90}
\end{align*}
$$

The mean and variance of $X \sim \mathrm{GH}_{d}$

$$
\begin{align*}
\mathrm{E}[X] & =\mu+\delta R_{\lambda}(\zeta) \Pi \Delta^{\frac{1}{2}}  \tag{4.91}\\
\operatorname{Var}[X] & =\delta^{2}\left\{\zeta^{-1} R_{\lambda}(\zeta) \Delta+S_{\lambda}(\zeta)\left(\Pi \Delta^{\frac{1}{2}}\right)^{\top}\left(\Pi \Delta^{\frac{1}{2}}\right)\right\} \tag{4.92}
\end{align*}
$$

where

$$
\begin{align*}
R_{\lambda}(x) & =\frac{K_{\lambda+1}(x)}{K_{\lambda}(x)}  \tag{4.93}\\
S_{\lambda}(x) & =\frac{K_{\lambda+2}(x) K_{\lambda}(x)-K_{\lambda+1}^{2}(x)}{K_{\lambda}^{2}(x)} \tag{4.94}
\end{align*}
$$

Theorem 4.12 Suppose that $X$ is a d-dimensional variate distributed according to the generalised hyperbolic distribution $\mathrm{GH}_{d}$. Let $\left(X_{1}, X_{2}\right)$ be a partitioning of $X$, let $r$ and $k$ denote the dimensions of $X_{1}$ and $X_{2}$, respectively, and let $\left(\beta_{1}, \beta_{2}\right)$ and $\left(\mu_{1}, \mu_{2}\right)$ be similar partitions of $\beta$ and $\mu$, let

$$
\Delta=\left(\begin{array}{ll}
\Delta_{11} & \Delta_{12}  \tag{4.95}\\
\Delta_{21} & \Delta_{22}
\end{array}\right)
$$

be a partition of $\Delta$ such that $\Delta_{11}$ is a $r \times r$ matrix. Then one has the following

1. The distribution of $X_{1}$ is the $r$-dimensional generalised hyperbolic distribution, $\mathrm{GH}_{r}\left(\lambda^{*}, \alpha^{*}, \beta^{*}, \delta^{*}, \mu^{*}, \Delta^{*}\right)$, where

$$
\begin{aligned}
& \lambda^{*}=\lambda \\
& \alpha^{*}=\left|\Delta_{11}\right|^{-\frac{1}{2 r}}\left\{\alpha^{2}-\beta_{2}\left(\Delta_{22}-\Delta_{21} \Delta_{11}^{-1} \Delta_{12}\right) \beta_{2}^{\top}\right\}^{\frac{1}{2}} \\
& \beta^{*}=\beta_{1}+\beta_{2} \Delta_{21} \Delta_{11}^{-1} \\
& \delta^{*}=\delta\left|\Delta_{11}\right|^{\frac{1}{2 \rho}} \\
& \mu^{*}=\mu_{1} \\
& \Delta^{*}=|\Delta|^{-\frac{1}{r}} \Delta_{11}
\end{aligned}
$$

2. The conditional distribution of $X_{2}$ given $X_{1}=x_{1}$ is the $k$-dimensional generalised hyperbolic distribution $\mathrm{GH}_{k}(\tilde{\lambda}, \tilde{\alpha}, \tilde{\beta}, \tilde{\delta}, \tilde{\mu}, \tilde{\Delta})$, where

$$
\begin{aligned}
\tilde{\lambda} & =\lambda-\frac{r}{2} \\
\tilde{\alpha} & =\alpha\left|\Delta_{11}\right|^{\frac{1}{2 k}} \\
\tilde{\beta} & =\beta_{2} \\
\tilde{\delta} & =\left|\Delta_{11}\right|^{-\frac{1}{2 k}}\left\{\delta^{2}+\left(x_{1}-\mu_{1}\right) \Delta_{11}^{-1}\left(x_{1}-\mu_{1}\right)^{\top}\right\}^{\frac{1}{2}} \\
\tilde{\mu} & =\mu_{2}+\left(x_{1}-\mu_{1}\right) \Delta_{11}^{-1} \Delta_{12} \\
\tilde{\Delta} & =\left|\Delta_{11}\right|^{\frac{1}{k}}\left(\Delta_{22}-\Delta_{21} \Delta_{11}^{-1} \Delta_{12}\right)
\end{aligned}
$$

3. Let $Y=X A+B$ be a regular affine transformation of $X$ and let $\|A\|$ denote the absolute value of the determinant of $A$. The distribution of $Y$ is the d-dimensional generalised hyperbolic distribution $\mathrm{GH}_{d}\left(\lambda^{+}, \alpha^{+}, \beta^{+}, \delta^{+}, \mu^{+}, \Delta^{+}\right)$, where

$$
\begin{aligned}
\lambda^{+} & =\lambda \\
\alpha^{+} & =\alpha\|A\|^{-\frac{1}{d}} \\
\beta^{+} & =\beta\left(A^{-1}\right)^{\top} \\
\delta^{+} & =\|A\|^{\frac{1}{d}} \\
\mu^{+} & =\mu A+B \\
\Delta^{+} & =\|A\|^{-\frac{2}{d}} A^{\top} \Delta A
\end{aligned}
$$

## Multivariate t-Distribution

If $X$ and $Y$ are independent and distributed as $N_{p}(\mu, \Sigma)$ and $\mathcal{X}_{n}^{2}$ respectively, and $X \sqrt{n / Y}=t-\mu$, then the pdf of $t$ is given by

$$
\begin{equation*}
f_{t}(t ; n, \Sigma, \mu)=\frac{\Gamma\{(n+p) / 2\}}{\Gamma(n / 2) n^{p / 2} \pi^{p / 2}|\Sigma|^{1 / 2}\left\{1+\frac{1}{n}(t-\mu)^{\top} \Sigma^{-1}(t-\mu)\right\}^{(n+p) / 2}} \tag{4.96}
\end{equation*}
$$

The distribution of $t$ is the noncentral $t$-distribution with $n$ degrees of freedom and the noncentrality parameter $\mu$, Giri (1996).

## Multivariate Laplace Distribution

Let $g$ and $G$ be the pdf and cdf of a $d$-dimensional Gaussian distribution $N_{d}(0, \Sigma)$, the pdf and cdf of a multivariate Laplace distribution can be written as

$$
\begin{align*}
& f_{M \text { Laplace }_{d}}(x ; m, \Sigma)=\int_{0}^{\infty} g\left(z^{-\frac{1}{2}} x-z^{\frac{1}{2}} m\right) z^{-\frac{d}{2}} e^{-z} d z  \tag{4.97}\\
& F_{M \text { Laplace }_{d}}(x, m, \Sigma)=\int_{0}^{\infty} G\left(z^{-\frac{1}{2}} x-z^{\frac{1}{2}} m\right) e^{-z} d z \tag{4.98}
\end{align*}
$$

the pdf can also be described as

$$
\begin{align*}
f_{M \text { Laplace }_{d}}(x ; m, \Sigma)= & \frac{2 e^{x^{\top} \Sigma^{-1} m}}{(2 \pi)^{\frac{d}{2}}|\Sigma|^{\frac{1}{2}}}\left(\frac{x^{\top} \Sigma^{-1} x}{2+m^{\top} \Sigma^{-1} m}\right)^{\frac{\lambda}{2}} \\
& \times K_{\lambda}\left(\sqrt{\left(2+m^{\top} \Sigma^{-1} m\right)\left(x^{\top} \Sigma^{-1} x\right)}\right) \tag{4.99}
\end{align*}
$$

where $\lambda=\frac{2-d}{2}$ and $K_{\lambda}(x)$ is the modified Bessel function of the third kind

$$
\begin{equation*}
K_{\lambda}(x)=\frac{1}{2}\left(\frac{x}{2}\right)^{\lambda} \int_{0}^{\infty} t^{-\lambda-1} e^{-t-\frac{x^{2}}{4 t}} d t, \quad x>0 \tag{4.100}
\end{equation*}
$$

Multivariate Laplace distribution has mean and variance

$$
\begin{align*}
\mathrm{E}[X] & =m  \tag{4.101}\\
\operatorname{Cov}[X] & =\Sigma+m m^{\top} \tag{4.102}
\end{align*}
$$



Fig. 4.13 Tail comparison of $t$-distribution Q MVAtdistail

## Multivariate Mixture Model

A multivariate mixture model comprises multivariate distributions, e.g. the pdf of a multivariate Gaussian distribution can be written as

$$
\begin{equation*}
f(x)=\sum_{l=1}^{L} \frac{w_{l}}{\left|2 \pi \Sigma_{l}\right|^{\frac{1}{2}}} e^{-\frac{1}{2}\left(x-\mu_{l}\right)^{\top} \Sigma^{-1}\left(x-\mu_{l}\right)} \tag{4.103}
\end{equation*}
$$

## Generalised Hyperbolic Distribution

The GH distribution has an exponential decaying speed

$$
\begin{equation*}
f_{\mathrm{GH}}(x ; \lambda, \alpha, \beta, \delta, \mu=0) \sim x^{\lambda-1} e^{-(\alpha-\beta) x} \quad \text { as } \quad x \rightarrow \infty, \tag{4.104}
\end{equation*}
$$

As a comparison to tail behaviour of $t$-distribution depicted in Fig.4.13, the Fig. 4.14 illustrates the tail behaviour of GH distributions with different value of $\lambda$ with $\alpha=1, \beta=0, \delta=1, \mu=0$. It is clear that among the four distributions, GH with $\lambda=1.5$ has the lowest decaying speed, while NIG decays faster.

In Fig. 4.15, Chen, Härdle, and Jeong (2008), four distributions and especially their tail-behaviour are compared. In order to keep the comparability of these distributions, we specified the means to 0 and standardised the variances to 1 . Furthermore we used one important subclass of the GH distribution: the NIG


Fig. 4.14 Tail comparison of GH distribution (pdf) $\mathrm{Q}_{\text {MVAghdistail }}$

Distribution comparison


Tail comparison


Fig. 4.15 Graphical comparison of the NIG distribution (line), standard normal distribution $\mathbf{Q}$ MVAghadatail
distribution with $\lambda=-\frac{1}{2}$ introduced above. On the left panel, the complete forms of these distributions are revealed. The Cauchy (dots) distribution has the lowest peak and the fattest tails. In other words, it has the flattest distribution. The NIG distribution decays second fast in the tails although it has the highest peak, which is more clearly displayed on the right panel.

### 4.7 Copulae

The cumulative distribution function (cdf) of a two-dimensional vector $\left(X_{1}, X_{2}\right)$ is given by

$$
\begin{equation*}
F\left(x_{1}, x_{2}\right)=\mathrm{P}\left(X_{1} \leq x_{1}, X_{2} \leq x_{2}\right) . \tag{4.105}
\end{equation*}
$$

For the case that $X_{1}$ and $X_{2}$ are independent, their joint cumulative distribution function $F\left(x_{1}, x_{2}\right)$ can be written as a product of their one-dimensional marginals:

$$
\begin{equation*}
F\left(x_{1}, x_{2}\right)=F_{X_{1}}\left(x_{1}\right) F_{X_{2}}\left(x_{2}\right)=\mathrm{P}\left(X_{1} \leq x_{1}\right) \mathrm{P}\left(X_{2} \leq x_{2}\right) . \tag{4.106}
\end{equation*}
$$

But how can we model dependence of $X_{1}$ and $X_{2}$ ? Most people would suggest linear correlation. Correlation is though an appropriate measure of dependence only when the random variables have an elliptical or spherical distribution, which include the normal multivariate distribution. Although the terms "correlation" and "dependency" are often used interchangeably, correlation is actually a rather imperfect measure of dependency, and there are many circumstances where correlation should not be used.

Copulae represent an elegant concept of connecting marginals with joint cumulative distribution functions. Copulae are functions that join or "couple" multivariate distribution functions to their one-dimensional marginal distribution functions. Let us consider a $d$-dimensional vector $X=\left(X_{1}, \ldots, X_{d}\right)^{\top}$. Using copulae, the marginal distribution functions $F_{X_{i}}(i=1, \ldots, d)$ can be separately modelled from their dependence structure and then coupled together to form the multivariate distribution $F_{X}$. Copula functions have a long history in probability theory and statistics. Their application in finance is very recent. Copulae are important in Value-at-Risk calculations and constitute an essential tool in quantitative finance (Härdle et al., 2009).

First let us concentrate on the two-dimensional case, then we will extend this concept to the $d$-dimensional case, for a random variable in $\mathbb{R}^{d}$ with $d \geq 1$. To be able to define a copula function, first we need to represent a concept of the volume of a rectangle, a 2 -increading function and a grounded function.

Let $U_{1}$ and $U_{2}$ be two sets in $\overline{\mathbb{R}}=\mathbb{R} \cup\{+\infty\} \cup\{-\infty\}$ and consider the function $F: U_{1} \times U_{2} \longrightarrow \overline{\mathbb{R}}$.

Definition 4.2 The $F$-volume of a rectangle $B=\left[x_{1}, x_{2}\right] \times\left[y_{1}, y_{2}\right] \subset U_{1} \times U_{2}$ is defined as:

$$
\begin{equation*}
V_{F}(B)=F\left(x_{2}, y_{2}\right)-F\left(x_{1}, y_{2}\right)-F\left(x_{2}, y_{1}\right)+F\left(x_{1}, y_{1}\right) \tag{4.107}
\end{equation*}
$$

Definition 4.3 $F$ is said to be a 2-increasing function if for every $B=\left[x_{1}, x_{2}\right] \times$ $\left[y_{1}, y_{2}\right] \subset U_{1} \times U_{2}$,

$$
\begin{equation*}
V_{F}(B) \geq 0 \tag{4.108}
\end{equation*}
$$

Remark 4.2 Note that "to be 2-increasing function" neither implies nor is implied by "to be increasing in each argument".

The following lemmas (Nelsen, 1999) will be very useful later for establishing the continuity of copulae.
Lemma 4.1 Let $U_{1}$ and $U_{2}$ be non-empty sets in $\overline{\mathbb{R}}$ and let $F: U_{1} \times U_{2} \longrightarrow \overline{\mathbb{R}}$ be a two-increasing function. Let $x_{1}, x_{2}$ be in $U_{1}$ with $x_{1} \leq x_{2}$, and $y_{1}, y_{2}$ be in $U_{2}$ with $y_{1} \leq y_{2}$. Then the function $t \mapsto F\left(t, y_{2}\right)-F\left(t, y_{1}\right)$ is non-decreasing on $U_{1}$ and the function $t \mapsto F\left(x_{2}, t\right)-F\left(x_{1}, t\right)$ is non-decreasing on $U_{2}$.

Definition 4.4 If $U_{1}$ and $U_{2}$ have a smallest element $\min U_{1}$ and $\min U_{2}$ respectively, then we say that a function $F: U_{1} \times U_{2} \longrightarrow \mathbb{R}$ is grounded if :

$$
\begin{align*}
& \text { for all } x \in U_{1}: F\left(x, \min U_{2}\right)=0 \text { and }  \tag{4.109}\\
& \text { for all } y \in U_{2}: F\left(\min U_{1}, y\right)=0 \tag{4.110}
\end{align*}
$$

In the following, we will refer to this definition of a cdf.
Definition 4.5 A cdf is a function from $\overline{\mathbb{R}}^{2} \mapsto[0,1]$ which
(i) is grounded
(ii) is 2-increasing
(iii) satisfies $F(\infty, \infty)=1$

Lemma 4.2 Let $U_{1}$ and $U_{2}$ be non-empty sets in $\overline{\mathbb{R}}$ and let $F: U_{1} \times U_{2} \longrightarrow \overline{\mathbb{R}}$ be a grounded two-increasing function. Then $F$ is non-decreasing in each argument.

Definition 4.6 If $U_{1}$ and $U_{2}$ have a greatest element $\max U_{1}$ and max $U_{2}$ respectively, then we say that a function $F: U_{1} \times U_{2} \longrightarrow \mathbb{R}$ has margins and that the margins of $F$ are given by:

$$
\begin{align*}
& F(x)=F\left(x, \max U_{2}\right) \text { for all } x \in U_{1}  \tag{4.111}\\
& F(y)=F\left(\max U_{1}, y\right) \text { for all } y \in U_{2} \tag{4.112}
\end{align*}
$$

Lemma 4.3 Let $U_{1}$ and $U_{2}$ be non-empty sets in $\overline{\mathbb{R}}$ and let $F: U_{1} \times U_{2} \longrightarrow \overline{\mathbb{R}}$ be a grounded two-increasing function which has margins. Let $\left(x_{1}, y_{1}\right),\left(x_{2}, y_{2}\right) \in$ $S_{1} \times S_{2}$. Then

$$
\begin{equation*}
\left|F\left(x_{2}, y_{2}\right)-F\left(x_{1}, y_{1}\right)\right| \leq\left|F\left(x_{2}\right)-F\left(x_{1}\right)\right|+\left|F\left(y_{2}\right)-F\left(y_{1}\right)\right| \tag{4.113}
\end{equation*}
$$

Definition 4.7 A two-dimensional copula is a function $C$ defined on the unit square $I^{2}=I \times I$ with $I=[0,1]$ such that
(i) for every $u \in I$ holds: $C(u, 0)=C(0, v)=0$, i.e. $C$ is grounded.
(ii) for every $u_{1}, u_{2}, v_{1}, v_{2} \in I$ with $u_{1} \leq u_{2}$ and $v_{1} \leq v_{2}$ holds:

$$
\begin{equation*}
C\left(u_{2}, v_{2}\right)-C\left(u_{2}, v_{1}\right)-C\left(u_{1}, v_{2}\right)+C\left(u_{1}, v_{1}\right) \geq 0 \tag{4.114}
\end{equation*}
$$

i.e. $C$ is 2-increasing.
(iii) for every $u \in I$ holds $C(u, 1)=u$ and $C(1, v)=v$.

Informally, a copula is a joint distribution function defined on the unit square $[0,1]^{2}$ which has uniform marginals. That means that if $F_{X_{1}}\left(x_{1}\right)$ and $F_{X_{2}}\left(x_{2}\right)$ are univariate distribution functions, then $C\left\{F_{X_{1}}\left(x_{1}\right), F_{X_{2}}\left(x_{2}\right)\right\}$ is a two-dimensional distribution function with marginals $F_{X_{1}}\left(x_{1}\right)$ and $F_{X_{2}}\left(x_{2}\right)$.

Example 4.25 The functions $\max (u+v-1,0), u v, \min (u, v)$ can be easily checked to be copula functions. They are called respectively the minimum, product and maximum copula.

Example 4.26 Consider the function

$$
\begin{align*}
C_{\rho}^{\text {Gauss }}(u, v) & =\Phi_{\rho}\left\{\Phi^{-1}(u), \Phi^{-1}(v)\right\}  \tag{4.115}\\
& =\int_{-\infty}^{\Phi_{1}^{-1}(u)} \int_{-\infty}^{\Phi_{2}^{-1}(v)} f_{\rho}\left(x_{1}, x_{2}\right) d x_{2} d x_{1}
\end{align*}
$$

where $\Phi_{\rho}$ is the joint two-dimensional standard normal distribution function with correlation coefficient $\rho$, while $\Phi_{1}$ and $\Phi_{2}$ refer to standard normal cdfs and

$$
\begin{equation*}
f_{\rho}\left(x_{1}, x_{2}\right)=\frac{1}{2 \pi \sqrt{1-\rho^{2}}} \exp \left\{-\frac{x_{1}^{2}-2 \rho x_{1} x_{2}+x_{2}^{2}}{2\left(1-\rho^{2}\right)}\right\} \tag{4.116}
\end{equation*}
$$

denotes the bivariate normal pdf.
It is easy to see that $C^{\text {Gauss }}$ is a copula, the so-called Gaussian or normal copula, since it is 2-increasing and

$$
\begin{gather*}
\Phi_{\rho}\left\{\Phi^{-1}(u), \Phi^{-1}(0)\right\}=\Phi_{\rho}\left\{\Phi^{-1}(0), \Phi^{-1}(v)\right\}=0  \tag{4.117}\\
\Phi_{\rho}\left\{\Phi^{-1}(u), \Phi^{-1}(1)\right\}=u \text { and } \Phi_{\rho}\left\{\Phi^{-1}(1), \Phi^{-1}(v)\right\}=v \tag{4.118}
\end{gather*}
$$



Fig. 4.16 Surface plot of the Gumbel-Hougaard copula, $\theta=30$ MVAghsurface

A simple and useful way to represent the graph of a copula is the contour diagram that is, graphs of its level sets-the sets in $I^{2}$ given by $C(u, v)=$ a constant. In Figs. 4.16 and 4.17 we present the countour diagrams of the Gumbel-Hougard copula (Example 4.4) for different values of the copula parameter $\theta$.

For $\theta=1$ the Gumbel-Hougaard copula reduces to the product copula, i.e.

$$
\begin{equation*}
C_{1}(u, v)=\Pi(u, v)=u v \tag{4.119}
\end{equation*}
$$

For $\theta \rightarrow \infty$, one finds for the Gumbel-Hougaard copula:

$$
\begin{equation*}
C_{\theta}(u, v) \longrightarrow \min (u, v)=M(u, v) \tag{4.120}
\end{equation*}
$$

where $M$ is also a copula such that $C(u, v) \leq M(u, v)$ for an arbitrary copula $C$. The copula $M$ is called the Fréchet-Hoeffding upper bound.

The two-dimensional function $W(u, v)=\max (u+v-1,0)$ defines a copula with $W(u, v) \leq C(u, v)$ for any other copula $C$. $W$ is called the Fréchet-Hoeffding lower bound.

In Fig. 4.18 we show an example of Gumbel-Hougaard copula sampling for fixed parameters $\sigma_{1}=1, \sigma_{2}=1$ and $\theta=3$.

One can demonstrate the so-called Fréchet-Hoeffding inequality, which we have already used in Example 1.3, and which states that each copula function is bounded by the minimum and maximum one:

$$
\begin{equation*}
W(u, v)=\max (u+v-1,0) \leq C(u, v) \leq \min (u, v)=M(u, v) \tag{4.121}
\end{equation*}
$$



Fig. 4.17 Contour plots of the Gumbel-Hougard copula $Q$ mVAghcontour

The full relationship between copula and joint cdf depends on Sklar theorem.
Example 4.27 Let us verify that the Gaussian copula satisfies Sklar's theorem in both directions. On the one side, let

$$
\begin{equation*}
F\left(x_{1}, x_{2}\right)=\int_{-\infty}^{x_{1}} \int_{-\infty}^{x_{2}} \frac{1}{2 \pi \sqrt{1-\rho^{2}}} \exp \left\{-\frac{u_{1}^{2}-2 \rho u_{1} u_{2}+u_{2}^{2}}{2\left(1-\rho^{2}\right)}\right\} d u_{2} d u_{1} . \tag{4.122}
\end{equation*}
$$

be a two-dimensional normal distribution function with standard normal cdf's $F_{X_{1}}\left(x_{1}\right)$ and $F_{X_{2}}\left(x_{2}\right)$. Since $F_{X_{1}}\left(x_{1}\right)$ and $F_{X_{2}}\left(x_{2}\right)$ are continuous, a unique copula

Fig. 4.18 10,000-sample output for $\sigma_{1}=1, \sigma_{2}=1$, $\theta=3$ Q MVAsample1000

$C$ exists such that for all $x_{1}, x_{2} \in \overline{\mathbb{R}}^{2}$ a two-dimensional distribution function can be written as a copula in $F_{X_{1}}\left(x_{1}\right)$ and $F_{X_{2}}\left(x_{2}\right)$ :

$$
\begin{equation*}
F\left(x_{1}, x_{2}\right)=C\left\{\Phi_{X_{1}}\left(x_{1}\right), \Phi_{X_{2}}\left(x_{2}\right)\right\} \tag{4.123}
\end{equation*}
$$

The Gaussian copula satisfies the above equality, therefore it is the unique copula mentioned in Sklar's theorem. This proves that the Gaussian copula, together with Gaussian marginals, gives the two-dimensional normal distribution.

Conversely, if $C$ is a copula and $F_{X_{1}}$ and $F_{X_{2}}$ are standard normal distribution functions, then

$$
\begin{align*}
C\left\{F_{X_{1}}\left(x_{1}\right), F_{X_{2}}\left(x_{2}\right)\right\}= & \int_{-\infty}^{\phi_{1}^{-1}\left\{F_{X_{1}}\left(x_{1}\right)\right\}} \int_{-\infty}^{\phi_{2}^{-1}\left\{F_{x_{2}}\left(x_{2}\right)\right\}} \frac{1}{2 \pi \sqrt{1-\rho^{2}}} \\
& \times \exp \left\{-\frac{x_{1}^{2}-2 \rho x_{1} x_{2}+x_{2}^{2}}{2\left(1-\rho^{2}\right)}\right\} d x_{2} d x_{1} \tag{4.124}
\end{align*}
$$

is evidently a joint (two-dimensional) distribution function. Its margins are

$$
\begin{align*}
& C\left\{F_{X_{1}}\left(x_{1}\right), F_{X_{2}}(+\infty)\right\}=\Phi_{\rho}\left[\Phi^{-1}\left\{F_{X_{1}}\left(x_{1}\right)\right\},+\infty\right]=F_{X_{1}}\left(x_{1}\right)  \tag{4.125}\\
& C\left\{F_{X_{1}}(+\infty), F_{X_{2}}\left(x_{2}\right)\right\}=\Phi_{\rho}\left[+\infty, \Phi^{-1}\left\{F_{X_{2}}\left(x_{2}\right)\right\}\right]=F_{X_{2}}\left(x_{2}\right) \tag{4.126}
\end{align*}
$$

The following proposition shows one attractive feature of the copula representation of dependence, i.e. that the dependence structure described by a copula
is invariant under increasing and continuous transformations of the marginal distributions.

Theorem 4.13 If $\left(X_{1}, X_{2}\right)$ have copula $C$ and set $g_{1}, g_{2}$ two continuously increasing functions, then $\left\{g_{1}\left(X_{1}\right), g_{2}\left(X_{2}\right)\right\}$ have the copula $C$, too.

Example 4.28 Independence implies that the product of the cdf's $F_{X_{1}}$ and $F_{X_{2}}$ equals the joint distribution function $F$, i.e.:

$$
\begin{equation*}
F\left(x_{1}, x_{2}\right)=F_{X_{1}}\left(x_{1}\right) F_{X_{2}}\left(x_{2}\right) . \tag{4.127}
\end{equation*}
$$

Thus, we obtain the independence or product copula $C=\Pi(u, v)=u v$.
While it is easily understood how a product copula describes an independence relationship, the converse is also true. Namely, the joint distribution function of two independent random variables can be interpreted as a product copula. This concept is formalised in the following theorem:

Theorem 4.14 Let $X_{1}$ and $X_{2}$ be random variables with continuous distribution functions $F_{X_{1}}$ and $F_{X_{2}}$ and the joint distribution function $F$. Then $X_{1}$ and $X_{2}$ are independent if and only if $C_{X_{1}, X_{2}}=\Pi$.
Example 4.29 Let us consider the Gaussian copula for the case $\rho=0$, i.e. vanishing correlation. In this case the Gaussian copula becomes

$$
\begin{align*}
C_{0}^{\text {Gauss }}(u, v) & =\int_{-\infty}^{\Phi_{1}^{-1}(u)} \varphi\left(x_{1}\right) d x_{1} \int_{-\infty}^{\Phi_{2}^{-1}(v)} \varphi\left(x_{2}\right) d x_{2} \\
& =u v  \tag{4.128}\\
& =\Pi(u, v)
\end{align*}
$$

The following theorem, which follows directly from Lemma 4.3, establishes the continuity of copulae .

Theorem 4.15 Let $C$ be a copula. Then for any $u_{1}, v_{1}, u_{2}, v_{2} \in I$ holds

$$
\begin{equation*}
\left|C\left(u_{2}, v_{2}\right)-C\left(u_{1}, v_{1}\right)\right| \leq\left|u_{2}-u_{1}\right|+\left|v_{2}-v_{1}\right| \tag{4.129}
\end{equation*}
$$

From (4.129) it follows that every copula $C$ is uniformly continuous on its domain.

A further important property of copulae concerns the partial derivatives of a copula with respect to its variables:
Theorem 4.16 Let $C(u, v)$ be a copula. For any $u \in I$, the partial derivative $\frac{\partial C(u, v)}{\partial v}$ exists for almost all $u \in I$. For such $u$ and $v$ one has:

$$
\begin{equation*}
\frac{\partial C(u, v)}{\partial v} \in I \tag{4.130}
\end{equation*}
$$

The analogous statement is true for the partial derivative $\frac{\partial C(u, v)}{\partial u}$ :

$$
\begin{equation*}
\frac{\partial C(u, v)}{\partial u} \in I \tag{4.131}
\end{equation*}
$$

Moreover, the functions

$$
\begin{aligned}
u & \mapsto C_{v}(u) \stackrel{\text { def }}{=} \partial C(u, v) / \partial v \text { and } \\
v & \mapsto C_{u}(v) \stackrel{\text { def }}{=} \partial C(u, v) / \partial u
\end{aligned}
$$

are defined and non-increasing almost everywhere on I.
Until now, we have considered copulae only in a two-dimensional setting. Let us now extend this concept to the $d$-dimensional case, for a random variable in $\mathbb{R}^{d}$ with $d \geq 1$.

Let $U_{1}, U_{2}, \ldots, U_{d}$ be non-empty sets in $\overline{\mathbb{R}}$ and consider the function $F: U_{1} \times$ $U_{2} \times \cdots \times U_{d} \longrightarrow \overline{\mathbb{R}}$. For $a=\left(a_{1}, a_{2}, \ldots, a_{d}\right)$ and $b=\left(b_{1}, b_{2}, \ldots, b_{d}\right)$ with $a \leq b$ (i.e. $a_{k} \leq b_{k}$ for all $k$ ) let $B=[a, b]=\left[a_{1}, b_{1}\right] \times\left[a_{2}, b_{2}\right] \times \cdots \times\left[a_{n}, b_{n}\right]$ be the $d$-box with vertices $c=\left(c_{1}, c_{2}, \ldots, c_{d}\right)$. It is obvious that each $c_{k}$ is either equal to $a_{k}$ or to $b_{k}$.

Definition 4.8 The $F$-volume of a $d$-box $B=[a, b]=\left[a_{1}, b_{1}\right] \times\left[a_{2}, b_{2}\right] \times \cdots \times$ $\left[a_{d}, b_{d}\right] \subset U_{1} \times U_{2} \times \cdots \times U_{d}$ is defined as follows:

$$
\begin{equation*}
V_{F}(B)=\sum_{k=1}^{d} \operatorname{sign}\left(c_{k}\right) F\left(c_{k}\right) \tag{4.132}
\end{equation*}
$$

where $\operatorname{sign}\left(c_{k}\right)=1$, if $c_{k}=a_{k}$ for even $k$ and $\operatorname{sign}\left(c_{k}\right)=-1$, if $c_{k}=a_{k}$ for odd $k$.
Example 4.30 For the case $d=3$, the $F$-volume of a 3-box $B=[a, b]=\left[x_{1}, x_{2}\right] \times$ $\left[y_{1}, y_{2}\right] \times\left[z_{1}, z_{2}\right]$ is defined as:

$$
\begin{aligned}
V_{F}(B)= & F\left(x_{2}, y_{2}, z_{2}\right)-F\left(x_{2}, y_{2}, z_{1}\right)-F\left(x_{2}, y_{1}, z_{2}\right)-F\left(x_{1}, y_{2}, z_{2}\right) \\
& +F\left(x_{2}, y_{1}, z_{1}\right)+F\left(x_{1}, y_{2}, z_{1}\right)+F\left(x_{1}, y_{1}, z_{2}\right)-F\left(x_{1}, y_{1}, z_{1}\right)
\end{aligned}
$$

Definition 4.9 $F$ is said to be a $d$-increasing function if for all $d$-boxes $B$ with vertices in $U_{1} \times U_{2} \times \cdots \times U_{d}$ holds:

$$
\begin{equation*}
V_{F}(B) \geq 0 \tag{4.133}
\end{equation*}
$$

Definition 4.10 If $U_{1}, U_{2}, \ldots, U_{d}$ have a smallest element $\min U_{1}, \min U_{2}, \ldots$, $\min U_{d}$ respectively, then we say that a function $F: U_{1} \times U_{2} \times \cdots \times U_{d} \longrightarrow \overline{\mathbb{R}}$ is grounded if :

$$
\begin{equation*}
F(x)=0 \text { for all } x \in U_{1} \times U_{2} \times \cdots \times U_{d} \tag{4.134}
\end{equation*}
$$

such that $x_{k}=\min U_{k}$ for at least one k .
The lemmas, which we presented for the two-dimensional case, have analogous multivariate versions, see Nelsen (1999).

Definition 4.11 A $d$-dimensional copula (or $d$-copula) is a function $C$ defined on the unit $d$-cube $I^{d}=I \times I \times \cdots \times I$ such that
(i) for every $u \in I^{d}$ holds: $C(u)=0$, if at least one coordinate of $u$ is equal to 0 ; i.e. $C$ is grounded.
(ii) for every $a, b \in I^{d}$ with $a \leq b$ holds:

$$
\begin{equation*}
V_{C}([a, b]) \geq 0 ; \tag{4.135}
\end{equation*}
$$

i.e. $C$ is 2 -increasing.
(iii) for every $u \in I^{d}$ holds: $C(u)=u_{k}$, if all coordinates of $u$ are 1 except $u_{k}$.

Analogously to the two-dimensional setting, let us state the Sklar's theorem for the $d$-dimensional case.

Theorem 4.17 (Sklar's Theorem in $d$-Dimensional Case) Let $F$ be a ddimensional distribution function with marginal distribution functions
$F_{X_{1}}, F_{X_{2}}, \ldots, F_{X_{d}}$. Then a d-copula $C$ exists such that for all $x_{1}, \ldots, x_{d} \in \overline{\mathbb{R}}^{d}$ :

$$
\begin{equation*}
F\left(x_{1}, x_{2}, \ldots, x_{d}\right)=C\left\{F_{X_{1}}\left(x_{1}\right), F_{X_{2}}\left(x_{2}\right), \ldots, F_{X_{d}}\left(x_{d}\right)\right\} \tag{4.136}
\end{equation*}
$$

Moreover, if $F_{X_{1}}, F_{X_{2}}, \ldots, F_{X_{d}}$ are continuous then $C$ is unique. Otherwise $C$ is uniquely determined on the Cartesian product $\operatorname{Im}\left(F_{X_{1}}\right) \times \operatorname{Im}\left(F_{X_{2}}\right) \times \cdots \times \operatorname{Im}\left(F_{X_{d}}\right)$.

Conversely, if $C$ is a copula and $F_{X_{1}}, F_{X_{2}}, \ldots, F_{X_{d}}$ are distribution functions then $F$ defined by (4.136) is a d-dimensional distribution function with marginals $F_{X_{1}}, F_{X_{2}}, \ldots, F_{X_{d}}$.

In order to illustrate the $d$-copulae we present the following examples:
Example 4.31 Let $\Phi$ denote the univariate standard normal distribution function and $\Phi_{\Sigma, d}$ the $d$-dimensional standard normal distribution function with correlation matrix $\Sigma$. Then the function

$$
\begin{align*}
C_{\rho}^{\text {Gauss }}(u, \Sigma) & =\Phi_{\Sigma, d}\left\{\Phi^{-1}\left(u_{1}\right), \ldots, \Phi^{-1}\left(u_{d}\right)\right\} \\
& =\int_{-\infty}^{\phi_{1}^{-1}\left(u_{d}\right)} \cdots \int_{-\infty}^{\phi_{2}^{-1}\left(u_{1}\right)} f_{\Sigma}\left(x_{1}, \ldots, x_{n}\right) d x_{1} \ldots d x_{d} \tag{4.137}
\end{align*}
$$

is the $d$-dimensional Gaussian or normal copula with correlation matrix $\Sigma$. The function

$$
\begin{align*}
& f_{\rho}\left(x_{1}, \ldots, x_{d}\right)=\frac{1}{\sqrt{\operatorname{det}(\Sigma)}}  \tag{4.138}\\
& \quad \times \exp \left\{-\frac{\left(\Phi^{-1}\left(u_{1}\right), \ldots, \Phi^{-1}\left(u_{d}\right)\right)^{\top}\left(\Sigma^{-1}-\mathcal{I}_{d}\right)\left(\Phi^{-1}\left(u_{1}\right), \ldots, \Phi^{-1}\left(u_{d}\right)\right)}{2}\right\}
\end{align*}
$$

is a copula density function. The copula dependence parameter $\alpha$ is the collection of all unknown correlation coefficients in $\Sigma$. If $\alpha \neq 0$, then the corresponding normal copula allows to generate joint symmetric dependence. However, it is not possible to model a tail dependence, i.e. joint extreme events have a zero probability.

Example 4.32 Let us consider the following function

$$
\begin{equation*}
C_{\theta}^{\mathrm{GH}}\left(u_{1}, \ldots, u_{d}\right)=\exp \left[-\left\{\sum_{j=1}^{d}\left(-\log u_{j}\right)^{\theta}\right\}^{1 / \theta}\right] \tag{4.139}
\end{equation*}
$$

One recognise this function is as the $d$-dimensional Gumbel-Hougaard copula function. Unlike the Gaussian copula, the copula (4.139) can generate an upper tail dependence.

Example 4.33 As in the two-dimensional setting, let us consider the $d$-dimensional Gumbel-Hougaard copula for the case $\theta=1$. In this case the Gumbel-Hougaard copula reduces to the $d$-dimensional product copula, i.e.

$$
\begin{equation*}
C_{1}\left(u_{1}, \ldots, u_{d}\right)=\prod_{j=1}^{d} u_{j}=\Pi^{d}(u) \tag{4.140}
\end{equation*}
$$

The extension of the two-dimensional copula $M$, which one gets from the $d$ dimensional Gumbel-Hougaard copula for $\theta \rightarrow \infty$ is denoted $M^{d}(u)$ :

$$
\begin{equation*}
C_{\theta}\left(u_{1}, \ldots u_{d}\right) \longrightarrow \min \left(u_{1}, \ldots, u_{d}\right)=M^{d}(u) \tag{4.141}
\end{equation*}
$$

The $d$-dimensional function

$$
\begin{equation*}
W^{d}(u)=\max \left(u_{1}+u_{2}+\cdots+u_{d}-d+1,0\right) \tag{4.142}
\end{equation*}
$$

defines a copula with $W(u) \leq C(u)$ for any other $d$-dimensional copula function $C(u) . W^{d}(u)$ is the Fréchet-Hoeffding lower bound in the $d$-dimensional case.

The functions $M^{d}$ and $\Pi^{d}$ are $d$-copulae for all $d \geq 2$, whereas the function $W^{d}$ fails to be a $d$-copula for any $d>2$ (Nelsen, 1999). However, the $d$-dimensional version of the Fréchet-Hoeffding inequality can be written as follows:

$$
\begin{equation*}
W^{d}(u) \leq C(u) \leq M^{d}(u) \tag{4.143}
\end{equation*}
$$

As we have already mentioned, copula functions have been widely applied in empirical finance.

|  | $\hookrightarrow$ |
| :--- | :--- |
|  | The cumulative distribution function (cdf) is defined as $F(x)=$ |
|  | $\mathrm{P}(X<x)$. |

### 4.8 Bootstrap

Recall that we need large sample sizes in order to sufficiently approximate the critical values computable by the CLT. Here large means $n>50$ for onedimensional data. How can we construct confidence intervals in the case of smaller sample sizes? One way is to use a method called the Bootstrap. The Bootstrap algorithm uses the data twice:

1. estimate the parameter of interest,
2. simulate from an estimated distribution to approximate the asymptotic distribution of the statistics of interest.

In detail, bootstrap works as follows. Consider the observations $x_{1}, \ldots, x_{n}$ of the sample $X_{1}, \ldots, X_{n}$ and estimate the empirical distribution function (EDF) $F_{n}$. In the case of one-dimensional data

$$
\begin{equation*}
F_{n}(x)=\frac{1}{n} \sum_{i=1}^{n} \boldsymbol{I}\left(X_{i} \leq x\right) \tag{4.144}
\end{equation*}
$$

This is a step function which is constant between neighbouring data points.


Fig. 4.19 The standard normal cdf (thick line) and the empirical distribution function (thin line) for $n=100$ Q MVAedfnormal


Fig. 4.20 The standard normal cdf (thick line) and the empirical distribution function (thin line) for $n=1,000$ O MVAedfnormal

Example 4.34 Suppose that we have $n=100$ standard normal $N(0,1)$ data points $X_{i}, i=1, \ldots, n$. The cdf of $X$ is $\Phi(x)=\int_{-\infty}^{x} \varphi(u) d u$ and is shown in Fig. 4.19 as the thin, solid line. The EDF is displayed as a thick step function line. Figure 4.20 shows the same setup for $n=1,000$ observations.

Now draw with replacement a new sample from this empirical distribution. That is we sample with replacement $n^{*}$ observations $X_{1}^{*}, \ldots, X_{n^{*}}^{*}$ from the original sample. This is called a Bootstrap sample. Usually one takes $n^{*}=n$.

Since we sample with replacement, a single observation from the original sample may appear several times in the Bootstrap sample. For instance, if the original sample consists of the three observations $x_{1}, x_{2}, x_{3}$, then a Bootstrap sample might look like $X_{1}^{*}=x_{3}, X_{2}^{*}=x_{2}, X_{3}^{*}=x_{3}$. Computationally, we find the Bootstrap sample by using a uniform random number generator to draw from the indices $1,2, \ldots, n$ of the original samples.

The Bootstrap observations are drawn randomly from the empirical distribution, i.e. the probability for each original observation to be selected into the Bootstrap sample is $1 / n$ for each draw. It is easy to compute that

$$
\mathrm{E}_{F_{n}}\left(X_{i}^{*}\right)=\frac{1}{n} \sum_{i=1}^{n} x_{i}=\bar{x} .
$$

This is the expected value given that the cdf is the original mean of the sample $x_{1}, \ldots, x_{n}$. The same holds for the variance, i.e.

$$
V_{F_{n}}\left(X_{i}^{*}\right)=\widehat{\sigma}^{2},
$$

where $\widehat{\sigma}^{2}=n^{-1} \sum\left(x_{i}-\bar{x}\right)^{2}$. The cdf of the bootstrap observations is defined as in (4.144). Figure 4.21 shows the cdf of the $n=100$ original observations as a solid line and two bootstrap cdf's as thin lines.

The CLT holds for the bootstrap sample. Analogously to Corollary 4.1 we have the following corollary.

Fig. 4.21 The $\operatorname{cdf} F_{n}$ (thick line) and two bootstrap cdf's $F_{n}^{*}$ (thin lines) O MVAedfbootstrap


Corollary 4.2 If $X_{1}^{*}, \ldots, X_{n}^{*}$ is a bootstrap sample from $X_{1}, \ldots, X_{n}$, then the distribution of

$$
\sqrt{n}\left(\frac{\bar{x}^{*}-\bar{x}}{\widehat{\sigma}^{*}}\right)
$$

also becomes $N(0,1)$ asymptotically, where $\bar{x}^{*}=n^{-1} \sum_{i=1}^{n} X_{i}^{*}$ and $\left(\widehat{\sigma}^{*}\right)^{2}=$ $n^{-1} \sum_{i=1}^{n}\left(X_{i}^{*}-\bar{x}^{*}\right)^{2}$.

How do we find a confidence interval for $\mu$ using the Bootstrap method? Recall that the quantile $u_{1-\alpha / 2}$ might be bad for small sample sizes because the true distribution of $\sqrt{n}\left(\frac{\bar{\alpha}-\mu}{\hat{\sigma}}\right)$ might be far away from the limit distribution $N(0,1)$. The Bootstrap idea enables us to "simulate" this distribution by computing $\sqrt{n}\left(\frac{\bar{x}^{*}-\bar{x}}{\widehat{\sigma}^{*}}\right)$ for many Bootstrap samples. In this way we can estimate an empirical ( $1-\alpha / 2$ )quantile $u_{1-\alpha / 2}^{*}$. The bootstrap improved confidence interval is then

$$
C_{1-\alpha}^{*}=\left[\bar{x}-\frac{\hat{\sigma}}{\sqrt{n}} u_{1-\alpha / 2}^{*}, \bar{x}+\frac{\hat{\sigma}}{\sqrt{n}} u_{1-\alpha / 2}^{*}\right] .
$$

By Corollary 4.2 we have

$$
P\left(\mu \in C_{1-\alpha}^{*}\right) \longrightarrow 1-\alpha \quad \text { as } n \rightarrow \infty,
$$

but with an improved speed of convergence, see Hall (1992).

|  | Summary |
| :--- | :--- |$\quad$| For small sample sizes the bootstrap improves the precision of the |
| :--- |
| confidence interval. |

### 4.9 Exercises

Exercise 4.1 Assume that the random vector $Y$ has the following normal distribution: $Y \sim N_{p}(0, \mathcal{I})$. Transform it according to (4.49) to create $X \sim N(\mu, \Sigma)$ with mean $\mu=(3,2)^{\top}$ and $\Sigma=\left(\begin{array}{cc}1 & -1.5 \\ -1.5 & 4\end{array}\right)$. How would you implement the resulting formula on a computer?

Exercise 4.2 Prove Theorem 4.7 using Theorem 4.5.
Exercise 4.3 Suppose that $X$ has mean zero and covariance $\Sigma=\left(\begin{array}{ll}1 & 0 \\ 0 & 2\end{array}\right)$. Let $Y=$ $X_{1}+X_{2}$. Write $Y$ as a linear transformation, i.e. find the transformation matrix $\mathcal{A}$. Then compute $\operatorname{Var}(Y)$ via (4.26). Can you obtain the result in another fashion?
Exercise 4.4 Calculate the mean and the variance of the estimate $\hat{\beta}$ in (3.50).
Exercise 4.5 Compute the conditional moments $\mathrm{E}\left(X_{2} \mid x_{1}\right)$ and $\mathrm{E}\left(X_{1} \mid x_{2}\right)$ for the pdf of Example 4.5.

Exercise 4.6 Prove the relation (4.28).
Exercise 4.7 Prove the relation (4.29).
Hint: Note that $\left.\operatorname{Var}\left(\mathrm{E}\left(X_{2} \mid X_{1}\right)\right)=\mathrm{E}\left(\mathrm{E}\left(X_{2} \mid X_{1}\right) \mathrm{E}\left(X_{2}^{\top} \mid X_{1}\right)\right)-\mathrm{E}\left(X_{2}\right) \mathrm{E}\left(X_{2}^{\top}\right)\right)$ and that
$\mathrm{E}\left(\operatorname{Var}\left(X_{2} \mid X_{1}\right)\right)=\mathrm{E}\left[\mathrm{E}\left(X_{2} X_{2}^{\top} \mid X_{1}\right)-\mathrm{E}\left(X_{2} \mid X_{1}\right) \mathrm{E}\left(X_{2}^{\top} \mid X_{1}\right)\right]$.
Exercise 4.8 Compute (4.46) for the pdf of Example 4.5.

## Exercise 4.9

Show that $\quad f_{Y}(y)=\left\{\begin{array}{ll}\frac{1}{2} y_{1}-\frac{1}{4} y_{2} & 0 \leq y_{1} \leq 2, \quad\left|y_{2}\right| \leq 1-\left|1-y_{1}\right| \\ 0 & \text { otherwise }\end{array} \quad\right.$ is a $p d f$.
Exercise 4.10 Compute (4.46) for a two-dimensional standard normal distribution. Show that the transformed random variables $Y_{1}$ and $Y_{2}$ are independent. Give a geometrical interpretation of this result based on iso-distance curves.

Exercise 4.11 Consider the Cauchy distribution which has no moment, so that the CLT cannot be applied. Simulate the distribution of $\bar{x}$ (for different $n$ 's). What can you expect for $n \rightarrow \infty$ ?

Hint: The Cauchy distribution can be simulated by the quotient of two independent standard normally distributed random variables.

Exercise 4.12 A European car company has tested a new model and reports the consumption of petrol $\left(X_{1}\right)$ and oil $\left(X_{2}\right)$. The expected consumption of petrol is $8 l$ per $100 \mathrm{~km}\left(\mu_{1}\right)$ and the expected consumption of oil is 11 per $10,000 \mathrm{~km}\left(\mu_{2}\right)$. The measured consumption of petrol is $8.11 \mathrm{per} 100 \mathrm{~km}\left(\bar{x}_{1}\right)$ and the measured consumption of oil is 1.11 per $10,000 \mathrm{~km}\left(\bar{x}_{2}\right)$. The asymptotic distribution of $\sqrt{n}\left\{\binom{\overline{x_{1}}}{\overline{x_{2}}}-\binom{\mu_{1}}{\mu_{2}}\right\}$ is $\left.N\binom{0}{0},\left(\begin{array}{cc}0.1 & 0.05 \\ 0.05 & 0.1\end{array}\right)\right)$.

For the American market the basic measuring units are miles ( 1 mile $\approx 1.6 \mathrm{~km}$ ) and gallons ( 1 gallon $\approx 3.8 l$ ). The consumptions of petrol $\left(Y_{1}\right)$ and oil $\left(Y_{2}\right)$ are usually reported in miles per gallon. Can you express $\bar{y}_{1}$ and $\bar{y}_{2}$ in terms of $\bar{x}_{1}$ and $\bar{x}_{2}$ ? Recompute the asymptotic distribution for the American market.
Exercise 4.13 Consider the pdf $f\left(x_{1}, x_{2}\right)=e^{-\left(x_{1}+x_{2}\right)}, x_{1}, x_{2}>0$ and let $U_{1}=$ $X_{1}+X_{2}$ and $U_{2}=X_{1}-X_{2}$. Compute $f\left(u_{1}, u_{2}\right)$.

## Exercise 4.14 Consider the pdf's

$$
\begin{array}{ll}
f\left(x_{1}, x_{2}\right)=4 x_{1} x_{2} e^{-x_{1}^{2}} & x_{1}, x_{2}>0 \\
f\left(x_{1}, x_{2}\right)=1 & 0<x_{1}, x_{2}<1 \text { and } x_{1}+x_{2}<1 \\
f\left(x_{1}, x_{2}\right)=\frac{1}{2} e^{-x_{1}} & x_{1}>\left|x_{2}\right| .
\end{array}
$$

For each of these pdf's compute $\mathrm{E}(X), \operatorname{Var}(X), \mathrm{E}\left(X_{1} \mid X_{2}\right), \mathrm{E}\left(X_{2} \mid X_{1}\right), \operatorname{Var}\left(X_{1} \mid X_{2}\right)$ and $\operatorname{Var}\left(X_{2} \mid X_{1}\right)$.
Exercise 4.15 Consider the pdf $f\left(x_{1}, x_{2}\right)=\frac{3}{4} x_{1}^{-\frac{1}{2}}, 0<x_{1}<x_{2}<1$. Compute $P\left(X_{1}<0.25\right), P\left(X_{2}<0.25\right)$ and $P\left(X_{2}<0.25 \mid X_{1}<0.25\right)$.
Exercise 4.16 Consider the pdf $f\left(x_{1}, x_{2}\right)=\frac{1}{2 \pi}, 0<x_{1}<2 \pi, 0<x_{2}<1$.
Let $U_{1}=\sin X_{1} \sqrt{-2 \log X_{2}}$ and $U_{2}=\cos X_{1} \sqrt{-2 \log X_{2}}$. Compute $f\left(u_{1}, u_{2}\right)$.
Exercise 4.17 Consider $f\left(x_{1}, x_{2}, x_{3}\right)=k\left(x_{1}+x_{2} x_{3}\right) ; 0<x_{1}, x_{2}, x_{3}<1$.
(a) Determine $k$ so that $f$ is a valid pdf of $\left(X_{1}, X_{2}, X_{3}\right)=X$.
(b) Compute the $(3 \times 3)$ matrix $\Sigma_{X}$.
(c) Compute the $(2 \times 2)$ matrix of the conditional variance of $\left(X_{2}, X_{3}\right)$ given $X_{1}=x_{1}$.
Exercise 4.18 Let $X \sim N_{2}\left(\binom{1}{2},\left(\begin{array}{cc}2 & a \\ a & 2\end{array}\right)\right)$.
(a) Represent the contour ellipses for $a=0 ;-\frac{1}{2} ;+\frac{1}{2} ; 1$.
(b) For $a=\frac{1}{2}$ find the regions of $X$ centred on $\mu$ which cover the area of the true parameter with probability 0.90 and 0.95 .

Exercise 4.19 Consider the pdf

$$
f\left(x_{1}, x_{2}\right)=\frac{1}{8 x_{2}} e^{-\left(\frac{x_{1}}{2 x_{2}}+\frac{x_{2}}{4}\right)} \quad x_{1}, x_{2}>0 .
$$

Compute $f\left(x_{2}\right)$ and $f\left(x_{1} \mid x_{2}\right)$. Also give the best approximation of $X_{1}$ by a function of $X_{2}$. Compute the variance of the error of the approximation.

## Exercise 4.20 Prove Theorem 4.6.

## Chapter 5 <br> Theory of the Multinormal

In the preceding chapter we saw how the multivariate normal distribution comes into play in many applications. It is useful to know more about this distribution, since it is often a good approximate distribution in many situations. Another reason for considering the multinormal distribution relies on the fact that it has many appealing properties: it is stable under linear transforms, zero correlation corresponds to independence, the marginals and all the conditionals are also multivariate normal variates, etc. The mathematical properties of the multinormal make analyses much simpler.

In this chapter we will first concentrate on the probabilistic properties of the multinormal, then we will introduce two "companion" distributions of the multinormal which naturally appear when sampling from a multivariate normal population: the Wishart and the Hotelling distributions. The latter is particularly important for most of the testing procedures proposed in Chap. 7.

### 5.1 Elementary Properties of the Multinormal

Let us first summarise some properties which were already derived in the previous chapter.

- The pdf of $X \sim N_{p}(\mu, \Sigma)$ is

$$
\begin{equation*}
f(x)=|2 \pi \Sigma|^{-1 / 2} \exp \left\{-\frac{1}{2}(x-\mu)^{\top} \Sigma^{-1}(x-\mu)\right\} . \tag{5.1}
\end{equation*}
$$

The expectation is $\mathrm{E}(X)=\mu$, the covariance can be calculated as
$\operatorname{Var}(X)=\mathrm{E}(X-\mu)(X-\mu)^{\top}=\Sigma$.

- Linear transformations turn normal random variables into normal random variables. If $X \sim N_{p}(\mu, \Sigma)$ and $\mathcal{A}(p \times p), c \in \mathbb{R}^{p}$, then $Y=\mathcal{A} X+c$ is $p$-variate Normal, i.e.

$$
\begin{equation*}
Y \sim N_{p}\left(\mathcal{A} \mu+c, \mathcal{A} \Sigma \mathcal{A}^{\top}\right) \tag{5.2}
\end{equation*}
$$

- If $X \sim N_{p}(\mu, \Sigma)$, then the Mahalanobis transformation is

$$
\begin{equation*}
Y=\Sigma^{-1 / 2}(X-\mu) \sim N_{p}\left(0, \mathcal{I}_{p}\right) \tag{5.3}
\end{equation*}
$$

and it holds that

$$
\begin{equation*}
Y^{\top} Y=(X-\mu)^{\top} \Sigma^{-1}(X-\mu) \sim \chi_{p}^{2} \tag{5.4}
\end{equation*}
$$

Often it is interesting to partition $X$ into sub-vectors $X_{1}$ and $X_{2}$. The following theorem tells us how to correct $X_{2}$ to obtain a vector which is independent of $X_{1}$.
Theorem 5.1 Let $X=\binom{X_{1}}{X_{2}} \sim N_{p}(\mu, \Sigma), X_{1} \in \mathbb{R}^{r}, X_{2} \in \mathbb{R}^{p-r}$. Define $X_{2.1}=$ $X_{2}-\Sigma_{21} \Sigma_{11}^{-1} X_{1}$ from the partitioned covariance matrix

$$
\Sigma=\left(\begin{array}{ll}
\Sigma_{11} & \Sigma_{12} \\
\Sigma_{21} & \Sigma_{22}
\end{array}\right)
$$

Then

$$
\begin{gather*}
X_{1} \sim N_{r}\left(\mu_{1}, \Sigma_{11}\right)  \tag{5.5}\\
X_{2.1} \sim N_{p-r}\left(\mu_{2.1}, \Sigma_{22.1}\right) \tag{5.6}
\end{gather*}
$$

are independent with

$$
\begin{equation*}
\mu_{2.1}=\mu_{2}-\Sigma_{21} \Sigma_{11}^{-1} \mu_{1}, \quad \Sigma_{22.1}=\Sigma_{22}-\Sigma_{21} \Sigma_{11}^{-1} \Sigma_{12} \tag{5.7}
\end{equation*}
$$

Proof

$$
\begin{aligned}
X_{1} & =\mathcal{A} X \quad \text { with } \quad \mathcal{A}=\left(\mathcal{I}_{r}, 0\right) \\
X_{2.1} & =\mathcal{B} X \quad \text { with } \quad \mathcal{B}=\left(-\Sigma_{21} \Sigma_{11}^{-1}, \mathcal{I}_{p-r}\right)
\end{aligned}
$$

Then, by (5.2) $X_{1}$ and $X_{2.1}$ are both normal. Note that

$$
\begin{aligned}
\operatorname{Cov}\left(X_{1}, X_{2.1}\right)=\mathcal{A} \Sigma \mathcal{B}^{\top} & =\left(\begin{array}{cc}
\begin{array}{|cc|}
\hline 1 & 0 \\
& \ddots \\
0 & \\
\hline
\end{array} & 1
\end{array}\right)\left(\begin{array}{cc}
\Sigma_{11} & \Sigma_{12} \\
\Sigma_{21} & \Sigma_{22}
\end{array}\right)\binom{\left(-\Sigma_{21} \Sigma_{11}^{-1}\right)^{\top}}{\begin{array}{|cc}
1 & \\
& 0 \\
0 & 1
\end{array}}, \\
\mathcal{A} \Sigma & =\left(\begin{array}{ll}
\mathcal{I}_{r} & 0
\end{array}\right)\left(\begin{array}{cc}
\Sigma_{11} & \Sigma_{12} \\
\Sigma_{21} & \Sigma_{22}
\end{array}\right)=\left(\begin{array}{ll}
\Sigma_{11} & \Sigma_{12}
\end{array}\right) \\
\text { hence, } \mathcal{A} \Sigma \mathcal{B}^{\top} & =\left(\begin{array}{ll}
\Sigma_{11} & \Sigma_{12}
\end{array}\right)\binom{\left(-\Sigma_{21} \Sigma_{11}^{-1}\right.}{\mathcal{I}_{p-r}} \\
& =\left(\begin{array}{ll}
-\Sigma_{11}\left(\Sigma_{21} \Sigma_{11}^{-1}\right)^{\top}+\Sigma_{12}
\end{array}\right) .
\end{aligned}
$$

Recall that $\Sigma_{21}=\left(\Sigma_{12}\right)^{\top}$. Hence $\mathcal{A} \Sigma \mathcal{B}^{\top}=-\Sigma_{11} \Sigma_{11}^{-1} \Sigma_{12}+\Sigma_{12} \equiv 0$.
Using (5.2) again we also have the joint distribution of ( $X_{1}, X_{2.1}$ ), namely

$$
\binom{X_{1}}{X_{2.1}}=\binom{\mathcal{A}}{\mathcal{B}} X \sim N_{p}\left(\binom{\mu_{1}}{\mu_{2.1}},\left(\begin{array}{cc}
\Sigma_{11} & 0 \\
0 & \Sigma_{22.1}
\end{array}\right)\right)
$$

With this block diagonal structure of the covariance matrix, the joint pdf of ( $X_{1}, X_{2.1}$ ) can easily be factorised into

$$
\begin{aligned}
f\left(x_{1}, x_{2.1}\right)= & \left|2 \pi \Sigma_{11}\right|^{-\frac{1}{2}} \exp \left\{-\frac{1}{2}\left(x_{1}-\mu_{1}\right)^{\top} \Sigma_{11}^{-1}\left(x_{1}-\mu_{1}\right)\right\} \\
& \times\left|2 \pi \Sigma_{22.1}\right|^{-\frac{1}{2}} \exp \left\{-\frac{1}{2}\left(x_{2.1}-\mu_{2.1}\right)^{\top} \Sigma_{22.1}^{-1}\left(x_{2.1}-\mu_{2.1}\right)\right\}
\end{aligned}
$$

from which the independence between $X_{1}$ and $X_{2.1}$ follows.
The next two corollaries are direct consequences of Theorem 5.1.
Corollary 5.1 Let $X=\binom{X_{1}}{X_{2}} \sim N_{p}(\mu, \Sigma), \Sigma=\left(\begin{array}{cc}\Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22}\end{array}\right) . \Sigma_{12}=0$ if and only if $X_{1}$ is independent of $X_{2}$.

The independence of two linear transforms of a multinormal $X$ can be shown via the following corollary.
Corollary 5.2 If $X \sim N_{p}(\mu, \Sigma)$ and given some matrices $\mathcal{A}$ and $\mathcal{B}$, then $\mathcal{A} X$ and $\mathcal{B} X$ are independent if and only if $\mathcal{A} \Sigma \mathcal{B}^{\top}=0$.

The following theorem is also useful. It generalises Theorem 4.6. The proof is left as an exercise.

Theorem 5.2 If $X \sim N_{p}(\mu, \Sigma), \mathcal{A}(q \times p), c \in \mathbb{R}^{q}$ and $q \leq p$, then $Y=\mathcal{A} X+c$ is a q-variate Normal, i.e.

$$
Y \sim N_{q}\left(\mathcal{A} \mu+c, \mathcal{A} \Sigma \mathcal{A}^{\top}\right) .
$$

The conditional distribution of $X_{2}$ given $X_{1}$ is given by the next theorem.
Theorem 5.3 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 $\Sigma_{22.1}$, i.e.

$$
\begin{equation*}
\left(X_{2} \mid X_{1}=x_{1}\right) \sim N_{p-r}\left(\mu_{2}+\Sigma_{21} \Sigma_{11}^{-1}\left(x_{1}-\mu_{1}\right), \Sigma_{22.1}\right) . \tag{5.8}
\end{equation*}
$$

Proof Since $X_{2}=X_{2.1}+\Sigma_{21} \Sigma_{11}^{-1} X_{1}$, for a fixed value of $X_{1}=x_{1}, X_{2}$ is equivalent to $X_{2.1}$ plus a constant term:

$$
\left(X_{2} \mid X_{1}=x_{1}\right)=\left(X_{2.1}+\Sigma_{21} \Sigma_{11}^{-1} x_{1}\right),
$$

which has the normal distribution $N\left(\mu_{2.1}+\Sigma_{21} \Sigma_{11}^{-1} x_{1}, \Sigma_{22.1}\right)$.
Note that the conditional mean of $\left(X_{2} \mid X_{1}\right)$ is a linear function of $X_{1}$ and that the conditional variance does not depend on the particular value of $X_{1}$. In the following example we consider a specific distribution.
Example 5.1 Suppose that $p=2, r=1, \mu=\binom{0}{0}$ and $\Sigma=\left(\begin{array}{cc}1 & -0.8 \\ -0.8 & 2\end{array}\right)$. Then $\Sigma_{11}=1, \Sigma_{21}=-0.8$ and $\Sigma_{22.1}=\Sigma_{22}-\Sigma_{21} \Sigma_{11}^{-1} \Sigma_{12}=2-(0.8)^{2}=1.36$. Hence the marginal pdf of $X_{1}$ is

$$
f_{X_{1}}\left(x_{1}\right)=\frac{1}{\sqrt{2 \pi}} \exp \left(-\frac{x_{1}^{2}}{2}\right)
$$

and the conditional pdf of $\left(X_{2} \mid X_{1}=x_{1}\right)$ is given by

$$
f\left(x_{2} \mid x_{1}\right)=\frac{1}{\sqrt{2 \pi(1.36)}} \exp \left\{-\frac{\left(x_{2}+0.8 x_{1}\right)^{2}}{2 \times(1.36)}\right\} .
$$

As mentioned above, the conditional mean of $\left(X_{2} \mid X_{1}\right)$ is linear in $X_{1}$. The shift in the density of $\left(X_{2} \mid X_{1}\right)$ can be seen in Fig. 5.1.

Sometimes it will be useful to reconstruct a joint distribution from the marginal distribution of $X_{1}$ and the conditional distribution $\left(X_{2} \mid X_{1}\right)$. The following theorem shows under which conditions this can be easily done in the multinormal framework.

Conditional Normal Densities $f(\mathbf{X 2 | X 1})$


Fig. 5.1 Shifts in the conditional density $\mathbf{a}$ MVAcondnorm

Theorem 5.4 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

$$
\begin{aligned}
\mu & =\binom{\mu_{1}}{\mathcal{A} \mu_{1}+b} \\
\Sigma & =\left(\begin{array}{cc}
\Sigma_{11} & \Sigma_{11} \mathcal{A}^{\top} \\
\mathcal{A} \Sigma_{11} \Omega+\mathcal{A} \Sigma_{11} \mathcal{A}^{\top}
\end{array}\right)
\end{aligned}
$$

Example 5.2 Consider the following random variables

$$
\begin{aligned}
& X_{1} \sim N_{1}(0,1), \\
& X_{2} \left\lvert\, X_{1}=x_{1} \sim N_{2}\left(\binom{2 x_{1}}{x_{1}+1},\left(\begin{array}{ll}
1 & 0 \\
0 & 1
\end{array}\right)\right) .\right.
\end{aligned}
$$

Using Theorem (5.4), where $\mathcal{A}=\left(\begin{array}{ll}2 & 1\end{array}\right)^{\top}, b=\left(\begin{array}{ll}0 & 1\end{array}\right)^{\top}$ and $\Omega=\mathcal{I}_{2}$, we easily obtain the following result:

$$
X=\binom{X_{1}}{X_{2}} \sim N_{3}\left(\left(\begin{array}{l}
0 \\
0 \\
1
\end{array}\right),\left(\begin{array}{lll}
1 & 2 & 1 \\
2 & 5 & 2 \\
1 & 2 & 2
\end{array}\right)\right)
$$

In particular, the marginal distribution of $X_{2}$ is

$$
X_{2} \sim N_{2}\left(\binom{0}{1},\left(\begin{array}{ll}
5 & 2 \\
2 & 2
\end{array}\right)\right)
$$

thus conditional on $X_{1}$, the two components of $X_{2}$ are independent but marginally they are not.

Note that the marginal mean vector and covariance matrix of $X_{2}$ could have also been computed directly by using (4.28)-(4.29). Using the derivation above, however, provides us with useful properties: we have multinormality.

## Conditional Approximations

As we saw in Chap. 4 (Theorem 4.3), the conditional expectation $\mathrm{E}\left(X_{2} \mid X_{1}\right)$ is the mean squared error (MSE) best approximation of $X_{2}$ by a function of $X_{1}$. We have in this case

$$
\begin{equation*}
X_{2}=\mathrm{E}\left(X_{2} \mid X_{1}\right)+U=\mu_{2}+\Sigma_{21} \Sigma_{11}^{-1}\left(X_{1}-\mu_{1}\right)+U \tag{5.9}
\end{equation*}
$$

Hence, the best approximation of $X_{2} \in \mathbb{R}^{p-r}$ by $X_{1} \in \mathbb{R}^{r}$ is the linear approximation that can be written as:

$$
\begin{equation*}
X_{2}=\beta_{0}+\mathcal{B} X_{1}+U \tag{5.10}
\end{equation*}
$$

with $\mathcal{B}=\Sigma_{21} \Sigma_{11}^{-1}, \beta_{0}=\mu_{2}-B \mu_{1}$ and $U \sim N\left(0, \Sigma_{22.1}\right)$.
Consider now the particular case where $r=p-1$. Now $X_{2} \in \mathbb{R}$ and $\mathcal{B}$ is a row vector $\beta^{\top}$ of dimension $(1 \times r)$

$$
\begin{equation*}
X_{2}=\beta_{0}+\beta^{\top} X_{1}+U \tag{5.11}
\end{equation*}
$$

This means, geometrically speaking, that the best MSE approximation of $X_{2}$ by a function of $X_{1}$ is a hyperplane. The marginal variance of $X_{2}$ can be decomposed via (5.11):

$$
\begin{equation*}
\sigma_{22}=\beta^{\top} \Sigma_{11} \beta+\sigma_{22.1}=\sigma_{21} \Sigma_{11}^{-1} \sigma_{12}+\sigma_{22.1} \tag{5.12}
\end{equation*}
$$

The ratio

$$
\begin{equation*}
\rho_{2.1 \ldots r}^{2}=\frac{\sigma_{21} \Sigma_{11}^{-1} \sigma_{12}}{\sigma_{22}} \tag{5.13}
\end{equation*}
$$

is known as the square of the multiple correlation between $X_{2}$ and the $r$ variables $X_{1}$. It is the percentage of the variance of $X_{2}$ which is explained by the linear
approximation $\beta_{0}+\beta^{\top} X_{1}$. The last term in (5.12) is the residual variance of $X_{2}$. The square of the multiple correlation corresponds to the coefficient of determination introduced in Sect. 3.4, see (3.39), but here it is defined in terms of the r.v. $X_{1}$ and $X_{2}$. It can be shown that $\rho_{2.1 \ldots r}$ is also the maximum correlation attainable between $X_{2}$ and a linear combination of the elements of $X_{1}$, the optimal linear combination being precisely given by $\beta^{\top} X_{1}$. Note that when $r=1$, the multiple correlation $\rho_{2.1}$ coincides with the usual simple correlation $\rho_{X_{2} X_{1}}$ between $X_{2}$ and $X_{1}$.

Example 5.3 Consider the "classic blue" pullover example (Example 3.15) and suppose that $X_{1}$ (sales), $X_{2}$ (price), $X_{3}$ (advertisement) and $X_{4}$ (sales assistants) are normally distributed with

$$
\mu=\left(\begin{array}{r}
172.7 \\
104.6 \\
104.0 \\
93.8
\end{array}\right) \text { and } \Sigma=\left(\begin{array}{rrrr}
1037.21 & & \\
-80.02 & 219.84 & \\
1430.70 & 92.10 & 2624.00 & \\
271.44 & -91.58 & 210.30 & 177.36
\end{array}\right)
$$

(These are in fact the sample mean and the sample covariance matrix but in this example we pretend that they are the true parameter values.)

The conditional distribution of $X_{1}$ given $\left(X_{2}, X_{3}, X_{4}\right)$ is thus an univariate normal with mean

$$
\mu_{1}+\sigma_{12} \Sigma_{22}^{-1}\left(\begin{array}{l}
X_{2}-\mu_{2} \\
X_{3}-\mu_{3} \\
X_{4}-\mu_{4}
\end{array}\right)=65.670-0.216 X_{2}+0.485 X_{3}+0.844 X_{4}
$$

and variance

$$
\sigma_{11.2}=\sigma_{11}-\sigma_{12} \Sigma_{22}^{-1} \sigma_{21}=96.761
$$

The linear approximation of the sales $\left(X_{1}\right)$ by the price $\left(X_{2}\right)$, advertisement $\left(X_{3}\right)$ and sales assistants $\left(X_{4}\right)$ is provided by the conditional mean above. (Note that this coincides with the results of Example 3.15 due to the particular choice of $\mu$ and $\Sigma$.) The quality of the approximation is given by the multiple correlation $\rho_{1.234}^{2}=\frac{\sigma_{12} \Sigma_{22}^{-1} \sigma_{21}}{\sigma_{11}}=0.907$. (Note again that this coincides with the coefficient of determination $r^{2}$ found in Example 3.15.)

This example also illustrates the concept of partial correlation. The correlation matrix between the four variables is given by

$$
P=\left(\begin{array}{rrrr}
1 & -0.168 & 0.867 & 0.633 \\
-0.168 & 1 & 0.121 & -0.464 \\
0.867 & 0.121 & 1 & 0.308 \\
0.633 & -0.464 & 0.308 & 1
\end{array}\right),
$$

so that the correlation between $X_{1}$ (sales) and $X_{2}$ (price) is -0.168 . We can compute the conditional distribution of $\left(X_{1}, X_{2}\right)$ given $\left(X_{3}, X_{4}\right)$, which is a bivariate normal with mean:

$$
\binom{\mu_{1}}{\mu_{2}}+\left(\begin{array}{ll}
\sigma_{13} & \sigma_{14} \\
\sigma_{23} & \sigma_{24}
\end{array}\right)\left(\begin{array}{ll}
\sigma_{33} & \sigma_{34} \\
\sigma_{43} & \sigma_{44}
\end{array}\right)^{-1}\binom{X_{3}-\mu_{3}}{X_{4}-\mu_{4}}=\binom{32.516+0.467 X_{3}+0.977 X_{4}}{153.644+0.085 X_{3}-0.617 X_{4}}
$$

and covariance matrix:

$$
\left(\begin{array}{ll}
\sigma_{11} & \sigma_{12} \\
\sigma_{21} & \sigma_{22}
\end{array}\right)-\left(\begin{array}{ll}
\sigma_{13} & \sigma_{14} \\
\sigma_{23} & \sigma_{24}
\end{array}\right)\left(\begin{array}{ll}
\sigma_{33} & \sigma_{34} \\
\sigma_{43} & \sigma_{44}
\end{array}\right)^{-1}\left(\begin{array}{ll}
\sigma_{31} & \sigma_{32} \\
\sigma_{41} & \sigma_{42}
\end{array}\right)=\left(\begin{array}{l}
104.006 \\
-33.574 \\
155.592
\end{array}\right)
$$

In particular, the last covariance matrix allows the partial correlation between $X_{1}$ and $X_{2}$ to be computed for a fixed level of $X_{3}$ and $X_{4}$ :

$$
\rho_{X_{1} X_{2} \mid X_{3} X_{4}}=\frac{-33.574}{\sqrt{104.006 \cdot 155.592}}=-0.264
$$

so that in this particular example with a fixed level of advertisement and sales assistance, the negative correlation between price and sales is more important than the marginal one.
Q MVAbluepullover

|  | Summary |
| :--- | :--- | :--- |$\quad$| If $X \sim N_{p}(\mu, \Sigma)$, then a linear transformation $\mathcal{A} X+c, \mathcal{A}(q \times p)$, |
| :--- |
| where $c \in \mathbb{R}^{q}$, has distribution $N_{q}\left(\mathcal{A} \mu+c, \mathcal{A} \Sigma \mathcal{A}^{\top}\right)$. |

### 5.2 The Wishart Distribution

The Wishart distribution (named after its discoverer) plays a prominent role in the analysis of estimated covariance matrices. If the mean of $X \sim N_{p}(\mu, \Sigma)$ is known to be $\mu=0$, then for a data matrix $\mathcal{X}(n \times p)$ the estimated covariance matrix is proportional to $\mathcal{X}^{\top} \mathcal{X}$. This is the point where the Wishart distribution comes in, because $\mathcal{M}(p \times p)=\mathcal{X}^{\top} \mathcal{X}=\sum_{i=1}^{n} x_{i} x_{i}^{\top}$ has a Wishart distribution $W_{p}(\Sigma, n)$.
Example 5.4 Set $p=1$, then for $X \sim N_{1}\left(0, \sigma^{2}\right)$ the data matrix of the observations

$$
\mathcal{X}=\left(x_{1}, \ldots, x_{n}\right)^{\top} \quad \text { with } \quad \mathcal{M}=\mathcal{X}^{\top} \mathcal{X}=\sum_{i=1}^{n} x_{i} x_{i}
$$

leads to the Wishart distribution $W_{1}\left(\sigma^{2}, n\right)=\sigma^{2} \chi_{n}^{2}$. The one-dimensional Wishart distribution is thus in fact a $\chi^{2}$ distribution.

When we talk about the distribution of a matrix, we mean of course the joint distribution of all its elements. More exactly: since $\mathcal{M}=\mathcal{X}^{\top} \mathcal{X}$ is symmetric we only need to consider the elements of the lower triangular matrix

$$
\mathcal{M}=\left(\begin{array}{cccc}
m_{11} & & &  \tag{5.14}\\
m_{21} & m_{22} & & \\
\vdots & \vdots & \ddots & \\
m_{p 1} & m_{p 2} & \ldots & m_{p p}
\end{array}\right)
$$

Hence the Wishart distribution is defined by the distribution of the vector

$$
\begin{equation*}
\left(m_{11}, \ldots, m_{p 1}, m_{22}, \ldots, m_{p 2}, \ldots, m_{p p}\right)^{\top} . \tag{5.15}
\end{equation*}
$$

Linear transformations of the data matrix $\mathcal{X}$ also lead to Wishart matrices.
Theorem 5.5 If $\mathcal{M} \sim W_{p}(\Sigma, n)$ and $\mathcal{B}(p \times q)$, then the distribution of $\mathcal{B}^{\top} \mathcal{M B}$ is Wishart $W_{q}\left(\mathcal{B}^{\top} \Sigma \mathcal{B}, n\right)$.

With this theorem we can standardise Wishart matrices since with $\mathcal{B}=\Sigma^{-1 / 2}$ the distribution of $\Sigma^{-1 / 2} \mathcal{M} \Sigma^{-1 / 2}$ is $W_{p}(\mathcal{I}, n)$. Another connection to the $\chi^{2}$-distribution is given by the following theorem.
Theorem 5.6 If $\mathcal{M} \sim W_{p}(\Sigma, m)$, and $a \in \mathbb{R}^{p}$ with $a^{\top} \Sigma a \neq 0$, then the distribution of $\frac{a^{\top} \mathcal{M} a}{a^{\top} \Sigma a}$ is $\chi_{m}^{2}$.

This theorem is an immediate consequence of Theorem 5.5 if we apply the linear transformation $x \mapsto a^{\top} x$. Central to the analysis of covariance matrices is the next theorem.

Theorem 5.7 (Cochran) Let $\mathcal{X}(n \times p)$ be a data matrix from a $N_{p}(0, \Sigma)$ distribution and let $\mathcal{C}(n \times n)$ be a symmetric matrix.
(a) $\mathcal{X}^{\top} \mathcal{C X}$ has the distribution of weighted Wishart random variables, i.e.

$$
\mathcal{X}^{\top} \mathcal{C} \mathcal{X}=\sum_{i=1}^{n} \lambda_{i} W_{p}(\Sigma, 1)
$$

where $\lambda_{i}, i=1, \ldots, n$, are the eigenvalues of $\mathcal{C}$.
(b) $\mathcal{X}^{\top} \mathcal{C} \mathcal{X}$ is Wishart if and only if $\mathcal{C}^{2}=\mathcal{C}$. In this case

$$
\mathcal{X}^{\top} \mathcal{C} \mathcal{X} \sim W_{p}(\Sigma, r)
$$

and $r=\operatorname{rank}(\mathcal{C})=\operatorname{tr}(\mathcal{C})$.
(c) $n \mathcal{S}=\mathcal{X}^{\top} \mathcal{H} \mathcal{X}$ is distributed as $W_{p}(\Sigma, n-1)$ (note that $\mathcal{S}$ is the sample covariance matrix).
(d) $\bar{x}$ and $\mathcal{S}$ are independent.

The following properties are useful:

1. If $\mathcal{M} \sim W_{p}(\Sigma, n)$, then $\mathrm{E}(\mathcal{M})=n \Sigma$.
2. If $\mathcal{M}_{i}$ are independent Wishart $W_{p}\left(\Sigma, n_{i}\right) i=1, \ldots, k$, then $\mathcal{M}=\sum_{i=1}^{k} \mathcal{M}_{i} \sim$ $W_{p}(\Sigma, n)$ where $n=\sum_{i=1}^{k} n_{i}$.
3. The density of $W_{p}(\Sigma, n-1)$ for a positive definite $\mathcal{M}$ is given by:

$$
\begin{equation*}
f_{\Sigma, n-1}(\mathcal{M})=\frac{|\mathcal{M}|^{\frac{1}{2}(n-p-2)} e^{-\frac{1}{2} \operatorname{tr}\left(\mathcal{M} \Sigma^{-1}\right)}}{2^{\frac{1}{2} p(n-1)} \pi^{\frac{1}{4} p(p-1)}|\Sigma|^{\frac{1}{2}(n-1)} \prod_{i=1}^{p} \Gamma\left\{\frac{n-i}{2}\right\}}, \tag{5.16}
\end{equation*}
$$

where $\Gamma$ is the gamma function: $\Gamma(z)=\int_{0}^{\infty} t^{z-1} e^{-t} d t$.
For further details on the Wishart distribution, see Mardia, Kent, and Bibby (1979).


## Summary

$\hookrightarrow$ The Wishart distribution is a generalisation of the $\chi^{2}$-distribution. In particular $W_{1}\left(\sigma^{2}, n\right)=\sigma^{2} \chi_{n}^{2}$.
$\hookrightarrow$ The empirical covariance matrix $\mathcal{S}$ has a $\frac{1}{n} W_{p}(\Sigma, n-1)$ distribution.
$\hookrightarrow \quad$ In the normal case, $\bar{x}$ and $\mathcal{S}$ are independent.
$\hookrightarrow \quad$ For $\mathcal{M} \sim W_{p}(\Sigma, m)$, $a^{\top} \mathcal{M} a / a^{\top} \Sigma a \sim \chi_{m}^{2}$.

### 5.3 Hotelling's $\boldsymbol{T}^{\mathbf{2}}$-Distribution

Suppose that $Y \in \mathbb{R}^{p}$ is a standard normal random vector, i.e. $Y \sim N_{p}(0, \mathcal{I})$, independent of the random matrix $\mathcal{M} \sim W_{p}(\mathcal{I}, n)$. What is the distribution of $Y^{\top} \mathcal{M}^{-1} Y$ ? The answer is provided by the Hotelling $T^{2}$-distribution: $n Y^{\top} \mathcal{M}^{-1} Y$ is Hotelling $T_{p, n}^{2}$ distributed.

The Hotelling $T^{2}$-distribution is a generalisation of the Student $t$-distribution. The general multinormal distribution $N(\mu, \Sigma)$ is considered in Theorem 5.8. The Hotelling $T^{2}$-distribution will play a central role in hypothesis testing in Chap. 7.

Theorem 5.8 If $X \sim N_{p}(\mu, \Sigma)$ is independent of $\mathcal{M} \sim W_{p}(\Sigma, n)$, then

$$
n(X-\mu)^{\top} \mathcal{M}^{-1}(X-\mu) \sim T_{p, n}^{2}
$$

Corollary 5.3 If $\bar{x}$ is the mean of a sample drawn from a normal population $N_{p}(\mu, \Sigma)$ and $\mathcal{S}$ is the sample covariance matrix, then

$$
\begin{equation*}
(n-1)(\bar{x}-\mu)^{\top} \mathcal{S}^{-1}(\bar{x}-\mu)=n(\bar{x}-\mu)^{\top} \mathcal{S}_{u}^{-1}(\bar{x}-\mu) \sim T_{p, n-1}^{2} \tag{5.17}
\end{equation*}
$$

Recall that $\mathcal{S}_{u}=\frac{n}{n-1} \mathcal{S}$ is an unbiased estimator of the covariance matrix. A connection between the Hotelling $T^{2}$ - and the $F$-distribution is given by the next theorem.

Theorem 5.9

$$
T_{p, n}^{2}=\frac{n p}{n-p+1} F_{p, n-p+1}
$$

Example 5.5 In the univariate case $(p=1)$, this theorem boils down to the well-known result:

$$
\left(\frac{\bar{x}-\mu}{\sqrt{\mathcal{S}_{u}} / \sqrt{n}}\right)^{2} \sim T_{1, n-1}^{2}=F_{1, n-1}=t_{n-1}^{2}
$$

For further details on Hotelling $T^{2}$-distribution see Mardia et al. (1979). The next corollary follows immediately from (3.23), (3.24) and from Theorem 5.8. It will be useful for testing linear restrictions in multinormal populations.

Corollary 5.4 Consider a linear transform of $X \sim N_{p}(\mu, \Sigma), Y=\mathcal{A} X$ where $\mathcal{A}(q \times p)$ with $(q \leq p)$. If $\bar{x}$ and $\mathcal{S}_{X}$ are the sample mean and the covariance matrix, we have

$$
\begin{gathered}
\bar{y}=\mathcal{A} \bar{x} \sim N_{q}\left(\mathcal{A} \mu, \frac{1}{n} \mathcal{A} \Sigma \mathcal{A}^{\top}\right) \\
n \mathcal{S}_{Y}=n \mathcal{A} \mathcal{S}_{X} \mathcal{A}^{\top} \sim W_{q}\left(\mathcal{A} \Sigma \mathcal{A}^{\top}, n-1\right) \\
(n-1)(\mathcal{A} \bar{x}-\mathcal{A} \mu)^{\top}\left(\mathcal{A} \mathcal{S}_{X} \mathcal{A}^{\top}\right)^{-1}(\mathcal{A} \bar{x}-\mathcal{A} \mu) \sim T_{q, n-1}^{2}
\end{gathered}
$$

The $T^{2}$ distribution is closely connected to the univariate $t$-statistic. In Example 5.4 we described the manner in which the Wishart distribution generalises the $\chi^{2}$-distribution. We can write (5.17) as:

$$
T^{2}=\sqrt{n}(\bar{x}-\mu)^{\top}\left(\frac{\sum_{j=1}^{n}\left(x_{j}-\bar{x}\right)\left(x_{j}-\bar{x}\right)^{\top}}{n-1}\right)^{-1} \sqrt{n}(\bar{x}-\mu)
$$

which is of the form

$$
\binom{\text { multivariate normal }}{\text { random vector }}^{\top}\left(\frac{\begin{array}{c}
\text { Wishart random } \\
\text { matrix }
\end{array}}{\text { degrees of freedom }}\right)^{-1}\binom{\text { multivariate normal }}{\text { random vector }}
$$

This is analogous to

$$
t^{2}=\sqrt{n}(\bar{x}-\mu)\left(s^{2}\right)^{-1} \sqrt{n}(\bar{x}-\mu)
$$

or

$$
\binom{\text { normal }}{\text { random variable }}\binom{\chi^{2} \text {-random }}{\frac{\text { variable }}{\text { degrees of freedom }}}^{-1}\binom{\text { normal }}{\text { random variable }}
$$

for the univariate case. Since the multivariate normal and Wishart random variables are independently distributed, their joint distribution is the product of the marginal normal and Wishart distributions. Using calculus, the distribution of $T^{2}$ as given above can be derived from this joint distribution.

|  | Summary |
| :--- | :--- |$\quad$| Hotelling's $T^{2}$-distribution is a generalisation of the $t$-distribution. |
| :--- |
|  | | In particular $T_{1, n}^{2}=t_{n}$. |
| :--- |

### 5.4 Spherical and Elliptical Distributions

The multinormal distribution belongs to the large family of elliptical distributions which has recently gained a lot of attention in financial mathematics. Elliptical distributions are often used, particularly in risk management.

Definition 5.1 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 $\phi($.$) which is then called the characteristic generator of the spherical$ distribution $S_{p}(\phi)$. We will write $Y \sim S_{p}(\phi)$.

This is only one of several possible ways to define spherical distributions. We can see spherical distributions as an extension of the standard multinormal distribution $N_{p}\left(0, \mathcal{I}_{p}\right)$.

Theorem 5.10 Spherical random variables have the following properties:

1. All marginal distributions of a spherically distributed random vector are spherical.
2. All the marginal characteristic functions have the same generator.
3. Let $X \sim S_{p}(\phi)$, then $X$ has the same distribution as $r u^{(p)}$ where $u^{(p)}$ is a random vector distributed uniformly on the unit sphere surface in $\mathbb{R}^{p}$ and $r \geq 0$ is a random variable independent of $u^{(p)}$. If $\mathrm{E}\left(r^{2}\right)<\infty$, then

$$
\mathrm{E}(X)=0, \quad \operatorname{Cov}(X)=\frac{\mathrm{E}\left(r^{2}\right)}{p} \mathcal{I}_{p}
$$

The random radius $r$ is related to the generator $\phi$ by a relation described in Fang, Kotz, and $\operatorname{Ng}$ (1990, p. 29). The moments of $X \sim S_{p}(\phi)$, provided that they exist, can be expressed in terms of one-dimensional integral.

A spherically distributed random vector does not, in general, necessarily possess a density. However, if it does, the marginal densities of dimension smaller than $p-1$ are continuous and the marginal densities of dimension smaller than $p-2$
are differentiable (except possibly at the origin in both cases). Univariate marginal densities for $p$ greater than 2 are non-decreasing on $(-\infty, 0)$ and non-increasing on $(0, \infty)$.

Definition 5.2 A $(p \times 1)$ random vector $X$ is said to have 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 $X \sim E C_{p}(\mu, \Sigma, \phi)$.

Remark 5.1 The elliptical distribution can be seen as an extension of $N_{p}(\mu, \Sigma)$.
Example 5.6 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. Moreover the $t$-distribution belongs to the family of $p$-dimensional spherical distributions.
Example 5.7 The multinormal distribution. Let $X \sim N_{p}(\mu, \Sigma)$. Then $X \sim E C_{p}(\mu, \Sigma, \phi)$ and $\phi(u)=\exp (-u / 2)$. Figure 4.3 shows a density surface of the multivariate normal distribution: $f(x)=\operatorname{det}(2 \pi \Sigma)^{-\frac{1}{2}} \exp \left\{-\frac{1}{2}(x-\mu)^{\top} \Sigma^{-1}\right.$ $(x-\mu)\}$ with $\Sigma=\left(\begin{array}{cc}1 & 0.6 \\ 0.6 & 1\end{array}\right)$ and $\mu=\binom{0}{0}$ Notice that the density is constant on ellipses. This is the reason for calling this family of distributions "elliptical".

Theorem 5.11 Elliptical random vectors $X$ have the following properties:

1. Any linear combination of elliptically distributed variables are elliptical.
2. Marginal distributions of elliptically distributed variables are elliptical.
3. A scalar function $\phi$ (.) can determine an elliptical distribution $E C_{p}(\mu, \Sigma, \phi)$ for every $\mu \in \mathbb{R}^{p}$ and $\Sigma \geq 0$ with $\operatorname{rank}(\Sigma)=k$ iff $\phi\left(t^{\top} t\right)$ is a $p$-dimensional characteristic function.
4. Assume that $X$ is non-degenerate. If $X \sim E C_{p}(\mu, \Sigma, \phi)$ and $X \sim E C_{p}$ $\left(\mu^{*}, \Sigma^{*}, \phi^{*}\right)$, then a constant $c>0$ exists that

$$
\mu=\mu^{*}, \quad \Sigma=c \Sigma^{*}, \quad \phi^{*}(.)=\phi\left(c^{-1} .\right)
$$

In other words $\Sigma, \phi, \mathcal{A}$ are not unique, unless we impose the condition that $\operatorname{det}(\Sigma)=1$.
5. The characteristic function of $X, \psi(t)=\mathrm{E}\left(e^{\mathrm{i} t^{\top} X}\right)$ is of the form

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

for a scalar function $\phi$.
6. $X \sim E C_{p}(\mu, \Sigma, \phi)$ with $\operatorname{rank}(\Sigma)=k$ iff $X$ has the same distribution as:

$$
\begin{equation*}
\mu+r \mathcal{A}^{\top} u^{(k)} \tag{5.18}
\end{equation*}
$$

where $r \geq 0$ is independent of $u^{(k)}$ which is a random vector distributed uniformly on the unit sphere surface in $\mathbb{R}^{k}$ and $\mathcal{A}$ is a $(k \times p)$ matrix such that $\mathcal{A}^{\top} \mathcal{A}=\Sigma$.
7. Assume that $X \sim E C_{p}(\mu, \Sigma, \phi)$ and $\mathrm{E}\left(r^{2}\right)<\infty$. Then

$$
\mathrm{E}(X)=\mu \quad \operatorname{Cov}(X)=\frac{\mathrm{E}\left(r^{2}\right)}{\operatorname{rank}(\Sigma)} \Sigma=-2 \phi^{\top}(0) \Sigma
$$

8. Assume that $X \sim E C_{p}(\mu, \Sigma, \phi)$ with $\operatorname{rank}(\Sigma)=k$. Then

$$
Q(X)=(X-\mu)^{\top} \Sigma^{-1}(X-\mu)
$$

has the same distribution as $r^{2}$ in Eq. (5.18).

### 5.5 Exercises

Exercise 5.1 Consider $X \sim N_{2}(\mu, \Sigma)$ with $\mu=(2,2)^{\top}$ and $\Sigma=\left(\begin{array}{ll}1 & 0 \\ 0 & 1\end{array}\right)$ and the matrices $\mathcal{A}=\binom{1}{1}^{\top}, \mathcal{B}=\binom{1}{-1}^{\top}$. Show that $\mathcal{A} X$ and $\mathcal{B} X$ are independent.
Exercise 5.2 Prove Theorem 5.4.
Exercise 5.3 Prove proposition (c) of Theorem 5.7.
Exercise 5.4 Let

$$
X \sim N_{2}\left(\binom{1}{2},\left(\begin{array}{ll}
2 & 1 \\
1 & 2
\end{array}\right)\right)
$$

and

$$
Y \left\lvert\, X \sim N_{2}\left(\binom{X_{1}}{X_{1}+X_{2}},\left(\begin{array}{ll}
1 & 0 \\
0 & 1
\end{array}\right)\right)\right.
$$

(a) Determine the distribution of $Y_{2} \mid Y_{1}$.
(b) Determine the distribution of $W=X-Y$.

Exercise 5.5 Consider $\left(\begin{array}{c}X \\ Y \\ Z\end{array}\right) \sim N_{3}(\mu, \Sigma)$. Compute $\mu$ and $\Sigma$ knowing that

$$
\begin{aligned}
Y \mid Z & \sim N_{1}(-Z, 1) \\
\mu_{Z \mid Y} & =-\frac{1}{3}-\frac{1}{3} Y \\
X \mid Y, Z & \sim N_{1}(2+2 Y+3 Z, 1)
\end{aligned}
$$

Determine the distributions of $X \mid Y$ and of $X \mid Y+Z$.
Exercise 5.6 Knowing that

$$
\begin{aligned}
Z & \sim N_{1}(0,1) \\
Y \mid Z & \sim N_{1}(1+Z, 1) \\
X \mid Y, Z & \sim N_{1}(1-Y, 1)
\end{aligned}
$$

(a) find the distribution of $\left(\begin{array}{l}X \\ Y \\ Z\end{array}\right)$ and of $Y \mid X, Z$.
(b) find the distribution of

$$
\binom{U}{V}=\binom{1+Z}{1-Y}
$$

(c) compute $\mathrm{E}(Y \mid U=2)$.

Exercise 5.7 Suppose $\binom{X}{Y} \sim N_{2}(\mu, \Sigma)$ with $\Sigma$ positive definite. Is it possible that
(a) $\mu_{X \mid Y}=3 Y^{2}$,
(b) $\sigma_{X X \mid Y}=2+Y^{2}$,
(c) $\mu_{X \mid Y}=3-Y$, and
(d) $\sigma_{X X \mid Y}=5$ ?

Exercise 5.8 Let $X \sim N_{3}\left(\left(\begin{array}{l}1 \\ 2 \\ 3\end{array}\right),\left(\begin{array}{rrr}11 & -6 & 2 \\ -6 & 10 & -4 \\ 2 & -4 & 6\end{array}\right)\right)$.
(a) Find the best linear approximation of $X_{3}$ by a linear function of $X_{1}$ and $X_{2}$ and compute the multiple correlation between $X_{3}$ and $\left(X_{1}, X_{2}\right)$.
(b) Let $Z_{1}=X_{2}-X_{3}, Z_{2}=X_{2}+X_{3}$ and $\left(Z_{3} \mid Z_{1}, Z_{2}\right) \sim N_{1}\left(Z_{1}+Z_{2}, 10\right)$. Compute the distribution of $\left(\begin{array}{l}Z_{1} \\ Z_{2} \\ Z_{3}\end{array}\right)$.

Exercise 5.9 Let $(X, Y, Z)^{\top}$ be a trivariate normal r.v. with

$$
\begin{aligned}
Y \mid Z & \sim N_{1}(2 Z, 24) \\
Z \mid X & \sim N_{1}(2 X+3,14) \\
X & \sim N_{1}(1,4) \\
\text { and } \rho_{X Y} & =0.5 .
\end{aligned}
$$

Find the distribution of $(X, Y, Z)^{\top}$ and compute the partial correlation between $X$ and $Y$ for fixed $Z$. Do you think it is reasonable to approximate $X$ by a linear function of $Y$ and $Z$ ?
Exercise 5.10 Let $X \sim N_{4}\left(\left(\begin{array}{l}1 \\ 2 \\ 3 \\ 4\end{array}\right),\left(\begin{array}{llrl}4 & 1 & 2 & 4 \\ 1 & 4 & 2 & 1 \\ 2 & 2 & 16 & 1 \\ 4 & 1 & 1 & 9\end{array}\right)\right)$.
(a) Give the best linear approximation of $X_{2}$ as a function of $\left(X_{1}, X_{4}\right)$ and evaluate the quality of the approximation.
(b) Give the best linear approximation of $X_{2}$ as a function of $\left(X_{1}, X_{3}, X_{4}\right)$ and compare your answer with part (a).

Exercise 5.11 Prove Theorem 5.2.
(Hint: complete the linear transformation $Z=\binom{\mathcal{A}}{\mathcal{I}_{p-q}} X+\binom{c}{0_{p-q}}$ and then use Theorem 5.1 to get the marginal of the first $q$ components of $Z$.)
Exercise 5.12 Prove Corollaries 5.1 and 5.2.

## Chapter 6 Theory of Estimation

We know from our basic knowledge of statistics that one of the objectives in statistics is to better understand and model the underlying process which generates data. This is known as statistical inference: we infer from information contained in sample properties of the population from which the observations are taken. In multivariate statistical inference, we do exactly the same. The basic ideas were introduced in Sect. 4.5 on sampling theory: we observed the values of a multivariate random variable $X$ and obtained a sample $\mathcal{X}=\left\{x_{i}\right\}_{i=1}^{n}$. Under random sampling, these observations are considered to be realisations of a sequence of i.i.d. random variables $X_{1}, \ldots, X_{n}$ where each $X_{i}$ is a $p$-variate random variable which replicates the parent or population random variable $X$. In this chapter, for notational convenience, we will no longer differentiate between a random variable $X_{i}$ and an observation of it, $x_{i}$, in our notation. We will simply write $x_{i}$ and it should be clear from the context whether a random variable or an observed value is meant.

Statistical inference infers from the i.i.d. random sample $\mathcal{X}$ the properties of the population: typically, some unknown characteristic $\theta$ of its distribution. In parametric statistics, $\theta$ is a $k$-variate vector $\theta \in \mathbb{R}^{k}$ characterising the unknown properties of the population pdf $f(x ; \theta)$ : this could be the mean, the covariance matrix, kurtosis, etc.

The aim will be to estimate $\theta$ from the sample $\mathcal{X}$ through estimators $\hat{\theta}$ which are functions of the sample: $\hat{\theta}=\hat{\theta}(\mathcal{X})$. When an estimator $\hat{\theta}$ is proposed, we must derive its sampling distribution to analyse its properties.

In this chapter the basic theoretical tools are developed which are needed to derive estimators and to determine their properties in general situations. We will basically rely on the maximum likelihood theory in our presentation. In many situations, the maximum likelihood estimators (MLEs) indeed share asymptotic optimal properties which make their use easy and appealing.

We will illustrate the multivariate normal population and also the linear regression model where the applications are numerous and the derivations are easy to do. In multivariate setups, the MLE is at times too complicated to be derived
analytically. In such cases, the estimators are obtained using numerical methods (nonlinear optimisation). The general theory and the asymptotic properties of these estimators remain simple and valid. The following Chap. 7 concentrates on hypothesis testing and confidence interval issues.

### 6.1 The Likelihood Function

Suppose that $\left\{x_{i}\right\}_{i=1}^{n}$ is an i.i.d. sample from a population with pdf $f(x ; \theta)$. The aim is to estimate $\theta \in \mathbb{R}^{k}$ which is a vector of unknown parameters. The likelihood function is defined as the joint density $L(\mathcal{X} ; \theta)$ of the observations $x_{i}$ considered as a function of $\theta$ :

$$
\begin{equation*}
L(\mathcal{X} ; \theta)=\prod_{i=1}^{n} f\left(x_{i} ; \theta\right) \tag{6.1}
\end{equation*}
$$

where $\mathcal{X}$ denotes the sample of the data matrix with the observations $x_{1}^{\top}, \ldots, x_{n}^{\top}$ in each row. The MLE of $\theta$ is defined as

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

Often it is easier to maximise the log-likelihood function

$$
\begin{equation*}
\ell(\mathcal{X} ; \theta)=\log L(\mathcal{X} ; \theta), \tag{6.2}
\end{equation*}
$$

which is equivalent since the logarithm is a monotone one-to-one function. Hence

$$
\hat{\theta}=\arg \max _{\theta} L(\mathcal{X} ; \theta)=\arg \max _{\theta} \ell(\mathcal{X} ; \theta) .
$$

The following examples illustrate cases where the maximisation process can be performed analytically, i.e., we will obtain an explicit analytical expression for $\hat{\theta}$. Unfortunately, in other situations, the maximisation process can be more intricate, involving nonlinear optimisation techniques. In the latter case, given a sample $\mathcal{X}$ and the likelihood function, numerical methods will be used to determine the value of $\theta$ maximising $L(\mathcal{X} ; \theta)$ or $\ell(\mathcal{X} ; \theta)$. These numerical methods are typically based on Newton-Raphson iterative techniques.

Example 6.1 Consider a sample $\left\{x_{i}\right\}_{i=1}^{n}$ from $N_{p}(\mu, \mathcal{I})$, i.e., from the pdf

$$
f(x ; \theta)=(2 \pi)^{-p / 2} \exp \left\{-\frac{1}{2}(x-\theta)^{\top}(x-\theta)\right\},
$$

where $\theta=\mu \in \mathbb{R}^{p}$ is the mean vector parameter. The log-likelihood is in this case

$$
\begin{equation*}
\ell(\mathcal{X} ; \theta)=\sum_{i=1}^{n} \log \left\{f\left(x_{i} ; \theta\right)\right\}=\log (2 \pi)^{-n p / 2}-\frac{1}{2} \sum_{i=1}^{n}\left(x_{i}-\theta\right)^{\top}\left(x_{i}-\theta\right) . \tag{6.3}
\end{equation*}
$$

The term $\left(x_{i}-\theta\right)^{\top}\left(x_{i}-\theta\right)$ equals

$$
\left(x_{i}-\bar{x}\right)^{\top}\left(x_{i}-\bar{x}\right)+(\bar{x}-\theta)^{\top}(\bar{x}-\theta)+2(\bar{x}-\theta)^{\top}\left(x_{i}-\bar{x}\right) .
$$

Summing this term over $i=1, \ldots, n$ we see that

$$
\sum_{i=1}^{n}\left(x_{i}-\theta\right)^{\top}\left(x_{i}-\theta\right)=\sum_{i=1}^{n}\left(x_{i}-\bar{x}\right)^{\top}\left(x_{i}-\bar{x}\right)+n(\bar{x}-\theta)^{\top}(\bar{x}-\theta) .
$$

Hence

$$
\ell(\mathcal{X} ; \theta)=\log (2 \pi)^{-n p / 2}-\frac{1}{2} \sum_{i=1}^{n}\left(x_{i}-\bar{x}\right)^{\top}\left(x_{i}-\bar{x}\right)-\frac{n}{2}(\bar{x}-\theta)^{\top}(\bar{x}-\theta) .
$$

Only the last term depends on $\theta$ and is obviously maximised for

$$
\hat{\theta}=\hat{\mu}=\bar{x}
$$

Thus $\bar{x}$ is the MLE of $\theta$ for this family of pdfs $f(x, \theta)$.
A more complex example is the following one where we derive the MLEs for $\mu$ and $\Sigma$.

Example 6.2 Suppose $\left\{x_{i}\right\}_{i=1}^{n}$ is a sample from a normal distribution $N_{p}(\mu, \Sigma)$. Here $\theta=(\mu, \Sigma)$ with $\Sigma$ interpreted as a vector. Due to the symmetry of $\Sigma$ the unknown parameter $\theta$ is in fact $\left\{p+\frac{1}{2} p(p+1)\right\}$-dimensional. Then

$$
\begin{equation*}
L(\mathcal{X} ; \theta)=|2 \pi \Sigma|^{-n / 2} \exp \left\{-\frac{1}{2} \sum_{i=1}^{n}\left(x_{i}-\mu\right)^{\top} \Sigma^{-1}\left(x_{i}-\mu\right)\right\} \tag{6.4}
\end{equation*}
$$

and

$$
\begin{equation*}
\ell(\mathcal{X} ; \theta)=-\frac{n}{2} \log |2 \pi \Sigma|-\frac{1}{2} \sum_{i=1}^{n}\left(x_{i}-\mu\right)^{\top} \Sigma^{-1}\left(x_{i}-\mu\right) . \tag{6.5}
\end{equation*}
$$

The term $\left(x_{i}-\mu\right)^{\top} \Sigma^{-1}\left(x_{i}-\mu\right)$ equals

$$
\left(x_{i}-\bar{x}\right)^{\top} \Sigma^{-1}\left(x_{i}-\bar{x}\right)+(\bar{x}-\mu)^{\top} \Sigma^{-1}(\bar{x}-\mu)+2(\bar{x}-\mu)^{\top} \Sigma^{-1}\left(x_{i}-\bar{x}\right)
$$

Summing this term over $i=1, \ldots, n$ we see that
$\sum_{i=1}^{n}\left(x_{i}-\mu\right)^{\top} \Sigma^{-1}\left(x_{i}-\mu\right)=\sum_{i=1}^{n}\left(x_{i}-\bar{x}\right)^{\top} \Sigma^{-1}\left(x_{i}-\bar{x}\right)+n(\bar{x}-\mu)^{\top} \Sigma^{-1}(\bar{x}-\mu)$.

Note that from (2.14)

$$
\begin{aligned}
\left(x_{i}-\bar{x}\right)^{\top} \Sigma^{-1}\left(x_{i}-\bar{x}\right) & =\operatorname{tr}\left\{\left(x_{i}-\bar{x}\right)^{\top} \Sigma^{-1}\left(x_{i}-\bar{x}\right)\right\} \\
& =\operatorname{tr}\left\{\Sigma^{-1}\left(x_{i}-\bar{x}\right)\left(x_{i}-\bar{x}\right)^{\top}\right\}
\end{aligned}
$$

Therefore, by summing over the index $i$ we finally arrive at

$$
\begin{aligned}
\sum_{i=1}^{n}\left(x_{i}-\mu\right)^{\top} \Sigma^{-1}\left(x_{i}-\mu\right)= & \operatorname{tr}\left\{\Sigma^{-1} \sum_{i=1}^{n}\left(x_{i}-\bar{x}\right)\left(x_{i}-\bar{x}\right)^{\top}\right\} \\
& +n(\bar{x}-\mu)^{\top} \Sigma^{-1}(\bar{x}-\mu) \\
= & \operatorname{tr}\left\{\Sigma^{-1} n \mathcal{S}\right\}+n(\bar{x}-\mu)^{\top} \Sigma^{-1}(\bar{x}-\mu)
\end{aligned}
$$

Thus the log-likelihood function for $N_{p}(\mu, \Sigma)$ is

$$
\begin{equation*}
\ell(\mathcal{X} ; \theta)=-\frac{n}{2} \log |2 \pi \Sigma|-\frac{n}{2} \operatorname{tr}\left\{\Sigma^{-1} \mathcal{S}\right\}-\frac{n}{2}(\bar{x}-\mu)^{\top} \Sigma^{-1}(\bar{x}-\mu) \tag{6.6}
\end{equation*}
$$

We can easily see that the third term is maximised by $\mu=\bar{x}$. In fact the MLEs are given by

$$
\hat{\mu}=\bar{x}, \quad \hat{\Sigma}=\mathcal{S}
$$

The derivation of $\hat{\Sigma}$ is a lot more complicated. It involves derivatives with respect to matrices with their notational complexities and will not be presented here; for more elaborate proof, see Mardia, Kent and Bibby (1979, pp. 103-104). Note that the unbiased covariance estimator $\mathcal{S}_{u}=\frac{n}{n-1} \mathcal{S}$ is not the MLE of $\Sigma$ !
Example 6.3 Consider the linear regression model $y_{i}=\beta^{\top} x_{i}+\varepsilon_{i}$ for $i=1, \ldots, n$, where $\varepsilon_{i}$ is i.i.d. and $N\left(0, \sigma^{2}\right)$ and where $x_{i} \in \mathbb{R}^{p}$. Here $\theta=\left(\beta^{\top}, \sigma\right)$ is a $(p+1)$ dimensional parameter vector. Denote

$$
y=\left(\begin{array}{c}
y_{1} \\
\vdots \\
y_{n}
\end{array}\right), \quad \mathcal{X}=\left(\begin{array}{c}
x_{1}^{\top} \\
\vdots \\
x_{n}^{\top}
\end{array}\right)
$$

Then

$$
L(y, \mathcal{X} ; \theta)=\prod_{i=1}^{n} \frac{1}{\sqrt{2 \pi} \sigma} \exp \left\{-\frac{1}{2 \sigma^{2}}\left(y_{i}-\beta^{\top} x_{i}\right)^{2}\right\}
$$

and

$$
\begin{aligned}
\ell(y, \mathcal{X} ; \theta) & =\log \left\{\frac{1}{(2 \pi)^{n / 2} \sigma^{n}}\right\}-\frac{1}{2 \sigma^{2}} \sum_{i=1}^{n}\left(y_{i}-\beta^{\top} x_{i}\right)^{2} \\
& =-\frac{n}{2} \log (2 \pi)-n \log \sigma-\frac{1}{2 \sigma^{2}}(y-\mathcal{X} \beta)^{\top}(y-\mathcal{X} \beta) \\
& =-\frac{n}{2} \log (2 \pi)-n \log \sigma-\frac{1}{2 \sigma^{2}}\left(y^{\top} y+\beta^{\top} \mathcal{X}^{\top} \mathcal{X} \beta-2 \beta^{\top} \mathcal{X}^{\top} y\right)
\end{aligned}
$$

Differentiating w.r.t. the parameters yields

$$
\begin{aligned}
\frac{\partial}{\partial \beta} \ell & =-\frac{1}{2 \sigma^{2}}\left(2 \mathcal{X}^{\top} \mathcal{X} \beta-2 \mathcal{X}^{\top} y\right) \\
\frac{\partial}{\partial \sigma} \ell & =-\frac{n}{\sigma}+\frac{1}{\sigma^{3}}\left\{(y-\mathcal{X} \beta)^{\top}(y-\mathcal{X} \beta)\right\} .
\end{aligned}
$$

Note that $\frac{\partial}{\partial \beta} \ell$ denotes the vector of the derivatives w.r.t. all components of $\beta$ (the gradient). Since the first equation only depends on $\beta$, we start with deriving $\hat{\beta}$.

$$
\mathcal{X}^{\top} \mathcal{X} \hat{\beta}=\mathcal{X}^{\top} y, \quad \text { hence } \quad \hat{\beta}=\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1} \mathcal{X}^{\top} y
$$

Plugging $\hat{\beta}$ into the second equation gives

$$
\frac{n}{\hat{\sigma}}=\frac{1}{\hat{\sigma}^{3}}(y-\mathcal{X} \hat{\beta})^{\top}(y-\mathcal{X} \hat{\beta}), \quad \text { hence } \quad \hat{\sigma}^{2}=\frac{1}{n}\|y-\mathcal{X} \hat{\beta}\|^{2},
$$

where $\|\cdot\|^{2}$ denotes the Euclidean vector norm from Sect. 2.6. We see that the MLE $\hat{\beta}$ is identical with the least squares estimator (3.52). The variance estimator

$$
\hat{\sigma}^{2}=\frac{1}{n} \sum_{i=1}^{n}\left(y_{i}-\hat{\beta}^{\top} x_{i}\right)^{2}
$$

is nothing else than the residual sum of squares (RSS) from (3.37) generalised to the case of multivariate $x_{i}$. Note that when the $x_{i}$ are considered to be fixed, we have

$$
\mathrm{E}(y)=\mathcal{X} \beta \text { and } \operatorname{Var}(y)=\sigma^{2} \mathcal{I}_{n}
$$

Then, using the properties of moments from Sect. 4.2 we have

$$
\begin{gather*}
\mathrm{E}(\hat{\beta})=\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1} \mathcal{X}^{\top} \mathrm{E}(y)=\beta,  \tag{6.7}\\
\operatorname{Var}(\hat{\beta})=\sigma^{2}\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1} . \tag{6.8}
\end{gather*}
$$

 Summary

### 6.2 The Cramer-Rao Lower Bound

As pointed out above, an important question in estimation theory is whether an estimator $\hat{\theta}$ has certain desired properties, in particular, if it converges to the unknown parameter $\theta$ it is supposed to estimate. One typical property we want for an estimator is unbiasedness, meaning that on the average, the estimator hits its target: $\mathrm{E}(\hat{\theta})=\theta$. We have seen for instance (see Example 6.2) that $\bar{x}$ is an unbiased estimator of $\mu$ and $\mathcal{S}$ is a biased estimator of $\Sigma$ in finite samples. If we restrict ourselves to unbiased estimation, then the natural question is whether the estimator shares some optimality properties in terms of its sampling variance. Since we focus on unbiasedness, we look for an estimator with the smallest possible variance.

In this context, the Cramer-Rao lower bound will give the minimal achievable variance for any unbiased estimator. This result is valid under very general regularity conditions (discussed below). One of the most important applications of the Cramer-Rao lower bound is that it provides the asymptotic optimality property of MLEs. The Cramer-Rao theorem involves the score function and its properties which will be derived first.

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

$$
\begin{equation*}
s(\mathcal{X} ; \theta)=\frac{\partial}{\partial \theta} \ell(\mathcal{X} ; \theta)=\frac{1}{L(\mathcal{X} ; \theta)} \frac{\partial}{\partial \theta} L(\mathcal{X} ; \theta) . \tag{6.9}
\end{equation*}
$$

The covariance matrix $\mathcal{F}_{n}=\operatorname{Var}\{s(\mathcal{X} ; \theta)\}$ is called the Fisher information matrix. In what follows, we will give some interesting properties of score functions.
Theorem 6.1 If $s=s(\mathcal{X} ; \theta)$ is the score function and if $\hat{\theta}=t=t(\mathcal{X}, \theta)$ is any function of $\mathcal{X}$ and $\theta$, then under regularity conditions

$$
\begin{equation*}
\mathrm{E}\left(s t^{\top}\right)=\frac{\partial}{\partial \theta} \mathrm{E}\left(t^{\top}\right)-\mathrm{E}\left(\frac{\partial t^{\top}}{\partial \theta}\right) \tag{6.10}
\end{equation*}
$$

The proof is left as an exercise (see Exercise 6.9). The regularity conditions required for this theorem are rather technical and ensure that the expressions (expectations and derivations) appearing in (6.10) are well defined. In particular, the support of the density $f(x ; \theta)$ should not depend on $\theta$. The next corollary is a direct consequence.
Corollary 6.1 If $s=s(\mathcal{X} ; \theta)$ is the score function, and $\hat{\theta}=t=t(\mathcal{X})$ is any unbiased estimator of $\theta$ (i.e., $\mathrm{E}(t)=\theta$ ), then

$$
\begin{equation*}
\mathrm{E}\left(s t^{\top}\right)=\operatorname{Cov}(s, t)=\mathcal{I}_{k} \tag{6.11}
\end{equation*}
$$

Note that the score function has mean zero (see Exercise 6.10).

$$
\begin{equation*}
\mathrm{E}\{s(\mathcal{X} ; \theta)\}=0 . \tag{6.12}
\end{equation*}
$$

Hence, $\mathrm{E}\left(s s^{\top}\right)=\operatorname{Var}(s)=\mathcal{F}_{n}$ and by setting $s=t$ in Theorem 6.1 it follows that

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

Remark 6.1 If $x_{1}, \ldots, x_{n}$ are i.i.d., $\mathcal{F}_{n}=n \mathcal{F}_{1}$ where $\mathcal{F}_{1}$ is the Fisher information matrix for sample size $n=1$.

Example 6.4 Consider an i.i.d. sample $\left\{x_{i}\right\}_{i=1}^{n}$ from $N_{p}(\theta, \mathcal{I})$. In this case the parameter $\theta$ is the mean $\mu$. It follows from (6.3) that

$$
\begin{aligned}
s(\mathcal{X} ; \theta) & =\frac{\partial}{\partial \theta} \ell(\mathcal{X} ; \theta) \\
& =-\frac{1}{2} \frac{\partial}{\partial \theta}\left\{\sum_{i=1}^{n}\left(x_{i}-\theta\right)^{\top}\left(x_{i}-\theta\right)\right\} \\
& =n(\bar{x}-\theta)
\end{aligned}
$$

Hence, the information matrix is

$$
\mathcal{F}_{n}=\operatorname{Var}\{n(\bar{x}-\theta)\}=n \mathcal{I}_{p} .
$$

How well can we estimate $\theta$ ? The answer is given in the following theorem which is from Cramer and Rao. As pointed out above, this theorem gives a lower bound for unbiased estimators. Hence, all estimators, which are unbiased and attain this lower bound, are minimum variance estimators.

Theorem 6.2 (Cramer-Rao) If $\hat{\theta}=t=t(\mathcal{X})$ is any unbiased estimator for $\theta$, then under regularity conditions

$$
\begin{equation*}
\operatorname{Var}(t) \geq \mathcal{F}_{n}^{-1}, \tag{6.13}
\end{equation*}
$$

where

$$
\begin{equation*}
\mathcal{F}_{n}=\mathrm{E}\left\{s(\mathcal{X} ; \theta) s(\mathcal{X} ; \theta)^{\top}\right\}=\operatorname{Var}\{s(\mathcal{X} ; \theta)\} \tag{6.14}
\end{equation*}
$$

is the Fisher information matrix.
Proof Consider the correlation $\rho_{Y, Z}$ between $Y$ and $Z$ where $Y=a^{\top} t, Z=c^{\top} s$. Here $s$ is the score and the vectors $a, c \in \mathbb{R}^{p}$. By Corollary 6.1 $\operatorname{Cov}(s, t)=\mathcal{I}$ and thus

$$
\begin{aligned}
\operatorname{Cov}(Y, Z) & =a^{\top} \operatorname{Cov}(t, s) c=a^{\top} c \\
\operatorname{Var}(Z) & =c^{\top} \operatorname{Var}(s) c=c^{\top} \mathcal{F}_{n} c .
\end{aligned}
$$

Hence,

$$
\begin{equation*}
\rho_{Y, Z}^{2}=\frac{\operatorname{Cov}^{2}(Y, Z)}{\operatorname{Var}(Y) \operatorname{Var}(Z)}=\frac{\left(a^{\top} c\right)^{2}}{a^{\top} \operatorname{Var}(t) a \cdot c^{\top} \mathcal{F}_{n} c} \leq 1 \tag{6.15}
\end{equation*}
$$

In particular, this holds for any $c \neq 0$. Therefore it holds also for the maximum of the left-hand side of (6.15) with respect to $c$. Since

$$
\max _{c} \frac{c^{\top} a a^{\top} c}{c^{\top} \mathcal{F}_{n} c}=\max _{c^{\top} \mathcal{F}_{n} c=1} c^{\top} a a^{\top} c
$$

and

$$
\max _{c^{\top} \mathcal{F}_{n} c=1} c^{\top} a a^{\top} c=a^{\top} \mathcal{F}_{n}^{-1} a
$$

by our maximisation Theorem 2.5, we have

$$
\frac{a^{\top} \mathcal{F}_{n}^{-1} a}{a^{\top} \operatorname{Var}(t) a} \leq 1 \quad \forall a \in \mathbb{R}^{p}, \quad a \neq 0
$$

i.e.,

$$
a^{\top}\left\{\operatorname{Var}(t)-\mathcal{F}_{n}^{-1}\right\} a \geq 0 \quad \forall a \in \mathbb{R}^{p}, \quad a \neq 0
$$

which is equivalent to $\operatorname{Var}(t) \geq \mathcal{F}_{n}^{-1}$.
MLEs attain the lower bound if the sample size $n$ goes to infinity. The next Theorem 6.3 states this and, in addition, gives the asymptotic sampling distribution of the maximum likelihood estimation, which turns out to be multinormal.
Theorem 6.3 Suppose that the sample $\left\{x_{i}\right\}_{i=1}^{n}$ is i.i.d. If $\hat{\theta}$ is the MLE for $\theta \in \mathbb{R}^{k}$, i.e., $\hat{\theta}=\arg \max _{\theta} L(\mathcal{X} ; \theta)$, then under some regularity conditions, as $n \rightarrow \infty$ :

$$
\begin{equation*}
\sqrt{n}(\hat{\theta}-\theta) \xrightarrow{\mathcal{L}} N_{k}\left(0, \mathcal{F}_{1}^{-1}\right) \tag{6.16}
\end{equation*}
$$

where $\mathcal{F}_{1}$ denotes the Fisher information for sample size $n=1$.
As a consequence of Theorem 6.3 we see that under regularity conditions the MLE is asymptotically unbiased, efficient (minimum variance) and normally distributed. Also it is a consistent estimator of $\theta$.

Note that from property (5.4) of the multinormal it follows that asymptotically

$$
\begin{equation*}
n(\hat{\theta}-\theta)^{\top} \mathcal{F}_{1}(\hat{\theta}-\theta) \xrightarrow{\mathcal{L}} \chi_{p}^{2} . \tag{6.17}
\end{equation*}
$$

If $\hat{\mathcal{F}}_{1}$ is a consistent estimator of $\mathcal{F}_{1}$ (e.g. $\hat{\mathcal{F}}_{1}=\mathcal{F}_{1}(\hat{\theta})$ ), we have equivalently

$$
\begin{equation*}
n(\hat{\theta}-\theta)^{\top} \hat{\mathcal{F}}_{1}(\hat{\theta}-\theta) \xrightarrow{\mathcal{L}} \chi_{p}^{2} . \tag{6.18}
\end{equation*}
$$

This expression is sometimes useful in testing hypotheses about $\theta$ and in constructing confidence regions for $\theta$ in a very general setup. These issues will be raised in more detail in the next chapter, but from (6.18) it can be seen, for instance, that when $n$ is large,

$$
\mathrm{P}\left\{n(\hat{\theta}-\theta)^{\top} \hat{\mathcal{F}}_{1}(\hat{\theta}-\theta) \leq \chi_{1-\alpha ; p}^{2}\right\} \approx 1-\alpha,
$$

where $\chi_{\nu ; p}^{2}$ denotes the $v$-quantile of a $\chi_{p}^{2}$ random variable. So, the ellipsoid $n(\hat{\theta}-$ $\theta)^{\top} \hat{\mathcal{F}}_{1}(\hat{\theta}-\theta) \leq \chi_{1-\alpha ; p}^{2}$ provides in $\mathbb{R}^{p}$ an asymptotic $(1-\alpha)$-confidence region for $\theta$.


### 6.3 Exercises

Exercise 6.1 Consider a uniform distribution on the interval $[0, \theta]$. What is the MLE of $\theta$ ? (Hint: the maximisation here cannot be performed by means of derivatives. Here the support of $x$ depends on $\theta$.)

Exercise 6.2 Consider an i.i.d. sample of size $n$ from the bivariate population with pdf $f\left(x_{1}, x_{2}\right)=\left(\theta_{1} \theta_{2}\right)^{-1} \exp \left(-x_{1} / \theta_{1}-x_{2} / \theta_{2}\right), x_{1}, x_{2}>0$. Compute the MLE of $\theta=\left(\theta_{1}, \theta_{2}\right)$. Find the Cramer-Rao lower bound. Is it possible to derive a minimal variance unbiased estimator of $\theta$ ?

Exercise 6.3 Show that the MLE of Example 6.1, $\hat{\mu}=\bar{x}$, is a minimal variance estimator for any finite sample size $n$ (i.e., without applying Theorem 6.3).

Exercise 6.4 We know from Example 6.4 that the MLE of Example 6.1 has $\mathcal{F}_{1}=$ $\mathcal{I}_{p}$. This leads to

$$
\sqrt{n}(\bar{x}-\mu) \xrightarrow{\mathcal{L}} N_{p}(0, \mathcal{I})
$$

by Theorem 6.3. Can you give an analogous result for the square $\bar{x}^{2}$ for the case $p=1$ ?

Exercise 6.5 Consider an i.i.d. sample of size $n$ from the bivariate population with pdf $f\left(x_{1}, x_{2}\right)=\left(\theta_{1}^{2} \theta_{2} x_{2}\right)^{-1} \exp \left(-x_{1} / \theta_{1} x_{2}-x_{2} / \theta_{1} \theta_{2}\right), x_{1}, x_{2}>0$. Compute the MLE of $\theta=\left(\theta_{1}, \theta_{2}\right)$. Find the Cramer-Rao lower bound and the asymptotic variance of $\hat{\theta}$.

Exercise 6.6 Consider a sample $\left\{x_{i}\right\}_{i=1}^{n}$ from $N_{p}\left(\mu, \Sigma_{0}\right)$ where $\Sigma_{0}$ is known. Compute the Cramer-Rao lower bound for $\mu$. Can you derive a minimal unbiased estimator for $\mu$ ?

Exercise 6.7 Let $X \sim N_{p}(\mu, \Sigma)$ where $\Sigma$ is unknown but we know $\Sigma=\operatorname{diag}\left(\sigma_{11}, \sigma_{22}, \ldots, \sigma_{p p}\right)$. From an i.i.d. sample of size $n$, find the MLE of $\mu$ and of $\Sigma$.

Exercise 6.8 Reconsider the setup of the previous exercise. Suppose that

$$
\Sigma=\operatorname{diag}\left(\sigma_{11}, \sigma_{22}, \ldots, \sigma_{p p}\right)
$$

Can you derive in this case the Cramer-Rao lower bound for $\theta^{\top}=$ $\left(\mu_{1} \ldots \mu_{p}, \sigma_{11} \ldots \sigma_{p p}\right)$ ?

Exercise 6.9 Prove Theorem 6.1. Hint: start from $\frac{\partial}{\partial \theta} \mathrm{E}\left(t^{\top}\right)=\frac{\partial}{\partial \theta} \int t^{\top}(\mathcal{X} ; \theta)$ $L(\mathcal{X} ; \theta) d \mathcal{X}$, then permute integral and derivatives and note that $s(\mathcal{X} ; \theta)=$ $\frac{1}{L(\mathcal{X} ; \theta)} \frac{\partial}{\partial \theta} L(\mathcal{X} ; \theta)$.
Exercise 6.10 Prove expression (6.12).
(Hint: start from $\mathrm{E}\{s(\mathcal{X} ; \theta)\}=\int \frac{1}{L(\mathcal{X} ; \theta)} \frac{\partial}{\partial \theta} L(\mathcal{X} ; \theta) L(\mathcal{X} ; \theta) \partial \mathcal{X}$ and then permute integral and derivative.)

## Chapter 7 <br> Hypothesis Testing

In the preceding chapter, the theoretical basis of estimation theory was presented. Now we turn our interest towards testing issues: we want to test the hypothesis $H_{0}$ that the unknown parameter $\theta$ belongs to some subspace of $\mathbb{R}^{q}$. This subspace is called the null set and will be denoted by $\Omega_{0} \subset \mathbb{R}^{q}$.

In many cases, this null set corresponds to restrictions which are imposed on the parameter space: $H_{0}$ corresponds to a "reduced model". As we have already seen in Chap. 3, the solution to a testing problem is in terms of a rejection region $R$ which is a set of values in the sample space which leads to the decision of rejecting the null hypothesis $H_{0}$ in favour of an alternative $H_{1}$, which is called the "full model".

In general, we want to construct a rejection region $R$ which controls the size of the type I error, i.e. the probability of rejecting the null hypothesis when it is true. More formally, a solution to a testing problem is of predetermined size $\alpha$ if:

$$
\mathrm{P}\left(\text { Rejecting } H_{0} \mid H_{0} \text { is true }\right)=\alpha .
$$

In fact, since $H_{0}$ is often a composite hypothesis, it is achieved by finding $R$ such that

$$
\sup _{\theta \in \Omega_{0}} \mathrm{P}(\mathcal{X} \in R \mid \theta)=\alpha
$$

In this chapter we will introduce a tool which allows us to build a rejection region in general situations; it is based on the likelihood ratio principle. This is a very useful technique because it allows us to derive a rejection region with an asymptotically appropriate size $\alpha$. The technique will be illustrated through various testing problems and examples. We concentrate on multinormal populations and linear models where the size of the test will often be exact even for finite sample sizes $n$.

Section 7.1 gives the basic ideas and Sect. 7.2 presents the general problem of testing linear restrictions. This allows us to propose solutions to frequent types
of analyses (including comparisons of several means, repeated measurements and profile analysis). Each case can be viewed as a simple specific case of testing linear restrictions. Special attention is devoted to confidence intervals and confidence regions for means and for linear restrictions on means in a multinormal setup.

### 7.1 Likelihood Ratio Test

Suppose that the distribution of $\left\{x_{i}\right\}_{i=1}^{n}, x_{i} \in \mathbb{R}^{p}$, depends on a parameter vector $\theta$. We will consider two hypotheses:

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

The hypothesis $H_{0}$ corresponds to the "reduced model" and $H_{1}$ to the "full model". This notation was already used in Chap. 3.

Example 7.1 Consider a multinormal $N_{p}(\theta, \mathcal{I})$. To test if $\theta$ equals a certain fixed value $\theta_{0}$ we construct the test problem:

$$
\begin{aligned}
& H_{0}: \theta=\theta_{0} \\
& H_{1}: \text { no constraints on } \theta
\end{aligned}
$$

or, equivalently, $\Omega_{0}=\left\{\theta_{0}\right\}, \Omega_{1}=\mathbb{R}^{p}$.
Define $L_{j}^{*}=\max _{\theta \in \Omega_{j}} L(\mathcal{X} ; \theta)$, the maxima of the likelihood for each of the hypotheses. Consider the likelihood ratio (LR)

$$
\begin{equation*}
\lambda(\mathcal{X})=\frac{L_{0}^{*}}{L_{1}^{*}} \tag{7.1}
\end{equation*}
$$

One tends to favour $H_{0}$ if the LR is high and $H_{1}$ if the LR is low. The likelihood ratio test (LRT) tells us when exactly to favour $H_{0}$ over $H_{1}$. A LRT of size $\alpha$ for testing $H_{0}$ against $H_{1}$ has the rejection region

$$
R=\{\mathcal{X}: \lambda(\mathcal{X})<c\}
$$

where $c$ is determined so that $\sup _{\theta \in \Omega_{0}} \mathrm{P}_{\theta}(\mathcal{X} \in R)=\alpha$. The difficulty here is to express $c$ as a function of $\alpha$, because $\lambda(\mathcal{X})$ might be a complicated function of $\mathcal{X}$.

Instead of $\lambda$ we may equivalently use the log-likelihood

$$
-2 \log \lambda=2\left(\ell_{1}^{*}-\ell_{0}^{*}\right)
$$

In this case the rejection region will be $R=\{\mathcal{X}:-2 \log \lambda(\mathcal{X})>k\}$. What is the distribution of $\lambda$ or of $-2 \log \lambda$ from which we need to compute $c$ or $k$ ?

Theorem 7.1 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

$$
\forall \theta \in \Omega_{0}:-2 \log \lambda \xrightarrow{\mathcal{L}} \chi_{q-r}^{2} \quad \text { as } \quad n \rightarrow \infty .
$$

An asymptotic rejection region can now be given by simply computing the $1-\alpha$ quantile $k=\chi_{1-\alpha ; q-r}^{2}$. The LRT rejection region is therefore

$$
R=\left\{\mathcal{X}:-2 \log \lambda(\mathcal{X})>\chi_{1-\alpha ; q-r}^{2}\right\}
$$

Theorem 7.1 is thus very helpful: it gives a general way of building rejection regions into many problems. Unfortunately, it is only an asymptotic result, meaning that the size of the test is only approximately equal to $\alpha$, although the approximation becomes better when the sample size $n$ increases. The question is "how large should $n$ be?". There is no definite rule: we encounter here the same problem that was already discussed with respect to the Central Limit Theorem in Chap. 4.

Fortunately, in many standard circumstances, we can derive exact tests even for finite samples because the test statistic $-2 \log \lambda(\mathcal{X})$ or a simple transformation of it turns out to have a simple form. This is the case in most of the following standard testing problems. All of them can be viewed as an illustration of the likelihood ratio principle.

Test Problem 1 is an amuse-bouche: in testing the mean of a multinormal population with a known covariance matrix the likelihood ratio statistic has a very simple quadratic form with a known distribution under $H_{0}$.

Test Problem 1. Suppose that $X_{1}, \ldots, X_{n}$ is an i.i.d. random sample from a $N_{p}(\mu, \Sigma)$ population.

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

In this case $H_{0}$ is a simple hypothesis, i.e. $\Omega_{0}=\left\{\mu_{0}\right\}$ and therefore the dimension $r$ of $\Omega_{0}$ equals 0 . Since we have imposed no constraints in $H_{1}$, the space $\Omega_{1}$ is the whole $\mathbb{R}^{p}$ which leads to $q=p$. From (6.6) we know that

$$
\ell_{0}^{*}=\ell\left(\mu_{0}, \Sigma\right)=-\frac{n}{2} \log |2 \pi \Sigma|-\frac{1}{2} n \operatorname{tr}\left(\Sigma^{-1} \mathcal{S}\right)-\frac{1}{2} n\left(\bar{x}-\mu_{0}\right)^{\top} \Sigma^{-1}\left(\bar{x}-\mu_{0}\right) .
$$

Under $H_{1}$ the maximum of $\ell(\mu, \Sigma)$ is

$$
\ell_{1}^{*}=\ell(\bar{x}, \Sigma)=-\frac{n}{2} \log |2 \pi \Sigma|-\frac{1}{2} n \operatorname{tr}\left(\Sigma^{-1} \mathcal{S}\right)
$$

Therefore,

$$
\begin{equation*}
-2 \log \lambda=2\left(\ell_{1}^{*}-\ell_{0}^{*}\right)=n\left(\bar{x}-\mu_{0}\right)^{\top} \Sigma^{-1}\left(\bar{x}-\mu_{0}\right) \tag{7.2}
\end{equation*}
$$

which, by Theorem 4.7, has a $\chi_{p}^{2}$-distribution under $H_{0}$.
Example 7.2 Consider the bank data again. Let us test whether the population mean of the forged bank notes is equal to

$$
\mu_{0}=(214.9,129.9,129.7,8.3,10.1,141.5)^{\top}
$$

(This is in fact the sample mean of the genuine bank notes.) The sample mean of the forged bank notes is

$$
\bar{x}=(214.8,130.3,130.2,10.5,11.1,139.4)^{\top} .
$$

Suppose for the moment that the estimated covariance matrix $\mathcal{S}_{f}$ given in (3.5) is the true covariance matrix $\Sigma$. We construct the LRT and obtain

$$
\begin{aligned}
-2 \log \lambda & =2\left(\ell_{1}^{*}-\ell_{0}^{*}\right)=n\left(\bar{x}-\mu_{0}\right)^{\top} \Sigma^{-1}\left(\bar{x}-\mu_{0}\right) \\
& =7362.32,
\end{aligned}
$$

the quantile $k=\chi_{0.95 ; 6}^{2}$ equals 12.592. The rejection consists of all values in the sample space which lead to values of the LRT statistic larger than 12.592. Under $H_{0}$ the value of $-2 \log \lambda$ is therefore highly significant. Hence, the true mean of the forged bank notes is significantly different from $\mu_{0}$ !

Test Problem 2 is the same as the preceding one but in a more realistic situation where the covariance matrix is unknown; here the Hotelling's $T^{2}$-distribution will be useful to determine an exact test and a confidence region for the unknown $\mu$.

Test Problem 2. Suppose that $X_{1}, \ldots, X_{n}$ is an i.i.d. random sample from a $N_{p}(\mu, \Sigma)$ population.

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

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

$$
\begin{align*}
\mathcal{S}_{0} & =\frac{1}{n}\left[x-1_{n} \mu_{0}^{\top}-1_{n} \bar{x}^{\top}+1_{n} \bar{x}^{\top}\right]^{\top}\left[x-1_{n} \mu_{0}^{\top}-1_{n} \bar{x}^{\top}+1_{n} \bar{x}^{\top}\right] \\
& =\mathcal{S}+\left(\bar{x}-\mu_{0}\right)\left(\bar{x}-\mu_{0}\right)^{\top} \\
\ell_{0}^{*} & =\ell\left(\mu_{0}, \mathcal{S}+d d^{\top}\right), \quad d=\left(\bar{x}-\mu_{0}\right) \tag{7.3}
\end{align*}
$$

and under $H_{1}$ we have

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

This leads after some calculation to

$$
\begin{aligned}
-2 \log \lambda= & 2\left(\ell_{1}^{*}-\ell_{0}^{*}\right) \\
= & -n \log |\mathcal{S}|-n \operatorname{tr}\left(\mathcal{S}^{-1} \mathcal{S}\right)-n(\bar{x}-\bar{x})^{\top} \mathcal{S}^{-1}(\bar{x}-\bar{x})+n \log \left|\mathcal{S}+d d^{\top}\right| \\
& +n \operatorname{tr}\left[\left(\mathcal{S}+d d^{\top}\right)^{-1} \mathcal{S}\right]+n\left(\bar{x}-\mu_{0}\right)^{\top}\left(\mathcal{S}+d d^{\top}\right)^{-1}\left(\bar{x}-\mu_{0}\right) \\
= & n \log \left|\frac{\mathcal{S}+d d^{\top}}{\mathcal{S}}\right|+n \operatorname{tr}\left[\left(\mathcal{S}+d d^{\top}\right)^{-1} \mathcal{S}\right]+n d^{\top}\left(\mathcal{S}+d d^{\top}\right)^{-1} d-n p \\
= & n \log \left|\frac{\mathcal{S}+d d^{\top}}{\mathcal{S}}\right|+n \operatorname{tr}\left[\left(\mathcal{S}+d d^{\top}\right)^{-1}\left(d d^{\top}+\mathcal{S}\right)\right]-n p \\
= & n \log \left|\frac{\mathcal{S}+d d^{\top}}{\mathcal{S}}\right| \\
= & n \log \left|1+\mathcal{S}^{-1 / 2} d d^{\top} \mathcal{S}^{-1 / 2}\right| .
\end{aligned}
$$

By using the result for the determinant of a partitioned matrix, it equals to

$$
\begin{aligned}
& n \log \left|\begin{array}{cc}
1 & -d^{\top} \mathcal{S}^{-1 / 2} \\
\mathcal{S}^{-1 / 2} d & I
\end{array}\right| \\
& \quad=n \log \left|\begin{array}{ccccc}
1 & -d^{\top} \mathcal{S}^{-1 / 2}{ }_{1}-d^{\top} \mathcal{S}^{-1 / 2}{ }_{2} \ldots & \ldots d^{\top} \mathcal{S}^{-1 / 2}{ }_{p} \\
\mathcal{S}^{-1 / 2} d_{1} & 1 & 0 & \ldots & 0 \\
\mathcal{S}^{-1 / 2} d_{2} & 0 & 1 & & 0 \\
\vdots & \vdots & & \ddots & \\
\mathcal{S}^{-1 / 2} d_{p} & 0 & 0 & \ldots & 1
\end{array}\right|
\end{aligned}
$$

$$
\begin{align*}
& =n \log 1+n \log \sum_{i=1}^{p}-d^{\top} \mathcal{S}^{-1 / 2}{ }_{i}(-1)^{1+(i+1)}\left|\begin{array}{ccccc}
\mathcal{S}^{-1 / 2} d_{1} & 1 & 0 & \ldots & 0 \\
\mathcal{S}^{-1 / 2} d_{2} & 0 & 1 & \ldots & 0 \\
\vdots & & \ddots & \\
\mathcal{S}^{-1 / 2} d_{i} & 0 & 0 & \ldots & 0 \\
\vdots & & & & \\
\mathcal{S}^{-1 / 2} d_{p} & 0 & 0 & \ldots & 1
\end{array}\right| \\
& =n \log 1+\sum_{i=1}^{p}-d^{\top} \mathcal{S}^{-1 / 2}{ }_{i}(-1)^{2+i} \mathcal{S}^{-1 / 2} d_{i}(-1)^{i+1} \\
& =n \log \left(1+d^{\top} \mathcal{S}^{-1} d\right) . \tag{7.4}
\end{align*}
$$

This statistic is a monotone function of $(n-1) d^{\top} \mathcal{S}^{-1} d$. This means that $-2 \log \lambda>$ $k$ if and only if $(n-1) d^{\top} \mathcal{S}^{-1} d>k^{\prime}$. The latter statistic has by Corollary 5.3, under $H_{0}$, a Hotelling's $T^{2}$-distribution. Therefore,

$$
\begin{equation*}
(n-1)\left(\bar{x}-\mu_{0}\right)^{\top} \mathcal{S}^{-1}\left(\bar{x}-\mu_{0}\right) \sim T_{p, n-1}^{2} \tag{7.5}
\end{equation*}
$$

or equivalently

$$
\begin{equation*}
\left(\frac{n-p}{p}\right)\left(\bar{x}-\mu_{0}\right)^{\top} \mathcal{S}^{-1}\left(\bar{x}-\mu_{0}\right) \sim F_{p, n-p} \tag{7.6}
\end{equation*}
$$

In this case an exact rejection region may be defined as

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

Alternatively, we have from Theorem 7.1 that under $H_{0}$ the asymptotic distribution of the test statistic is

$$
-2 \log \lambda \xrightarrow{\mathcal{L}} \chi_{p}^{2}, \quad \text { as } n \rightarrow \infty
$$

which leads to the (asymptotically valid) rejection region

$$
n \log \left\{1+\left(\bar{x}-\mu_{0}\right)^{\top} S^{-1}\left(\bar{x}-\mu_{0}\right)\right\}>\chi_{1-\alpha ; p}^{2}
$$

but of course, in this case, we would prefer to use the exact $F$-test provided just above.
Example 7.3 Consider the problem of Example 7.2 again. We know that $\mathcal{S}_{f}$ is the empirical analogue for $\Sigma_{f}$, the covariance matrix for the forged banknotes. The test statistic (7.5) has the value $1,153.4$ or its equivalent for the $F$ distribution in (7.6)
is 182.5 which is highly significant $\left(F_{0.95 ; 6,94}=2.1966\right)$ so that we conclude that $\mu_{f} \neq \mu_{0}$.

## Confidence Region for $\mu$

When estimating a multidimensional parameter $\theta \in \mathbb{R}^{k}$ from a sample, we saw in Chap. 6 how to determine the estimator $\hat{\theta}=\hat{\theta}(\mathcal{X})$. For the observed data we end up with a point estimate, which is the corresponding observed value of $\hat{\theta}$. We know $\hat{\theta}(\mathcal{X})$ is a random variable and we often prefer to determine a confidence region for $\theta$. A confidence region (CR) is a random subset of $\mathbb{R}^{k}$ (determined by appropriate statistics) such that we are "confident", at a certain given level $1-\alpha$, that this region contains $\theta$ :

$$
\mathrm{P}(\theta \in \mathrm{CR})=1-\alpha .
$$

This is just a multidimensional generalisation of the basic univariate confidence interval. Confidence regions are particularly useful when a hypothesis $H_{0}$ on $\theta$ is rejected, because they eventually help in identifying which component of $\theta$ is responsible for the rejection.

There are only a few cases where confidence regions can be easily assessed, and include most of the testing problems on mean presented in this section.

Corollary 5.3 provides a pivotal quantity which allows confidence regions for $\mu$ to be constructed. Since $\left(\frac{n-p}{p}\right)(\bar{x}-\mu)^{\top} \mathcal{S}^{-1}(\bar{x}-\mu) \sim F_{p, n-p}$, we have

$$
\mathrm{P}\left\{\left(\frac{n-p}{p}\right)(\mu-\bar{x})^{\top} \mathcal{S}^{-1}(\mu-\bar{x})<F_{1-\alpha ; p, n-p}\right\}=1-\alpha .
$$

Then,

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

is a confidence region at level $(1-\alpha)$ for $\mu$. It is the interior of an iso-distance ellipsoid in $\mathbb{R}^{p}$ centred at $\bar{x}$, with a scaling matrix $\mathcal{S}^{-1}$ and a distance constant $\left(\frac{p}{n-p}\right) F_{1-\alpha ; p, n-p}$. 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 of say, $1-\alpha$.

Below, we consider a more general problem. We construct simultaneous confidence intervals for all possible linear combinations $a^{\top} \mu, a \in \mathbb{R}^{p}$ of the elements of $\mu$.

Suppose for a moment that we fix a particular projection vector $a$. We are back to a standard univariate problem of finding a confidence interval for the mean $a^{\top} \mu$ of a univariate random variable $a^{\top} X$. We can use the $t$-statistics and an obvious confidence interval for $a^{\top} \mu$ is given by the values $a^{\top} \mu$ such that

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

or equivalently

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

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

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

Now it is easy to prove (using Theorem 2.5) that:

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

Therefore, simultaneously for all $a \in \mathbb{R}^{p}$, the interval

$$
\begin{equation*}
\left(a^{\top} \bar{x}-\sqrt{K_{\alpha} a^{\top} \mathcal{S} a}, a^{\top} \bar{x}+\sqrt{K_{\alpha} a^{\top} \mathcal{S} a}\right) \tag{7.7}
\end{equation*}
$$

where $K_{\alpha}=\frac{p}{n-p} F_{1-\alpha ; p, n-p}$, will contain $a^{\top} \mu$ with probability $(1-\alpha)$.
A particular choice of $a$ are the columns of the identity matrix $\mathcal{I}_{p}$, providing simultaneous confidence intervals for $\mu_{1}, \ldots, \mu_{p}$. We therefore have with probability $(1-\alpha)$ for $j=1, \ldots, p$

$$
\begin{equation*}
\bar{x}_{j}-\sqrt{\frac{p}{n-p} F_{1-\alpha ; p, n-p} s_{j j}} \leq \mu_{j} \leq \bar{x}_{j}+\sqrt{\frac{p}{n-p} F_{1-\alpha ; p, n-p} s_{j j}} . \tag{7.8}
\end{equation*}
$$

It should be noted that these intervals define a rectangle inscribing the confidence ellipsoid for $\mu$ given above. They are particularly useful when a null hypothesis $H_{0}$ of the type described above is rejected and one would like to see which component(s) are mainly responsible for the rejection.

Example 7.4 The $95 \%$ confidence region for $\mu_{f}$, the mean of the forged banknotes, is given by the ellipsoid:

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

The $95 \%$ simultaneous confidence intervals are given by (we use $F_{0.95 ; 6,94}=$ 2.1966)

$$
\begin{aligned}
& 214.692 \leq \mu_{1} \leq \quad 214.954 \\
& 130.205 \leq \mu_{2} \leq 130.395 \\
& 130.082 \leq \mu_{3} \leq 130.304 \\
& 10.108 \leq \mu_{4} \leq 10.952 \\
& 10.896 \leq \mu_{5} \leq \\
& \leq 11.370 \\
& 139.242 \leq \mu_{6} \leq \quad 139.658 .
\end{aligned}
$$

Comparing the inequalities 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 Examples 7.2 and 7.3.

In addition, the method can provide other confidence intervals. We have at the same level of confidence (choosing $a^{\top}=(0,0,0,1,-1,0)$ )

$$
-1.211 \leq \mu_{4}-\mu_{5} \leq 0.005
$$

showing that for the forged bills, the lower border is essentially smaller than the upper border.

Remark 7.1 It should be noted that the confidence region is an ellipsoid whose characteristics depend on the whole matrix $\mathcal{S}$. In particular, the slope of the axis depends on the eigenvectors of $S$ and therefore on the covariances $s_{i j}$. However, the rectangle inscribing the confidence ellipsoid provides the simultaneous confidence intervals for $\mu_{j}, j=1, \ldots, p$. They do not depend on the covariances $s_{i j}$, but only on the variances $s_{i j}$ [see (7.8)]. In particular, it may happen that a tested value $\mu_{0}$ is covered by the confidence ellipsoid but not covered by the intervals (7.8). In this case, $\mu_{0}$ is rejected by a test based on the simultaneous confidence intervals but not rejected by a test based on the confidence ellipsoid. The simultaneous confidence intervals are easier to handle than the full ellipsoid but we have lost some information, namely the covariance between the components (see Exercise 7.14).

The following problem concerns the covariance matrix in a multinormal population: in this situation the test statistic has a slightly more complicated distribution. We will therefore invoke the approximation of Theorem 7.1 in order to derive a test of approximate size $\alpha$.

Test Problem 3. Suppose that $X_{1}, \ldots, X_{n}$ is an i.i.d. random sample from a $N_{p}(\mu, \Sigma)$ population.

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

Under $H_{0}$ we have $\hat{\mu}=\bar{x}$, and $\Sigma=\Sigma_{0}$, whereas under $H_{1}$ we have $\hat{\mu}=\bar{x}$, and $\hat{\Sigma}=\mathcal{S}$. Hence

$$
\begin{aligned}
& \ell_{0}^{*}=\ell\left(\bar{x}, \Sigma_{0}\right)=-\frac{1}{2} n \log \left|2 \pi \Sigma_{0}\right|-\frac{1}{2} n \operatorname{tr}\left(\Sigma_{0}^{-1} \mathcal{S}\right) \\
& \ell_{1}^{*}=\ell(\bar{x}, \mathcal{S})=-\frac{1}{2} n \log |2 \pi \mathcal{S}|-\frac{1}{2} n p
\end{aligned}
$$

and thus

$$
\begin{aligned}
-2 \log \lambda & =2\left(\ell_{1}^{*}-\ell_{0}^{*}\right) \\
& =n \operatorname{tr}\left(\Sigma_{0}^{-1} \mathcal{S}\right)-n \log \left|\Sigma_{0}^{-1} \mathcal{S}\right|-n p
\end{aligned}
$$

Note that this statistic is a function of the eigenvalues of $\Sigma_{0}^{-1} \mathcal{S}$. Unfortunately, the exact finite sample distribution of $-2 \log \lambda$ is very complicated. Asymptotically, we have under $H_{0}$

$$
-2 \log \lambda \xrightarrow{\mathcal{L}} \chi_{m}^{2} \quad \text { as } \quad n \rightarrow \infty
$$

with $m=\frac{1}{2}\{p(p+1)\}$, since a $(p \times p)$ covariance matrix has only these $m$ parameters as a consequence of its symmetry.

Example 7.5 Consider the US companies data set (Table 22.5) and suppose we are interested in the companies of the energy sector, analysing their assets $\left(X_{1}\right)$ and sales $\left(X_{2}\right)$. The sample is of size 15 and provides the value of $\mathcal{S}=10^{7} \times$ $\left[\begin{array}{ll}1.6635 & 1.2410 \\ 1.2410 & 1.3747\end{array}\right]$. We want to test if $\operatorname{Var}\binom{X_{1}}{X_{2}}=10^{7} \times\left[\begin{array}{ll}1.2248 & 1.1425 \\ 1.1425 & 1.5112\end{array}\right]=\Sigma_{0}$. ( $\Sigma_{0}$ is in fact the empirical variance matrix for $X_{1}$ and $X_{2}$ for the manufacturing sector). The test statistic (a MVAusenergy) turns out to be $-2 \log \lambda=5.4046$ which is not significant for $\chi_{3}^{2}$ ( $p$-value $=0.1445$ ). So we cannot conclude that $\Sigma \neq \Sigma_{0}$.

In the next testing problem, we address a question that was already stated in Chap. 3, Sect. 3.6: testing a particular value of the coefficients $\beta$ in a linear model. The presentation is carried out in general terms so that it can be built on in the next section where we will test linear restrictions on $\beta$.

Test Problem 4. Suppose that $Y_{1}, \ldots, Y_{n}$ are independent r.v.'s with $Y_{i} \sim N_{1}\left(\beta^{\top} x_{i}, \sigma^{2}\right), x_{i} \in \mathbb{R}^{p}$.
$H_{0}: \beta=\beta_{0}, \sigma^{2}$ unknown versus $H_{1}:$ no constraints.

Under $H_{0}$ we have $\beta=\beta_{0}, \hat{\sigma}_{0}^{2}=\frac{1}{n}\left\|y-\mathcal{X} \beta_{0}\right\|^{2}$ and under $H_{1}$ we have $\hat{\beta}=$ $\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1} \mathcal{X}^{\top} y, \hat{\sigma}^{2}=\frac{1}{n}\|y-\mathcal{X} \hat{\beta}\|^{2}$ (see Example 6.3). Hence by Theorem 7.1

$$
\begin{aligned}
-2 \log \lambda & =2\left(\ell_{1}^{*}-\ell_{0}^{*}\right) \\
& =n \log \left(\frac{\left\|y-\mathcal{X} \beta_{0}\right\|^{2}}{\|y-\mathcal{X} \hat{\beta}\|^{2}}\right) \\
& \xrightarrow{\mathcal{L}} \chi_{p}^{2}
\end{aligned}
$$

We draw upon the result (3.45) which gives us

$$
F=\frac{(n-p)}{p}\left(\frac{\left\|y-\mathcal{X} \beta_{0}\right\|^{2}}{\|y-\mathcal{X} \hat{\beta}\|^{2}}-1\right) \sim F_{p, n-p}
$$

so that in this case we again have an exact distribution.
Example 7.6 Let us consider our "classic blue" pullovers again. In Example 3.11 we tried to model the dependency of sales on prices. As we have seen in Fig. 3.5 the slope of the regression curve is rather small, hence we might ask if $\binom{\alpha}{\beta}=\binom{211}{0}$. Here

$$
y=\left(\begin{array}{c}
y_{1} \\
\vdots \\
y_{10}
\end{array}\right)=\left(\begin{array}{c}
x_{1,1} \\
\vdots \\
x_{10,1}
\end{array}\right), \quad \mathcal{X}=\left(\begin{array}{cc}
1 & x_{1,2} \\
\vdots & \vdots \\
1 & x_{10,2}
\end{array}\right)
$$

The test statistic for the LR test is

$$
-2 \log \lambda=9.10
$$

which under the $\chi_{2}^{2}$ distribution is significant. The exact $F$-test statistic

$$
F=5.93
$$

is also significant under the $F_{2,8}$ distribution $\left(F_{2,8 ; 0.95}=4.46\right)$.


## Summary

$\hookrightarrow$ The hypotheses $H_{0}: \theta \in \Omega_{0}$ against $H_{1}: \theta \in \Omega_{1}$ can be tested using the LRT. The likelihood ratio (LR) is the quotient $\lambda(\mathcal{X})=$ $L_{0}^{*} / L_{1}^{*}$ where the $L_{j}^{*}$ are the maxima of the likelihood for each of the hypotheses.

| Summary (continued) |
| :---: |
| $\hookrightarrow$ The test statistic in the LRT is $\lambda(\mathcal{X})$ or equivalently its logarithm $\log \lambda(\mathcal{X})$. If $\Omega_{1}$ is $q$-dimensional and $\Omega_{0} \subset \Omega_{1} r$-dimensional, then the asymptotic distribution of $-2 \log \lambda$ is $\chi_{q-r}^{2}$. This allows $H_{0}$ to be tested against $H_{1}$ by calculating the test statistic $-2 \log \lambda=$ $2\left(\ell_{1}^{*}-\ell_{0}^{*}\right)$ where $\ell_{j}^{*}=\log L_{j}^{*}$. |
| $\hookrightarrow$ The hypothesis $H_{0}: \mu=\mu_{0}$ for $X \sim N_{p}(\mu, \Sigma)$, where $\Sigma$ is known, leads to $-2 \log \lambda=n\left(\bar{x}-\mu_{0}\right)^{\top} \Sigma^{-1}\left(\bar{x}-\mu_{0}\right) \sim \chi_{p}^{2}$. |
| $\hookrightarrow$ The hypothesis $H_{0}: \mu=\mu_{0}$ for $X \sim N_{p}(\mu, \Sigma)$, where $\Sigma$ is unknown, leads to $-2 \log \lambda=n \log \left\{1+\left(\bar{x}-\mu_{0}\right)^{\top} \mathcal{S}^{-1}(\bar{x}-\right.$ $\left.\left.\mu_{0}\right)\right\} \longrightarrow \chi_{p}^{2}$, and $(n-1)\left(\bar{x}-\mu_{0}\right)^{\top} \mathcal{S}^{-1}\left(\bar{x}-\mu_{0}\right) \sim T_{p, n-1}^{2}$ |
| $\begin{aligned} & \hookrightarrow \text { The hypothesis } H_{0}: \Sigma=\Sigma_{0} \text { for } X \sim N_{p}(\mu, \Sigma) \text {, where } \mu \text { is } \\ & \text { unknown, leads to }-2 \log \lambda=n \operatorname{tr}\left(\Sigma_{0}^{-1} \mathcal{S}\right)-n \log \left\|\Sigma_{0}^{-1} \mathcal{S}\right\|-n p \longrightarrow \\ & \\ & \chi_{m}^{2}, m=\frac{1}{2} p(p+1) . \end{aligned}$ |
| $\hookrightarrow$ The hypothesis $H_{0}: \beta=\beta_{0}$ for $Y_{i} \sim N_{1}\left(\beta^{\top} x_{i}, \sigma^{2}\right)$, where $\sigma^{2}$ is unknown, leads to $-2 \log \lambda=n \log \left(\frac{\left\\|y-\mathcal{X} \beta_{0}\right\\|^{2}}{\\|y-\mathcal{X}\\|^{2}}\right) \longrightarrow \chi_{p}^{2}$. |

### 7.2 Linear Hypothesis

In this section, we present a very general procedure which allows a linear hypothesis to be tested, i.e. a linear restriction, either on a vector mean $\mu$ or on the coefficient $\beta$ of a linear model. The presented technique covers many of the practical testing problems on means or regression coefficients.

Linear hypotheses are of the form $\mathcal{A} \mu=a$ with known matrices $\mathcal{A}(q \times p)$ and $a(q \times 1)$ with $q \leq p$.
Example 7.7 Let $\mu=\left(\mu_{1}, \mu_{2}\right)^{\top}$. The hypothesis that $\mu_{1}=\mu_{2}$ can be equivalently written as:

$$
\mathcal{A} \mu=\left(\begin{array}{ll}
1 & -1
\end{array}\right)\binom{\mu_{1}}{\mu_{2}}=0=a .
$$

The general idea is to test a normal population $H_{0}: \mathcal{A} \mu=a$ (restricted model) against the full model $H_{1}$ where no restrictions are put on $\mu$. Due to the properties of the multinormal, we can easily adapt the Test Problems 1 and 2 to this new situation. Indeed we know, from Theorem 5.2, that $y_{i}=\mathcal{A} x_{i} \sim N_{q}\left(\mu_{y}, \Sigma_{y}\right)$, where $\mu_{y}=$ $\mathcal{A} \mu$ and $\Sigma_{y}=\mathcal{A} \Sigma \mathcal{A}^{\top}$.

Testing the null $H_{0}: \mathcal{A} \mu=a$, is the same as testing $H_{0}: \quad \mu_{y}=a$. The appropriate statistics are $\bar{y}$ and $\mathcal{S}_{y}$ which can be derived from the original statistics $\bar{x}$ and $\mathcal{S}$ available from $\mathcal{X}$ :

$$
\bar{y}=\mathcal{A} \bar{x}, \quad \mathcal{S}_{y}=\mathcal{A S} \mathcal{A}^{\top} .
$$

Here the difference between the translated sample mean and the tested value is $d=$ $\mathcal{A} \bar{x}-a$. We are now in the situation to proceed to Test Problems 5 and 6.

Test Problem 5. Suppose $X_{1}, \ldots, X_{n}$ is an i.i.d. random sample from a $N_{p}(\mu, \Sigma)$ population.

$$
H_{0}: \mathcal{A} \mu=a, \Sigma \text { known versus } H_{1}: \text { no constraints. }
$$

By (7.2) we have that, under $H_{0}$ :

$$
n(\mathcal{A} \bar{x}-a)^{\top}\left(\mathcal{A} \Sigma \mathcal{A}^{\top}\right)^{-1}(\mathcal{A} \bar{x}-a) \sim \mathcal{X}_{q}^{2}
$$

and we reject $H_{0}$ if this test statistic is too large at the desired significance level.
Example 7.8 We consider hypotheses on partitioned mean vectors $\mu=\binom{\mu_{1}}{\mu_{2}}$. Let us first look at

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

for $N_{2 p}\left(\binom{\mu_{1}}{\mu_{2}},\left(\begin{array}{cc}\Sigma & 0 \\ 0 & \Sigma\end{array}\right)\right)$ with known $\Sigma$. This is equivalent to $\mathcal{A}=(\mathcal{I},-\mathcal{I}), a=$ $(0, \ldots, 0)^{\top} \in \mathbb{R}^{p}$ and leads to

$$
-2 \log \lambda=n\left(\bar{x}_{1}-\bar{x}_{2}\right)(2 \Sigma)^{-1}\left(\bar{x}_{1}-\bar{x}_{2}\right) \sim \chi_{p}^{2}
$$

Another example is the test whether $\mu_{1}=0$, i.e.

$$
H_{0}: \mu_{1}=0, \text { versus } H_{1}: \text { no constraints, }
$$

for $N_{2 p}\left(\binom{\mu_{1}}{\mu_{2}},\left(\begin{array}{cc}\Sigma & 0 \\ 0 & \Sigma\end{array}\right)\right)$ with known $\Sigma$. This is equivalent to $\mathcal{A} \mu=a$ with $\mathcal{A}=$ $(\mathcal{I}, 0)$, and $a=(0, \ldots, 0)^{\top} \in \mathbb{R}^{p}$. Hence

$$
-2 \log \lambda=n \bar{x}_{1} \Sigma^{-1} \bar{x}_{1} \sim \chi_{p}^{2} .
$$

Test Problem 6. Suppose $X_{1}, \ldots, X_{n}$ is an i.i.d. random sample from a $N_{p}(\mu, \Sigma)$ population.
$H_{0}: \mathcal{A} \mu=a, \Sigma$ unknown versus $H_{1}:$ no constraints.

From Corollary (5.4) and under $H_{0}$ it follows immediately that

$$
\begin{equation*}
(n-1)(\mathcal{A} \bar{x}-a)^{\top}\left(\mathcal{A S} \mathcal{A}^{\top}\right)^{-1}(\mathcal{A} \bar{x}-a) \sim T_{q, n-1}^{2} \tag{7.9}
\end{equation*}
$$

since indeed under $H_{0}$,

$$
\mathcal{A} \bar{x} \sim N_{q}\left(a, n^{-1} \mathcal{A} \Sigma \mathcal{A}^{\top}\right)
$$

is independent of

$$
n \mathcal{A S} \mathcal{A}^{\top} \sim W_{q}\left(\mathcal{A} \Sigma \mathcal{A}^{\top}, n-1\right)
$$

Example 7.9 Let's come back again to the bank data set and suppose that we want to test if $\mu_{4}=\mu_{5}$, i.e. the hypothesis that the lower border mean equals the larger border mean for the forged bills. In this case:

$$
\left.\begin{array}{rl}
\mathcal{A} & =\left(\begin{array}{lllll}
0 & 0 & 0 & 1
\end{array}-10\right.
\end{array}\right)
$$

The test statistic is:

$$
99(\mathcal{A} \bar{x})^{\top}\left(\mathcal{A} S_{f} \mathcal{A}^{\top}\right)^{-1}(\mathcal{A} \bar{x}) \sim T_{1,99}^{2}=F_{1,99 .} .
$$

The observed value is 13.638 which is significant at the $5 \%$ level.

## Repeated Measurements

In many situations, $n$ independent sampling units are observed at $p$ different times or under $p$ different experimental conditions (different treatments, etc.). So here we repeat $p$ one-dimensional measurements on $n$ different subjects. For instance, we observe the results from $n$ students taking $p$ different exams. We end up with a $(n \times p)$ matrix. We can thus consider the situation where we have $X_{1}, \ldots, X_{n}$ i.i.d. from a normal distribution $N_{p}(\mu, \Sigma)$ when there are $p$ repeated measurements. The hypothesis of interest in this case is that there are no treatment effects, $H_{0}: \mu_{1}=$
$\mu_{2}=\cdots=\mu_{p}$. This hypothesis is a direct application of Test Problem 6. Indeed, introducing an appropriate matrix transform on $\mu$ we have

$$
H_{0}: \mathcal{C} \mu=0 \text { where } \mathcal{C}((p-1) \times p)=\left(\begin{array}{ccccc}
1 & -1 & 0 & \cdots & 0  \tag{7.10}\\
0 & 1 & -1 & \cdots & 0 \\
\vdots & \vdots & \vdots & \vdots & \vdots \\
0 & \cdots & 0 & 1 & -1
\end{array}\right)
$$

Note that in many cases one of the experimental conditions is the "control" (a placebo, standard drug or reference condition). Suppose it is the first component. In that case one is interested in studying differences to the control variable. The matrix $\mathcal{C}$ has therefore a different form

$$
\mathcal{C}((p-1) \times p)=\left(\begin{array}{ccccc}
1 & -1 & 0 & \cdots & 0 \\
1 & 0 & -1 & \cdots & 0 \\
\vdots & \vdots & \vdots & \vdots & \vdots \\
1 & 0 & 0 & \cdots & -1
\end{array}\right)
$$

By (7.9) the null hypothesis will be rejected if:

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

As a matter of fact, $\mathcal{C} \mu$ is the mean of the random variable $y_{i}=\mathcal{C} x_{i}$

$$
y_{i} \sim N_{p-1}\left(\mathcal{C} \mu, \mathcal{C} \Sigma \mathcal{C}^{\top}\right)
$$

Simultaneous confidence intervals for linear combinations of the mean of $y_{i}$ have been derived above in (7.7). For all $a \in \mathbb{R}^{p-1}$, with probability $(1-\alpha)$ we have

$$
a^{\top} \mathcal{C} \mu \in a^{\top} \mathcal{C} \bar{x} \pm \sqrt{\frac{(p-1)}{n-p+1} F_{1-\alpha ; p-1, n-p+1} a^{\top} \mathcal{C S C}^{\top} a}
$$

Due to the nature of the problem here, the row sums of the elements in $\mathcal{C}$ are zero: $\mathcal{C} 1_{p}=0$, therefore $a^{\top} \mathcal{C}$ is a vector having sum of elements equals to 0 . This is called a contrast. Let $b=\mathcal{C}^{\top} a$. We have $b^{\top} 1_{p}=\sum_{j=1}^{p} b_{j}=0$. The result above thus provides for all contrasts of $\mu$, and $b^{\top} \mu$ simultaneous confidence intervals at level $(1-\alpha)$

$$
b^{\top} \mu \in b^{\top} \bar{x} \pm \sqrt{\frac{(p-1)}{n-p+1} F_{1-\alpha ; p-1, n-p+1} b^{\top} \mathcal{S} b}
$$

Examples of contrasts for $p=4$ are $b^{\top}=(1-100)$ or $(100-1)$ or even ( $1-\frac{1}{3}-\frac{1}{3}-\frac{1}{3}$ ) when the control is to be compared with the mean of three different treatments.

Example 7.10 Bock (1975) considers the evolution of the vocabulary of children from the eighth through eleventh grade. The data set contains the scores of a vocabulary test of 40 randomly chosen children. This is a repeated measurement situation, ( $n=40, p=4$ ), since the same children were observed from grades 8 to 11. The statistics of interest are:

$$
\left.\begin{array}{rl}
\bar{x} & =(1.086,2.544,2.851,3.420
\end{array}\right)^{\top} .
$$

Suppose we are interested in the yearly evolution of the children. Then the matrix $\mathcal{C}$ providing successive differences of $\mu_{j}$ is:

$$
\mathcal{C}=\left(\begin{array}{rrrr}
1 & -1 & 0 & 0 \\
0 & 1 & -1 & 0 \\
0 & 0 & 1 & -1
\end{array}\right)
$$

The value of the test statistic is $F_{\text {obs }}=53.134$ which is highly significant for $F_{3.37}$. There are significant differences between the successive means. However, the analysis of the contrasts shows the following simultaneous $95 \%$ confidence intervals

$$
\begin{aligned}
& -1.958 \leq \mu_{1}-\mu_{2} \leq \quad-0.959 \\
& -0.949 \leq \mu_{2}-\mu_{3} \leq \quad 0.335 \\
& -1.171 \leq \mu_{3}-\mu_{4} \leq 0.036
\end{aligned}
$$

Thus, the rejection of $H_{0}$ is mainly due to the difference between the childrens' performances in the first and second year. The confidence intervals for the following contrasts may also be of interest:

$$
\begin{array}{ll}
-2.283 \leq \mu_{1}-\frac{1}{3}\left(\mu_{2}+\mu_{3}+\mu_{4}\right) & \leq-1.423 \\
-1.777 \leq \frac{1}{3}\left(\mu_{1}+\mu_{2}+\mu_{3}\right)-\mu_{4} & \leq-0.742 \\
-1.479 \leq \quad \mu_{2}-\mu_{4} & \leq-0.272 .
\end{array}
$$

They show that $\mu_{1}$ is different from the average of the 3 other years (the same being true for $\mu_{4}$ ) and $\mu_{4}$ turns out to be higher than $\mu_{2}$ (and of course higher than $\mu_{1}$ ).

Test Problem 7 illustrates how the likelihood ratio can be applied to testing a linear restriction on the coefficient $\beta$ of a linear model. It is also shown how a transformation of the test statistic leads to an exact $F$ test as presented in Chap. 3.

Test Problem 7. Suppose $Y_{1}, \ldots, Y_{n}$, are independent with $Y_{i} \sim N_{1}\left(\beta^{\top} x_{i}, \sigma^{2}\right)$, and $x_{i} \in \mathbb{R}^{p}$.
$H_{0}: \mathcal{A} \beta=a, \sigma^{2}$ unknown versus $H_{1}:$ no constraints.

To get the constrained maximum likelihood estimators under $H_{0}$, let $f(\beta, \lambda)=$ $(y-x \beta)^{\top}(y-x \beta)-\lambda^{\top}(\mathcal{A} \beta-a)$ where $\lambda \in \mathbb{R}^{q}$ and solve $\frac{\partial f(\beta, \lambda)}{\partial \beta}=0$ and $\frac{\partial f(\beta, \lambda)}{\partial \lambda}=0$ (Exercise 3.24), thus we obtain:

$$
\tilde{\beta}=\hat{\beta}-\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1} \mathcal{A}^{\top}\left\{\mathcal{A}\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1} \mathcal{A}^{\top}\right\}^{-1}(\mathcal{A} \hat{\beta}-a)
$$

for $\beta$ and $\tilde{\sigma}^{2}=\frac{1}{n}(y-\mathcal{X} \tilde{\beta})^{\top}(y-\mathcal{X} \tilde{\beta})$. The estimate $\hat{\beta}$ denotes the unconstrained MLE as before. Hence, the LR statistic is

$$
\begin{aligned}
-2 \log \lambda & =2\left(\ell_{1}^{*}-\ell_{0}^{*}\right) \\
& =n \log \left(\frac{\|y-\mathcal{X} \tilde{\beta}\|^{2}}{\|y-\mathcal{X} \hat{\beta}\|^{2}}\right) \\
& \xrightarrow{\mathcal{L}} \chi_{q}^{2}
\end{aligned}
$$

where $q$ is the number of elements of $a$. This problem also has an exact $F$-test since

$$
\begin{aligned}
& \frac{n-p}{q}\left(\frac{\|y-\mathcal{X} \tilde{\beta}\|^{2}}{\|y-\mathcal{X} \hat{\beta}\|^{2}}-1\right) \\
& \quad=\frac{n-p}{q} \frac{(\mathcal{A} \hat{\beta}-a)^{\top}\left\{\mathcal{A}\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1} \mathcal{A}^{\top}\right\}^{-1}(\mathcal{A} \hat{\beta}-a)}{(y-\mathcal{X} \hat{\beta})^{\top}(y-\mathcal{X} \hat{\beta})} \sim F_{q, n-p} .
\end{aligned}
$$

Example 7.11 Let us continue with the "classic blue" pullovers. We can once more test if $\beta=0$ in the regression of sales on prices. It holds that

$$
\beta=0 \quad \text { iff } \quad\left(\begin{array}{ll}
0 & 1
\end{array}\right)\binom{\alpha}{\beta}=0
$$

The LR statistic here is

$$
-2 \log \lambda=0.284
$$

which is not significant for the $\chi_{1}^{2}$ distribution. The $F$-test statistic

$$
F=0.231
$$

is also not significant. Hence, we can assume independence of sales and prices (alone). Recall that this conclusion has to be revised if we consider the prices together with advertising costs and hours sales manager hours.

Recall the different conclusion that was made in Example 7.6 when we rejected $H_{0}: \alpha=211$ and $\beta=0$. The rejection there came from the fact that the pair of values was rejected. Indeed, if $\beta=0$ the estimator of $\alpha$ would be $\bar{y}=172.70$ and this is too far from 211.

Example 7.12 Let us now consider the multivariate regression in the "classic blue" pullovers example. From Example 3.15 we know that the estimated parameters in the model

$$
X_{1}=\alpha+\beta_{1} X_{2}+\beta_{2} X_{3}+\beta_{3} X_{4}+\varepsilon
$$

are

$$
\hat{\alpha}=65.670, \quad \hat{\beta}_{1}=-0.216, \quad \hat{\beta}_{2}=0.485, \quad \hat{\beta}_{3}=0.844
$$

Hence, we could postulate the approximate relation:

$$
\beta_{1} \approx-\frac{1}{2} \beta_{2}
$$

which means in practice that augmenting the price by 20 EUR requires the advertising costs to increase by 10 EUR in order to keep the number of pullovers sold constant. Vice versa, reducing the price by 20 EUR yields the same result as before if we reduced the advertising costs by 10 EUR. Let us now test whether the hypothesis

$$
H_{0}: \quad \beta_{1}=-\frac{1}{2} \beta_{2}
$$

is valid. This is equivalent to

$$
\left(\begin{array}{llll}
0 & 1 & \frac{1}{2} & 0
\end{array}\right)\left(\begin{array}{c}
\alpha \\
\beta_{1} \\
\beta_{2} \\
\beta_{3}
\end{array}\right)=0
$$

The LR statistic in this case is equal to (Q MVAlrtest)

$$
-2 \log \lambda=0.012
$$

the $F$ statistic is

$$
F=0.007
$$

Hence, in both cases we will not reject the null hypothesis.

## Comparison of Two Mean Vectors

In many situations, we want to compare two groups of individuals for whom a set of $p$ characteristics has been observed. We have two random samples $\left\{x_{i 1}\right\}_{i=1}^{n_{1}}$ and $\left\{x_{j}\right\}_{j=1}^{n_{2}}$ from two distinct $p$-variate normal populations. Several testing issues can be addressed in this framework. In Test Problem 8 we will first test the hypothesis of equal mean vectors in the two groups under the assumption of equality of the two covariance matrices. This task can be solved by adapting Test Problem 2.

In Test Problem 9 a procedure for testing the equality of the two covariance matrices is presented. If the covariance matrices differ, the procedure of Test Problem 8 is no longer valid. If the equality of the covariance matrices is rejected, an easy rule for comparing two means with no restrictions on the covariance matrices is provided in Test Problem 10.

Test Problem 8. Assume that $X_{i 1} \sim N_{p}\left(\mu_{1}, \Sigma\right)$, with $i=1, \ldots, n_{1}$ and $X_{j 2} \sim N_{p}\left(\mu_{2}, \Sigma\right)$, with $j=1, \ldots, n_{2}$, where all the variables are independent.

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

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

$$
\begin{align*}
\left(\bar{x}_{1}-\bar{x}_{2}\right) & \sim N_{p}\left(\delta, \frac{n_{1}+n_{2}}{n_{1} n_{2}} \Sigma\right)  \tag{7.11}\\
n_{1} S_{1}+n_{2} S_{2} & \sim W_{p}\left(\Sigma, n_{1}+n_{2}-2\right) . \tag{7.12}
\end{align*}
$$

Let $\mathcal{S}=\left(n_{1}+n_{2}\right)^{-1}\left(n_{1} S_{1}+n_{2} S_{2}\right)$ be the weighted mean of $\mathcal{S}_{1}$ and $\mathcal{S}_{2}$. Since the two samples are independent and since $\mathcal{S}_{k}$ is independent of $\bar{x}_{k}$ (for $k=1,2$ ) it follows that $\mathcal{S}$ is independent of $\left(\bar{x}_{1}-\bar{x}_{2}\right)$. Hence, Theorem 5.8 applies and leads to a $T^{2}$-distribution:

$$
\begin{equation*}
\left.\frac{n_{1} n_{2}\left(n_{1}+n_{2}-2\right)}{\left(n_{1}+n_{2}\right)^{2}}\left\{\left(\bar{x}_{1}-\bar{x}_{2}\right)-\delta\right\}^{\top} \mathcal{S}^{-1}\left\{\left(\bar{x}_{1}-\bar{x}_{2}\right)-\delta\right\}\right) \sim T_{p, n_{1}+n_{2}-2}^{2} \tag{7.13}
\end{equation*}
$$

or

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

This result, as in Test Problem 2, can be used to test $H_{0}: \delta=0$ or to construct a confidence region for $\delta \in \mathbb{R}^{p}$. The rejection region is given by:

$$
\begin{equation*}
\frac{n_{1} n_{2}\left(n_{1}+n_{2}-p-1\right)}{p\left(n_{1}+n_{2}\right)^{2}}\left(\bar{x}_{1}-\bar{x}_{2}\right)^{\top} \mathcal{S}^{-1}\left(\bar{x}_{1}-\bar{x}_{2}\right) \geq F_{1-\alpha ; p, n_{1}+n_{2}-p-1} \tag{7.14}
\end{equation*}
$$

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

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

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

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

In particular we have at the $(1-\alpha)$ level, for $j=1, \ldots, p$,

$$
\begin{equation*}
\delta_{j} \in\left(\bar{x}_{1 j}-\bar{x}_{2 j}\right) \pm \sqrt{\frac{p\left(n_{1}+n_{2}\right)^{2}}{\left(n_{1}+n_{2}-p-1\right)\left(n_{1} n_{2}\right)} F_{1-\alpha ; p, n_{1}+n_{2}-p-1} s_{j j}} \tag{7.15}
\end{equation*}
$$

Example 7.13 Let us come back to the questions raised in Example 7.5. We compare the means of assets $\left(X_{1}\right)$ and of sales $\left(X_{2}\right)$ for two sectors, energy (group 1 ) and manufacturing (group 2). With $n_{1}=15, n_{2}=10$, and $p=2$ we obtain the statistics:

$$
\bar{x}_{1}=\binom{4084.0}{2580.5}, \bar{x}_{2}=\binom{4307.2}{4925.2}
$$

and

$$
\mathcal{S}_{1}=10^{7}\left(\begin{array}{ll}
1.6635 & 1.2410 \\
1.2410 & 1.3747
\end{array}\right), \mathcal{S}_{2}=10^{7}\left(\begin{array}{ll}
1.2248 & 1.1425 \\
1.1425 & 1.5112
\end{array}\right),
$$

so that

$$
\mathcal{S}=10^{7}\left(\begin{array}{ll}
1.4880 & 1.2016 \\
1.2016 & 1.4293
\end{array}\right)
$$

The observed value of the test statistic (7.14) is $F=2.7036$. 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 ( $F>F_{0.90 ; 2,22}=2.5613$ ). By directly applying (7.15), the $95 \%$ simultaneous confidence intervals for the differences ( $a$ MVAsimcidif) are obtained as:

$$
\begin{aligned}
& -4628.6 \leq \mu_{1 a}-\mu_{2 a} \leq 4182.2 \\
& -6662.4 \leq \mu_{1 s}-\mu_{2 s} \leq 1973.0
\end{aligned}
$$

Example 7.14 In order to illustrate the presented test procedures it is interesting to analyse some simulated data. This simulation will point out the importance of the covariances in testing means. We created two independent normal samples in $\mathbb{R}^{4}$ of sizes $n_{1}=30$ and $n_{2}=20$ with:

$$
\begin{aligned}
& \mu_{1}=(8,6,10,10)^{\top} \\
& \mu_{2}=(6,6,10,13)^{\top} .
\end{aligned}
$$

One may consider this as an example of $X=\left(X_{1}, \ldots, X_{n}\right)^{\top}$ being the students' scores from four tests, where the two groups of students were subjected to two different methods of teaching. First we simulate the two samples with $\Sigma=\mathcal{I}_{4}$ and obtain the statistics:

$$
\begin{aligned}
& \bar{x}_{1}=(7.607,5.945,10.213,9.635)^{\top} \\
& \bar{x}_{2}=(6.222,6.444,9.560,13.041)^{\top} \\
& \mathcal{S}_{1}=\left(\begin{array}{rrrr}
0.812 & -0.229 & -0.034 & 0.073 \\
-0.229 & 1.001 & 0.010 & -0.059 \\
-0.034 & 0.010 & 1.078 & -0.098 \\
0.073 & -0.059 & -0.098 & 0.823
\end{array}\right) \\
& \mathcal{S}_{2}=\left(\begin{array}{rrrr}
0.559 & -0.057 & -0.271 & 0.306 \\
-0.057 & 1.237 & 0.181 & 0.021 \\
-0.271 & 0.181 & 1.159 & -0.130 \\
0.306 & 0.021 & -0.130 & 0.683
\end{array}\right) .
\end{aligned}
$$

The test statistic (7.14) takes the value $F=60.65$ which is highly significant: the small variance allows the difference to be detected even with these relatively moderate sample sizes. We conclude (at the $95 \%$ level) that:

$$
\begin{aligned}
& 0.6213 \leq \delta_{1} \leq 2.2691 \\
&-1.5217 \leq \delta_{2} \leq 0.5241 \\
&-0.3766 \leq \delta_{3} \leq 1.6830 \\
&-4.2614 \leq \delta_{4} \leq-2.5494
\end{aligned}
$$

which confirms that the means for $X_{1}$ and $X_{4}$ are different.

Consider now a different simulation scenario where the standard deviations are four times larger: $\Sigma=16 \mathcal{I}_{4}$. Here we obtain:

$$
\begin{aligned}
& \bar{x}_{1}=(7.312,6.304,10.840,10.902)^{\top} \\
& \bar{x}_{2}=(6.353,5.890,8.604,11.283)^{\top} \\
& \mathcal{S}_{1}=\left(\begin{array}{rrrr}
21.907 & 1.415 & -2.050 & 2.379 \\
1.415 & 11.853 & 2.104 & -1.864 \\
-2.050 & 2.104 & 17.230 & 0.905 \\
2.379 & -1.864 & 0.905 & 9.037
\end{array}\right) \\
& \mathcal{S}_{2}=\left(\begin{array}{rrrr}
20.349 & -9.463 & 0.958 & -6.507 \\
-9.463 & 15.502 & -3.383 & -2.551 \\
0.958 & -3.383 & 14.470 & -0.323 \\
-6.507 & -2.551 & -0.323 & 10.311
\end{array}\right) .
\end{aligned}
$$

Now the test statistic takes the value 1.54 which is no longer significant ( $F_{0.95,4,45}=$ 2.58). Now we cannot reject the null hypothesis (which we know to be false!) since the increase in variances prohibits the detection of differences of such magnitude.

The following situation illustrates once more the role of the covariances between covariates. Suppose that $\Sigma=16 \mathcal{I}_{4}$ as above but with $\sigma_{14}=\sigma_{41}=-3.999$ (this corresponds to a negative correlation $r_{41}=-0.9997$ ). We have:

$$
\begin{aligned}
& \bar{x}_{1}=(8.484,5.908,9.024,10.459)^{\top} \\
& \bar{x}_{2}=(4.959,7.307,9.057,13.803)^{\top} \\
& \mathcal{S}_{1}=\left(\begin{array}{rrrr}
14.649 & -0.024 & 1.248 & -3.961 \\
-0.024 & 15.825 & 0.746 & 4.301 \\
1.248 & 0.746 & 9.446 & 1.241 \\
-3.961 & 4.301 & 1.241 & 20.002
\end{array}\right) \\
& \mathcal{S}_{2}=\left(\begin{array}{rrrr}
14.035 & -2.372 & 5.596-1.601 \\
-2.372 & 9.173 & -2.027 & -2.954 \\
5.596 & -2.027 & 9.021 & -1.301 \\
-1.601 & -2.954 & -1.301 & 9.593
\end{array}\right) .
\end{aligned}
$$

The value of $F$ is 3.853 which is significant at the $5 \%$ level ( $p$-value $=0.0089$ ). So the null hypothesis $\delta=\mu_{1}-\mu_{2}=0$ is outside the $95 \%$ confidence ellipsoid. However, the simultaneous confidence intervals, which do not take the covariances into account are given by:

$$
\begin{aligned}
& -0.1837 \leq \delta_{1} \leq 7.2343 \\
& -4.9452 \leq \delta_{2} \leq 2.1466 \\
& -3.0091 \leq \delta_{3} \leq 2.9438 \\
& -7.2336 \leq \delta_{4} \leq 0.5450
\end{aligned}
$$

They contain the null value (see Remark 7.1 above) although they are very asymmetric for $\delta_{1}$ and $\delta_{4}$.

Example 7.15 Let us compare the vectors of means of the forged and the genuine bank notes. The matrices $\mathcal{S}_{f}$ and $\mathcal{S}_{g}$ were given in Example 3.1 and since here $n_{f}=n_{g}=100, \mathcal{S}$ is the simple average of $\mathcal{S}_{f}$ and $\mathcal{S}_{g}: \mathcal{S}=\frac{1}{2}\left(\mathcal{S}_{f}+\mathcal{S}_{g}\right)$.

$$
\begin{aligned}
& \bar{x}_{g}=(214.97,129.94,129.72,8.305,10.168,141.52)^{\top} \\
& \bar{x}_{f}=(214.82,130.3,130.19,10.53,11.133,139.45)^{\top}
\end{aligned}
$$

The test statistic is given by (7.14) and turns out to be $F=391.92$ which is highly significant for $F_{6,193}$. The $95 \%$ simultaneous confidence intervals for the differences $\delta_{j}=\mu_{g j}-\mu_{f j}, j=1, \ldots, p$ are:

$$
\begin{aligned}
&-0.0443 \leq \delta_{1} \leq 0.3363 \\
&-0.5186 \leq \delta_{2} \leq-0.1954 \\
&-0.6416 \leq \delta_{3} \leq-0.3044 \\
&-2.6981 \leq \delta_{4} \leq-1.7519 \\
&-1.2952 \leq \delta_{5} \leq-0.6348 \\
& 1.8072 \leq \delta_{6} \leq 2.3268 .
\end{aligned}
$$

All of the components (except for the first one) show significant differences in the means. The main effects are taken by the lower border $\left(X_{4}\right)$ and the diagonal $\left(X_{6}\right)$.

The preceding test implicitly uses the fact that the two samples are extracted from two different populations with common variance $\Sigma$. In this case, the test statistic (7.14) measures the distance between the two centers of gravity of the two groups w.r.t. the common metric given by the pooled variance matrix $\mathcal{S}$. If $\Sigma_{1} \neq \Sigma_{2}$ no such matrix exists. There are no satisfactory test procedures for testing the equality of variance matrices which are robust with respect to normality assumptions of the populations. The following test extends Bartlett's test for equality of variances in the univariate case. But this test is known to be very sensitive to departures from normality.

Test Problem 9 (Comparison of Covariance Matrices). Let $X_{i h} \sim N_{p}\left(\mu_{h}, \Sigma_{h}\right)$, $i=1, \ldots, n_{h}, h=1, \ldots, k$ be independent random variables,

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

Each sub-sample provides $\mathcal{S}_{h}$, an estimator of $\Sigma_{h}$, with

$$
n_{h} S_{h} \sim W_{p}\left(\Sigma_{h}, n_{h}-1\right)
$$

Under $H_{0}, \sum_{h=1}^{k} n_{h} \mathcal{S}_{h} \sim W_{p}(\Sigma, n-k)$ (Sect. 5.2), where $\Sigma$ is the common covariance matrix $X_{i h}$ and $n=\sum_{h=1}^{k} n_{h}$. Let $\mathcal{S}=\frac{n_{1} \mathcal{S}_{1}+\cdots+n_{k} \mathcal{S}_{k}}{n}$ be the weighted average of the $\mathcal{S}_{h}$ (this is in fact the MLE of $\Sigma$ when $H_{0}$ is true). The LRT leads to the statistic

$$
\begin{equation*}
-2 \log \lambda=n \log |\mathcal{S}|-\sum_{h=1}^{k} n_{h} \log \left|\mathcal{S}_{h}\right| \tag{7.16}
\end{equation*}
$$

which under $H_{0}$ is approximately distributed as a $\mathcal{X}_{m}^{2}$ where $m=\frac{1}{2}(k-1) p(p+1)$.
Example 7.16 Let's come back to Example 7.13, where the mean of assets and sales have been compared for companies from the energy and manufacturing sector assuming that $\Sigma_{1}=\Sigma_{2}$. The test of $\Sigma_{1}=\Sigma_{2}$ leads to the value of the test statistic

$$
\begin{equation*}
-2 \log \lambda=0.9076 \tag{7.17}
\end{equation*}
$$

which is not significant ( $p$-value for a $\chi_{3}^{2}=0.82$ ). We cannot reject $H_{0}$ and the comparison of the means performed above is valid.
Example 7.17 Let us compare the covariance matrices of the forged and the genuine bank notes (the matrices $S_{f}$ and $S_{g}$ are shown in Example 3.1). A first look seems to suggest that $\Sigma_{1} \neq \Sigma_{2}$. The pooled variance $\mathcal{S}$ is given by $\mathcal{S}=\frac{1}{2}\left(\mathcal{S}_{f}+\mathcal{S}_{g}\right)$ since here $n_{f}=n_{g}$. The test statistic here is $-2 \log \lambda=127.21$, which is highly significant $\chi^{2}$ with 21 degrees of freedom. As expected, we reject the hypothesis of equal covariance matrices, and as a result the procedure for comparing the two means in Example 7.15 is not valid.
What can we do with unequal covariance matrices? When both $n_{1}$ and $n_{2}$ are large, we have a simple solution:

Test Problem 10 (Comparison of Two Means, Unequal Covariance Matrices, Large Samples). Assume that $X_{i 1} \sim N_{p}\left(\mu_{1}, \Sigma_{1}\right)$, with $i=1, \ldots, n_{1}$ and $X_{j 2} \sim N_{p}\left(\mu_{2}, \Sigma_{2}\right)$, with $j=1, \ldots, n_{2}$ are independent random variables.

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

Letting $\delta=\mu_{1}-\mu_{2}$, we have

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

Therefore, by (5.4)

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

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

$$
\begin{equation*}
\left(\bar{x}_{1}-\bar{x}_{2}\right)^{\top}\left(\frac{\mathcal{S}_{1}}{n_{1}}+\frac{\mathcal{S}_{2}}{n_{2}}\right)^{-1}\left(\bar{x}_{1}-\bar{x}_{2}\right) \xrightarrow{\mathcal{L}} \chi_{p}^{2} . \tag{7.18}
\end{equation*}
$$

This can be used in place of (7.13) for testing $H_{0}$, defining a confidence region for $\delta$ or constructing simultaneous confidence intervals for $\delta_{j}, j=1, \ldots, p$.

For instance, the rejection region at the level $\alpha$ will be

$$
\begin{equation*}
\left(\bar{x}_{1}-\bar{x}_{2}\right)^{\top}\left(\frac{\mathcal{S}_{1}}{n_{1}}+\frac{\mathcal{S}_{2}}{n_{2}}\right)^{-1}\left(\bar{x}_{1}-\bar{x}_{2}\right)>\chi_{1-\alpha ; p}^{2} \tag{7.19}
\end{equation*}
$$

and the $(1-\alpha)$ simultaneous confidence intervals for $\delta_{j}, j=1, \ldots, p$ are:

$$
\begin{equation*}
\delta_{j} \in\left(\bar{x}_{1}-\bar{x}_{2}\right) \pm \sqrt{\chi_{1-\alpha ; p}^{2}\left(\frac{s_{j j}^{(1)}}{n_{1}}+\frac{s_{j j}^{(2)}}{n_{2}}\right)}, \tag{7.20}
\end{equation*}
$$

where $s_{j j}^{(i)}$ is the $(j, j)$ element of the matrix $\mathcal{S}_{i}$. This may be compared to (7.15) where the pooled variance was used.

Remark 7.2 We see, by comparing the statistics (7.19) with (7.14), that we measure here the distance between $\bar{x}_{1}$ and $\bar{x}_{2}$ using the metric $\left(\frac{\mathcal{S}_{1}}{n_{1}}+\frac{\mathcal{S}_{2}}{n_{2}}\right)$. It should be noted that when $n_{1}=n_{2}$, the two methods are essentially the same since then $\mathcal{S}=\frac{1}{2}\left(\mathcal{S}_{1}+\mathcal{S}_{2}\right)$. If the covariances are different but have the same eigenvectors (different eigenvalues), one can apply the common principal component (CPC) technique, see Chap. 11.

Example 7.18 Let us use the last test to compare the forged and the genuine bank notes again ( $n_{1}$ and $n_{2}$ are both large). The test statistic (7.19) turns out to be 2,436.8 which is again highly significant. The $95 \%$ simultaneous confidence intervals are:

$$
\begin{aligned}
&-0.0389 \leq \delta_{1} \leq 0.3309 \\
&-0.5140 \leq \delta_{2} \leq-0.2000 \\
&-0.6368 \leq \delta_{3} \leq-0.3092 \\
&-2.6846 \leq \delta_{4} \leq-1.7654 \\
&-1.2858 \leq \delta_{5} \leq-0.6442 \\
& 1.8146 \leq \delta_{6} \leq 2.3194
\end{aligned}
$$

showing that all the components except the first are different from zero, the largest difference coming from $X_{6}$ (length of the diagonal) and $X_{4}$ (lower border). The
results are very similar to those obtained in Example 7.15. This is due to the fact that here $n_{1}=n_{2}$ as we already mentioned in the remark above.

## Profile Analysis

Another useful application of Test Problem 6 is the repeated measurements problem applied to two independent groups. This problem arises in practice when we observe repeated measurements of characteristics (or measures of the same type under different experimental conditions) on the different groups which have to be compared. It is important that the $p$ measures (the "profile") are comparable, and, in particular, are reported in the same units. For instance, they may be measures of blood pressure at $p$ different points in time, one group being the control group and the other the group receiving a new treatment. The observations may be the scores obtained from $p$ different tests of two different experimental groups. One is then interested in comparing the profiles of each group: the profile being just the vectors of the means of the $p$ responses (the comparison may be visualised in a two-dimensional graph using the parallel coordinate plot introduced in Sect. 1.7).

We are thus in the same statistical situation as for the comparison of two means:

$$
\begin{array}{ll}
X_{i 1} \sim N_{p}\left(\mu_{1}, \Sigma\right) & i=1, \ldots, n_{1} \\
X_{i 2} \sim N_{p}\left(\mu_{2}, \Sigma\right) & i=1, \ldots, n_{2}
\end{array}
$$

where all variables are independent. Suppose the two population profiles look like in Fig. 7.1.

The following questions are of interest:

1. Are the profiles similar in the sense of being parallel (which means no interaction between the treatments and the groups)?
2. If the profiles are parallel, are they at the same level?
3. If the profiles are parallel, is there any treatment effect, i.e. are the profiles horizontal (profiles remain the same no matter which treatment received)?

The above questions are easily translated into linear constraints on the means and a test statistic can be obtained accordingly.

## Parallel Profiles

Let $\mathcal{C}$ be a $(p-1) \times p$ matrix defined as $\mathcal{C}=\left(\begin{array}{rrrrr}1 & -1 & 0 & \cdots & 0 \\ 0 & 1 & -1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & \cdots & 0 & 1 & -1\end{array}\right)$.


Fig. 7.1 Example of population profiles $Q_{\text {MVAprofil }}$

The hypothesis to be tested is

$$
H_{0}^{(1)}: \mathcal{C}\left(\mu_{1}-\mu_{2}\right)=0
$$

From (7.11), (7.12) and Corollary 5.4 we know that under $H_{0}$ :

$$
\begin{equation*}
\frac{n_{1} n_{2}}{\left(n_{1}+n_{2}\right)^{2}}\left(n_{1}+n_{2}-2\right)\left\{\mathcal{C}\left(\bar{x}_{1}-\bar{x}_{2}\right)\right\}^{\top}\left(\mathcal{C S C}^{\top}\right)^{-1} \mathcal{C}\left(\bar{x}_{1}-\bar{x}_{2}\right) \sim T_{p-1, n_{1}+n_{2}-2}^{2} \tag{7.21}
\end{equation*}
$$

where $\mathcal{S}$ is the pooled covariance matrix. The hypothesis is rejected if

$$
\frac{n_{1} n_{2}\left(n_{1}+n_{1}-p\right)}{\left(n_{1}+n_{2}\right)^{2}(p-1)}(\mathcal{C} \bar{x})^{\top}\left(\mathcal{C S C}^{\top}\right)^{-1} \mathcal{C} \bar{x}>F_{1-\alpha ; p-1, n_{1}+n_{2}-p}
$$

## Equality of Two Levels

The question of equality of the two levels is meaningful only if the two profiles are parallel. In the case of interactions (rejection of $H_{0}^{(1)}$ ), the two populations react differently to the treatments and the question of the level has no meaning. The equality of the two levels can be formalised as

$$
H_{0}^{(2)}: 1_{p}^{\top}\left(\mu_{1}-\mu_{2}\right)=0
$$

since

$$
1_{p}^{\top}\left(\bar{x}_{1}-\bar{x}_{2}\right) \sim N_{1}\left(1_{p}^{\top}\left(\mu_{1}-\mu_{2}\right), \frac{n_{1}+n_{2}}{n_{1} n_{2}} 1_{p}^{\top} \Sigma 1_{p}\right)
$$

and

$$
\left(n_{1}+n_{2}\right) 1_{p}^{\top} \mathcal{S} 1_{p} \sim W_{1}\left(1_{p}^{\top} \Sigma 1_{p}, n_{1}+n_{2}-2\right)
$$

Using Corollary 5.4 we have that:

$$
\begin{align*}
\frac{n_{1} n_{2}}{\left(n_{1}+n_{2}\right)^{2}}\left(n_{1}+n_{2}-2\right) \frac{\left\{1_{p}^{\top}\left(\bar{x}_{1}-\bar{x}_{2}\right)\right\}^{2}}{1_{p}^{\top} \mathcal{S} 1_{p}} & \sim T_{1, n_{1}+n_{2}-2}^{2}  \tag{7.22}\\
& =F_{1, n_{1}+n_{2}-2}
\end{align*}
$$

The rejection region is

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

## Treatment Effect

If it is rejected that the profiles are parallel, then two independent analyses should be done on the two groups using the repeated measurement approach. But if it is accepted that they are parallel, then we can exploit the information contained in both groups (possibly at different levels) to test a treatment effect, i.e. if the two profiles are horizontal. This may be written as:

$$
H_{0}^{(3)}: \mathcal{C}\left(\mu_{1}+\mu_{2}\right)=0
$$

Consider the average profile $\bar{x}$

$$
\bar{x}=\frac{n_{1} \bar{x}_{1}+n_{2} \bar{x}_{2}}{n_{1}+n_{2}} .
$$

Clearly,

$$
\bar{x} \sim N_{p}\left(\frac{n_{1} \mu_{1}+n_{2} \mu_{2}}{n_{1}+n_{2}}, \frac{1}{n_{1}+n_{2}} \Sigma\right) .
$$

Now it is not hard to prove that $H_{0}^{(3)}$ with $H_{0}^{(1)}$ implies that

$$
\mathcal{C}\left(\frac{n_{1} \mu_{1}+n_{2} \mu_{2}}{n_{1}+n_{2}}\right)=0
$$

So under parallel, horizontal profiles we have

$$
\sqrt{n_{1}+n_{2}} \mathcal{C} \bar{x} \sim N_{p}\left(0, \mathcal{C} \Sigma \mathcal{C}^{\top}\right) .
$$

From Corollary 5.4 we again obtain

$$
\begin{equation*}
\left(n_{1}+n_{2}-2\right)(\mathcal{C} \bar{x})^{\top}\left(\mathcal{C S C}^{\top}\right)^{-1} \mathcal{C} \bar{x} \sim T_{p-1, n_{1}+n_{2}-2}^{2} . \tag{7.23}
\end{equation*}
$$

This leads to the rejection region of $H_{0}^{(3)}$, namely

$$
\frac{n_{1}+n_{2}-p}{p-1}(\mathcal{C} \bar{x})^{\top}\left(\mathcal{C S C}^{\top}\right)^{-1} \mathcal{C} \bar{x}>F_{1-\alpha ; p-1, n_{1}+n_{2}-p}
$$

Example 7.19 Morrison (1990) proposed a test in which the results of four sub-tests of the Wechsler Adult Intelligence Scale (WAIS) are compared for two categories of people: group 1 contains $n_{1}=37$ people who do not have a senile factor and group 2 contains $n_{2}=12$ people who have a senile factor. The four WAIS sub-tests are $X_{1}$ (information), $X_{2}$ (similarities), $X_{3}$ (arithmetic) and $X_{4}$ (picture completion). The relevant statistics are

$$
\begin{aligned}
& \bar{x}_{1}=(12.57,9.57,11.49,7.97)^{\top} \\
& \bar{x}_{2}=(8.75,5.33,8.50,4.75)^{\top} \\
& \mathcal{S}_{1}=\left(\begin{array}{rrrr}
11.164 & 8.840 & 6.210 & 2.020 \\
8.840 & 11.759 & 5.778 & 0.529 \\
6.210 & 5.778 & 10.790 & 1.743 \\
2.020 & 0.529 & 1.743 & 3.594
\end{array}\right) \\
& \mathcal{S}_{2}=\left(\begin{array}{rrrr}
9.688 & 9.583 & 8.875 & 7.021 \\
9.583 & 16.722 & 11.083 & 8.167 \\
8.875 & 11.083 & 12.083 & 4.875 \\
7.021 & 8.167 & 4.875 & 11.688
\end{array}\right) .
\end{aligned}
$$

The test statistic for testing if the two profiles are parallel is $F=0.4634$, which is not significant ( $p$-value $=0.71$ ). Thus it is accepted that the two are parallel. The second test statistic (testing the equality of the levels of the two profiles) is $F=17.21$, which is highly significant ( $p$-value $\approx 10^{-4}$ ). The global level of the test for the non-senile people is superior to the senile group. The final test (testing the horizontality of the average profile) has the test statistic $F=53.32$, which is
also highly significant ( $p$-value $\approx 10^{-14}$ ). This implies that there are substantial differences among the means of the different subtests.

|  | Summary |
| :--- | :--- |

### 7.3 Boston Housing

Returning to the Boston Housing data set, we are now in a position to test if the means of the variables vary according to their location, for example, when they are located in a district with high valued houses. In Chap. 1, we built two groups of observations according to the value of $X_{14}$ being less than or equal to the median of $X_{14}$ (a group of 256 districts) and greater than the median (a group of 250 districts). In what follows, we use the transformed variables motivated in Sect. 1.9.

Testing the equality of the means from the two groups was proposed in a multivariate setup, so we restrict the analysis to the variables $X_{1}, X_{5}, X_{8}, X_{11}$, and $X_{13}$ to see if the differences between the two groups that were identified in Chap. 1 can be confirmed by a formal test. As in Test Problem 8, the hypothesis to be tested is

$$
H_{0}: \mu_{1}=\mu_{2}, \text { where } \mu_{1} \in \mathbb{R}^{5}, n_{1}=256, \text { and } n_{2}=250
$$

$\Sigma$ is not known. The $F$-statistic given in (7.13) is equal to 126.30 , which is much higher than the critical value $F_{0.95 ; 5,500}=2.23$. Therefore, we reject the hypothesis of equal means.

To see which component, $X_{1}, X_{5}, X_{8}, X_{11}$, or $X_{13}$, is responsible for this rejection, take a look at the simultaneous confidence intervals defined in (7.14):

$$
\begin{aligned}
\delta_{1} & \in\left(\begin{array}{ll}
1.4020, & 2.5499
\end{array}\right) \\
\delta_{5} & \in\left(\begin{array}{ll}
0.1315, & 0.2383
\end{array}\right) \\
\delta_{8} & \in(-0.5344,-0.2222) \\
\delta_{11} & \in\left(\begin{array}{ll}
1.0375, & 1.7384
\end{array}\right) \\
\delta_{13} & \in\left(\begin{array}{ll}
1.1577, & 1.5818
\end{array}\right)
\end{aligned}
$$

These confidence intervals confirm that all of the $\delta_{j}$ are significantly different from zero (note there is a negative effect for $X_{8}$ : weighted distances to employment centers) Q MVAsimcibh.

We could also check if the factor "being bounded by the river" (variable $X_{4}$ ) has some effect on the other variables. To do this compare the means of $\left(X_{5}, X_{8}, X_{9}, X_{12}, X_{13}, X_{14}\right)^{\top}$. There are two groups: $n_{1}=35$ districts bounded by the river and $n_{2}=471$ districts not bounded by the river. Test Problem 8 ( $H_{0}: \mu_{1}=\mu_{2}$ ) is applied again with $p=6$. The resulting test statistic, $F=5.81$, is highly significant $\left(F_{0.95 ; 6,499}=2.12\right)$. The simultaneous confidence intervals indicate that only $X_{14}$ (the value of the houses) is responsible for the hypothesis being rejected. At a significance level of 0.95

$$
\begin{aligned}
& \delta_{5} \in(-0.0603,0.1919) \\
& \delta_{8} \in(-0.5225,0.1527) \\
& \delta_{9} \in(-0.5051,0.5938) \\
& \delta_{12} \in(-0.3974,0.7481) \\
& \delta_{13} \in(-0.8595,0.3782) \\
& \delta_{14} \in(0.0014,0.5084)
\end{aligned}
$$

## Testing Linear Restrictions

In Chap. 3 a linear model was proposed that explained the variations of the price $X_{14}$ by the variations of the other variables. Using the same procedure that was shown in Testing Problem 7, we are in a position to test a set of linear restrictions on the vector of regression coefficients $\beta$.

The model we estimated in Sect. 3.7 provides the following (a MVAlinregbh):

| Variable | $\hat{\beta}_{j}$ | $S E\left(\hat{\beta}_{j}\right)$ | $t$ | $p$-Value |
| :--- | ---: | :--- | ---: | :--- |
| Constant | 4.1769 | 0.3790 | 11.020 | 0.0000 |
| $X_{1}$ | -0.0146 | 0.0117 | -1.254 | 0.2105 |
| $X_{2}$ | 0.0014 | 0.0056 | 0.247 | 0.8051 |
| $X_{3}$ | -0.0127 | 0.0223 | -0.570 | 0.5692 |
| $X_{4}$ | 0.1100 | 0.0366 | 3.002 | 0.0028 |
| $X_{5}$ | -0.2831 | 0.1053 | -2.688 | 0.0074 |
| $X_{6}$ | 0.4211 | 0.1102 | 3.822 | 0.0001 |
| $X_{7}$ | 0.0064 | 0.0049 | 1.317 | 0.1885 |
| $X_{8}$ | -0.1832 | 0.0368 | -4.977 | 0.0000 |
| $X_{9}$ | 0.0684 | 0.0225 | 3.042 | 0.0025 |
| $X_{10}$ | -0.2018 | 0.0484 | -4.167 | 0.0000 |
| $X_{11}$ | -0.0400 | 0.0081 | -4.946 | 0.0000 |
| $X_{12}$ | 0.0445 | 0.0115 | 3.882 | 0.0001 |
| $X_{13}$ | -0.2626 | 0.0161 | -16.320 | 0.0000 |

Recall that the estimated residuals $Y-\mathcal{X} \hat{\beta}$ did not show a big departure from normality, which means that the testing procedure developed above can be used.

1. First a global test of significance for the regression coefficients is performed,

$$
H_{0}:\left(\beta_{1}, \ldots, \beta_{13}\right)=0
$$

This is obtained by defining $\mathcal{A}=\left(0_{13}, \mathcal{I}_{13}\right)$ and $a=0_{13}$ so that $H_{0}$ is equivalent to $\mathcal{A} \beta=a$ where $\beta=\left(\beta_{0}, \beta_{1}, \ldots, \beta_{13}\right)^{\top}$. Based on the observed values $F=$ 123.20. This is highly significant $\left(F_{0.95 ; 13,492}=1.7401\right)$, thus we reject $H_{0}$. Note that under $H_{0} \hat{\beta}_{H_{0}}=(3.0345,0, \ldots, 0)$ where $3.0345=\bar{y}$.
2. Since we are interested in the effect that being located close to the river has on the value of the houses, the second test is $H_{0}: \beta_{4}=0$. This is done by fixing

$$
\mathcal{A}=(0,0,0,0,1,0,0,0,0,0,0,0,0,0)^{\top}
$$

and $a=0$ to obtain the equivalent hypothesis $H_{0}: \mathcal{A} \beta=a$. The result is again significant: $F=9.0125\left(F_{0.95 ; 1,492}=3.8604\right)$ with a $p$-value of 0.0028 . Note that this is the same $p$-value obtained in the individual test $\beta_{4}=0$ in Chap.3, computed using a different setup.
3. A third test notices the fact that some of the regressors in the full model (3.57) appear to be insignificant (that is they have high individual $p$-values). It can be confirmed from a joint test if the corresponding reduced model, formulated

Table 7.1 Linear regression for Boston housing data set $\mathbf{Q}$ MVAlinreg2bh

| Variable | $\hat{\beta}_{j}$ | $S E$ | $t$ | $p$-Value |
| :--- | ---: | :--- | :--- | :--- |
| Const | 4.1582 | 0.3628 | 11.462 | 0.0000 |
| $X_{4}$ | 0.1087 | 0.0362 | 2.999 | 0.0028 |
| $X_{5}$ | -0.3055 | 0.0973 | -3.140 | 0.0018 |
| $X_{6}$ | 0.4668 | 0.1059 | 4.407 | 0.0000 |
| $X_{8}$ | -0.1855 | 0.0327 | -5.679 | 0.0000 |
| $X_{9}$ | 0.0492 | 0.0183 | 2.690 | 0.0074 |
| $X_{10}$ | -0.2096 | 0.0446 | -4.705 | 0.0000 |
| $X_{11}$ | -0.0410 | 0.0078 | -5.280 | 0.0000 |
| $X_{12}$ | 0.0481 | 0.0112 | 4.306 | 0.0000 |
| $X_{13}$ | -0.2588 | 0.0149 | -17.396 | 0.0000 |

by deleting the insignificant variables, is rejected by the data. We want to test $H_{0}: \beta_{1}=\beta_{2}=\beta_{3}=\beta_{7}=0$. Hence,

$$
\mathcal{A}=\left(\begin{array}{llllllllllllll}
0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0
\end{array}\right)
$$

and $a=0_{4}$. The test statistic is 0.9344 , which is not significant for $F_{4,492}$. Given that the $p$-value is equal to 0.44 , we cannot reject the null hypothesis nor the corresponding reduced model. The value of $\hat{\beta}$ under the null hypothesis is
$\hat{\beta}_{H_{0}}=(4.16,0,0,0,0.11,-0.31,0.47,0,-0.19,0.05,-0.20,-0.04,0.05,-0.26)^{\top}$.
A possible reduced model is

$$
X_{14}=\beta_{0}+\beta_{4} X_{4}+\beta_{5} X_{5}+\beta_{6} X_{6}+\beta_{8} X_{8}+\cdots+\beta_{13} X_{13}+\varepsilon
$$

Estimating this reduced model using OLS, as was done in Chap. 3, provides the results shown in Table 7.1.

Note that the reduced model has $r^{2}=0.763$ which is very close to $r^{2}=0.765$ obtained from the full model. Clearly, including variables $X_{1}, X_{2}, X_{3}$, and $X_{7}$ does not provide valuable information in explaining the variation of $X_{14}$, the price of the houses.

### 7.4 Exercises

Exercise 7.1 Use Theorem 7.1 to derive a test for testing the hypothesis that a dice is balanced, based on $n$ tosses of that dice. (Hint: use the multinomial probability function.)
Exercise 7.2 Consider $N_{3}(\mu, \Sigma)$. Formulate the hypothesis $H_{0}: \mu_{1}=\mu_{2}=\mu_{3}$ in terms of $\mathcal{A} \mu=a$.
Exercise 7.3 Simulate a normal sample with $\mu=\binom{1}{2}$ and $\Sigma=\left(\begin{array}{cc}1 & 0.5 \\ 0.5 & 2\end{array}\right)$ and test $H_{0}: 2 \mu_{1}-\mu_{2}=0.2$ first with $\Sigma$ known and then with $\Sigma$ unknown. Compare the results.

Exercise 7.4 Derive expression (7.3) for the LRT statistic in Test Problem 2.
Exercise 7.5 With the simulated data set of Example 7.14, test the hypothesis of equality of the covariance matrices.
Exercise 7.6 In the US companies data set, test the equality of means between the energy and manufacturing sectors, taking the full vector of observations $X_{1}$ to $X_{6}$. Derive the simultaneous confidence intervals for the differences.
Exercise 7.7 Let $X \sim N_{2}(\mu, \Sigma)$ where $\Sigma$ is known to be $\Sigma=\left(\begin{array}{rr}2 & -1 \\ -1 & 2\end{array}\right)$. We have an i.i.d. sample of size $n=6$ providing $\bar{x}^{\top}=\left(1 \frac{1}{2}\right)$. Solve the following test problems ( $\alpha=0.05$ ):
(a) $H_{0}: \quad \mu=\left(2, \frac{2}{3}\right)^{\top} \quad H_{1}: \quad \mu \neq\left(2, \frac{2}{3}\right)^{\top}$
(b) $H_{0}: \quad \mu_{1}+\mu_{2}=\frac{7}{2} \quad H_{1}: \quad \mu_{1}+\mu_{2} \neq \frac{7}{2}$
(c) $H_{0}: \quad \mu_{1}-\mu_{2}=\frac{1}{2} \quad H_{1}: \quad \mu_{1}-\mu_{2} \neq \frac{1}{2}$
(d) $H_{0}: \quad \mu_{1}=2 \quad H_{1}: \quad \mu_{1} \neq 2$

For each case, represent the rejection region graphically (comment).
Exercise 7.8 Repeat the preceding exercise with $\Sigma$ unknown and $S=\left(\begin{array}{rr}2 & -1 \\ -1 & 2\end{array}\right)$. Compare the results.

Exercise 7.9 Consider $X \sim N_{3}(\mu, \Sigma)$. An i.i.d. sample of size $n=10$ provides:

$$
\begin{aligned}
& \bar{x}=(1,0,2)^{\top} \\
& S=\left(\begin{array}{lll}
3 & 2 & 1 \\
2 & 3 & 1 \\
1 & 1 & 4
\end{array}\right) .
\end{aligned}
$$

(a) Knowing that the eigenvalues of $S$ are integers, describe a $95 \%$ confidence region for $\mu$. (Hint: to compute eigenvalues use $|S|=\prod_{j=1}^{3} \lambda_{j}$ and $\operatorname{tr}(S)=$ $\left.\sum_{j=1}^{3} \lambda_{j}.\right)$
(b) Calculate the simultaneous confidence intervals for $\mu_{1}, \mu_{2}$ and $\mu_{3}$.
(c) Can we assert that $\mu_{1}$ is an average of $\mu_{2}$ and $\mu_{3}$ ?

Exercise 7.10 Consider two independent i.i.d. samples, each of size 10, from two bivariate normal populations. The results are summarised below:

$$
\begin{gathered}
\bar{x}_{1}=(3,1)^{\top} ; \bar{x}_{2}=(1,1)^{\top} \\
S_{1}=\left(\begin{array}{rr}
4 & -1 \\
-1 & 2
\end{array}\right) ; S_{2}=\left(\begin{array}{rr}
2 & -2 \\
-2 & 4
\end{array}\right) .
\end{gathered}
$$

Provide a solution to the following tests:
(a) $H_{0}: \quad \mu_{1}=\mu_{2} \quad H_{1}: \quad \mu_{1} \neq \mu_{2}$
(b) $H_{0}: \quad \mu_{11}=\mu_{21} \quad H_{1}: \quad \mu_{11} \neq \mu_{21}$
(c) $H_{0}: \quad \mu_{12}=\mu_{22} \quad H_{1}: \quad \mu_{12} \neq \mu_{22}$

Compare the solutions and comment.
Exercise 7.11 Prove expression (7.4) in the Test Problem 2 with log-likelihoods $\ell_{0}^{*}$ and $\ell_{1}^{*}$. [Hint: use (2.29).]

Exercise 7.12 Assume that $X \sim N_{p}(\mu, \Sigma)$ where $\Sigma$ is unknown.
(a) Derive the $\log L R T$ for testing the independence of the $p$ components, that is $H_{0}: \Sigma$ is a diagonal matrix. (Solution: $-2 \log \lambda=-n \log |\mathcal{R}|$ where $\mathcal{R}$ is the correlation matrix, which is asymptotically a $\chi_{\frac{1}{2} p(p-1)}^{2}$ under $H_{0}$.)
(b) Assume that $\Sigma$ is a diagonal matrix (all the variables are independent). Can an asymptotic test for $H_{0}: \mu=\mu_{o}$ against $H_{1}: \mu \neq \mu_{o}$ be derived? How would this compare to $p$ independent univariate $t$-tests on each $\mu_{j}$ ?
(c) Show an easy derivation of an asymptotic test for testing the equality of the $p$ means [Hint: use $(C \bar{X})^{\top}\left(C S C^{\top}\right)^{-1} C \bar{X} \rightarrow \chi_{p-1}^{2}$ where $\mathcal{S}=\operatorname{diag}\left(s_{11}, \ldots, s_{p p}\right)$ and $\mathcal{C}$ is defined as in (7.10)]. Compare this to the simple ANOVA procedure used in Sect. 3.5.

Exercise 7.13 The yields of wheat have been measured in 30 parcels that have been randomly attributed to three lots prepared by one of three different fertiliser $A, B$ and C. The data are

| Fertiliser yield | A | B | C |
| :---: | :--- | :--- | :--- |
| 1 | 4 | 6 | 2 |
| 2 | 3 | 7 | 1 |
| 3 | 2 | 7 | 1 |
| 4 | 5 | 5 | 1 |
| 5 | 4 | 5 | 3 |
| 6 | 4 | 5 | 4 |
| 7 | 3 | 8 | 3 |
| 8 | 3 | 9 | 3 |
| 9 | 3 | 9 | 2 |
| 10 | 1 | 6 | 2 |

Using Exercise 7.12,
(a) test the independence between the three variables.
(b) test whether $\mu^{\top}=\left[\begin{array}{ll}2 & 6\end{array}\right]$ and compare this to the three univariate $t$-tests.
(c) test whether $\mu_{1}=\mu_{2}=\mu_{3}$ using simple ANOVA and the $\chi^{2}$ approximation.

Exercise 7.14 Consider an i.i.d. sample of size $n=5$ from a bivariate normal distribution

$$
X \sim N_{2}\left(\mu,\left(\begin{array}{ll}
3 & \rho \\
\rho & 1
\end{array}\right)\right)
$$

where $\rho$ is a known parameter. Suppose $\bar{x}^{\top}=(10)$. For what value of $\rho$ would the hypothesis $H_{0}: \mu^{\top}=(00)$ be rejected in favour of $H_{1}: \mu^{\top} \neq(00)$ (at the $5 \%$ level)?

Exercise 7.15 Using Example 7.14, test the last two cases described there and test the sample number one ( $n_{1}=30$ ), to see if they are from a normal population with $\Sigma=4 \mathcal{I}_{4}$ (the sample covariance matrix to be used is given by $S_{1}$ ).

Exercise 7.16 Consider the bank data set. For the counterfeit bank notes, we want to know if the length of the diagonal $\left(X_{6}\right)$ can be predicted by a linear model in $X_{1}$ to $X_{5}$. Estimate the linear model and test if the coefficients are significantly different from zero.

Exercise 7.17 In Example 7.10, can you predict the vocabulary score of the children in eleventh grade, by knowing the results from grades 8-9 and 10? Estimate a linear model and test its significance.

Exercise 7.18 Test the equality of the covariance matrices from the two groups in the WAIS subtest (Example 7.19).

Exercise 7.19 Prove expressions (7.21)-(7.23).

Exercise 7.20 Using Theorem 6.3 and expression (7.16), construct an asymptotic rejection region of size $\alpha$ for testing, in a general model $f(x, \theta)$, with $\theta \in \mathbb{R}^{k}$, $H_{0}: \theta=\theta_{0}$ against $H_{1}: \theta \neq \theta_{0}$.
Exercise 7.21 Exercise 6.5 considered the pdf $f\left(x_{1}, x_{2}\right)=\frac{1}{\theta_{1}^{2} \theta_{2}^{2} x_{2}} e^{-\left(\frac{x_{1}}{\theta_{1} x_{2}}+\frac{x_{2}}{\theta_{1} \theta_{2}}\right)}$
$x_{1}, x_{2}>0$. Solve the problem of testing $H_{0}: \theta^{\top}=\left(\theta_{01}, \theta_{02}\right)$ from an iid sample of size $n$ on $x=\left(x_{1}, x_{2}\right)^{\top}$, where $n$ is large.

Exercise 7.22 In Olkin and Veath (1980), the evolution of citrate concentrations in plasma is observed at three different times of day, $X_{1}(8 \mathrm{am}), X_{2}$ (11 am) and $X_{3}(3 \mathrm{pm})$, for two groups of patients who follow different diets. (The patients were randomly attributed to each group under a balanced design $n_{1}=n_{2}=5$.)
The data are:

| Group | $X_{1}(8 \mathrm{am})$ | $X_{2}(11 \mathrm{am})$ | $X_{3}(3 \mathrm{pm})$ |
| :--- | :--- | :--- | :--- |
| I | 125 | 137 | 121 |
|  | 144 | 173 | 147 |
|  | 105 | 119 | 125 |
|  | 151 | 149 | 128 |
| II | 137 | 139 | 109 |
|  | 93 | 121 | 107 |
|  | 116 | 135 | 106 |
|  | 109 | 83 | 100 |
|  | 89 | 95 | 83 |
|  | 116 | 128 | 100 |

Test if the profiles of the groups are parallel, if they are at the same level and if they are horizontal.

## Part III <br> Multivariate Techniques

## Chapter 8 <br> Regression Models

The aim of regression models is to model the variation of a quantitative response variable $y$ in terms of the variation of one or several explanatory variables $\left(x_{1}, \ldots, x_{p}\right)^{\top}$. We have already introduced such models in Chaps. 3 and 7 where linear models were written in (3.50) as

$$
y=\mathcal{X} \beta+\varepsilon,
$$

where $y(n \times 1)$ is the vector of observation for the response variable, $\mathcal{X}(n \times p)$ is the data matrix of the $p$ explanatory variables and $\varepsilon$ are the errors. Linear models are not restricted to handle only linear relationships between $y$ and $x$. Curvature is allowed by including appropriate higher order terms in the design matrix $\mathcal{X}$.

Example 8.1 If $y$ represents response and $x_{1}, x_{2}$ are two factors that explain the variation of $y$ via the quadratic response model:

$$
\begin{equation*}
y_{i}=\beta_{0}+\beta_{1} x_{i 1}+\beta_{2} x_{i 2}+\beta_{3} x_{i 1}^{2}+\beta_{4} x_{i 2}^{2}+\beta_{5} x_{i 1} x_{i 2}+\varepsilon_{i}, i=1, \ldots, n . \tag{8.1}
\end{equation*}
$$

This model (8.1) belongs to the class of linear models because it is linear in $\beta$. The data matrix $\mathcal{X}$ is:

$$
\mathcal{X}=\left(\begin{array}{cccccc}
1 & x_{11} & x_{12} & x_{11}^{2} & x_{12}^{2} & x_{11} x_{12} \\
1 & x_{21} & x_{22} & x_{21}^{2} & x_{22}^{2} & x_{21} x_{22} \\
\vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\
1 & x_{n 1} & x_{n 2} & x_{n 1}^{2} & x_{n 2}^{2} & x_{n 1} x_{n 2}
\end{array}\right)
$$

For a given value of $\beta$, the response surface can be represented in a threedimensional plot as in Fig. 8.1 where we display $y=20+1 x_{1}+2 x_{2}-8 x_{1}^{2}-$ $6 x_{2}^{2}+6 x_{1} x_{2}$, i.e. $\beta=(20,1,2,-8,-6,+6)^{\top}$.


Fig. 8.1 A 3-D response surface $Q$ MVAresponsesurface

Note also that pure non-linear models can sometimes be rewritten as a linear model by choosing an appropriate transformation of the coordinates of the variables. For instance the Cobb-Douglas production function

$$
y_{i}=k x_{i 1}^{\beta_{1}} x_{i 2}^{\beta_{2}} x_{i 3}^{\beta_{3}},
$$

where $y$ is the level of the production of a plant and $\left(x_{1}, x_{2}, x_{3}\right)^{\top}$ are three factors of production (e.g. labour, capital and energy), can be transformed into a linear model in the log scale. We have indeed

$$
\log y_{i}=\beta_{0}+\beta_{1} \log x_{i 1}+\beta_{2} \log x_{i 2}+\beta_{3} \log x_{i 3},
$$

where $\beta_{0}=\log k$ and the $\beta_{j}, j=1, \ldots, 3$ are the elasticities $\left(\beta_{j}=\right.$ $\left.\partial \log y / \partial \log x_{j}\right)$.

Linear models are flexible and cover a wide class of models. If $\mathcal{X}$ has full rank, they can easily be estimated by least squares $\hat{\beta}=\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1} \mathcal{X}^{\top} y$ and linear restrictions on the $\beta$ 's can be tested using the tools developed in Chap. 7.

In Chap. 3, we saw that even qualitative explanatory variables can be used by defining appropriate coding of the nominal values of $x$. In this chapter, we will extend our toolbox by showing how to code these qualitative factors in a way which allows the introduction of several qualitative factors including the possibility of interactions. This covers more general ANOVA models than those introduced in Chap. 3. This includes the ANCOVA models where qualitative and quantitative variables are both present in the explanatory variables.

When the response variable is qualitative or categorical (for instance, an individual can be employed or unemployed, a company may be bankrupt or not, the opinion of one person relative to a particular issue can be "in favour", "against" or "indifferent to", etc.), linear models have to be adapted to this particular situation. The most useful models for these cases will be presented in the second part of the chapter; this covers the log-linear models for contingency tables (where we analyse the relations between several categorical variables) and the logit model for quantal or binomial responses where we analyse the probability of being in one state as a function of explanatory variables.

### 8.1 General ANOVA and ANCOVA Models

### 8.1.1 ANOVA Models

## One-Factor Models

In Sect. 3.5, we introduced the example of analysing the effect of one factor (three possible marketing strategies) on the sales of a product (a pullover), see Table 3.2. The standard way to present one factor ANOVA models with $p$ levels is as follows

$$
\begin{equation*}
y_{k \ell}=\mu+\alpha_{\ell}+\varepsilon_{k \ell}, k=1, \ldots, n_{\ell}, \text { and } \ell=1, \ldots, p, \tag{8.2}
\end{equation*}
$$

all the $\varepsilon_{k \ell}$ being independent. Here $\ell$ is the label which indicates the level of the factor and $\alpha_{\ell}$ is the effect of the $\ell$ th level: it measures the deviation from $\mu$, the global mean of $y$, due to this level of the factor. In this notation, we need to impose the restriction $\sum_{\ell=1}^{p} \alpha_{\ell}=0$ in order to identify $\mu$ as the mean of $y$. This presentation is equivalent, but slightly different, to the one presented in Chap. 3 (compare with Eq. (3.41)), but it allows for easier extension to the multiple factors case. Note also that here we allow different sample sizes for each level of the factor (an unbalanced design, more general than the balanced design presented in Chap. 3).

To simplify the presentation, assume as in the pullover example that $p=3$. In this case, one could be tempted to write the model (8.2) under the general form of a linear model by using three indicator variables

$$
y_{i}=\mu+\alpha_{1} x_{i 1}+\alpha_{2} x_{i 2}+\alpha_{3} x_{i 3}+\varepsilon_{i},
$$

where $x_{i \ell}$ is equal to 1 or 0 according to the $i$ th observation and belongs (or not) to the level $\ell$ of the factor. In matrix notation and letting, for simplicity, $n_{1}=n_{2}=$ $n_{3}=2$ we have with $\beta=\left(\mu, \alpha_{1}, \alpha_{2}, \alpha_{3}\right)^{\top}$

$$
\begin{equation*}
y=\mathcal{X} \beta+\varepsilon, \tag{8.3}
\end{equation*}
$$

where the design matrix $\mathcal{X}$ is given by:

$$
\mathcal{X}=\left(\begin{array}{llll}
1 & 1 & 0 & 0 \\
1 & 1 & 0 & 0 \\
1 & 0 & 1 & 0 \\
1 & 0 & 1 & 0 \\
1 & 0 & 0 & 1 \\
1 & 0 & 0 & 1
\end{array}\right)
$$

Unfortunately, this type of coding is not useful because the matrix $\mathcal{X}$ is not of full rank (the sum of each row is equal to the same constant 2) and therefore the matrix $\mathcal{X}^{\top} \mathcal{X}$ is not invertible. One way to overcome this problem is to change the coding by introducing the additional constraint that the effects add up to zero. There are many ways to achieve this. Noting that $\alpha_{3}=-\alpha_{1}-\alpha_{2}$, we do not need to introduce $\alpha_{3}$ explicitly in the model. The linear model could indeed be written as

$$
y_{i}=\mu+\alpha_{1} x_{i 1}+\alpha_{2} x_{i 2}+\varepsilon_{i}
$$

with a design matrix defined as

$$
\mathcal{X}=\left(\begin{array}{rrr}
1 & 1 & 0 \\
1 & 1 & 0 \\
1 & 0 & 1 \\
1 & 0 & 1 \\
1 & -1 & -1 \\
1 & -1 & -1
\end{array}\right),
$$

which automatically implies that $\alpha_{3}=-\left(\alpha_{1}+\alpha_{2}\right)$. The linear model (8.3) is now correct with $\beta=\left(\mu, \alpha_{1}, \alpha_{2}\right)^{\top}$. The least squares estimator $\hat{\beta}=\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1} \mathcal{X}^{\top} y$ can be computed providing the estimator of the ANOVA parameters $\mu$ and $\alpha_{\ell}, \ell=$ $1, \ldots, 3$. Any linear constraint on $\beta$ can be tested by using the techniques described in Chap. 7. For instance, the null hypothesis of no factor effect $H_{0}: \alpha_{1}=\alpha_{2}=$ $\alpha_{3}=0$ can be written as $H_{0}: \mathcal{A} \beta=a$, where $\mathcal{A}=\left(\begin{array}{lll}0 & 1 & 0 \\ 0 & 0 & 1\end{array}\right)$ and $a=\left(\begin{array}{ll}0 & 0\end{array}\right)^{\top}$.

## Multiple-Factors Models

The coding above can be extended to more general situations with many qualitative variables (factors) and with the possibility of interactions between the factors. Suppose that in a marketing example, the sales of a product can be explained by two factors: the marketing strategy with three levels (as in the pullover example) but also the location of the shop that may be either in a big shopping centre or in a less commercial location (two levels for this factor). We might also think that there is an

Table 8.1 A two factor ANOVA data set, factor $A$, three levels of the marketing strategy and factor $B$, two levels for the location

|  | $B_{1}$ | $B_{2}$ |
| :--- | :--- | :--- |
| $A_{1}$ | 18 | 15 |
|  | 15 | 20 |
|  |  | 25 |
|  |  | 30 |
| $A_{2}$ | 5 | 10 |
|  | 8 | 12 |
|  | 8 |  |
| $A_{3}$ | 10 | 20 |
|  | 14 | 25 |

The figures represent the resulting sales during the same period
interaction between the two factors: the marketing strategy might have a different effect in a shopping centre than in a small quiet area. To fix the idea the data are collected as in Table 8.1.

The general two factor model with interactions can be written as

$$
\begin{equation*}
y_{i j k}=\mu+\alpha_{i}+\gamma_{j}+(\alpha \gamma)_{i j}+\varepsilon_{i j k} ; i=1, \ldots, r, j=1, \ldots, s, k=1, \ldots, n_{i j} \tag{8.4}
\end{equation*}
$$

where the identification constraints are:

$$
\begin{align*}
& \sum_{i=1}^{r} \alpha_{i}=0 \text { and } \sum_{j=1}^{s} \gamma_{j}=0 \\
& \sum_{i=1}^{r}(\alpha \gamma)_{i j}=0, j=1, \ldots, s  \tag{8.5}\\
& \sum_{j=1}^{s}(\alpha \gamma)_{i j}=0, i=1, \ldots, r .
\end{align*}
$$

In our example of Table 8.1 we have $r=3$ and $s=2$. The $\alpha$ 's measure the effect of the marketing strategy (three levels) and the $\gamma$ 's the effect of the location (two levels). A positive (negative) value of one of these parameters would indicate a favourable (unfavourable) effect on the expected sales; the global average of sales being represented by the parameter $\mu$. The interactions are measured by the parameters $(\alpha \gamma)_{i j}, i=1, \ldots, r, j=1, \ldots, s$, again identification constraints implies the $(r+s)$ constraints in (8.5) on the interactions terms.

For example, a positive value of $(\alpha \gamma)_{11}$ would indicate that the effect of the sale strategy $A_{1}$ (advertisement in local newspaper), if any, is more favourable on the sales in the location $B_{1}$ (in a big commercial centre) than in the location $B_{2}$ (not a commercial centre) with the relation $(\alpha \gamma)_{11}=-(\alpha \gamma)_{12}$. As another example, a negative value of $(\alpha \gamma)_{31}$ would indicate that the marketing strategy $A_{3}$ (luxury
presentation in shop windows) has less effect, if any, in location type $B_{1}$ than in $B_{2}$ : again $(\alpha \gamma)_{31}=-(\alpha \gamma)_{32}$, etc.

The nice thing is that it is easy to extend the coding rule for one-factor model to this general situation, in order to present the model a standard linear model with the appropriate design matrix $\mathcal{X}$. To build the columns of $\mathcal{X}$ for the effect of each factor, we will need, as above, $r-1$ (and $s-1$ ) variables for coding a qualitative variable with $r$ (and $s$, respectively) levels with the convention defined above in the one-factor case. For the interactions between a $r$ between a $r$ level factor and a $s$ level factor, we will need $(r-1) \times(s-1)$ additional columns that will be obtained by performing the product, element by element, of the corresponding main effect columns. So, at the end, for a full model with all the interactions, we have $\{1+r-1+s-1+(r-1)(s-1)\}=r s$ parameters where the first column of 1 's is for the intercept (the constant $\mu$ ). We illustrate this for our marketing example where $r=3$ and $s=2$. We first describe a model without interactions.

## 1. Model without interactions

Without the interactions (all the $(\alpha \gamma)_{i j}=0$ ) the model could be written with $3=(r-1)+(s-1)$ coded variables in a simple linear model form as in (8.3), with the matrices:

$$
y=\left(\begin{array}{c}
18 \\
15 \\
15 \\
20 \\
25 \\
30 \\
5 \\
8 \\
8 \\
10 \\
12 \\
10 \\
14 \\
20 \\
25
\end{array}\right), \mathcal{X}=\left(\begin{array}{rrrr}
1 & 1 & 0 & 1 \\
1 & 1 & 0 & 1 \\
1 & 1 & 0 & -1 \\
1 & 1 & 0 & -1 \\
1 & 1 & 0 & -1 \\
1 & 1 & 0 & -1 \\
1 & 0 & 1 & 1 \\
1 & 0 & 1 & 1 \\
1 & 0 & 1 & 1 \\
1 & 0 & 1 & -1 \\
1 & 0 & 1 & -1 \\
1 & -1 & -1 & 1 \\
1 & -1 & -1 & 1 \\
1 & -1 & -1 & -1 \\
1-1 & -1 & -1
\end{array}\right),
$$

and $\beta=\left(\mu, \alpha_{1}, \alpha_{2}, \gamma_{1}\right)^{\top}$. Then, $\alpha_{3}=-\left(\alpha_{1}+\alpha_{2}\right)$ and $\gamma_{2}=-\gamma_{1}$.
2. Model with interactions

A model with interaction between $A$ and $B$ is obtained by adding new columns to the design matrix. We need $2=(r-1) \times(s-1)$ new coding variables which are defined as the product, element-by-element, of the corresponding columns obtained for the main effects. For instance for the interaction parameter $(\alpha \gamma)_{11}$, we multiply the column used for coding $\alpha_{1}$ by the column defined for coding $\gamma_{1}$, where the product is element-by-element. The same is done for the parameter
$(\alpha \gamma)_{21}$. No other columns are necessary, since the remaining interactions are derived from the identification constraints (8.5). We obtain

$$
\mathcal{X}=\left(\begin{array}{rrrrrr}
1 & 1 & 0 & 1 & 1 & 0 \\
1 & 1 & 0 & 1 & 1 & 0 \\
1 & 1 & 0 & -1 & -1 & 0 \\
1 & 1 & 0 & -1 & -1 & 0 \\
1 & 1 & 0 & -1 & -1 & 0 \\
1 & 1 & 0 & -1 & -1 & 0 \\
1 & 0 & 1 & 1 & 0 & 1 \\
1 & 0 & 1 & 1 & 0 & 1 \\
1 & 0 & 1 & 1 & 0 & 1 \\
1 & 0 & 1 & -1 & 0 & -1 \\
1 & 0 & 1 & -1 & 0 & -1 \\
1 & -1 & -1 & 1 & -1 & -1 \\
1 & -1 & -1 & 1 & -1 & -1 \\
1 & -1 & -1 & -1 & 1 & 1 \\
1 & -1 & -1 & -1 & 1 & 1
\end{array}\right),
$$

with $\beta=\left(\mu, \alpha_{1}, \alpha_{2}, \gamma_{1},(\alpha \gamma)_{11},(\alpha \gamma)_{21}\right)^{\top}$. The other interactions can indeed be derived from (8.5)

$$
\begin{aligned}
& (\alpha \gamma)_{12}=-(\alpha \gamma)_{11} \\
& (\alpha \gamma)_{22}=-(\alpha \gamma)_{21} \\
& (\alpha \gamma)_{31}=-\left((\alpha \gamma)_{11}+(\alpha \gamma)_{21}\right) \\
& (\alpha \gamma)_{32}=-(\alpha \gamma)_{31} .
\end{aligned}
$$

The estimation of $\beta$ is again simply given by the least squares solution $\hat{\beta}=$ $\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1} \mathcal{X}^{\top} y$.

Example 8.2 Let us come back to the marketing data provided by the two-way Table 8.1. The values of $\hat{\beta}$ in the full model, with interactions, are given in Table 8.2. The $p$-values in the right column are for the individual tests: it appears that the interactions do not provide additional significant explanation of $y$, but the effect of the two factors seems significant.

Using the techniques of Chap. 7, we can test some reduced model corresponding to linear constraints on the $\beta$ 's. The full model is the model with all the parameters, including all the interactions. The overall fit test $H_{0}$ : all the parameters, except $\mu$, are equal to zero, gives the value $F_{\text {observed }}=6.5772$ with a $p$-value of 0.0077 for a $F_{5,9}$, so that $H_{0}$ is rejected. In this case, the $\mathrm{RSS}_{\text {reduced }}=735.3333$. So there is some effect by the factors.

Table 8.2 Estimation of the two factors ANOVA model with data from Table 8.1

|  | $\hat{\beta}$ | $p$-Values |
| :--- | ---: | :--- |
| $\mu$ | 15.25 |  |
| $\alpha_{1}$ | 4.25 | 0.0218 |
| $\alpha_{2}$ | -6.25 | 0.0033 |
| $\gamma_{1}$ | -3.42 | 0.0139 |
| $(\alpha \gamma)_{11}$ | 0.42 | 0.7922 |
| $(\alpha \gamma)_{21}$ | 1.42 | 0.8096 |
| RSS $_{\text {full }}$ | 158.00 |  |

We then test a less reduced model. We can test if the interaction terms are significantly different to zero. This is a linear constraint on $\beta$ with

$$
\mathcal{A}=\left(\begin{array}{cccccc}
0 & 0 & 0 & 0 & 1 & 0 \\
0 & 0 & 0 & 0 & 0 & 1
\end{array}\right) ; a=\binom{0}{0}
$$

Under the null we obtain:

$$
\hat{\beta}_{H_{0}}=\left(\begin{array}{r}
15.3035 \\
4.0975 \\
-6.0440 \\
-3.2972 \\
0 \\
0
\end{array}\right)
$$

and $\mathrm{RSS}_{\text {reduced }}=181.8019$. The observed value of $F=0.6779$ which is not significant $(r=11, f=9)$ the $p$-value $=\mathrm{P}\left(F_{2,9} \geq 0.6779\right)=0.5318$, confirming the absence of interactions.

Now taking the model without the interactions as the full model, we can test if one of the main effects $\alpha$ (marketing strategy) or $\gamma$ (location) or both are significantly different from zero. We leave this as an exercise for the reader.

### 8.1.2 ANCOVA Models

ANCOVA (ANalysis of COVAriances) are mixed models where some variables are qualitative and others are quantitative. The same coding of the ANOVA will be used for the qualitative variable. The design matrix $\mathcal{X}$ is completed by the columns for the quantitative explanatory variables $x$. Interactions between a qualitative variable (a factor with $r$ levels) and a quantitative one $x$ is also possible, this corresponds to situations where the effect of $x$ on the response $y$ is different according to the level of the factor. This is achieved by adding into the design matrix $\mathcal{X}$, a new column obtained by the product, element-by-element, of the quantitative variable with the
coded variables for the factor ( $r-1$ interaction variables if the categorical variable has $r$ levels).

For instance consider a simple model where a response $y$ is explained by one explanatory variable $x$ and one factor with two levels (for instance the gender level 1 for men and level 2 for women), we would have in the case $n_{1}=n_{2}=3$

$$
\mathcal{X}=\left(\begin{array}{rrrr}
1 & x_{1} & 1 & x_{1} \\
1 & x_{2} & 1 & x_{2} \\
1 & x_{3} & 1 & x_{3} \\
1 & x_{4} & -1 & -x_{4} \\
1 & x_{5} & -1 & -x_{5} \\
1 & x_{6} & -1 & -x_{6}
\end{array}\right),
$$

with $\beta=\left(\beta_{1}, \beta_{2}, \beta_{3}, \beta_{4}\right)^{\top}$. The intercept and the slope are $\left(\beta_{1}+\beta_{3}\right)$ and $\left(\beta_{1}+\beta_{4}\right)$ for men and $\left(\beta_{1}-\beta_{3}\right)$ and $\left(\beta_{1}-\beta_{4}\right)$ for women. This situation is displayed in Fig. 8.2.

Example 8.3 Consider the Car Data provided in Sect. 22.3. We want to analyse the effect of the weight ( $W$ ), the displacement $(D)$ on the mileage $(M)$. But we would like to test if the origin of the car (the factor $C$ ) has some effect on the response and if the effect of the continuous variables is different for the different levels of the factor.

From the regression results in Table 8.3, we observe that only the weight affects the mileage, while the displacement does not. We also consider the origin of the car, however, both the displacement and the factor are not significant. Table 8.4 is for different factor levels.


Fig. 8.2 A model with interaction

Table 8.3 Estimation of the effects of weight and displacement on the mileage Q MVAcareffect

|  | $\hat{\beta}$ | $p$-Values | $\tilde{\beta}$ | $p$-Values |
| :--- | ---: | :--- | :--- | :--- |
| $\mu$ | 41.0066 | 0.0000 | 43.4031 | 0.0000 |
| $W$ | -0.0073 | 0.0000 | -0.0074 | 0.0000 |
| $D$ | 0.0118 | 0.2250 | 0.0081 | 0.4140 |
| $C$ |  |  | -0.9675 | 0.1250 |

Table 8.4 Different factor levels on the response $\mathbf{Q}$ MVAcareffect

|  | $\mu$ | $p$-Values | $W$ | $p$-Values | $D$ | $p$-Values |
| :--- | :--- | :--- | ---: | :--- | :--- | :--- |
| $c=1$ | 40.043 | 0.0000 | -0.0065 | 0.0000 | 0.0058 | 0.3790 |
| $c=2$ | 47.557 | 0.0005 | 0.0081 | 0.3666 | -0.3582 | 0.0160 |
| $c=3$ | 44.174 | 0.0002 | 0.0039 | 0.7556 | -0.2650 | 0.3031 |

### 8.1.3 Boston Housing

In Chaps. 3 and 7, linear models were used to analyse if the variations of the price (the variables were transformed in Sect. 1.9) could be explained by other variables. A reduced model was obtained in Sect. 7.3 with the results shown in Table 7.1, with $r^{2}=0.763$. The model was:

$$
\begin{aligned}
X_{14}= & \beta_{0}+\beta_{4} X_{4}+\beta_{5} X_{5}+\beta_{6} X_{6}+\beta_{8} X_{8}+\beta_{9} X_{9}+\beta_{10} X_{10}+\beta_{11} X_{11} \\
& +\beta_{12} X_{12}+\beta_{13} X_{13}
\end{aligned}
$$

One factor $\left(X_{4}\right)$ was coded as a binary variable (1, if the house is close to the Charles River and 0 if it is not). Taking advantage of the ANCOVA models described above, we would like to add to a new factor built from the original quantitative variable $X_{9}=$ index of accessibility to radial highways. So we will transform $X_{4}$ as being 1 if close to the Charles River and -1 if not, and we will replace $X_{9}$ by a new factor coded $X_{15}=1$ if $X_{9} \geq$ median $\left(X_{9}\right)$ and $X_{15}=-1$ if $X_{9}<\operatorname{median}\left(X_{9}\right)$. We also want to consider the interaction of $X_{4}$ with $X_{12}$ (proportion of blacks) and the interaction of $X_{4}$ with the new factor $X_{15}$. The results are shown in Table 8.5.

|  | Summary |
| :--- | :--- |
|  | ANOVA models can be dividend into one-factor models and <br> multiple factor models. |
|  | Multiple factor models analyse many qualitative variables and the <br> interactions between them. |

Table 8.5 Estimation of the ANCOVA model using the Boston housing data Q MVAboshousing

|  | $\hat{\beta}$ | $p$-Values | $\tilde{\beta}$ | $p$-Values |
| :--- | ---: | :--- | ---: | :--- |
| $\beta_{0}$ | 32.27 | 0.00 | 27.65 | 0.00 |
| $\beta_{4}$ | 1.54 | 0.00 | -3.19 | 0.32 |
| $\beta_{5}$ | -17.59 | 0.00 | -16.50 | 0.00 |
| $\beta_{6}$ | 4.27 | 0.00 | 4.23 | 0.00 |
| $\beta_{8}$ | -1.13 | 0.00 | -1.10 | 0.00 |
| $\beta_{10}$ | 0.00 | 0.97 | 0.00 | 0.95 |
| $\beta_{11}$ | -0.97 | 0.00 | -0.97 | 0.00 |
| $\beta_{12}$ | 0.01 | 0.00 | 0.02 | 0.01 |
| $\beta_{13}$ | -0.54 | 0.00 | -0.54 | 0.00 |
| $\beta_{15}$ | 0.21 | 0.46 | 0.23 | 0.66 |
| $\beta_{4 * 14}$ |  |  | 0.01 | 0.13 |
| $\beta_{4 * 15}$ |  |  | 0.03 | 0.95 |

## Summary (continued)

$\hookrightarrow$ ANCOVA models are mixed models with qualitative and quantitative variables, and can also incorporate the interaction between a qualitative and a quantitative variable.

### 8.2 Categorical Responses

### 8.2.1 Multinomial Sampling and Contingency Tables

In many applications, the response variable of interest is qualitative or categorical, in the sense that the response can take its nominal value in one of, say, $K$ classes or categories. Often we observe counts $y_{k}$, the number of observations in category $k=1, \ldots, K$. If the total number of observations $n=\sum_{k=1}^{K} y_{k}$ is fixed and we may assume independence of the observations, we obtain a multinomial sampling process.

If we denote by $p_{k}$ the probability of observing the $k$ th category with $\sum_{k=1}^{K} p_{k}=$ 1, we have $\mathrm{E}\left(Y_{k}\right)=m_{k}=n p_{k}$. The likelihood of the sample can then be written as:

$$
\begin{equation*}
L=\frac{n!}{\prod_{k=1}^{K} y_{k}!} \prod_{k=1}^{K}\left(\frac{m_{k}}{n}\right)^{y_{k}} . \tag{8.6}
\end{equation*}
$$

In contingency tables, the categories are defined by several qualitative variables. For example in a $(J \times K)$ two-way table, the observations (counts) $y_{j k}, j=1, \ldots, J$ and $k=1, \ldots, K$ are reported for row $j$ and column $k$. Here $n=\sum_{j=1}^{J} \sum_{k=1}^{K} y_{j k}$. Log-linear models introduce a linear structure on the logarithms of the expected frequencies $m_{j k}=\mathrm{E}\left(y_{j k}\right)=n p_{j k}$, with $\sum_{j=1}^{J} \sum_{k=1}^{K} p_{j k}=1$. Log-linear structures on $m_{j k}$ will impose the same structure for the $p_{j k}$, the estimation of the model will then be obtained by constrained maximum likelihood. Three-way tables ( $J \times K \times L$ ) may be analysed in the same way.

Sometimes additional information is available on explanatory variables $x$. In this case, the logit model will be appropriate when the categorical response is binary ( $K=2$ ). We will introduce these models when the main response of interest is binary (for instance tables $(2 \times K)$ or $(2 \times K \times L))$. Further, we will show how they can be adapted to the case of contingency tables. Contingency tables are also analysed by multivariate descriptive tools in Chap. 15.

### 8.2.2 Log-Linear Models for Contingency Tables

## Two-Way Tables

Consider a $(J \times K)$ two-way table, where $y_{j k}$ is the number of observations having the nominal value $j$ for the first qualitative character and nominal value $k$ for the second character. Since the total number of observations is fixed $n=$ $\sum_{j=1}^{J} \sum_{k=1}^{K} y_{j k}$, there are $J K-1$ free cells in the table. The multinomial likelihood can be written as in (8.6)

$$
\begin{equation*}
L=\frac{n!}{\prod_{j=1}^{J} \prod_{k=1}^{K} y_{j k}!} \prod_{j=1}^{J} \prod_{k=1}^{K}\left(\frac{m_{j k}}{n}\right)^{y_{j k}} \tag{8.7}
\end{equation*}
$$

where we now introduce a log-linear structure to analyse the role of the rows and the columns to determine the parameters $m_{j k}=\mathrm{E}\left(y_{j k}\right)$ (or $\left.p_{j k}\right)$.

1. Model without interaction

Suppose that there is no interaction between the rows and the columns: this corresponds to the hypothesis of independence between the two qualitative characters. In other words, $p_{j k}=p_{j} p_{k}$ for all $j, k$. This implies the log-linear model:

$$
\begin{equation*}
\log m_{j k}=\mu+\alpha_{j}+\gamma_{k} \text { for } j=1, \ldots, J, k=1, \ldots, K \tag{8.8}
\end{equation*}
$$

where, as in ANOVA models for identification purposes $\sum_{j=1}^{J} \alpha_{j}=\sum_{k=1}^{K} \gamma_{k}=$ 0 . Using the same coding devices as above, the model can be written as

$$
\begin{equation*}
\log m=\mathcal{X} \beta \tag{8.9}
\end{equation*}
$$

For a $(2 \times 3)$ table we have:

$$
\log m=\left(\begin{array}{l}
\log m_{11} \\
\log m_{12} \\
\log m_{13} \\
\log m_{21} \\
\log m_{22} \\
\log m_{23}
\end{array}\right), \mathcal{X}=\left(\begin{array}{rrrr}
1 & 1 & 1 & 0 \\
1 & 1 & 0 & 1 \\
1 & 1 & -1 & -1 \\
1 & -1 & 1 & 0 \\
1 & -1 & 0 & 1 \\
1 & -1 & -1 & -1
\end{array}\right), \beta=\left(\begin{array}{c}
\beta_{0} \\
\beta_{1} \\
\beta_{2} \\
\beta_{3}
\end{array}\right)
$$

where the first column of $\mathcal{X}$ is for the constant term, the second column is the coded column for the 2-levels row effect and the two last columns are the coded columns for the 3-levels column effect. The estimation is obtained by maximising the log-likelihood which is equivalent to maximising the function $L(\beta)$ in $\beta$ :

$$
\begin{equation*}
L(\beta)=\sum_{j=1}^{J} \sum_{k=1}^{K} y_{j k} \log m_{j k} . \tag{8.10}
\end{equation*}
$$

The maximisation is under the constraint $\sum_{j, k} m_{j k}=n$. In summary we have $1+(J-1)+(K-1)-1$ free parameters for $J K-1$ free cells. The number of degrees of freedom in the model is the number of free cells minus the number of free parameters. It is given by

$$
r=J K-1-(J-1)-(K-1)=(J-1)(K-1) .
$$

In the example above, we have therefore $(3-1) \times(2-1)=2$ degrees of freedom.
The original parameters of the model can then be estimated as:

$$
\begin{align*}
& \alpha_{1}=\beta_{1} \\
& \alpha_{2}=-\beta_{1} \\
& \gamma_{1}=\beta_{2} \\
& \gamma_{2}=\beta_{3} \\
& \gamma_{3}=-\left(\beta_{2}+\beta_{3}\right) . \tag{8.11}
\end{align*}
$$

## 2. Model with interactions

In two-way tables the interactions between the two variables are of interest. This corresponds to the general (full) model

$$
\begin{equation*}
\log m_{j k}=\mu+\alpha_{j}+\gamma_{k}+(\alpha \gamma)_{j k}, j=1, \ldots, J, k=1, \ldots, K \tag{8.12}
\end{equation*}
$$

where in addition, we have the $J+K$ restrictions

$$
\begin{align*}
& \sum_{k=1}^{K}(\alpha \gamma)_{j k}=0, \text { for } j=1, \ldots, J \\
& \sum_{j=1}^{J}(\alpha \gamma)_{j k}=0, \text { for } k=1, \ldots, K \tag{8.13}
\end{align*}
$$

As in the ANOVA model, the interactions may be coded by adding $(J-1)(K-1)$ columns to $\mathcal{X}$, obtained by the product of the corresponding coded variables. In our example for the $(2 \times 3)$ table the design matrix $\mathcal{X}$ is completed with two more columns:

$$
\mathcal{X}=\left(\begin{array}{rrrrrr}
1 & 1 & 1 & 0 & 1 & 0 \\
1 & 1 & 0 & 1 & 0 & 1 \\
1 & 1 & -1 & -1 & -1 & -1 \\
1 & -1 & 1 & 0 & -1 & 0 \\
1 & -1 & 0 & 1 & 0 & -1 \\
1 & -1 & -1 & -1 & 1 & 1
\end{array}\right), \beta=\left(\begin{array}{l}
\beta_{0} \\
\beta_{1} \\
\beta_{2} \\
\beta_{3} \\
\beta_{4} \\
\beta_{5}
\end{array}\right) .
$$

Now the interactions are determined by using (8.13):

$$
\begin{aligned}
& (\alpha \gamma)_{11}=\beta_{4} \\
& (\alpha \gamma)_{12}=\beta_{5} \\
& (\alpha \gamma)_{13}=-\left\{(\alpha \gamma)_{11}+(\alpha \gamma)_{12}\right\}=-\left(\beta_{4}+\beta_{5}\right) \\
& (\alpha \gamma)_{21}=-(\alpha \gamma)_{11}=-\beta_{4} \\
& (\alpha \gamma)_{22}=-(\alpha \gamma)_{12}=-\beta_{5} \\
& (\alpha \gamma)_{23}=-(\alpha \gamma)_{13}=\beta_{4}+\beta_{5}
\end{aligned}
$$

We have again a log-linear model as in (8.9) and the estimation of $\beta$ goes through the maximisation in $\beta$ of $L(\beta)$ given by (8.10) under the same constraint.

The model with all the interaction terms is called the saturated model. In this model there are no degrees of freedom, the number of free parameters to be estimated equals the number of free cells. The parameters of interest are the interactions. In particular, we are interested in testing their significance. These issues will be addressed below.

## Three-Way Tables

The models presented above for two-way tables can be extended to higher order tables but at a cost of notational complexity. We show how to adapt to threeway tables. This deserves special attention due to the presence of higher-order interactions in the saturated model.

A $(J \times K \times L)$ three-way table may be constructed under multinomial sampling as follows: each of the $n$ observations falls in one, and only one, category of each of three categorical variables having $J, K$ and $L$ modalities respectively. We end up with a three-dimensional table with $J K L$ cells containing the counts $y_{j k \ell}$ where $n=\sum_{j, k, \ell} y_{j k \ell}$. The expected counts depend on the unknown probabilities $p_{j k \ell}$ in the usual way:

$$
m_{j k \ell}=n p_{j k \ell}, j=1, \ldots, J, k=1, \ldots, K, \ell=1, \ldots, L .
$$

1. The saturated model

A full saturated log-linear model reads as follows:

$$
\begin{align*}
\log m_{j k \ell}= & \mu+\alpha_{j}+\beta_{k}+\gamma_{\ell}+(\alpha \beta)_{j k}+(\alpha \gamma)_{j \ell}+(\beta \gamma)_{k \ell}+(\alpha \beta \gamma)_{j k \ell} \\
& j=1, \ldots, J, k=1, \ldots, K, \ell=1, \ldots, L \tag{8.14}
\end{align*}
$$

The restrictions are the following (using the "dot" notation for summation on the corresponding indices):

$$
\begin{aligned}
& \alpha_{(\bullet)}=\beta_{(\bullet)}=\gamma_{(\bullet)}=0 \\
& (\alpha \beta)_{j \bullet}=(\alpha \gamma)_{j \bullet}=(\beta \gamma)_{k \bullet}=0 \\
& (\alpha \beta)_{\bullet k}=(\alpha \gamma)_{\bullet \ell}=(\beta \gamma)_{\bullet \ell}=0 \\
& (\alpha \beta \gamma)_{j k \bullet}=(\alpha \beta \gamma)_{j \bullet \ell}=(\alpha \beta \gamma)_{\bullet k \ell}=0
\end{aligned}
$$

The parameters $(\alpha \beta)_{j k},(\alpha \gamma)_{j \ell},(\beta \gamma)_{k \ell}$ are called first-order interactions. The second-order interactions are the parameters $(\alpha \beta \gamma)_{j k \ell}$, they allow to take into account heterogeneities in the interactions between two of the three variables. For instance, let $\ell$ stand for the two gender categories ( $L=2$ ), if we suppose that $(\alpha \beta \gamma)_{j k 1}=-(\alpha \beta \gamma)_{j k 2} \neq 0$, we mean that the interactions between the variable $J$ and $K$ are not the same for both gender categories.

The estimation of the parameters of the saturated model are obtained through maximisation of the log-likelihood. In the multinomial sampling scheme, it corresponds to maximising the function:

$$
L=\sum_{j, k, \ell} y_{j k \ell} \log m_{j k \ell},
$$

under the constraint $\sum_{j, k, \ell} m_{j k \ell}=n$.

The number of degrees of freedom in the saturated model is again zero. Indeed, the number of free parameters in the model is

$$
\begin{aligned}
1 & +(J-1)+(K-1)+(L-1)+(J-1)(K-1)+(J-1)(L-1) \\
& +(K-1)(L-1)+(J-1)(K-1)(L-1)-1=J K L-1 .
\end{aligned}
$$

This is indeed equal to the number of free cells in the table and so, there is no degree of freedom.
2. Hierarchical non-saturated models

As illustrated above, a saturated model has no degrees of freedom. Non-saturated models correspond to reduced models where some parameters are fixed to be equal to zero. They are thus particular cases of the saturated model (8.14). The hierarchical non-saturated models that we will consider here, are models where once a set of parameters is set equal to zero, all the parameters of higher-order containing the same indices are also set equal to zero.

For instance if we suppose $\alpha_{1}=0$, we only consider non-saturated models where also $(\alpha \gamma)_{1 \ell}=(\alpha \beta)_{1 k}=(\alpha \beta \gamma)_{1 k \ell}=0$ for all values of $k$ and $\ell$. If we only suppose that $(\alpha \beta)_{12}=0$, we also assume that $(\alpha \beta \gamma)_{12 \ell}=0$ for all $\ell$.

Hierarchical models have the advantage of being more easily interpretable. Indeed without this hierarchy, the models would be difficult to interpret. What would be, for instance, the meaning of the parameter $(\alpha \beta \gamma)_{12 \ell}$, if we know that $(\alpha \beta)_{12}=0$ ? The estimation of the non-saturated models will be achieved by the usual way i.e. by maximising the $\log$-likelihood function $L$ as above but under the new constraints of the reduced model.

### 8.2.3 Testing Issues with Count Data

One of the main practical interests in regression models for contingency tables is to test restrictions on the parameters of a more complete model. These testing ideas are created in the same spirit as in Sect. 3.5 where we tested restrictions in ANOVA models.

In linear models, the test statistics is based on the comparison of the goodness of fit for the full model and for the reduced model. Goodness of fit is measured by the residual sum of squares (RSS). The idea here will be the same here but with a more appropriate measure for goodness of fit. Once a model has been estimated, we can compute the predicted value under that model for each cell of the table. We will denote, as above, the observed value in a cell by $y_{k}$ and $\hat{m}_{k}$ will denote the expected value predicted by the model. The goodness of fit may be appreciated by measuring, in some way, the distance between the series of observed and of predicted values.

Two statistics are proposed: the Pearson chi-square $X^{2}$ and the Deviance noted $G^{2}$. They are defined as follows:

$$
\begin{align*}
& X^{2}=\sum_{k=1}^{K} \frac{\left(y_{k}-\hat{m}_{k}\right)^{2}}{\hat{m}_{k}}  \tag{8.15}\\
& G^{2}=2 \sum_{k=1}^{K} y_{k} \log \left(\frac{y_{k}}{\hat{m}_{k}}\right) \tag{8.16}
\end{align*}
$$

where $K$ is the total number of cells of the table. The deviance is directly related to the log-likelihood ratio statistic and is usually preferred because it can be used to compare nested models as we usually do in this context.

Under the hypothesis that the model used to compute the predicted value is true, both statistics (for large samples) are approximately distributed as a $\chi^{2}$ variable with degrees of freedom $d . f$. depending on the model. The $d . f$. can be computed as follows:

$$
\begin{equation*}
d . f .=\# \text { free cells }-\# \text { free parameters estimated. } \tag{8.17}
\end{equation*}
$$

For saturated models, the fit is perfect: $X^{2}=G^{2}=0$ with $d . f .=0$.
Suppose now that we want to test a reduced model which is a restricted version of a full model. The deviance can then be used as the $F$ statistics in linear regression. The test procedure is straightforward:
$H_{0}:$ reduced model with $r$ degrees of freedom
$H_{1}:$ full model with $f$ degrees of freedom.

Since, the full model contains more parameters, we expect the deviance to be smaller. We reject the $H_{0}$ if this reduction is significant, i.e. if $G_{H_{0}}^{2}-G_{H_{1}}^{2}$ is large enough. Under $H_{0}$ one has:

$$
G_{H_{0}}^{2}-G_{H_{1}}^{2} \sim \chi_{r-f}^{2}
$$

We reject $H_{0}$ if the $p$-value:

$$
\mathrm{P}\left\{\chi_{r-f}^{2}>\left(G_{H_{0}}^{2}-G_{H_{1}}^{2}\right)\right\} .
$$

is small. Suppose we want to test the independence in a $(J \times K)$ two-way table (no interaction). Here the full model is the saturated one with no degrees of freedom $(f=0)$ and the restricted model has $r=(J-1)(K-1)$ degrees of freedom. We reject $H_{0}$ if the $p$-value of $H_{0} \mathrm{P}\left\{\chi_{r}^{2}>\left(G_{H_{0}}^{2}\right)\right\}$ is too small.

This test is equivalent to the Pearson chi-square test for independence in two-way tables ( $G_{H_{0}}^{2} \approx X_{H_{0}}^{2}$ when $n$ is large).

Table 8.6 A three-way contingency table: top table for men and bottom table for women $\mathbf{a}$ MVAdrug

| M | A1 | A2 | A3 | A4 | A5 |
| :--- | :--- | :--- | :--- | :--- | :--- |
| DY | 21 | 32 | 70 | 43 | 19 |
| DN | 683 | 596 | 705 | 295 | 99 |
| F | A1 | A2 | A3 | A4 | A5 |
| DY | 46 | 89 | 169 | 98 | 51 |
| DN | 738 | 700 | 847 | 336 | 196 |

Table 8.7 Coefficient estimates based on the saturated model Q mVAdrug

|  |  |  |  |
| :--- | ---: | :--- | ---: |
| $\hat{\beta}_{0}$ intercept | 5.0089 | $\hat{\beta}_{10}$ | $\hat{\beta}$ |
| $\hat{\beta}_{1}$ gender: M | -0.2867 | $\hat{\beta}_{11}$ | 0.0205 |
| $\hat{\beta}_{2}$ drug: DY | -1.0660 | $\hat{\beta}_{12}$ drug*age | 0.0482 |
| $\hat{\beta}_{3}$ age | -0.0080 | $\hat{\beta}_{13}$ | -0.4983 |
| $\hat{\beta}_{4}$ | 0.2151 | $\hat{\beta}_{14}$ | -0.1807 |
| $\hat{\beta}_{5}$ | 0.6607 | $\hat{\beta}_{15}$ | 0.0857 |
| $\hat{\beta}_{6}$ | -0.0463 | $\hat{\beta}_{16}$ gender*drug*age | 0.2766 |
| $\hat{\beta}_{7}$ gender*drug | -0.1632 | $\hat{\beta}_{17}$ | -0.0134 |
| $\hat{\beta}_{8}$ gender*age | 0.0713 | $\hat{\beta}_{18}$ | -0.0523 |
| $\hat{\beta}_{9}$ | -0.0092 | $\hat{\beta}_{19}$ | -0.0112 |

Example 8.4 Everitt and Dunn (1998) provide a three-dimensional $(2 \times 2 \times 5)$ count table of $n=5,833$ interviewed people. The count were on prescribed psychotropic drugs in the fortnight prior to the interview as a function of age and gender. The data are summarised in Table 8.6, where the categories for the three factors are M for male, F for female, DY for "yes" having taken drugs, DN for "no" not having taking drugs and the five age categories: A1 (16-29), A2 (30-44), A3 (45-64), A4 (65-74), A5 for over 74. The table provides the observed frequencies $y_{j k \ell}$ in each of the cells of the three-way table: where $j$ stands for gender, $k$ for drug and $\ell$ for age categories. The design matrix $\mathcal{X}$ for the full saturated model can be found in the quantlet Q MVAdrug.

The saturated model gives the estimates displayed in Table 8.7.
We see for instance that $\hat{\beta}_{1}<0$, so there are fewer men than women in the study, since $\hat{\beta}_{7}$ is also negative it seems that the tendency of men taking the drug is less important than for women. Also, note that $\hat{\beta}_{12}$ to $\hat{\beta}_{15}$ forms an increasing sequence, so that the age factor seems to increase the tendency to take the drug. Note that in this saturated model, there are no degrees of freedom and the fit is perfect, $\hat{m}_{j k \ell}=y_{j k \ell}$ for all the cells of the table.

The second order interactions have a lower order of magnitude, so we want to test if they are significantly different to zero. We consider a restricted model where

Table 8.8 Coefficients estimates based on the maximum likelihood method Q MVAdrug3waysTab

|  | $\hat{\beta}$ |  | $\hat{\beta}$ |
| :--- | ---: | :--- | ---: |
| $\hat{\beta}_{0}$ intercept | 5.0051 | $\hat{\beta}_{8}$ gender*age | 0.0795 |
| $\hat{\beta}_{1}$ gender: M | -0.2919 | $\hat{\beta}_{9}$ | 0.0321 |
| $\hat{\beta}_{2}$ drug: DY | -1.0717 | $\hat{\beta}_{10}$ | 0.0265 |
| $\hat{\beta}_{3}$ age | -0.0030 | $\hat{\beta}_{11}$ | 0.0534 |
| $\hat{\beta}_{4}$ | 0.2358 | $\hat{\beta}_{12}$ drug*age | -0.4915 |
| $\hat{\beta}_{5}$ | 0.6649 | $\hat{\beta}_{13}$ | -0.1576 |
| $\hat{\beta}_{6}$ | -0.0425 | $\hat{\beta}_{14}$ | 0.0917 |
| $\hat{\beta}_{7}$ gender*drug | -0.1734 | $\hat{\beta}_{15}$ | 0.2822 |

$(\alpha \beta \gamma)_{j k \ell}$ are all set to zero. This can be achieved by testing $H_{0}: \beta_{16}=\beta_{17}=$ $\beta_{18}=\beta_{19}=0$. The maximum likelihood estimators of the restricted model are obtained by deleting the last four columns in the design matrix $\mathcal{X}$. The results are given in Table 8.8.

We have $J=2, K=2$ and $L=5$, this makes $J K L-1=19$ free cells. The full model has $f=0$ degrees of freedom and the reduced model has $r=4$ degrees of freedom. The $G^{2}$ deviance is given by 2.3004 ; it has 4 degrees of freedom (the chisquare statistics is 2.3745 ). The $p$-value of the restricted model is 0.6807 , so we do not reject the null hypothesis (the restricted model without 2 nd order interaction). In others words, age does not interfere with the interactions between gender and drugs, or equivalently, gender does not interfere in the interactions between age and drugs. The reader can verify that the first order interactions are significant, by taking, for instance, the model without interactions of the second order as the new full model and testing a reduced model where all the first order interactions are all set to zero. Q MVAdrug3waysTab

### 8.2.4 Logit Models

Logit models are useful to analyse how explanatory variables influence a binary response $y$. The response $y$ may take the two values 1 and 0 to denote the presence or absence of a certain qualitative trait (a person can be employed or unemployed, a firm can be bankrupt or not, a patient can be affected by a certain disease or not, etc.). Logit models are designed to estimate the probability of $y=1$ as a logistic function of linear combinations of $x$. Logit models can be adapted to the analysis of contingency tables where one of the qualitative variables is binary. One obtains the probability of being in one of the two states of this binary variable as a function of the other variables. We concentrate here on $(2 \times K)$ and $(2 \times K \times L)$ tables.

## Logit Models for Binary Response

Consider the vector $y(n \times 1)$ of observations on a binary response variable (a value of " 1 " indicating the presence of a particular qualitative trait and a value of " 0 ", its absence). The logit model makes the assumption that the probability for observing $y_{i}=1$ given a particular value of $x_{i}=\left(x_{i 1}, \ldots, x_{i p}\right)^{\top}$ is given by the logistic function of a "score", a linear combination of $x$ :

$$
\begin{equation*}
p\left(x_{i}\right)=\mathrm{P}\left(y_{i}=1 \mid x_{i}\right)=\frac{\exp \left(\beta_{0}+\sum_{j=1}^{p} \beta_{j} x_{i j}\right)}{1+\exp \left(\beta_{0}+\sum_{j=1}^{p} \beta_{j} x_{i j}\right)} \tag{8.19}
\end{equation*}
$$

This entails the probability of the absence of the trait:

$$
1-p\left(x_{i}\right)=\mathrm{P}\left(y_{i}=0 \mid x_{i}\right)=\frac{1}{1+\exp \left(\beta_{0}+\sum_{j=1}^{p} \beta_{j} x_{i j}\right)}
$$

which implies

$$
\begin{equation*}
\log \left\{\frac{p\left(x_{i}\right)}{1-p\left(x_{i}\right)}\right\}=\beta_{0}+\sum_{j=1}^{p} \beta_{j} x_{i j} \tag{8.20}
\end{equation*}
$$

This indicates that the logit model is equivalent to a log-linear model for the odds ratio $p\left(x_{i}\right) /\left\{1-p\left(x_{i}\right)\right\}$. A positive value of $\beta_{j}$ indicates an explanatory variable $x_{j}$ that will favour the presence of the trait since it improves the odds. A zero value of $\beta_{j}$ corresponds to the absence of an effect of this variable on the appearance of the qualitative trait.

For i.i.d observations the likelihood function is:

$$
L\left(\beta_{0}, \beta\right)=\prod_{i=1}^{n} p\left(x_{i}\right)^{y_{i}}\left\{1-p\left(x_{i}\right)\right\}^{1-y_{i}} .
$$

The maximum likelihood estimators of the $\beta$ 's are obtained as the solution of the non-linear maximisation problem $\left(\hat{\beta}_{0}, \hat{\beta}\right)=\arg \max _{\beta_{0}, \beta} \log L\left(\beta_{0}, \beta\right)$ where

$$
\log L\left(\beta_{0}, \beta\right)=\sum_{i=1}^{n}\left[y_{i} \log p\left(x_{i}\right)+\left(1-y_{i}\right) \log \left\{1-p\left(x_{i}\right)\right\}\right]
$$

The asymptotic theory of the MLE of Chap. 6 (see Theorem 6.3) applies and thus asymptotic inference on $\beta$ is available (test of hypothesis or confidence intervals).

Example 8.5 In the bankruptcy data set (see Sect. 22.22), we have measures on 5 financial characteristics on 66 banks, 33 among them being bankrupt and the other 33 still being solvent. The logit model can be used to evaluate the probability of

Table 8.9 Probabilities of the bankruptcies with the logit model Q MVAbankrupt

|  | $\hat{\beta}$ | $p$-Values |
| :--- | ---: | :---: |
| $\beta_{0}$ | 3.6042 | 0.0660 |
| $\beta_{3}$ | -0.2031 | 0.0037 |
| $\beta_{4}$ | -0.0205 | 0.0183 |

Table 8.10 A $(2 \times K)$
contingency table

|  | 1 | $\cdots$ | $k$ | $\cdots$ | $K$ | Total |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 1 | $y_{11}$ | $\cdots$ | $y_{1 k}$ | $\cdots$ | $y_{1 K}$ | $y_{1}$ |
| 2 | $y_{21}$ | $\cdots$ | $y_{2 k}$ | $\cdots$ | $y_{2 K}$ | $y_{2}$ |
| Total | $y_{\bullet}$ | $\cdots$ | $y_{\bullet} k$ | $\cdots$ | $y_{\bullet K}$ | $y_{\bullet}=n$ |

bankruptcy as a function of these financial ratios. We obtain the results summarised in Table 8.9. We observe that only $\beta_{3}$ and $\beta_{4}$ are significant.

## Logit Models for Contingency Tables

The logit model may contain quantitative and qualitative explanatory variables. In the latter case, the variable may be coded according to the rules described in the ANOVA/ANCOVA sections above. This enables a revisit to the contingency tables where one of the variables is binary and is the variable of interest. How can the probability of taking one of the two nominal values be evaluated as a function of the other variables? We keep the notations of Sect. 8.1 and suppose, without loss of generality, that the first variable with $J=2$ is the binary variable of interest. In the drug Example 8.4, we have a $(2 \times 2 \times 5)$ table and one is interested in the probability of taking a drug as a function of age and gender.
$(2 \times K)$ Tables with Binomial Sampling
In Table 8.10 we have displayed the situation. Let $p_{k}$ be the probability of falling into the first row for the $k$-th column, $k=1, \ldots, K$. Since we are mainly interested in the probabilities $p_{k}$ as a function of $k$, we suppose here that $y_{\bullet k}$ are fixed for $k=$ $1, \ldots, K$ (or we work conditionally on the observed value of these column totals), where $y_{\bullet k}=\sum_{j=1}^{J} y_{j k}$. Therefore, we have $K$ independent binomial processes with parameters $\left(y_{\bullet k}, p_{k}\right)$. Since the column variable is nominal we can use an ANOVA model to analyse the effect of the column variable on $p_{k}$ through the logs of the odds

$$
\begin{equation*}
\log \left(\frac{p_{k}}{1-p_{k}}\right)=\eta_{0}+\eta_{k}, k=1, \ldots, K \tag{8.21}
\end{equation*}
$$

where $\sum_{k=1}^{K} \eta_{k}=0$. As in the ANOVA models, one of the interests will be to test $H_{0}: \eta_{1}=\cdots=\eta_{K}=0$. The log-linear model for the odds has its equivalent in a logit formulation for $p_{k}$

$$
\begin{equation*}
p_{k}=\frac{\exp \left(\eta_{0}+\eta_{k}\right)}{1+\exp \left(\eta_{0}+\eta_{k}\right)}, k=1, \ldots, K \tag{8.22}
\end{equation*}
$$

Note that we can code the RHS of (8.21) as a linear model $\mathcal{X} \theta$, where for instance, for a $(2 \times 4)$ table $(K=4)$ we have:

$$
\mathcal{X}=\left(\begin{array}{rrrr}
1 & 1 & 0 & 0 \\
1 & 0 & 1 & 0 \\
1 & 0 & 0 & 1 \\
1 & -1 & -1 & -1
\end{array}\right), \theta=\left(\begin{array}{c}
\beta_{0} \\
\beta_{1} \\
\beta_{2} \\
\beta_{3}
\end{array}\right)
$$

where $\eta_{0}=\beta_{0}, \eta_{1}=\beta_{1}, \eta_{2}=\beta_{2}, \eta_{3}=\beta_{3}$ and $\eta_{4}=-\left(\beta_{1}+\beta_{2}+\beta_{3}\right)$. The logit model for $p_{k}, k=1, \ldots, K$ can now be written, with some abuse of notation, as the $K$-vector

$$
p=\frac{\exp (\mathcal{X} \theta)}{1+\exp (\mathcal{X} \theta)},
$$

where the division has to be understood as being element-by-element. The MLE of $\theta$ is obtained by maximising the log-likelihood

$$
\begin{equation*}
L(\theta)=\sum_{k=1}^{K}\left\{y_{1 k} \log p_{k}+y_{2 k} \log \left(1-p_{k}\right)\right\} \tag{8.23}
\end{equation*}
$$

where the $p_{k}$ are elements of the $K$-vector $p$.
This logit model is a saturated model. Indeed the number of free parameters is $K$, the dimension of $\theta$, and the number of free cells is also equal to $K$ since we consider the column totals $y_{\bullet k}$ as being fixed. So, there are no degrees of freedom in this model. It can be proven that this logit model is equivalent to the saturated model for a table $(2 \times K)$ presented in Sect. 8.2.2 where all the interactions are present in the model. The hypothesis of all interactions $(\alpha \gamma)_{j k}$ being equal to zero (independence case) is equivalent to the hypothesis that the $\eta_{k}, k=1, \ldots, K$ are all equal to zero (no column effect on the probabilities $p_{k}$ ).

The main interest of the logit presentation is its flexibility when the variable defining the column categories is a quantitative variable (age group, number of children, etc.). Indeed, when this is the case, the logit model allows to quantify the effect of the column category by using less parameters and a more flexible relationship than a linear relation. Suppose that we could attach a representative value $x_{k}$ to each column category for this class (for instance, it could be the median
value, or the average value of the class category). We can then choose the following logit model for $p_{k}$

$$
\begin{equation*}
p_{k}=\frac{\exp \left(\eta_{0}+\eta_{1} x_{k}\right)}{1+\exp \left(\eta_{0}+\eta_{1} x_{k}\right)}, k=1, \ldots, K \tag{8.24}
\end{equation*}
$$

where we now have only two free parameters for $K$ free cells, so we have $K-2$ degrees of freedom. We could even introduce a quadratic term to allow some curvature effect of $x$ on the odds

$$
p_{k}=\frac{\exp \left(\eta_{0}+\eta_{1} x_{k}+\eta_{2} x_{k}^{2}\right)}{1+\exp \left(\eta_{0}+\eta_{1} x_{k}+\eta_{2} x_{k}^{2}\right)}, k=1, \ldots, K .
$$

In this latter case, we would still have $K-3$ degrees of freedom.
We can follow the same idea for a three-way table when we want to model the behaviour of the first binary variable as a function of the two other variables defining the table. In the drug example, one is interested in analysing the tendency of taking a psychotropic drug as a function of the gender category and of the age. Fix the number of observations in each cell $k \ell$ (i.e. $y_{\bullet k \ell}$ ), so that we have a binomial sampling process with an unknown parameter $p_{k \ell}$ for each cell. As for the two-way case above, we can either use ANOVA-like models for the logarithm of the odds and ANCOVA-like models when one (or both) of the two qualitative variables defining the $K$ and/or $L$ categories is a quantitative variable.

One may study the following ANOVA model for the logarithms of the odds

$$
\log \left(\frac{p_{k \ell}}{1-p_{k \ell}}\right)=\mu+\eta_{k}+\zeta_{\ell}, k=1, \ldots, K, \ell=1, \ldots, L
$$

with $\eta=\zeta=0$. As another example, if $x_{\ell}$ is a representative value (like the average age of the group) of the $\ell$ th level of the third categorical variable, one might think of:

$$
\begin{equation*}
\log \left(\frac{p_{k \ell}}{1-p_{k \ell}}\right)=\mu+\eta_{k}+\zeta x_{\ell}, k=1, \ldots, K, \ell=1, \ldots, L \tag{8.25}
\end{equation*}
$$

with the constraint $\eta=0$. Here also, interactions and the curvature effect for $x_{\ell}$ can be introduced, as shown in the following example. Since the cell totals $y_{\bullet k \ell}$ are considered as fixed, the log-likelihood to be maximised is:

$$
\begin{equation*}
\sum_{k=1}^{K} \sum_{\ell=1}^{L}\left\{y_{1 k \ell} \log p_{k \ell}+y_{2 k \ell} \log \left(1-p_{k \ell}\right)\right\} \tag{8.26}
\end{equation*}
$$

where $p_{k \ell}$ follows the appropriate logistic model.

Example 8.6 Consider again Example 8.4. One is interested in the influence of gender and age on drug prescription. Take the number of observations for each "gender-age group" combination, $y_{\bullet k \ell}$ as fixed. A logit model (8.25) can be used for the odds-ratios of the probability of taking drugs, where the value $x_{\ell}$ is the average age of the group. In the linear form it may be written as one of the two following equivalent forms:

$$
\begin{aligned}
& \log \left(\frac{p}{1-p}\right)=\mathcal{X} \theta \\
& p=\frac{\exp (\mathcal{X} \theta)}{1+\exp (\mathcal{X} \theta)}
\end{aligned}
$$

where $\theta=\left(\beta_{0}, \beta_{1}, \beta_{2}\right)^{\top}$ and the design matrix $\mathcal{X}$ is given by

$$
\mathcal{X}=\left(\begin{array}{rrr}
1.0 & 1.0 & 23.2 \\
1.0 & 1.0 & 36.5 \\
1.0 & 1.0 & 54.3 \\
1.0 & 1.0 & 69.2 \\
1.0 & 1.0 & 79.5 \\
1.0 & -1.0 & 23.2 \\
1.0 & -1.0 & 36.5 \\
1.0 & -1.0 & 54.3 \\
1.0 & -1.0 & 69.2 \\
1.0 & -1.0 & 79.5
\end{array}\right)
$$

The first column of $\mathcal{X}$ is for the intercept, the second is the coded variable for the two gender categories and the last column is the average of the ages for the corresponding age-group. Then we estimate $\beta$ by maximising the log-likelihood function (8.26). We obtain:

$$
\begin{aligned}
& \hat{\beta}_{0}=-3.5612 \\
& \hat{\beta}_{1}=-0.3426 \\
& \hat{\beta}_{2}=0.0280
\end{aligned}
$$

the intercept for men is $\hat{\beta}_{0}+\hat{\beta}_{1}=-3.9038$ and for women is $\hat{\beta}_{0}-\hat{\beta}_{1}=-3.2186$, indicating a gender effect and the common slope for the positive age effect being $\hat{\beta}_{2}=0.0280$. The fit appears to be reasonably good. There are $K \times L=2 \times$ $5=10$ free cells in the table. A saturated "full" model with ten parameters and a zero degree of freedom would involve a constant (one parameter) plus an effect for gender (one parameter) plus an effect for age (four parameters) and finally the


Fig. 8.3 Fit of the log of the odds-ratios for taking drugs: linear model for age effect with a "gender" effect (no interaction). Men are the stars and women are the circles $\mathbf{Q}$ MVAdruglogistic
interactions between gender and age (four parameters). The model retained above is a "reduced model" with only three parameters that can be tested against the most general saturated model. We obtain the value of the deviance $G_{H_{0}}^{2}=11.5584$ with 7 degrees of freedom $(7=10-3)$, whereas, $G_{H_{1}}^{2}=0$ with no degree of freedom. This gives a $p$-value $=0.1160$, so we cannot reject the reduced model.

Figure 8.3 shows how well the model fits the data. It displays the fitted values of the log of the odds-ratios by the linear model for the men and the women along with the log of the odds-ratios computed from the observed corresponding frequencies. It seems that the age effect shows a curvature. So we fit a model introducing the square of the ages. This gives the following design matrix:

$$
\mathcal{X}=\left(\begin{array}{rrrr}
1.0 & 1.0 & 23.2 & 538.24 \\
1.0 & 1.0 & 36.5 & 1332.25 \\
1.0 & 1.0 & 54.3 & 2948.49 \\
1.0 & 1.0 & 69.2 & 4788.64 \\
1.0 & 1.0 & 79.5 & 6320.25 \\
1.0 & -1.0 & 23.2 & 538.24 \\
1.0 & -1.0 & 36.5 & 1332.25 \\
1.0 & -1.0 & 54.3 & 2948.49 \\
1.0 & -1.0 & 69.2 & 4788.64 \\
1.0 & -1.0 & 79.5 & 6320.25
\end{array}\right)
$$

The maximum likelihood estimators are:

$$
\begin{aligned}
& \hat{\beta}_{0}=-4.4996 \\
& \hat{\beta}_{1}=-0.3457 \\
& \hat{\beta}_{2}=0.0697 \\
& \hat{\beta}_{3}=-0.0004
\end{aligned}
$$

Q MVAdruglogistic
The fit is better for this more flexible alternative, giving a deviance $G_{H_{1}}^{2}=$ 3.3251 with 6 degrees of freedom $(6=10-4)$. If we test $H_{0}:$ no curvature for the age effect against $H_{1}$ : curvature for the age effect, the reduction of the deviance is $G_{H_{0}}^{2}-G_{H_{1}}^{2}=11.5584-3.3251=8.2333$ with one degree of freedom. The $p$-value $=0.0041$, so we reject the reduced model (no curvature) in favour of the more general model with a curvature term.

We know already from Example 8.4 that second order interactions are not significant for this data set (the influence of age on taking a drug is the same for both gender categories), so we can keep this model as a final reasonable model to analyse the probability of taking the drug as a function of the gender and of the age. To summarise this analysis we end up saying that the probability of taking a psychotropic drug can be modelled as (with some abuse of notation)

$$
\begin{equation*}
\log \left(\frac{p}{1-p}\right)=\beta_{0}+\beta_{1} * \operatorname{Sex}+\beta_{2} * \text { Age }+\beta_{3} * \mathrm{Age}^{2} \tag{8.27}
\end{equation*}
$$



## Summary

$\hookrightarrow$ In contingency tables, the categories are defined by the qualitative variables.
$\hookrightarrow$ The saturated model has all of the interaction terms, and 0 degree of freedom.
$\hookrightarrow$ The non-saturated model is a reduced model since it fixes some parameters to be zero.


### 8.3 Exercises

Exercise 8.1 For the one factor ANOVA model, show that if the model is "balanced" $\left(n_{1}=n_{2}=n_{3}\right)$, we have $\hat{\mu}=\bar{y}$. If the model is not balanced, show that $\bar{y}=\hat{\mu}+n_{1} \hat{\alpha}_{1}+n_{2} \hat{\alpha}_{2}+n_{3} \hat{\alpha}_{3}$.

Exercise 8.2 Redo the calculations of Example 8.2 and test if the main effects of the marketing strategy and of the location are significant.

Exercise 8.3 Redo the calculations of Example 8.3 with the Car Data set.
Exercise 8.4 Calculate the prediction interval for "classic blue" pullover sales (Example 3.2) corresponding to price $=120$.

Exercise 8.5 Redo the calculations of the Boston housing example in Sect. 8.1.3
Exercise 8.6 We want to analyse the variations in the consumption of packs of cigarettes per month as a function of the brand (A or B), of the price per pack and as a function of the gender of the smoker ( $M$ or $F$ ). The data are below.

| $y$ | Price | Gender | Brand |
| :--- | :--- | :--- | :--- |
| 30 | 3.5 | M | A |
| 4 | 4 | F | B |
| 20 | 4.1 | F | B |
| 15 | 3.75 | M | A |
| 24 | 3.25 | F | A |
| 11 | 5 | F | B |
| 8 | 4.1 | F | B |
| 9 | 3.5 | M | A |
| 17 | 4.5 | M | B |
| 1 | 4 | F | B |
| 23 | 3.65 | M | A |
| 13 | 3.5 | M | A |

1. In addition to the effects of brand, price and gender, test if there is an interaction between the brand and the price.
2. How would the design matrix of a full model with all the interactions between the variables appear? What would be the number of degrees of freedom of such a model?
3. We would like to introduce a curvature term for the price variable. How can we proceed? Test if this coefficient is significant.

Exercise 8.7 In the drug Example 8.4, test if the first order interactions are significant.

## Chapter 9 <br> Variable Selection

Variable selection is very important in statistical modelling. We are frequently not only interested in using a model for prediction but also need to correctly identify the relevant variables, that is, to recover the correct model under given assumptions. It is known that under certain conditions, the ordinary least squares (OLS) method produces poor prediction results and does not yield a parsimonious model causing overfitting. Therefore the objective of the variable selection methods is to find the variables which are the most relevant for prediction. Such methods are particularly important when the true underlying model has a sparse representation (many parameters close to zero). The identification of relevant variables will reduce the noise and therefore improve the prediction performance of the fitted model.

Some popular regularisation methods used are the ridge regression, subset selection, $L_{1}$ norm penalisation and their modifications and combinations. Ridge regression, for instance, which minimises a penalised residual sum of squares using the squared $L_{2}$ norm penalty, is employed to improve the OLS estimate through a bias-variance trade-off. However, ridge regression has a drawback that it cannot yield a parsimonious model since it keeps all predictors in the model and therefore creates an interpretability problem. It also gives prediction errors close to those from the OLS model.

Another approach proposed for variable selection is the so-called "least absolute shrinkage and selection operator" (Lasso), aims at combining the features of ridge regression and subset selection either retaining (and shrinking) the coefficients or setting them to zero. This method received several extensions such as the Elastic net, a combination of Lasso and ridge regression or the Group Lasso used when predictors are divided into groups. This chapter describes the application of Lasso, Group Lasso as well as the Elastic net in linear regression model with continuous and binary response (logit model) variables.

### 9.1 Lasso

Tibshirani (1996) first introduced Lasso for generalised linear models, where the response variable $y$ is continuous rather than categorical. Lasso has two important characteristics. First, it has an $L_{1}$-penalty term which performs shrinkage on coefficients in a way similar to ridge regression, where an $L_{2}$ penalty is used.

Second, unlike ridge regression, Lasso performs variable subset selection driving some coefficients to exactly zero due to the nature of the constraint, where the objective function may touch the quadratic constraint area at a corner. For this reason, the Lasso is able to produce sparse solutions and is therefore able to combine good features of both ridge regression and subset selection procedure. It yields interpretable models and has the stability of ridge regression.

### 9.1.1 Lasso in the Linear Regression Model

The linear regression model can be written as follows:

$$
y=\mathcal{X} \beta+\varepsilon
$$

where $y$ is an $(n \times 1)$ vector of observations for the response variable, $\mathcal{X}=$ $\left(x_{1}^{\top}, \ldots, x_{n}^{\top}\right)^{\top}, x_{i} \in \mathbb{R}^{p}, i=1, \ldots, n$ is a data matrix of $p$ explanatory variables, and $\varepsilon=\left(\varepsilon_{1}, \ldots, \varepsilon_{n}\right)^{\top}$ is a vector of errors where $\mathrm{E}\left(\varepsilon_{i}\right)=0$ and $\operatorname{Var}\left(\varepsilon_{i}\right)=\sigma^{2}$, $i=1, \ldots, n$.

In this framework, $\mathrm{E}(y \mid \mathcal{X})=\mathcal{X} \beta$ with $\beta=\left(\beta_{1}, \ldots, \beta_{p}\right)^{\top}$. Further assume that the columns of $\mathcal{X}$ are standardised such that $n^{-1} \sum_{i=1}^{n} x_{i j}=0$ and $n^{-1} \sum_{i=1}^{n} x_{i j}^{2}=$ 1. The Lasso estimate $\hat{\beta}$ can then be defined as follows

$$
\begin{equation*}
\hat{\beta}=\arg \min _{\beta}\left\{\sum_{i=1}^{n}\left(y_{i}-x_{i}^{\top} \beta\right)^{2}\right\}, \text { subject to } \sum_{j=1}^{p}\left|\beta_{j}\right| \leq s, \tag{9.1}
\end{equation*}
$$

where $s \geq 0$ is the tuning parameter which controls the amount of shrinkage. For the OLS estimate $\hat{\beta}^{0}=\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1} \mathcal{X}^{\top} y$ a choice of tuning parameter $s<s_{0}$, where $s_{0}=\sum_{j=1}^{p}\left|\hat{\beta}_{j}^{0}\right|$, will cause shrinkage of the solutions towards 0 , and ultimately some coefficients may be exactly equal to 0 . For values $s \geq s_{0}$ the Lasso coefficients are equal to the unpenalised OLS coefficients.

An alternative representation of (9.1) is:

$$
\begin{equation*}
\hat{\beta}=\arg \min _{\beta}\left\{\sum_{i=1}^{n}\left(y_{i}-x_{i}^{\top} \beta\right)^{2}+\lambda \sum_{j=1}^{p}\left|\beta_{j}\right|\right\} \tag{9.2}
\end{equation*}
$$

with a tuning parameter $\lambda \geq 0$. As $\lambda$ increases, the Lasso estimates are continuously shrunk toward zero. Then if $\lambda$ is quite large, some coefficients are exactly zero. For $\lambda=0$ the Lasso coefficients coincide with the OLS estimate. In fact, if the solution to (9.1) is denoted as $\hat{\beta}_{s}$ and the solution to (9.2) as $\hat{\beta}_{\lambda}$, then $\forall \lambda>0$ and the resulting solution $\hat{\beta}_{\lambda} \exists s_{\lambda}$ such that $\hat{\beta}_{\lambda}=\hat{\beta}_{s_{\lambda}}$ and vice versa which implies a one-toone correspondence between these parameters. However, this does not hold if it is required that $\lambda \geq 0$ only and not $\lambda>0$, because if, for instance, $\lambda=0$, then $\hat{\beta}_{\lambda}$ is the same for any $s \geq\|\hat{\beta}\|_{1}$ and the correspondence is no longer one-to-one.

## Geometrical Aspects in $\mathbb{R}^{\mathbf{2}}$

The Lasso estimate under the least squares loss function solves a quadratic programming problem with linear inequality constraints. The criterion $\sum_{i=1}^{n}\left(y_{i}-x_{i}^{\top} \beta\right)^{2}$ yields the quadratic form objective function

$$
\begin{equation*}
\left(\beta-\hat{\beta}^{0}\right)^{\top} \mathcal{W}\left(\beta-\hat{\beta}^{0}\right) \tag{9.3}
\end{equation*}
$$

with $\mathcal{W}=\mathcal{X}^{\top} \mathcal{X}$. For the special case when $p=2, \beta=\left(\beta_{1}, \beta_{2}\right)^{\top}$, the resulting elliptical contour lines are centred around the OLS estimate and the linear constraints are represented by square (shaded area) shown in Fig. 9.1. The Lasso solution is the first place that the contours touch the square, and this sometimes occurs at a corner, corresponding to a zero coefficient. The nature of the Lasso shrinkage may not occur completely obvious. In the work by Efron, Hastie, Johnstone, and Tibshirani (2004) the Least Angle Regression (LAR) algorithm

Fig. 9.1 Lasso in the general design case for $s=4$ and OLS estimate $\hat{\beta}^{0}=(6,7)^{\top}$ Q MVAlassocontour

with a Lasso modification was described which computes the whole path of Lasso solutions and gives a better understanding of the shrinkage nature.

## The LAR Algorithm and Lasso Solution Paths

The LAR algorithm may be introduced in the simple three-dimensional case as follows (assume that the number of covariates $p=3$ ):

- first, standardise all the covariates to have mean 0 and unit length as well as make the response variable have mean zero;
- start with $\hat{\beta}=0$;
- initialise the algorithm with the first two covariates: let $\mathcal{X}=\left(x_{1}, x_{2}\right)$ and calculate the prediction vector $\hat{y}_{0}=\mathcal{X} \hat{\beta}=0$;
- calculate $\bar{y}_{2}$ the projection of $y$ onto $\mathcal{L}\left(x_{1}, x_{2}\right)$, the linear space spanned by $x_{1}$ and $x_{2}$;
- compute the vector of current correlations between the covariates $\mathcal{X}$ and the two-dimensional current residual vector: $C^{\hat{y}_{0}}=\mathcal{X}^{\top}\left(\bar{y}_{2}-\hat{y}_{0}\right)=\left(c_{1}^{\hat{y}_{0}}, c_{2}^{\hat{y}_{0}}\right)^{\top}$. According to Fig. 9.2, the current residual $\bar{y}_{2}-\hat{y}_{0}$ makes a smaller angle with $x_{1}$, than with $x_{2}$, therefore $c_{1}^{\hat{y}_{0}}>c_{2}^{\hat{y}_{0}}$;
- augment $\hat{y}_{0}$ in the direction of $x_{1}$ so that $\hat{y}_{1}=\hat{y}_{0}+\hat{\gamma}_{1} x_{1}$ with $\hat{\gamma}_{1}$ chosen such that $c_{1}^{\hat{y}_{0}}=c_{2}^{\hat{y}_{0}}$ which means that the new current residual $\bar{y}_{2}-\hat{y}_{1}$ makes equal angles (is equiangular) with $x_{1}$ and $x_{2}$;
- suppose that another regressor $x_{3}$ enters the model: calculate a new projection $\bar{y}_{3}$ of $y$ onto $\mathcal{L}\left(x_{1}, x_{2}, x_{3}\right)$;
- recompute the current correlations vector $C^{\hat{y}_{1}}=\left(c_{1}^{\hat{y}_{1}}, c_{2}^{\hat{y}_{1}}, c_{3}^{\hat{y}_{1}}\right)^{\top}$ with $\mathcal{X}=$ $\left(x_{1}, x_{2}, x_{3}\right), \bar{y}_{3}$ and $\hat{y}_{1}$;
- augment $\hat{y}_{1}$ in the equiangular direction so that $\hat{y}_{2}=\hat{y}_{1}+\hat{\gamma}_{2} u_{2}$ with $\hat{\gamma}_{2}$ chosen such that $c_{1}^{\hat{y_{1}}}=c_{2}^{\hat{y}_{1}}=c_{3}^{\hat{y}_{1}}$, then the new current residual $\bar{y}_{3}-\hat{y}_{2}$ goes


Fig. 9.2 Illustration of LARS algorithm
equiangularly between $x_{1}, x_{2}$ and $x_{3}$ (here $u_{2}$ is the unit vector lying along the equiangular direction $\hat{y}_{2}$ );

- the three-dimensional algorithm is terminated with the calculation of the final prediction vector $\hat{y}_{3}=\hat{y}_{2}+\hat{\gamma}_{3} u_{3}$ with $\hat{\gamma}_{3}$ chosen such that $\hat{y}_{3}=\bar{y}_{3}$.

In the case of $p>3$ covariates, $\hat{y}_{3}$ would be smaller than $\bar{y}_{3}$ initiating another change of direction, as illustrated in Fig. 9.2.

In this setup, it is important that the covariate vectors $x_{1}, x_{2}, x_{3}$ are linearly independent. The LAR algorithm "moves" the variable coefficients to their least squares values. So the Lasso adjustment necessary for the sparse solution is that if a nonzero coefficient happens to return to zero, it should be dropped from the current ("active") set of variables and not be considered in further computations. The general LAR algorithm for $p$ predictors can be summarised as follows.

## Least Angle Regression Algorithm

1. The covariates are standardised to have mean 0 and unit length 1 and the response has mean 0 :

$$
\sum_{i=1}^{n} y_{i}=0, \quad \sum_{i=1}^{n} x_{i j}=0, \quad \sum_{i=1}^{n} x_{i j}^{2}=1 ; \quad j=1,2, \ldots, p
$$

The task is to construct the fit $\hat{\beta}=\left(\hat{\beta}_{1}, \ldots, \hat{\beta}_{p}\right)^{\top}$ by iteratively changing the prediction vector $\hat{y}=\sum_{j=1}^{p} x_{j} \hat{\beta}_{j}=\mathcal{X} \hat{\beta}$.
2. Denote $\mathcal{A}$ equal to a subset of the indices $\{1,2, \ldots, p\}$, begin with $\hat{y}_{\mathcal{A}}=$ $\hat{y}_{o}=0$ and calculate the vector of current correlations

$$
\hat{c}=\mathcal{X}^{\top}\left(y-\hat{y}_{\mathcal{A}}\right)
$$

3. Then review the current set $\mathcal{A}=\left\{j:\left|\hat{c}_{j}\right|=\hat{C}\right\}$ as the set of indices corresponding to the covariates with the greatest absolute current correlations, where $\hat{\hat{C}}=\max _{j}\left\{\left|\hat{c}_{j}\right|\right\}$; let $s_{j}=\operatorname{sign}\left\{\hat{c}_{j}\right\}$ for $j \in \mathcal{A}$ and compute the matrix $\mathcal{X}_{\mathcal{A}}=\left(s_{j} x_{j}\right)_{j \in \mathcal{A}}$, the scalar $A_{\mathcal{A}}=\left(1_{\mathcal{A}}^{\top} \mathcal{G}_{\mathcal{A}}^{-1} 1_{\mathcal{A}}\right)^{-\frac{1}{2}}$ with $\mathcal{G}_{\mathcal{A}}=\mathcal{X}_{\mathcal{A}}^{\top} \mathcal{X}_{\mathcal{A}}$ and $1_{\mathcal{A}}^{\top}$ being a vector of ones of length $|\mathcal{A}|$, and the so-called equiangular vector $u_{\mathcal{A}}=\mathcal{X}_{\mathcal{A}} w_{\mathcal{A}}$ with $w_{\mathcal{A}}=A_{\mathcal{A}} \mathcal{G}_{\mathcal{A}}^{-1} 1_{\mathcal{A}}$ which makes equal angles, each less than $90^{\circ}$, with the columns of $\mathcal{X}_{\mathcal{A}}$.
4. Calculate the inner product vector $a \stackrel{\text { def }}{=} \mathcal{X}^{\top} u_{\mathcal{A}}$ and the direction

$$
\hat{\gamma}=\min _{j \in \mathcal{A}^{c}}^{+}\left\{\frac{\hat{C}-\hat{c}_{j}}{A_{\mathcal{A}}-a_{j}}, \frac{\hat{C}+\hat{c}_{j}}{A_{\mathcal{A}}+a_{j}}\right\}
$$

5. Define $\hat{d}$ to be the $m$-vector equaling $s_{j} w_{\mathcal{A}_{j}}$ for $j \in \mathcal{A}$ and zero elsewhere and $\gamma_{j}=-\hat{\beta}_{j} / \hat{d}_{j}$ yielding $\tilde{\gamma}=\min _{\gamma_{j}>0}\left\{\gamma_{j}\right\}$
(a) If $\tilde{\gamma}<\hat{\gamma}$, calculate the next LARS step as

$$
\hat{y}_{\mathcal{A}_{+}}=\hat{y}_{\mathcal{A}}+\tilde{\gamma} u_{\mathcal{A}}
$$

with $\mathcal{A}_{+}=\mathcal{A}-\{\tilde{j}\}$.
(b) Else: calculate the next step as

$$
\hat{y}_{\mathcal{A}_{+}}=\hat{y}_{\mathcal{A}}+\hat{\gamma} u_{\mathcal{A}}
$$

6. Iterate until all $p$ predictors have been entered, some of which are ultimately dropped from the active set $\mathcal{A}$.

This algorithm can be implemented on a grid from 0 to 1 of the standardised coefficients constraint $s$ resulting in the complete paths of the Lasso coefficients and illustrating the nature of Lasso shrinkage.

Once the Lasso solution paths have been obtained, it is important to decide on a rule how to choose the "optimal" solution, or, equally, the regularisation parameter $\lambda$. There are several existing methods to do this and the most popular examples are the $K$-fold cross-validation, generalised cross-validation, Schwartz's (Bayesian) Information Criterion (BIC). All these methods can be viewed as degrees-offreedom adjustments to the residual squared error (RSE) which underestimates the true prediction error

$$
\mathrm{RSE} \stackrel{\text { def }}{=} \sum_{i=1}^{n}\left(y_{i}-\hat{y}_{i}\right)^{2}
$$

Consider the generalised cross-validation statistic:

$$
\begin{equation*}
\operatorname{GCV}(\lambda)=n^{-1} \operatorname{RSE}_{\lambda} /\{1-\operatorname{df}(\lambda) / n\}^{2}, \tag{9.4}
\end{equation*}
$$

where $\operatorname{RSE}_{\lambda}$ is the residual sum of squares for the constrained fit with a particular regularisation parameter $\lambda$. An alternative is the BIC

$$
\begin{equation*}
\mathrm{BIC}=n \log \left(\hat{\sigma}^{2}\right)+\log (n) \cdot \operatorname{df}(\lambda) \tag{9.5}
\end{equation*}
$$

with the estimation of error variance $\hat{\sigma}^{2}=n^{-1} \sum_{i=1}^{n}\left(y_{i}-\hat{y}_{i}\right)^{2}$.

The degrees of freedom of the predicted vector $\hat{y}$ in the Lasso problem with the linear Gaussian model with normally distributed errors having zero expectation and variance $\sigma^{2}$, written $\varepsilon_{i} \sim N\left(0, \sigma^{2}\right)$, can be defined as follows:

$$
\begin{equation*}
\operatorname{df}(\lambda) \stackrel{\text { def }}{=} \sigma^{-2} \sum_{i=1}^{n} \operatorname{Cov}\left(\hat{y}_{i}, y_{i}\right) \tag{9.6}
\end{equation*}
$$

which can actually be used for both linear and non-linear models. This expression for $\operatorname{df}(\lambda)$ can be viewed as a quantitative measure of the prediction error bias dependence on how much each $y_{i}$ affects its fitted value $\hat{y}_{i}$. The estimate $\hat{\beta}$ minimising the GCV statistic can then be chosen. The following example shows how to compute $\operatorname{df}(\lambda)$.

Example 9.1 (Calculation of $\mathrm{df}(\lambda)$ ) As no closed-form solution exists for the Lasso problem, an approximation should be calculated. The constraint $\sum\left|\beta_{j}\right| \leq s$ can be rewritten as $\sum \beta_{j}^{2} /\left|\beta_{j}\right| \leq s$. Using the duality between the constrained and unconstrained problems and one-to-one correspondence between $s$ and $\lambda$, the Lasso solution is computed as the ridge regression estimate

$$
\hat{\beta}=\left(\mathcal{X}^{\top} \mathcal{X}+\lambda B^{-1}\right)^{-1} \mathcal{X}^{\top} y
$$

where $B=\operatorname{diag}\left(\left|\hat{\beta}_{j}\right|\right)$. Then it follows that

$$
\begin{aligned}
\hat{y} & =\mathcal{X} \hat{\beta} \\
& =\mathcal{X}\left(\mathcal{X}^{\top} \mathcal{X}+\lambda B^{-1}\right)^{-1} \mathcal{X}^{\top} y .
\end{aligned}
$$

Then, to calculate $\operatorname{Cov}\left(\hat{y}_{i}, y_{i}\right)$, one could use $\operatorname{Cov}\left(\hat{y}_{i}, y_{i}\right)=\operatorname{Cov}\left(e_{i}^{\top} \hat{y}, e_{i}^{\top} y\right)=$ $e_{i}^{\top} \operatorname{Cov}(\hat{y}, y) e_{i}$, where $e_{i}$ is a vector where the $i^{\prime}$ th entry is 1 and the rest are zero. Furthermore, each entry in the sum of (9.6) can be calculated to be

$$
\begin{align*}
\operatorname{Cov}\left(\hat{y}_{i}, y_{i}\right) & =e_{i}^{\top} \operatorname{Cov}(\hat{y}, y) e_{i}  \tag{9.7}\\
& =e_{i}^{\top} \mathcal{X}\left(\mathcal{X}^{\top} \mathcal{X}+\lambda B^{-1}\right)^{-1} \mathcal{X}^{\top} \operatorname{Cov}(y, y) e_{i}  \tag{9.8}\\
& =\sigma^{2}\left(\mathcal{X}^{\top} e_{i}\right)^{\top}\left(\mathcal{X}^{\top} \mathcal{X}+\lambda B^{-1}\right)^{-1}\left(\mathcal{X}^{\top} e_{i}\right)  \tag{9.9}\\
& =\sigma^{2} x_{i}^{\top}\left(\mathcal{X}^{\top} \mathcal{X}+\lambda B^{-1}\right)^{-1} x_{i} . \tag{9.10}
\end{align*}
$$

Using the fact that (9.10) are scalars for all $i$ as well as the properties of the trace of a matrix and matrix multiplication rules mentioned in Chap. 2, one obtains the final closed-form expression for the effective degrees of freedom in the Lasso problem:

$$
\operatorname{df}(\lambda)=\frac{1}{\sigma^{2}} \sum_{i=1}^{n} \operatorname{tr}\left\{\sigma^{2} x_{i}^{\top}\left(\mathcal{X}^{\top} \mathcal{X}+\lambda B^{-1}\right)^{-1} x_{i}\right\}
$$

$$
\begin{aligned}
& =\sum_{i=1}^{n} \operatorname{tr}\left\{x_{i} x_{i}^{\top}\left(\mathcal{X}^{\top} \mathcal{X}+\lambda B^{-1}\right)^{-1}\right\} \\
& =\operatorname{tr}\left\{\left(\sum_{i=1}^{n} x_{i} x_{i}^{\top}\right)\left(\mathcal{X}^{\top} \mathcal{X}+\lambda B^{-1}\right)^{-1}\right\} \\
& =\operatorname{tr}\left\{\mathcal{X}^{\top} \mathcal{X}\left(\mathcal{X}^{\top} \mathcal{X}+\lambda B^{-1}\right)^{-1}\right\} \\
& =\operatorname{tr}\left\{\mathcal{X}\left(\mathcal{X}^{\top} \mathcal{X}+\lambda B^{-1}\right)^{-1} \mathcal{X}^{\top}\right\} .
\end{aligned}
$$

It should be noted that the formula for the effective degrees of freedom derived above is valid in the case of the underlying model with non-random regressors. When the random design is used and the set of nonzero predictors is not fixed, another estimator should be used.

## Orthonormal Design Case

A computationally convenient special case is the so-called orthonormal design framework. In the orthonormal design case $\mathcal{X}^{\top} \mathcal{X}$ is a diagonal matrix that $\mathcal{X}^{\top} \mathcal{X}=$ $\mathcal{I}$. Here the explicit Lasso estimate is

$$
\begin{align*}
\hat{\beta}_{j} & =\operatorname{sign}\left(\hat{\beta}_{j}^{0}\right)\left(\left|\hat{\beta}_{j}^{0}\right|-\gamma\right)^{+}  \tag{9.11}\\
\gamma & =\frac{\lambda}{2} \text { subject to } \sum_{j=1}^{p}\left|\hat{\beta}_{j}\right|=s . \tag{9.12}
\end{align*}
$$

The formula shows what was already mentioned in the beginning, namely that the Lasso estimate is a compromise between subset selection and ridge regression, the estimate is either shrunk by $\gamma$ or is set to zero. As a consequence Lasso coefficients can take values between zero and $\hat{\beta}_{j}^{0}$.
Example 9.2 (Orthonormal Design Case for $p=2$ ) Let $\hat{\beta}=\left(\widehat{\beta_{1}}, \widehat{\beta_{2}}\right)^{\top}$ w.l.o.g. be in the first quadrant, i.e. $\widehat{\beta_{1}} \geq 0$ and $\widehat{\beta_{2}} \geq 0$. This gives us the first condition. The orthonormal design ensures that the elliptical contour lines describe circles around the OLS estimate. Thus we get a linear function going through the point $\hat{\beta}^{0}$ and being orthogonal (if possible) to the first condition. Equalising both conditions

$$
\begin{align*}
\hat{\beta}_{1}+\hat{\beta}_{2} & =s  \tag{9.13}\\
\hat{\beta}_{2} & =\hat{\beta}_{1}+\left(\hat{\beta}_{2}^{0}-\hat{\beta}_{1}^{0}\right) \tag{9.14}
\end{align*}
$$

the Lasso estimate can now be accurately determined:

$$
\begin{align*}
& \hat{\beta}_{1}=\left(\frac{s}{2}+\frac{\hat{\beta}_{1}^{0}-\hat{\beta}_{2}^{0}}{2}\right)^{+}  \tag{9.15}\\
& \hat{\beta}_{2}=\left(\frac{s}{2}-\frac{\hat{\beta}_{1}^{0}-\hat{\beta}_{2}^{0}}{2}\right)^{+} . \tag{9.16}
\end{align*}
$$

For cases in which $\left(\frac{s}{2}+\frac{\hat{\beta}_{1}^{0}-\hat{\beta}_{2}^{0}}{2}\right) \leq 0$ or $\left(\frac{s}{2}-\frac{\hat{\beta}_{1}^{0}-\hat{\beta}_{2}^{0}}{2}\right) \leq 0$ the corresponding Lasso estimates will always be zero as the position of the $\hat{\beta}_{1}^{0}$ and corresponding contour lines do not make it possible to get the orthogonality condition mentioned above. Let $\hat{\beta}^{0}=(6,7)^{\top}$ and tuning parameter $s=4$. In this case the Lasso estimator is given by, as shown in Fig. 9.3:

$$
\begin{align*}
& \hat{\beta}_{1}=\frac{4}{2}+\frac{6-7}{2}=1.5,  \tag{9.17}\\
& \hat{\beta}_{2}=\frac{4}{2}-\frac{6-7}{2}=2.5 . \tag{9.18}
\end{align*}
$$



Fig. 9.3 Lasso in the orthonormal design case for $s=4$ and OLS estimate $\hat{\beta}^{0}=(6,7)^{\top} \mathrm{a}$ MVAlassocontour

In terms of $\lambda$, the Lasso solution (9.11) in the orthonormal design case can be calculated in a usual unconstrained minimisation problem. Note that in this case the least squares solution is given by

$$
\hat{\beta}^{0}=\left(\mathcal{X}^{\top} \mathcal{X}\right)^{-1} \mathcal{X}^{\top} y=\mathcal{X}^{\top} y
$$

Then the minimisation problem is written as

$$
\begin{aligned}
\hat{\beta} & =\arg \min _{\beta \in \mathbb{R}^{p}}\|y-\mathcal{X} \beta\|_{2}^{2}+\lambda\|\beta\|_{1} \\
& =\arg \min _{\beta \in \mathbb{R}^{p}}(y-\mathcal{X} \beta)^{\top}(y-\mathcal{X} \beta)+\lambda \sum_{j=1}^{p}\left|\beta_{j}\right| \\
& =\arg \min _{\beta \in \mathbb{R}^{p}}-2 y^{\top} \mathcal{X} \beta+\beta^{\top} \beta+\lambda \sum_{j=1}^{p}\left|\beta_{j}\right| \\
& =\arg \min _{\beta \in \mathbb{R}^{p}}-2 \hat{\beta}^{0 \top} \beta+\beta^{\top} \beta+\lambda \sum_{j=1}^{p}\left|\beta_{j}\right| \\
& =\arg \min _{\beta \in \mathbb{R}^{p}} \sum_{j=1}^{p}\left(-2 \hat{\beta}_{j}^{0} \beta_{j}+\beta_{j}^{2}+\lambda\left|\beta_{j}\right|\right) .
\end{aligned}
$$

The objective function can now be minimised by separate minimisation of its $j$ th element. To solve

$$
\begin{equation*}
\min _{\beta}\left(-2 \hat{\beta}^{0} \beta+\beta^{2}-\lambda|\beta|\right), \tag{9.19}
\end{equation*}
$$

where the index $j$ was dropped for simplicity, let's first assume that $\hat{\beta}^{0}>0$, then $\beta \geq 0$, because a lower value for the objective function may be obtained by changing the sign. Then the solution for the modified problem

$$
\begin{equation*}
\min _{\beta}\left(-2 \hat{\beta}^{0} \beta+\beta^{2}+\lambda \beta\right) \tag{9.20}
\end{equation*}
$$

is, obviously, $\hat{\beta}=\hat{\beta}^{0}-\gamma$, where $\gamma=\lambda / 2$, as in (9.11). To ensure the sign consistency for this case, one could see that the solution is

$$
\begin{equation*}
\hat{\beta}=\left(\hat{\beta}^{0}-\gamma\right)^{+}=\operatorname{sign}\left(\hat{\beta}^{0}\right)\left(\left|\hat{\beta}^{0}\right|-\gamma\right)^{+} \tag{9.21}
\end{equation*}
$$

Now let us take $\hat{\beta}^{0} \leq 0$, then $\beta \leq 0$ as well and the solution for the new problem

$$
\begin{equation*}
\min _{\beta}\left(-2 \hat{\beta}^{0} \beta+\beta^{2}-\lambda \beta\right) \tag{9.22}
\end{equation*}
$$

is $\hat{\beta}=\hat{\beta}^{0}+\gamma$, but the sign consistency requires that

$$
\begin{aligned}
\hat{\beta} & =\left(\hat{\beta}^{0}+\gamma\right)^{-} \\
& =-\left(-\hat{\beta}^{0}-\gamma\right)^{+} \\
& =\operatorname{sign}\left(\hat{\beta}^{0}\right)\left(\left|\hat{\beta}^{0}\right|-\gamma\right)^{+} .
\end{aligned}
$$

As the solutions are the same in both cases, the expression $\operatorname{sign}\left(\hat{\beta}^{0}\right)\left(\left|\hat{\beta}^{0}\right|-\gamma\right)^{+}$is indeed the solution to the original Lasso problem.

## General Lasso Solution

For a fixed $s \geq 0$ the Lasso estimation problem is a least squares problem subjected to $2^{p}$ linear inequality constraints as there are $2^{p}$ different possible signs for $\beta=\left(\beta_{1}, \ldots, \beta_{p}\right)^{\top}$. Lawson and Hansen (1974) suggested solving the least squares problem subject to a general linear inequality constraint $G \beta \leq h$ where $G(m \times p)$ corresponds to the $m=2^{p}$ constraints and $h=s 1_{m}$. As $m$ could be very large, this procedure is not very fast computationally. Therefore Lawson and Hansen (1974) introduced the inequality constraints sequentially in their algorithm, seeking a feasible solution.

Let $g(\beta)=\sum_{i=1}^{n}\left(y_{i}-x_{i}^{\top} \beta\right)^{2}$ and let $\delta_{k}, k=1, \ldots, 2^{p}$, be column vectors of $p$-tuples of the form $( \pm 1, \ldots, \pm 1)$. It follows that the linear inequality condition can be equivalently described as $\delta_{k}^{\top} \beta \leq s, k=1, \ldots, 2^{p}$. Now let $E=\left\{k \mid \delta_{k}^{\top} \beta=\right.$ $s\}$ the equality set, $m_{E}$ the number of elements of $E$ and $G_{E}=\left(\delta_{k}^{\top}\right)_{k \in E}$ a matrix whose rows are all $\delta_{k}$ 's for $k \in E$. Now the algorithm works as follows, see Tibshirani (1996):

1. Find OLS estimate $\hat{\beta}^{0}$ and let $\delta_{k_{0}}=\operatorname{sign}\left(\hat{\beta}^{0}\right), E=\left\{k_{0}\right\}$.
2. Find $\hat{\beta}$ to minimise $g(\beta)$ subject to $G_{E} \beta \leq s 1_{m_{E}}$.
3. If $\sum_{j=1}^{p}\left|\hat{\beta}_{j}\right| \leq s$ the computation is complete.
4. If $\sum_{j=1}^{p}\left|\hat{\beta}_{j}\right|>s$ add $k$ to the set $E$ where $\delta_{k}=\operatorname{sign}(\hat{\beta})$ and go back to step 2 .
5. The final iteration is a solution to the original problem.

As the number of steps is limited by $m=2^{p}$, the algorithm has to converge in finite time. The average number of iterations in practice is between $0.5 p$ and 0.75 p.

Example 9.3 Let us consider the car data set (Table 22.3) where $n=74$. We want to study in-what way the price $\left(X_{1}\right)$ depends on the 12 other variables $\left(X_{2}\right), \ldots,\left(X_{13}\right)$, which are represented by $j=1,2, \ldots, 12$, using Lasso regression. In Fig. 9.4 one can clearly see that coefficients become nonzero one at a time, that means the variables enter the regression equation sequentially as the scaled shrinkage parameter $\hat{s}=s /\left\|\hat{\beta}^{0}\right\|_{1}$ increases, in order $j=6,11,9,3, \ldots$


Fig. 9.4 Lasso estimates of standardised regression $\hat{\beta}_{j}$ for car data with $n=74$ and $p=12$ Q MVAlassoregress
(representing $X_{7}, X_{12}, X_{10}, X_{4}, \ldots$ ), hence the $L_{1}$ penalty results in variable selection and the variables which are most relevant are shrunk less. In this example, an optimal $\hat{s}$ can be found such that the fitted model gives the smallest residual (see Exercise 9.3).

### 9.1.2 Lasso in High Dimensions

The problem with the algorithm by Tibshirani to calculate the Lasso solutions is that it is initialised from an OLS solution of the unconstrained problem which does not correspond to the true model. Another problem is that for the case of $p>n$, this computation is infeasible. Therefore it may be optimal to start with a small initial guess for $\beta$ and iterate through a different kind of an algorithm to obtain the Lasso solutions. Such an algorithm is based on the properties of the Lasso problem as a convex programming one. Osborne et al. (2000) showed that the original Lasso estimate problem (9.1) can be rewritten as:

$$
\begin{equation*}
\hat{\beta}=\arg \min _{\beta \in \mathbb{R}^{p}} \frac{1}{2}(y-\mathcal{X} \beta)^{\top}(y-\mathcal{X} \beta) \stackrel{\text { def }}{=} \frac{1}{2} r^{\top} r, \quad \text { subject to } \quad s-\|\beta\|_{1} \geq 0, \tag{9.23}
\end{equation*}
$$

where $r \stackrel{\text { def }}{=}(y-\mathcal{X} \beta)$. Let $\mathcal{J}=\left\{i_{1}, \ldots, i_{p}\right\}$ be the set of indices such that $\left|\left(\mathcal{X}^{\top} r\right)_{i_{j}}\right|=\left\|\mathcal{X}^{\top} r\right\|_{\infty}$, for $j=1, \ldots, p$; so indices in $\mathcal{J}$ correspond to nonzero elements of $\beta$. Also let $P$ be the permutation matrix that permutes the elements of the coefficient vector $\beta$ so that the first elements are the nonzero elements:
$\beta=P^{\top}\left(\beta_{\mathcal{J}}, 0\right)^{\top}$. Denote $\theta_{\mathcal{J}}=\operatorname{sign}\left(\beta_{\mathcal{J}}\right)$ be equal to 1 if the corresponding element of $\beta_{\mathcal{J}}$ is positive and -1 otherwise. Further denoting $f(\beta)=$ $(y-\mathcal{X} \beta)^{\top}(y-\mathcal{X} \beta)$ the following optimisation algorithm is based on the local linearisation of (9.1) around $\beta$ :
$\hat{\beta}=\arg \min _{h} f(\beta+h), \quad$ subject to $\theta_{\mathcal{J}}^{\top}\left(\beta_{\mathcal{J}}+h_{\mathcal{J}}\right) \leq s$ and $h=P^{\top}\left(h_{\mathcal{J}}, 0\right)^{\top}$,
the solution for which can be shown to be equal to

$$
h_{\mathcal{J}}=\left(\mathcal{X}_{\mathcal{J}}^{\top} \mathcal{X}_{\mathcal{J}}\right)^{-1}\left\{\mathcal{X}_{\mathcal{J}}^{\top}\left(y-\mathcal{X}_{\mathcal{J}} \beta_{\mathcal{J}}\right)-\mu \theta_{\mathcal{J}}\right\}
$$

where

$$
\mu=\max \left\{0, \frac{\theta_{\mathcal{J}}^{\top}\left(\mathcal{X}_{\mathcal{J}}^{\top} \mathcal{X}_{\mathcal{J}}\right)^{-1} \mathcal{X}_{\mathcal{J}}^{\top} y-s}{\theta_{\mathcal{J}}^{\top}\left(\mathcal{X}_{\mathcal{J}}^{\top} \mathcal{X}_{\mathcal{J}}\right)^{-1} \theta_{\mathcal{J}}}\right\} .
$$

The procedure as a whole is implemented as shown in the "Lasso solution-path optimisation" algorithm. As shown in the algorithm, indices may enter and leave the set $\mathcal{J}$, which makes the Lasso problem similar to other subset selection techniques. Moreover, one can compute the whole path of Lasso solutions for $0 \leq s \leq s_{0}$, each time taking the solution for the previous $s$ as a starting point for the next one.

### 9.1.3 Lasso in Logit Model

The Lasso model can be extended to generalised linear models, one of the most common of which is the logistic regression (logit) model. Coefficients in the logit model have probabilistic interpretation. In the logit model, the linear predictor $\mathcal{X} \beta$ is related to the conditional mean $\mu$ of the response variable $y$ via the logit link $\log \{\mu /(1-\mu)\}$. As the response variable is binary, it is binomial-distributed and $\mu=p\left(x_{i}\right)$. Therefore, as defined in (9.25), the logit model for $y \in\{0,1\}$ of $(n \times 1)$ observations on a binary response variable and $x_{i}=\left(x_{i 1}, \ldots, x_{i p}\right)^{\top}$ is,

$$
\log \left\{\frac{p\left(x_{i}\right)}{1-p\left(x_{i}\right)}\right\}=\sum_{j=1}^{p} \beta_{j} x_{i j},
$$

where

$$
\begin{equation*}
p\left(x_{i}\right)=\mathrm{P}\left(y_{i}=1 \mid x_{i}\right)=\frac{\exp \left(\sum_{j=1}^{p} \beta_{j} x_{i j}\right)}{1+\exp \left(\sum_{j=1}^{p} \beta_{j} x_{i j}\right)} . \tag{9.25}
\end{equation*}
$$

```
Algorithm Lasso solution-path optimisation
    procedure FIND(optimal \(\beta\) )
        Choose initial \(\beta\) and \(\mathcal{J}\) (e.g. \(\beta \leftarrow 0, \mathcal{J} \leftarrow \emptyset\) )
        repeat
            Solve (9.23) to obtain \(h\)
            Set \(\hat{\beta} \leftarrow \beta+h\)
            if \(\operatorname{sign}\left(\hat{\beta}_{\mathcal{J}}\right)=\theta_{\mathcal{J}}\) then
                    Obtain the solution \(\beta=\hat{\beta}\)
            else
                    repeat
                        Find the smallest \(\gamma, 0<\gamma<1, k \in \mathcal{J}\) such that \(0=\beta_{k}+\gamma h_{k}\)
                    Set \(\beta=\beta+\gamma h\)
                                Set \(\theta_{k}=-\theta_{k}\)
                Solve (9.23) again to obtain a new \(h\)
                if \(\theta_{\mathcal{J}}^{\top}\left(\beta_{\mathcal{J}}+h_{\mathcal{J}}\right) \leq s\) then
                        \(\hat{\beta}=\beta+h\)
                else
                        Update \(\mathcal{J} \leftarrow \mathcal{J}_{-k}\)
                        Recompute \(\beta_{\mathcal{J}}, \theta_{\mathcal{J}}, h\)
                end if
                    until \(\operatorname{sign}\left(\hat{\beta}_{\mathcal{J}}\right)=\theta_{\mathcal{J}}\)
                end if
                Compute \(\hat{v} \leftarrow \mathcal{X}^{\top} \hat{r} /\left\|\mathcal{X}_{\mathcal{J}}^{\top} \hat{r}\right\|_{\infty}=P^{\top}\left(\hat{v}_{1}, \hat{v}_{2}\right)^{\top} \quad \triangleright\) here \(\hat{r}=y-\mathcal{X} \hat{\beta}\)
                if \(-1 \leq\left(\hat{v}_{2}\right)_{t} \leq 1\) for \(1 \leq l \leq p-|\mathcal{J}|\) then
                    \(\hat{\beta}\) is a solution
                else
                    Find \(J\) such that \(\left|\left(\hat{v}_{2}\right)_{J}\right|\) is maximised
                    Update \(\mathcal{J} \leftarrow(\mathcal{J}, J)\)
                    Update \(\hat{\beta}_{\mathcal{J}} \leftarrow\left(\hat{\beta}_{\mathcal{J}}, 0\right)^{\top}\)
                    Update \(\theta_{\mathcal{J}} \leftarrow\left(\theta_{\mathcal{J}}, \operatorname{sign}\left(\hat{v}_{2}\right)_{J}\right)^{\top}\)
                end if
                Set \(\beta \leftarrow \hat{\beta}\)
        until \(-1 \leq\left(\hat{v}_{2}\right)_{t} \leq 1\) for \(1 \leq t \leq p-|\mathcal{J}|\)
    end procedure
```

The Lasso estimate for the logit model is obtained by solving the following optimisation problem:

$$
\begin{equation*}
\hat{\beta}=\arg \min _{\beta}\left\{\sum_{i=1}^{n} g\left(-y_{i} x_{i}^{\top} \beta\right)\right\}, \text { subject to } \sum_{j=1}^{p}\left|\beta_{j}\right| \leq s, \tag{9.26}
\end{equation*}
$$

with tuning parameter $s \geq 0$ and log-loss function $g(u)=\log \{1+\exp (u)\}$. An alternative representation of the Lasso estimate $\hat{\beta}$ in the logit model is:

$$
\begin{equation*}
\arg \min _{\beta}\left\{\sum_{i=1}^{n} g\left(-y_{i} x_{i}^{\top} \beta\right)+\lambda \sum_{j=1}^{p}\left|\beta_{j}\right|\right\} . \tag{9.27}
\end{equation*}
$$

Shevade and Keerthi (2003) developed a simple and efficient algorithm to solve the optimisation in (9.27) based on the Gauss-Seidel method using coordinatewise descent approach. The algorithm is asymptotically convergent and easy to implement. Firstly, define the following terms,

$$
\begin{align*}
u_{i} & =-y_{i} x_{i}^{\top} \beta \\
F_{j} & =\sum_{i=1}^{n} \frac{\exp \left(u_{i}\right)}{\exp \left(1+u_{i}\right)} y_{i} x_{i j} . \tag{9.28}
\end{align*}
$$

The first order optimality conditions for (9.27) are:

$$
\begin{aligned}
F_{j}=0 \quad & \text { if } \quad j=0, \\
F_{j}=\lambda \quad & \text { if } \quad \beta_{j}>0, j>0, \\
F_{j}=-\lambda \quad & \text { if } \quad \beta_{j}<0, j>0, \\
-\lambda \leq F_{j} \leq \lambda \quad & \text { if } \quad \beta_{j}=0, j>0 .
\end{aligned}
$$

A new variable is defined

$$
\begin{aligned}
& v_{j}=\left|F_{j}\right| \quad \text { if } \quad j=0, \\
& =\left|\lambda-F_{j}\right| \quad \text { if } \quad \beta_{j}>0, j>0, \\
& =\left|\lambda+F_{j}\right| \quad \text { if } \quad \beta_{j}<0, j>0, \\
& =\psi_{j} \quad \text { if } \quad \beta_{j}=0, j>0 .
\end{aligned}
$$

where $\psi_{j}=\max \left\{\left(F_{j}-\lambda\right),\left(-\lambda-F_{j}\right), 0\right\}$. Thus, the first-order optimality conditions can be written as

$$
\begin{equation*}
v_{j}=0 \quad \forall j . \tag{9.29}
\end{equation*}
$$

It is difficult to obtain exact optimality condition, so the stopping criterion for (9.27) is defined as follows (for some small $\varepsilon$ ),

$$
\begin{equation*}
v_{j} \leq \varepsilon \quad \forall j \tag{9.30}
\end{equation*}
$$

To write the algorithm, let us define $I_{z}=\left\{j: \beta_{j}=0, j>0\right\}$ and $I_{n z}=$ $\left\{j: \beta_{j} \neq 0, j>0\right\}$ for sets of zero estimates and sets of nonzero estimates, respectively, and $I=I_{z} \cup I_{n z}$. The algorithm consists of two loops. The first loop runs over the variables in $I_{z}$ to choose the maximum violator, $v$. In the second loop $W$ is optimised with respect to $\beta_{v}$, therefore the set $I_{n z}$ is modified and maximum violator in $I_{n z}$ is obtained. The second loop is repeated until no violators are found in $I_{n z}$. The algorithm alternates between the first and second loop until no violators exist in both $I_{z}$ and $I_{n z}$.

```
Algorithm Lasso in logit model
    procedure FIND(optimal Lasso estimate \(\hat{\beta}\) )
        Set \(\beta_{j}=0\) for all \(j\)
        while an optimality violator exists in \(I_{z}\) do
            Find the maximum violator \((v)\) in \(I_{z}\)
            repeat
                Optimise \(W\) with respect to \(\beta_{v}\)
            Find the maximum violator (v) in \(I_{n z}\)
            until no violator exists in \(I_{n z}\)
        end while
    end procedure
```

Another way to obtain the lasso estimate in the logit model is by maximising the likelihood function of logit model with lasso constraint. The log-likelihood function of logit model is written as

$$
\begin{equation*}
\log L(\beta)=\sum_{i=1}^{n}\left[y_{i} \log p\left(x_{i}\right)+\left(1-y_{i}\right) \log \left\{1-p\left(x_{i}\right)\right\}\right] . \tag{9.31}
\end{equation*}
$$

Suppose $\ell(\beta)=\log L(\beta)$, with $\beta=\left(\beta_{1}, \ldots, \beta_{p}\right)^{\top}$, the Lasso estimates are obtained by maximising the penalised $\log$ likelihood for logit model as follows

$$
\begin{equation*}
\hat{\beta}=\arg \max _{\beta}\left\{n^{-1} \sum_{i=1}^{n} \ell(\beta)\right\} \text {, subject to } \sum_{j=1}^{p}\left|\beta_{j}\right| \leq s \tag{9.32}
\end{equation*}
$$

It can solved by a general non-linear programming procedure or by using iteratively reweighted least squares (IRLS). Friedman, Hastie, and Tibshirani (2010) developed an algorithm to solve the problem in (9.32). An alternative representation of the Lasso problem is defined as follows:

$$
\begin{equation*}
\hat{\beta}=\arg \max _{\beta}\left\{n^{-1} \sum_{i=1}^{n} \ell(\beta)-\lambda \sum_{j=1}^{p}\left|\beta_{j}\right|\right\} \tag{9.33}
\end{equation*}
$$

Example 9.4 Following Example 9.3, the price ( $X_{1}$ ) of car data set (Table 22.3) has average $6,192.28$. We now define a new categorical variable which takes the value 0 if $X_{1} \leq 6,000$ and otherwise is equal to 1 . We want to study in what way the price ( $X_{1}$ ) depends on the 12 other variables $\left(X_{2}, \ldots, X_{13}\right.$ ) using Lasso in logit model.

In Fig. 9.5 one can see that coefficients' dynamics depends on the shrinkage parameter $s=\|\hat{\beta}(\lambda)\|_{1}$, the $L_{1}$ norm of estimated coefficients. An optimal $s$ can be chosen such that the fitted model gives the smallest residual (see Exercise 9.4).


Fig. 9.5 Lasso estimates $\hat{\beta}_{j}$ of logit model for car data with $n=74$ and $p=12$ Q MVAlassologit

### 9.2 Elastic Net

Although the Lasso is widely used in variable selection, it has several drawbacks. Zou and Hastie (2005) stated that:

1. if $p>n$, the Lasso selects at most $n$ variables before it saturates;
2. if there is a group of variables which has very high correlation, then the Lasso tends to select only one variable from this group;
3. for usual $n>p$ condition, if there are high correlations between predictors, the prediction performance of the Lasso is dominated by ridge regression, see Tibshirani (1996).

Zou and Hastie (2005) introduced the Elastic net which combines good features of the $L_{1}$-norm and $L_{2}$-norm penalties. The Elastic net is a regularised regression method which overcomes the limitations of the Lasso. This method is very useful when $p \gg n$ or there are many correlated variables. The advantages are: (1) a group of correlated variables can be selected without arbitrary omissions, (2) the number of selected variables is no longer limited by the sample size.

### 9.2.1 Elastic Net in Linear Regression Model

We describe the Elastic net in linear regression model. For simplicity reason we assume that the $x_{i j}$ are standardised such that $\sum_{i=1}^{n} x_{i j}=0$ and $n^{-1} \sum_{i=1}^{n} x_{i j}^{2}=1$. The Elastic net penalty $P_{\alpha}(\beta)$ leads to the following modification of the problem to obtain the estimator $\hat{\beta}$

$$
\begin{equation*}
\arg \min _{\beta}\left\{(2 n)^{-1} \sum_{i=1}^{n}\left(y_{i}-x_{i}^{\top} \beta\right)^{2}+\lambda P_{\alpha}(\beta)\right\} \tag{9.34}
\end{equation*}
$$

where

$$
\begin{align*}
P_{\alpha}(\beta) & =\frac{1}{2}(1-\alpha)\|\beta\|_{2}^{2}+\alpha\|\beta\|_{1} \\
& =\sum_{j=1}^{p}\left\{\frac{1}{2}(1-\alpha) \beta_{j}^{2}+\alpha\left|\beta_{j}\right|\right\} . \tag{9.35}
\end{align*}
$$

The penalty $P_{\alpha}(\beta)$ is a compromise between ridge regression and the Lasso. If $\alpha=0$ then the criterion is the ridge regression and if $\alpha=1$ the method will be the Lasso. Practically, for small $\varepsilon>0$, the Elastic net with $\alpha=1-\varepsilon$ performs like the Lasso, but removes degeneracies and erratic variable selection behaviour caused by extreme correlation. Given a specific $\lambda$, as $\alpha$ increases from 0 to 1 , the sparsity of the Elastic net solutions increases monotonically from 0 to the sparsity of the Lasso solutions.

The Elastic net optimisation problem can be represented as the usual Lasso problem, using modified $\mathcal{X}$ and $y$ vectors, as shown in the following example.

Example 9.5 To turn the Elastic net optimisation problem into the usual Lasso one, one should first augment $y$ with $p$ additional zeros to obtain $\tilde{y}=(y, 0)^{\top}$. Then, augment $\mathcal{X}$ with the multiple of the $p \times p$ identity matrix $\sqrt{\lambda \alpha} \mathcal{I}$ to get $\tilde{\mathcal{X}}=\left(\mathcal{X}^{\top}, \sqrt{\lambda \alpha} I\right)^{\top}$. Next, define $\tilde{\lambda}=\lambda(1-\alpha)$ and solve the original Lasso minimisation problem in terms of the new input $\tilde{y}, \tilde{\mathcal{X}}$ and $\tilde{\lambda}$. This new problem is equivalent to the original Elastic net problem:

$$
\begin{aligned}
\|\tilde{y}-\tilde{\mathcal{X}} \beta\|_{2}^{2}+\tilde{\lambda}\|\beta\|_{1} & =\left\|\left[\begin{array}{c}
y \\
0
\end{array}\right]-\left[\begin{array}{c}
\mathcal{X} \beta \\
\sqrt{\lambda \alpha} \mathcal{I} \beta
\end{array}\right]\right\|_{2}^{2}+\lambda(1-\alpha)\|\beta\|_{1}, \\
& =\|y-\mathcal{X} \beta\|_{2}^{2}-\lambda \alpha\|\beta\|_{2}^{2}+\lambda\|\beta\|_{1}-\lambda \alpha\|\beta\|_{1}, \\
& =\|y-\mathcal{X} \beta\|_{2}^{2}+\lambda\left\{\alpha\|\beta\|_{2}^{2}+(1-\alpha)\|\beta\|_{1}\right\},
\end{aligned}
$$

which is equivalent to the original Elastic net problem.
We follow the idea of Friedman et al. (2010) who used a coordinate descent algorithm to solve the optimisation problem in (9.34). Let us suppose to have
estimates $\tilde{\beta}_{k}$ for $k \neq j$. Then we optimise (9.34) partially with respect to $\beta_{j}$ by computing the gradient at $\beta_{j}=\tilde{\beta}_{j}$, which only exists if $\tilde{\beta}_{j} \neq 0$. Having the softthresholding operator $S(z, \gamma)$ as

$$
\operatorname{sign}(z)(|z|-\gamma)^{+}=\left\{\begin{array}{llc}
z-\gamma & \text { if } \quad z>0 \text { and } \gamma<|z|,  \tag{9.36}\\
z+\gamma & \text { if } \quad z<0 \text { and } \gamma<|z|, \\
0 & \text { if } \quad \gamma \geq|z| .
\end{array}\right.
$$

it can be shown that the coordinate-wise update has the following form

$$
\begin{equation*}
\widetilde{\beta_{j}}=\frac{S\left\{n^{-1} \sum_{i=1}^{n} x_{i j}\left(y_{i}-\tilde{y}_{i}^{(j)}\right), \lambda \alpha\right\}}{1+\lambda(1-\alpha)} \tag{9.37}
\end{equation*}
$$

where $\tilde{y}_{i}^{(j)}=\sum_{k \neq j} x_{i k} \tilde{\beta}_{k}$ is a fitted value which excludes the contribution $x_{i j}$, therefore $y_{i}-\tilde{y}_{i}^{(j)}$ is partial residual for fitting $\beta_{j}$.

The algorithm computes the least square estimate for the partial residual $y_{i}-$ $\tilde{y}_{i}^{(j)}$, then applies the soft-thresholding rule to perform the Lasso contribution to the penalty $P_{\alpha}(\beta)$. Afterwards, a proportional shrinkage is applied to ridge penalty. There are several methods used to update the current estimate $\tilde{\beta}$. We describe the simplest updating method, the so-called naive update.

The partial residual can be rewritten as follows:

$$
\begin{align*}
y_{i}-\tilde{y}_{i}^{(j)} & =y_{i}-\hat{y}_{i}+x_{i j} \widetilde{\beta_{j}} \\
& =r_{i}+x_{i j} \widetilde{\beta_{j}}, \tag{9.38}
\end{align*}
$$

with $\widehat{y_{i}}$ being the current fit and $r_{i}$ the current residual. As $x_{j}$ is standardised, therefore

$$
\begin{equation*}
\frac{1}{n} \sum_{i=1}^{n} x_{i j}\left(y_{i}-\tilde{y}_{i}^{(j)}\right)=\frac{1}{n} \sum_{i=1}^{n} x_{i j} r_{i}+\widetilde{\beta_{j}} \tag{9.39}
\end{equation*}
$$

Note that the first term on the right-hand side of the new partial residual is the gradient of the loss with respect to $\beta_{j}$.

### 9.2.2 Elastic Net in Logit Model

The Elastic net penalty can similarly be applied to the logit model. Recall the loglikelihood function of the logit model in (9.31),

$$
\log L(\beta)=\sum_{i=1}^{n}\left[y_{i} \log p\left(x_{i}\right)+\left(1-y_{i}\right) \log \left\{1-p\left(x_{i}\right)\right\}\right] .
$$

Penalised log-likelihood for the logit model using Elastic net has the following form

$$
\begin{equation*}
\max _{\beta}\left\{n^{-1} \sum_{i=1}^{n} \ell(\beta)-\lambda P_{\alpha}(\beta)\right\}, \tag{9.40}
\end{equation*}
$$

with $\ell(\beta)=\log L(\beta)$. The solution of (9.40) can be found by means of a Newton algorithm. For a fixed $\lambda$ and given a current parameter $\tilde{\beta}$, the quadratic approximation (Taylor expansion) is updated about current estimates $\tilde{\beta}$ as follows:

$$
\begin{equation*}
\ell_{Q}(\beta)=-(2 n)^{-1} \sum_{i=1}^{n} w_{i}\left(z_{i}-x_{i}^{\top} \beta\right)^{2}+C(\tilde{\beta})^{2}, \tag{9.41}
\end{equation*}
$$

where working response and weight, respectively, are:

$$
\begin{aligned}
z_{i} & =x_{i}^{\top} \tilde{\beta}+\frac{y_{i}-\tilde{p}\left(x_{i}\right)}{\tilde{p}\left(x_{i}\right)\left\{1-\tilde{p}\left(x_{i}\right)\right\}}, \\
w_{i} & =\tilde{p}\left(x_{i}\right)\left\{1-\tilde{p}\left(x_{i}\right)\right\}
\end{aligned}
$$

A Newton update is obtained by minimising $\ell_{Q}(\beta)$.
Friedman et al. (2010) proposed similar approach creating an outer loop for each value of $\lambda$, which computes a quadratic approximation in (9.41) about current estimates $\tilde{\beta}$. Afterwards, a coordinate descent algorithm is used to solve the following penalised weighted least squares problem (PWLS)

$$
\begin{equation*}
\min _{\beta}\left\{-\ell_{Q}(\beta)+\lambda P_{\alpha}(\beta)\right\} . \tag{9.42}
\end{equation*}
$$

This inner coordinate descent loop continues until the maximum change in (9.42) is less than a very small threshold.

### 9.3 Group Lasso

The Group Lasso was first introduced by Yuan and Lin (2006) and was motivated by the fact that the predictor variables can occur in several groups and one could want a parsimonious model which uses only a few of these groups. That is, assume that there are $K$ groups and the vector of coefficients is structured as follows

$$
\beta^{G}=\left(\beta_{1}^{\top}, \ldots, \beta_{K}^{\top}\right)^{\top} \in \mathbb{R}^{\sum_{k} p_{k}},
$$

where $p_{k}$ is the coefficient vector dimension of the $k$ th group, $k=1, \ldots, K$. A sparse set of groups is produced, although within each group either all entries of $\beta_{k}$,
$k=1, \ldots, K$, a corresponding element of the whole vector $\beta^{G}$ are zero or all of them are nonzero. The Group Lasso problem can be formulated in general as

$$
\begin{equation*}
\arg \min _{\beta \in \mathbb{R}^{\Sigma_{k}} p_{k}} n^{-1}\left\|y-\sum_{k=1}^{K} \mathcal{X}_{k} \beta_{k}\right\|_{2}^{2}+\lambda \sum_{k=1}^{K} \sqrt{p_{k}}\left\|\beta_{k}\right\|_{2}, \tag{9.43}
\end{equation*}
$$

where $\mathcal{X}_{k}$ is the $k$ th component of the matrix $\mathcal{X}$ with columns corresponding to the predictors in the group $k, \beta_{k}$ is the coefficient vector for that group and $p_{k}$ is the cardinality of the group, i.e. the size of the coefficient vector which serves as a balancing weight in the case of widely differing group sizes. It is obvious that if groups consist of single elements, i.e. $p_{k}=1 \forall k$, then the Group Lasso problem is reduced to the usual Lasso one.

The computation of the Group Lasso solution involves calculating the necessary and sufficient subgradient KKT conditions for $\hat{\beta}^{G}=\left(\hat{\beta}_{1}^{\top}, \ldots, \hat{\beta}_{K}^{\top}\right)^{\top}$ to be a solution for (9.43)

$$
\begin{equation*}
-\mathcal{X}_{k}^{\top}\left(y-\sum_{k=1}^{K} \mathcal{X}_{k} \beta_{k}\right)+\frac{\lambda \beta_{k} \sqrt{p_{k}}}{\left\|\beta_{k}\right\|}=0 \tag{9.44}
\end{equation*}
$$

if $\beta_{k} \neq 0$; otherwise, for $\beta_{k}=0$, it holds that

$$
\begin{equation*}
\left\|\mathcal{X}_{k}^{\top}\left(y-\sum_{l \neq k} \mathcal{X}_{l} \hat{\beta}_{l}\right)\right\| \leq \lambda \sqrt{p_{k}} . \tag{9.45}
\end{equation*}
$$

Expressions (9.44) and (9.45) allow to calculate the solution, the so-called update step which can be used to implement an iterative algorithm to solve the problem (9.43). The solution resulting from the KKT conditions is readily shown to be the following:

$$
\begin{equation*}
\hat{\beta}_{k}=\left\{\left(\lambda \sqrt{p_{k}}\left\|\hat{\beta}_{k}\right\|^{-1}+\mathcal{X}_{k}^{\top} \mathcal{X}_{k}\right)^{-1}\right\}^{+} \mathcal{X}_{k}^{\top} \hat{r}_{k}, \tag{9.46}
\end{equation*}
$$

where the residual $\hat{r}_{k}$ is defined as $\hat{r}_{k} \stackrel{\text { def }}{=} y-\sum_{l \neq k} \mathcal{X}_{l} \hat{\beta}_{l}$. As a special (orthonormal) case, when $\mathcal{X}_{l}^{\top} \mathcal{X}_{l}=\mathcal{I}$, the solution is simplified to the $\hat{\beta}_{k}=\left(\lambda \sqrt{p_{k}}\left\|\hat{\beta}_{k}\right\|^{-1}+\right.$ 1) $\mathcal{X}_{k}^{\top} \hat{r}_{k}$. To obtain a full solution to this problem, Yuan and Lin (2006) suggest using a blockwise coordinate descent algorithm which iteratively applies the estimate (9.46) to $k=1, \ldots, K$.

Meier, van de Geer, and Bühlmann (2008) extended the Group Lasso to the case of logistic regression and demonstrated convergence of several algorithms for the computation of the solution as well as outlined consistency results for the Group Lasso logit estimator. The general setup for that model involves a binary response
variable $y_{i} \in\{0,1\}$ and $K$ groups predictor variable $x_{i}=\left(x_{i 1}^{\top}, \ldots, x_{i k}^{\top}\right)^{\top}$, both $x_{i}$ and $y_{i}$ are i.i.d., $i=1, \ldots, n$. Then the logistic linear regression model may be written as before:

$$
\begin{equation*}
\log \left\{\frac{p\left(x_{i}\right)}{1-p\left(x_{i}\right)}\right\}=\eta\left(x_{i}\right) \stackrel{\text { def }}{=} \beta_{0}+\sum_{k=1}^{K} x_{i k}^{\top} \beta_{k}, \tag{9.47}
\end{equation*}
$$

where the conditional probability $p\left(x_{i}\right)=\mathrm{P}\left(y_{i}=1 \mid x_{i}\right)$. The Group Lasso logit estimator $\hat{\beta}$ then minimises the objective function

$$
\begin{equation*}
\hat{\beta}=\arg \min _{\beta \in \mathbb{R}^{p+1}}\left\{-\ell(\beta)+\lambda \sum_{k=1}^{K} \sqrt{p_{k}}\left\|\beta_{k}\right\|_{2}\right\}, \tag{9.48}
\end{equation*}
$$

where $\ell(\cdot)$ is the log-likelihood function

$$
\ell(\beta)=\sum_{i=1}^{n} y_{i} \eta\left(x_{i}\right)-\log \left[1+\exp \left\{\eta\left(x_{i}\right)\right\}\right] .
$$

The problem is solved through a group-wise minimisation of the penalised objective function by, for example, the block-coordinate descent method.

Example 9.6 The Group Lasso results can be illustrated by an application to the MEMset Donor dataset of human donor splice sites with a sequence length of 7 base pairs. The full dataset (training and test parts) consists of 12.623 true ( $y_{i}=1$ ) and 269.155 false $\left(y_{i}=0\right)$ human donor sites. Each element of data represents a sequence of DNA within a window of the splice site which consists of the last three positions of the exon and first 4 positions of the intron; so the strings of length 7 are made up of 4 characters A, C, T, G and therefore the predictor variables are 7 factors, each having 4 levels. False splice sites are sequences on the DNA which match the consensus sequence at position four and five. Figure 9.6 shows how the Group Lasso does shrinkage on the level of groups built by DNA letters.

As is seen from Example 9.6, the solution to the Group Lasso problem yields a sparse solution only regarding the "between" case, that is, it excludes some of the groups from the model but then all coefficients in the remaining groups are nonzero. To ensure both the sparsity of groups and within each group, Simon, Friedman, Hastie, and Tibshirani (2013) proposed the so-called "sparse Group Lasso" which uses a more general penalty which yields sparsity an both inter- and intragroup level. The sparse Group Lasso estimate solves the problem

$$
\begin{equation*}
\hat{\beta}=\arg \min _{\beta \in \mathbb{R}^{p}}\left\|y-\sum_{k=1}^{K} \mathcal{X}_{k} \beta_{k}\right\|_{2}^{2}+\lambda_{1} \sum_{k=1}^{K}\left\|\beta_{k}\right\|_{2}+\lambda_{2}\|\beta\|_{1}, \tag{9.49}
\end{equation*}
$$

where $\beta=\left(\beta_{1}, \beta_{2}, \ldots, \beta_{K}\right)^{\top}$ is the entire parameter vector.


Fig. 9.6 Lasso estimates of standardised regression $\hat{\beta}_{j}$ for car data with $n=74$ and $p=12$ a MVAgrouplasso

| $\hookrightarrow$ | Lasso gives a sparse solution. Lasso estimate combines best of both <br> ridge regression and subset regression. |
| :--- | :--- |
|  | If there is a group of variables which has very high correlation, then <br> the Lasso tends to select only one variable from the group. |
|  | The LARS algorithm computes the whole path of Lasso solutions <br> and is feasible for the high-dimensional case $p>n$. |
|  | Elastic net combines good features of $L_{1}$-norm and $L_{2}$-norm <br> penalties. |
|  | The Elastic net is very useful when $p>n$ or there are many <br> correlated variables. |
| $\hookrightarrow$ | The Sparse Group Lasso can perform shrinkage both on inter- and <br> intragroup level. |

### 9.4 Exercises

Exercise 9.1 Derive the explicit Lasso estimate in (9.11) for the orthonormal design case.
Exercise 9.2 Compare Lasso orthonormal design case for $p=2$ graphically to ridge regression, i.e. to the problem $\hat{\beta}=\operatorname{argmin}\left\{\sum_{i=1}^{n}\left(y_{i}-x_{i}^{\top} \beta\right)^{2}\right\}$ subject to $\sum_{j=1}^{p} \beta_{j}{ }^{2} \leq s$.
Why does Lasso produce variable selection and ridge regression does not?
Exercise 9.3 Optimise the value of s such that the fitted model in Example 9.3 produces the smallest residual.

Exercise 9.4 Optimise the value of s such that the fitted model in Example 9.4 produces the smallest residual.

## Chapter 10 <br> Decomposition of Data Matrices by Factors

In Chap. 1 basic descriptive techniques were developed which provided tools for "looking" at multivariate data. They were based on adaptations of bivariate or univariate devices used to reduce the dimensions of the observations. In the following three chapters, issues of reducing the dimension of a multivariate data set will be discussed. The perspectives will be different but the tools will be related.

In this chapter, we take a descriptive perspective and show how using a geometrical approach provides a "best" way of reducing the dimension of a data matrix. It is derived with respect to a least-squares criterion. The result will be low dimensional graphical pictures of the data matrix. This involves the decomposition of the data matrix into "factors". These "factors" will be sorted in decreasing order of importance. The approach is very general and is the core idea of many multivariate techniques. We deliberately use the word "factor" here as a tool or transformation for structural interpretation in an exploratory analysis. In practice, the matrix to be decomposed will be some transformation of the original data matrix and as shown in the following chapters, these transformations provide easier interpretations of the obtained graphs in lower dimensional spaces.

Chapter 11 addresses the issue of reducing the dimensionality of a multivariate random variable by using linear combinations (the principal components). The identified principal components are ordered in decreasing order of importance. When applied in practice to a data matrix, the principal components will turn out to be the factors of a transformed data matrix (the data will be centred and eventually standardised).

Factor analysis is discussed in Chap. 12. The same problem of reducing the dimension of a multivariate random variable is addressed but in this case the number of factors is fixed from the start. Each factor is interpreted as a latent characteristic of the individuals revealed by the original variables. The nonuniqueness of the solutions is dealt with by searching for the representation with the easiest interpretation for the analysis.

Summarising, this chapter can be seen as a foundation since it develops a basic tool for reducing the dimension of a multivariate data matrix.

### 10.1 The Geometric Point of View

As a matter of introducing certain ideas, assume that the data matrix $\mathcal{X}(n \times p)$ is composed of $n$ observations (or individuals) of $p$ variables.

There are in fact two ways of looking at $\mathcal{X}$, row by row or column by column:

1. Each row (observation) is a vector $x_{i}^{\top}=\left(x_{i 1}, \ldots, x_{i p}\right) \in \mathbb{R}^{p}$.

From this point of view our data matrix $\mathcal{X}$ is representable as a cloud of $n$ points in $\mathbb{R}^{p}$ as shown in Fig. 10.1.
2. Each column (variable) is a vector $x_{[j]}=\left(x_{1 j}, \ldots, x_{n j}\right)^{\top} \in \mathbb{R}^{n}$.

From this point of view the data matrix $\mathcal{X}$ is a cloud of $p$ points in $\mathbb{R}^{n}$ as shown in Fig. 10.2.

When $n$ and/or $p$ are large (larger than 2 or 3 ), we cannot produce interpretable graphs of these clouds of points. Therefore, the aim of the factorial methods to be developed here is twofold. We shall try to simultaneously approximate the column space $C(\mathcal{X})$ and the row space $C\left(\mathcal{X}^{\top}\right)$ with smaller subspaces. The hope is of course that this can be done without losing too much information about the variation and structure of the point clouds in both spaces. Ideally, this will provide insights into the structure of $\mathcal{X}$ through graphs in $\mathbb{R}, \mathbb{R}^{2}$ or $\mathbb{R}^{3}$. The main focus then is to find the dimension reducing factors.


Fig. 10.1 Cloud of $n$ points in $\mathbb{R}^{p}$


Fig. 10.2 Cloud of $p$ points in $\mathbb{R}^{n}$


## Summary

$\hookrightarrow$ Each row (individual) of $\mathcal{X}$ is a $p$-dimensional vector. From this point of view $\mathcal{X}$ can be considered as a cloud of $n$ points in $\mathbb{R}^{p}$.
$\hookrightarrow$ Each column (variable) of $\mathcal{X}$ is a $n$-dimensional vector. From this point of view $\mathcal{X}$ can be considered as a cloud of $p$ points in $\mathbb{R}^{n}$.

### 10.2 Fitting the $p$-Dimensional Point Cloud

## Subspaces of Dimension 1

In this section $\mathcal{X}$ is represented by a cloud of $n$ points in $\mathbb{R}^{p}$ (considering each row). The question is how to project this point cloud onto a space of lower dimension. To begin consider the simplest problem, namely finding a subspace of dimension 1. The problem boils down to finding a straight line $F_{1}$ through the origin. The direction of this line can be defined by a unit vector $u_{1} \in \mathbb{R}^{p}$. Hence, we are searching for the vector $u_{1}$ which gives the "best" fit of the initial cloud of $n$ points. The situation is depicted in Fig. 10.3.

The representation of the $i$ th individual $x_{i} \in \mathbb{R}^{p}$ on this line is obtained by the projection of the corresponding point onto $u_{1}$, i.e. the projection point $p_{x_{i}}$. We know from (2.42) that the coordinate of $x_{i}$ on $F_{1}$ is given by

$$
\begin{equation*}
p_{x_{i}}=x_{i}^{\top} \frac{u_{1}}{\left\|u_{1}\right\|}=x_{i}^{\top} u_{1} . \tag{10.1}
\end{equation*}
$$



Fig. 10.3 Projection of point cloud onto $u$ space of lower dimension

We define the best line $F_{1}$ in the following "least-squares" sense: Find $u_{1} \in \mathbb{R}^{p}$ which minimises

$$
\begin{equation*}
\sum_{i=1}^{n}\left\|x_{i}-p_{x_{i}}\right\|^{2} \tag{10.2}
\end{equation*}
$$

Since $\left\|x_{i}-p_{x_{i}}\right\|^{2}=\left\|x_{i}\right\|^{2}-\left\|p_{x_{i}}\right\|^{2}$ by Pythagoras's theorem, the problem of minimising (10.2) is equivalent to maximising $\sum_{i=1}^{n}\left\|p_{x_{i}}\right\|^{2}$. Thus the problem is to find $u_{1} \in \mathbb{R}^{p}$ that maximises $\sum_{i=1}^{n}\left\|p_{x_{i}}\right\|^{2}$ under the constraint $\left\|u_{1}\right\|=1$. With (10.1) we can write

$$
\left(\begin{array}{c}
p_{x_{1}} \\
p_{x_{2}} \\
\vdots \\
p_{x_{n}}
\end{array}\right)=\left(\begin{array}{c}
x_{1}^{\top} u_{1} \\
x_{2}^{\top} u_{1} \\
\vdots \\
x_{n}^{\top} u_{1}
\end{array}\right)=\mathcal{X} u_{1}
$$

and the problem can finally be reformulated as: find $u_{1} \in \mathbb{R}^{p}$ with $\left\|u_{1}\right\|=1$ that maximises the quadratic form $\left(\mathcal{X} u_{1}\right)^{\top}\left(\mathcal{X} u_{1}\right)$ or

$$
\begin{equation*}
\max _{u_{1}^{\top} u_{1}=1} u_{1}^{\top}\left(\mathcal{X}^{\top} \mathcal{X}\right) u_{1} \tag{10.3}
\end{equation*}
$$

The solution is given by Theorem 2.5 (using $\mathcal{A}=\mathcal{X}^{\top} \mathcal{X}$ and $\mathcal{B}=\mathcal{I}$ in the theorem).

Theorem 10.1 The vector $u_{1}$ which minimises (10.2) is the eigenvector of $\mathcal{X}^{\top} \mathcal{X}$ associated with the largest eigenvalue $\lambda_{1}$ of $\mathcal{X}^{\top} \mathcal{X}$.

Note that if the data have been centred, i.e. $\bar{x}=0$, then $\mathcal{X}=\mathcal{X}_{c}$, where $\mathcal{X}_{c}$ is the centred data matrix, and $\frac{1}{n} \mathcal{X}^{\top} \mathcal{X}$ is the covariance matrix. Thus Theorem 10.1 says that we are searching for a maximum of the quadratic form (10.3) w.r.t. the covariance matrix $\mathcal{S}_{\mathcal{X}}=n^{-1} \mathcal{X}^{\top} \mathcal{X}$.

## Representation of the Cloud on $\boldsymbol{F}_{1}$

The coordinates of the $n$ individuals on $F_{1}$ are given by $\mathcal{X} u_{1}, \mathcal{X} u_{1}$ is called the first factorial variable or the first factor and $u_{1}$ the first factorial axis. The $n$ individuals, $x_{i}$, are now represented by a new factorial variable $z_{1}=\mathcal{X} u_{1}$. This factorial variable is a linear combination of the original variables $\left(x_{[1]}, \ldots, x_{[p]}\right)$ whose coefficients are given by the vector $u_{1}$, i.e.

$$
\begin{equation*}
z_{1}=u_{11} x_{[1]}+\cdots+u_{p 1} x_{[p]} . \tag{10.4}
\end{equation*}
$$



Fig. 10.4 Representation of the individuals $x_{1}, \ldots, x_{n}$ as a two-dimensional point cloud

## Subspaces of Dimension 2

If we approximate the $n$ individuals by a plane (dimension 2), it can be shown via Theorem 2.5 that this space contains $u_{1}$. The plane is determined by the best linear fit ( $u_{1}$ ) and a unit vector $u_{2}$ orthogonal to $u_{1}$ which maximises the quadratic form $u_{2}^{\top}\left(\mathcal{X}^{\top} \mathcal{X}\right) u_{2}$ under the constraints

$$
\left\|u_{2}\right\|=1, \text { and } u_{1}^{\top} u_{2}=0 .
$$

Theorem 10.2 The second factorial axis, $u_{2}$, is the eigenvector of $\mathcal{X}^{\top} \mathcal{X}$ corresponding to the second largest eigenvalue $\lambda_{2}$ of $\mathcal{X}^{\top} \mathcal{X}$.

The unit vector $u_{2}$ characterises a second line, $F_{2}$, on which the points are projected. The coordinates of the $n$ individuals on $F_{2}$ are given by $z_{2}=\mathcal{X} u_{2}$. The variable $z_{2}$ is called the second factorial variable or the second factor. The representation of the $n$ individuals in two-dimensional space ( $z_{1}=\mathcal{X} u_{1}$ vs. $\left.z_{2}=\mathcal{X} u_{2}\right)$ is shown in Fig. 10.4.

## Subspaces of Dimension $q(q \leq p)$

In the case of $q$ dimensions the task is again to minimise (10.2) but with projection points in a $q$-dimensional subspace. Following the same argument as above, it can be shown via Theorem 2.5 that this best subspace is generated by $u_{1}, u_{2}, \ldots, u_{q}$, the orthonormal eigenvectors of $\mathcal{X}^{\top} \mathcal{X}$ associated with the corresponding eigenvalues
$\lambda_{1} \geq \lambda_{2} \geq \cdots \geq \lambda_{q}$. The coordinates of the $n$ individuals on the $k$ th factorial axis, $u_{k}$, are given by the $k$ th factorial variable $z_{k}=\mathcal{X} u_{k}$ for $k=1, \ldots, q$. Each factorial variable $z_{k}=\left(z_{1 k}, z_{2 k}, \ldots, z_{n k}\right)^{\top}$ is a linear combination of the original variables $x_{[1]}, x_{[2]}, \ldots, x_{[p]}$ whose coefficients are given by the elements of the $k$ th vector $u_{k}: z_{i k}=\sum_{m=1}^{p} x_{i m} u_{m k}$.

| Su Summary |
| :---: |
| $\hookrightarrow$ The $p$-dimensional point cloud of individuals can be graphically represented by projecting each element into spaces of smaller dimensions. |
| $\hookrightarrow$ The first factorial axis is $u_{1}$ and defines a line $F_{1}$ through the origin. This line is found by minimising the orthogonal distances (10.2). The factor $u_{1}$ equals the eigenvector of $\mathcal{X}^{\top} \mathcal{X}$ corresponding to its largest eigenvalue. The coordinates for representing the point cloud on a straight line are given by $z_{1}=\mathcal{X} u_{1}$. |
| $\hookrightarrow$ The second factorial axis is $u_{2}$, where $u_{2}$ denotes the eigenvector of $\mathcal{X}^{\top} \mathcal{X}$ corresponding to its second largest eigenvalue. The coordinates for representing the point cloud on a plane are given by $z_{1}=\mathcal{X} u_{1}$ and $z_{2}=\mathcal{X} u_{2}$. |
| $\hookrightarrow$ The factor directions $1, \ldots, q$ are $u_{1}, \ldots, u_{q}$, which denote the eigenvectors of $\mathcal{X}^{\top} \mathcal{X}$ corresponding to the $q$ largest eigenvalues. The coordinates for representing the point cloud of individuals on a $q$-dimensional subspace are given by $z_{1}=\mathcal{X} u_{1}, \ldots, z_{q}=\mathcal{X} u_{q}$. |

### 10.3 Fitting the $n$-Dimensional Point Cloud

## Subspaces of Dimension 1

Suppose that $\mathcal{X}$ is represented by a cloud of $p$ points (variables) in $\mathbb{R}^{n}$ (considering each column). How can this cloud be projected into a lower dimensional space? We start as before with one dimension. In other words, we have to find a straight line $G_{1}$, which is defined by the unit vector $v_{1} \in \mathbb{R}^{n}$, and which gives the best fit of the initial cloud of $p$ points.

Algebraically, this is the same problem as above (replace $\mathcal{X}$ by $\mathcal{X}^{\top}$ and follow Sect. 10.2): the representation of the $j$ th variable $x_{[j]} \in \mathbb{R}^{n}$ is obtained by the projection of the corresponding point onto the straight line $G_{1}$ or the direction $v_{1}$. Hence we have to find $v_{1}$ such that $\sum_{j=1}^{p}\left\|p_{x_{[j]}}\right\|^{2}$ is maximised, or equivalently, we have to find the unit vector $v_{1}$ which maximises $\left(\mathcal{X}^{\top} v_{1}\right)^{\top}\left(\mathcal{X} v_{1}\right)=v_{1}^{\top}\left(\mathcal{X} \mathcal{X}^{\top}\right) v_{1}$. The solution is given by Theorem 2.5.

Theorem 10.3 $v_{1}$ is the eigenvector of $\mathcal{X} \mathcal{X}^{\top}$ corresponding to the largest eigenvalue $\mu_{1}$ of $\mathcal{X} \mathcal{X}^{\top}$.

## Representation of the Cloud on $\boldsymbol{G}_{1}$

The coordinates of the $p$ variables on $G_{1}$ are given by $w_{1}=\mathcal{X}^{\top} v_{1}$, the first factorial axis. The $p$ variables are now represented by a linear combination of the original individuals $x_{1}, \ldots, x_{n}$, whose coefficients are given by the vector $v_{1}$, i.e. for $j=$ $1, \ldots, p$

$$
\begin{equation*}
w_{1 j}=v_{11} x_{1 j}+\cdots+v_{1 n} x_{n j} . \tag{10.5}
\end{equation*}
$$

## Subspaces of Dimension q( $\boldsymbol{q} \leq \boldsymbol{n}$ )

The representation of the $p$ variables in a subspace of dimension $q$ is done in the same manner as for the $n$ individuals above. The best subspace is generated by the orthonormal eigenvectors $v_{1}, v_{2}, \ldots, v_{q}$ of $\mathcal{X} \mathcal{X}^{\top}$ associated with the eigenvalues $\mu_{1} \geq \mu_{2} \geq \cdots \geq \mu_{q}$. The coordinates of the $p$ variables on the $k$ th factorial axis are given by the factorial variables $w_{k}=\mathcal{X}^{\top} v_{k}, k=1, \ldots, q$. Each factorial variable $w_{k}=\left(w_{k 1}, w_{k 2}, \ldots, w_{k p}\right)^{\top}$ is a linear combination of the original individuals $x_{1}, x_{2}, \ldots, x_{n}$ whose coefficients are given by the elements of the $k$ th vector $v_{k}: w_{k j}=\sum_{m=1}^{n} v_{k m} x_{m j}$. The representation in a subspace of dimension $q=2$ is depicted in Fig. 10.5.


Fig. 10.5 Representation of the variables $x_{[1]}, \ldots, x_{[p]}$ as a two-dimensional point cloud

## 

## Summary

$\hookrightarrow$ The $n$-dimensional point cloud of variables can be graphically represented by projecting each element into spaces of smaller dimensions.
$\hookrightarrow$ The first factor direction is $v_{1}$ and defines a line $G_{1}$ through the origin. The vector $v_{1}$ equals the eigenvector of $\mathcal{X} \mathcal{X}^{\top}$ corresponding to the largest eigenvalue of $\mathcal{X} \mathcal{X}^{\top}$. The coordinates for representing the point cloud on a straight line are $w_{1}=\mathcal{X}^{\top} v_{1}$.
$\hookrightarrow$ The second factor direction is $v_{2}$, where $v_{2}$ denotes the eigenvector of $\mathcal{X} \mathcal{X}^{\top}$ corresponding to its second largest eigenvalue. The coordinates for representing the point cloud on a plane are given by $w_{1}=\mathcal{X}^{\top} v_{1}$ and $w_{2}=\mathcal{X}^{\top} v_{2}$.
$\hookrightarrow$ The factor directions $1, \ldots, q$ are $v_{1}, \ldots, v_{q}$, which denote the eigenvectors of $\mathcal{X} \mathcal{X}^{\top}$ corresponding to the $q$ largest eigenvalues. The coordinates for representing the point cloud of variables on a $q$ dimensional subspace are given by $w_{1}=\mathcal{X}^{\top} v_{1}, \ldots, w_{q}=\mathcal{X}^{\top} v_{q}$.

### 10.4 Relations Between Subspaces

The aim of this section is to present a duality relationship between the two approaches shown in Sects. 10.2 and 10.3. Consider the eigenvector equations in $\mathbb{R}^{n}$

$$
\begin{equation*}
\left(\mathcal{X} \mathcal{X}^{\top}\right) v_{k}=\mu_{k} v_{k} \tag{10.6}
\end{equation*}
$$

for $k \leq r$, where $r=\operatorname{rank}\left(\mathcal{X} \mathcal{X}^{\top}\right)=\operatorname{rank}(\mathcal{X}) \leq \min (p, n)$. Multiplying by $\mathcal{X}^{\top}$, we have

$$
\begin{align*}
\mathcal{X}^{\top}\left(\mathcal{X} \mathcal{X}^{\top}\right) v_{k} & =\mu_{k} \mathcal{X}^{\top} v_{k}  \tag{10.7}\\
\text { or } \quad\left(\mathcal{X}^{\top} \mathcal{X}\right)\left(\mathcal{X}^{\top} v_{k}\right) & =\mu_{k}\left(\mathcal{X}^{\top} v_{k}\right) \tag{10.8}
\end{align*}
$$

so that each eigenvector $v_{k}$ of $\mathcal{X} \mathcal{X}^{\top}$ corresponds to an eigenvector $\left(\mathcal{X}^{\top} v_{k}\right)$ of $\mathcal{X}^{\top} \mathcal{X}$ associated with the same eigenvalue $\mu_{k}$. This means that every nonzero eigenvalue of $\mathcal{X} \mathcal{X}^{\top}$ is an eigenvalue of $\mathcal{X}^{\top} \mathcal{X}$. The corresponding eigenvectors are related by

$$
u_{k}=c_{k} \mathcal{X}^{\top} v_{k}
$$

where $c_{k}$ is some constant.

Now consider the eigenvector equations in $\mathbb{R}^{p}$ :

$$
\begin{equation*}
\left(\mathcal{X}^{\top} \mathcal{X}\right) u_{k}=\lambda_{k} u_{k} \tag{10.9}
\end{equation*}
$$

for $k \leq r$. Multiplying by $\mathcal{X}$, we have

$$
\begin{equation*}
\left(\mathcal{X} \mathcal{X}^{\top}\right)\left(\mathcal{X} u_{k}\right)=\lambda_{k}\left(\mathcal{X} u_{k}\right) \tag{10.10}
\end{equation*}
$$

i.e. each eigenvector $u_{k}$ of $\mathcal{X}^{\top} \mathcal{X}$ corresponds to an eigenvector $\mathcal{X} u_{k}$ of $\mathcal{X} \mathcal{X}^{\top}$ associated with the same eigenvalue $\lambda_{k}$. Therefore, every nonzero eigenvalue of $\left(\mathcal{X}^{\top} \mathcal{X}\right)$ is an eigenvalue of $\mathcal{X} \mathcal{X}^{\top}$. The corresponding eigenvectors are related by

$$
v_{k}=d_{k} \mathcal{X} u_{k},
$$

where $d_{k}$ is some constant. Now, since $u_{k}^{\top} u_{k}=v_{k}^{\top} v_{k}=1$ we have $c_{k}=d_{k}=\frac{1}{\sqrt{\lambda_{k}}}$. This lead to the following result:

Theorem 10.4 (Duality Relations) Let $r$ be the rank of $\mathcal{X}$. For $k \leq r$, the eigenvalues $\lambda_{k}$ of $\mathcal{X}^{\top} \mathcal{X}$ and $\mathcal{X} \mathcal{X}^{\top}$ are the same and the eigenvectors ( $u_{k}$ and $v_{k}$, respectively) are related by

$$
\begin{align*}
u_{k} & =\frac{1}{\sqrt{\lambda_{k}}} \mathcal{X}^{\top} v_{k}  \tag{10.11}\\
v_{k} & =\frac{1}{\sqrt{\lambda_{k}}} \mathcal{X} u_{k} \tag{10.12}
\end{align*}
$$

Note that the projection of the $p$ variables on the factorial axis $v_{k}$ is given by

$$
\begin{equation*}
w_{k}=\mathcal{X}^{\top} v_{k}=\frac{1}{\sqrt{\lambda_{k}}} \mathcal{X}^{\top} \mathcal{X} u_{k}=\sqrt{\lambda_{k}} u_{k} . \tag{10.13}
\end{equation*}
$$

Therefore, the eigenvectors $v_{k}$ do not have to be explicitly recomputed to get $w_{k}$.
Note that $u_{k}$ and $v_{k}$ provide the SVD of $\mathcal{X}$ (see Theorem 2.2). Letting $U=\left[\begin{array}{llll}u_{1} & u_{2} & \ldots & u_{r}\end{array}\right], V=\left[\begin{array}{llll}v_{1} & v_{2} & \ldots & v_{r}\end{array}\right]$ and $\Lambda=\operatorname{diag}\left(\lambda_{1}, \ldots, \lambda_{r}\right)$ we have

$$
\mathcal{X}=V \Lambda^{1 / 2} U^{\top}
$$

so that

$$
\begin{equation*}
x_{i j}=\sum_{k=1}^{r} \lambda_{k}^{1 / 2} v_{i k} u_{j k} . \tag{10.14}
\end{equation*}
$$

In the following section this method is applied in analysing consumption behaviour across different household types.

| Summary |
| :---: |
| $\hookrightarrow$ The matrices $\mathcal{X}^{\top} \mathcal{X}$ and $\mathcal{X} \mathcal{X}^{\top}$ have the same nonzero eigenvalues $\lambda_{1}, \ldots, \lambda_{r}$, where $r=\operatorname{rank}(\mathcal{X})$. |
| $\hookrightarrow$ The eigenvectors of $\mathcal{X}^{\top} \mathcal{X}$ can be calculated from the eigenvectors of $\mathcal{X} \mathcal{X}^{\top}$ and vice versa: $u_{k}=\frac{1}{\sqrt{\lambda_{k}}} \mathcal{X}^{\top} v_{k} \quad \text { and } \quad v_{k}=\frac{1}{\sqrt{\lambda_{k}}} \mathcal{X} u_{k}$ |
| $\hookrightarrow$ The coordinates representing the variables (columns) of $\mathcal{X}$ in a $q$-dimensional subspace can be easily calculated by $w_{k}=\sqrt{\lambda_{k}} u_{k}$. |

### 10.5 Practical Computation

The practical implementation of the techniques introduced begins with the computation of the eigenvalues $\lambda_{1} \geq \lambda_{2} \geq \cdots \geq \lambda_{p}$ and the corresponding eigenvectors $u_{1}, \ldots, u_{p}$ of $\mathcal{X}^{\top} \mathcal{X}$. (Since $p$ is usually less than $n$, this is numerically less involved than computing $v_{k}$ directly for $k=1, \ldots, p$.) The representation of the $n$ individuals on a plane is then obtained by plotting $z_{1}=\mathcal{X} u_{1}$ versus $z_{2}=\mathcal{X} u_{2}$ ( $z_{3}=\mathcal{X} u_{3}$ may eventually be added if a third dimension is helpful). Using the Duality Relation (10.13) representations for the $p$ variables can easily be obtained. These representations can be visualised in a scatterplot of $w_{1}=\sqrt{\lambda_{1}} u_{1}$ against $w_{2}=\sqrt{\lambda_{2}} u_{2}$ (and eventually against $w_{3}=\sqrt{\lambda_{3}} u_{3}$ ). Higher dimensional factorial resolutions can be obtained (by computing $z_{k}$ and $w_{k}$ for $k>3$ ) but, of course, cannot be plotted.

A standard way of evaluating the quality of the factorial representations in a subspace of dimension $q$ is given by the ratio

$$
\begin{equation*}
\tau_{q}=\frac{\lambda_{1}+\lambda_{2}+\cdots+\lambda_{q}}{\lambda_{1}+\lambda_{2}+\cdots+\lambda_{p}}, \tag{10.15}
\end{equation*}
$$

where $0 \leq \tau_{q} \leq 1$. In general, the scalar product $y^{\top} y$ is called the inertia of $y \in \mathbb{R}^{n}$ w.r.t. the origin. Therefore, the ratio $\tau_{q}$ is usually interpreted as the percentage of the inertia explained by the first $q$ factors. Note that $\lambda_{j}=\left(\mathcal{X} u_{j}\right)^{\top}\left(\mathcal{X} u_{j}\right)=z_{j}^{\top} z_{j}$. Thus, $\lambda_{j}$ is the inertia of the $j$ th factorial variable w.r.t. the origin. The denominator in (10.15) is a measure of the total inertia of the $p$ variables, $x_{[j]}$. Indeed, by (2.3)

$$
\sum_{j=1}^{p} \lambda_{j}=\operatorname{tr}\left(\mathcal{X}^{\top} \mathcal{X}\right)=\sum_{j=1}^{p} \sum_{i=1}^{n} x_{i j}^{2}=\sum_{j=1}^{p} x_{[j]}^{\top} x_{[j]} .
$$

Remark 10.1 It is clear that the sum $\sum_{j=1}^{q} \lambda_{j}$ is the sum of the inertia of the first $q$ factorial variables $z_{1}, z_{2}, \ldots, z_{q}$.

Example 10.1 We consider the data set in Table 22.6 which gives the food expenditures of various French families (manual workers $=$ MA, employees $=E M$, managers $=\mathrm{CA}$ ) with varying numbers of children ( $2,3,4$ or 5 children). We are interested in investigating whether certain household types prefer certain food types. We can answer this question using the factorial approximations developed here.

The correlation matrix corresponding to the data is

$$
\mathcal{R}=\left(\begin{array}{rrrrrrr}
1.00 & 0.59 & 0.20 & 0.32 & 0.25 & 0.86 & 0.30 \\
0.59 & 1.00 & 0.86 & 0.88 & 0.83 & 0.66 & -0.36 \\
0.20 & 0.86 & 1.00 & 0.96 & 0.93 & 0.33 & -0.49 \\
0.32 & 0.88 & 0.96 & 1.00 & 0.98 & 0.37 & -0.44 \\
0.25 & 0.83 & 0.93 & 0.98 & 1.00 & 0.23 & -0.40 \\
0.86 & 0.66 & 0.33 & 0.37 & 0.23 & 1.00 & 0.01 \\
0.30 & -0.36 & -0.49 & -0.44 & -0.40 & 0.01 & 1.00
\end{array}\right) .
$$

We observe a rather high correlation (0.98) between meat and poultry, whereas the correlation for expenditure for milk and wine ( 0.01 ) is rather small. Are there household types that prefer, say, meat over bread?

We shall now represent food expenditures and households simultaneously using two factors. First, note that in this particular problem the origin has no specific meaning (it represents a "zero" consumer). So it makes sense to compare the consumption of any family to that of an "average family" rather than to the origin. Therefore, the data is first centred (the origin is translated to the centre of gravity, $\bar{x}$ ). Furthermore, since the dispersions of the seven variables are quite different each variable is standardised so that each has the same weight in the analysis (mean 0 and variance 1). Finally, for convenience, we divide each element in the matrix by $\sqrt{n}=\sqrt{12}$. (This will only change the scaling of the plots in the graphical representation.)

The data matrix to be analysed is

$$
\mathcal{X}_{*}=\frac{1}{\sqrt{n}} \mathcal{H X} \mathcal{D}^{-1 / 2}
$$

where $\mathcal{H}$ is the centering matrix and $\mathcal{D}=\operatorname{diag}\left(s_{X_{i} X_{i}}\right)$ (see Sect.3.3). Note that by standardising by $\sqrt{n}$, it follows that $\mathcal{X}_{*}^{\top} \mathcal{X}_{*}=\mathcal{R}$ where $\mathcal{R}$ is the correlation matrix of the original data. Calculating

$$
\lambda=(4.33,1.83,0.63,0.13,0.06,0.02,0.00)^{\top}
$$

shows that the directions of the first two eigenvectors play a dominant role $\left(\tau_{2}=\right.$ $88 \%$ ), whereas the other directions contribute less than $15 \%$ of inertia. A twodimensional plot should suffice for interpreting this data set.



| bread | -0.49873 | -0.84162 |
| :--- | ---: | ---: |
| veget | -0.96975 | -0.13310 |
| fruit | -0.92913 | 0.27791 |
| meat | -0.96210 | 0.19107 |
| poult | -0.91125 | 0.26590 |
| milk | -0.58434 | -0.70690 |
| wine | 0.42820 | -0.64815 |


| ma2 | 1.12853 | 0.14353 |
| :--- | ---: | ---: |
| em 2 | 0.74582 | 0.70752 |
| ca 2 | 0.04654 | 0.28641 |
| $\mathrm{ma3}$ | 0.80563 | -0.12763 |
| em 3 | 0.66887 | 0.06423 |
| $\mathrm{ca3}$ | -0.66886 | 0.53487 |
| $\mathrm{ma4}$ | 0.36826 | -0.54151 |
| em 4 | 0.09955 | -0.24969 |
| $\mathrm{ca4}$ | -0.63184 | 0.68533 |
| ma5 | -0.08726 | -1.09621 |
| em5 | -0.77026 | -0.44656 |
| ca5 | -1.70498 | 0.03971 |

Fig. 10.6 Representation of food expenditures and family types in two dimensions $a$ MVAdecofood

The coordinates of the projected data points are given in the two lower windows of Fig. 10.6. Let us first examine the food expenditure window. In this window we see the representation of the $p=7$ variables given by the first two factors. The plot shows the factorial variables $w_{1}$ and $w_{2}$ in the same fashion as Fig. 10.4. We see that the points for meat, poultry, vegetables and fruits are close to each other in the lower left of the graph. The expenditures for bread and milk can be found in the upper left, whereas wine stands alone in the upper right. The first factor, $w_{1}$, may be interpreted as the meat/fruit factor of consumption, the second factor, $w_{2}$, as the bread/wine component.

In the lower window on the right-hand side, we show the factorial variables $z_{1}$ and $z_{2}$ from the fit of the $n=12$ household types. Note that by the Duality Relations of Theorem 10.4, the factorial variables $z_{j}$ are linear combinations of the factors $w_{k}$ from the left window. The points displayed in the consumer window (graph on the right) are plotted relative to an average consumer represented by the origin. The manager families are located in the lower left corner of the graph whereas the manual workers and employees tend to be in the upper right. The factorial variables for CA5 (managers with five children) lie close to the meat/fruit factor. Relative to the average consumer this household type is a large consumer of meat/poultry and
fruits/vegetables. In Chap. 11, we will return to these plots interpreting them in a much deeper way. At this stage, it suffices to notice that the plots provide a graphical representation in $\mathbb{R}^{2}$ of the information contained in the original, high-dimensional $(12 \times 7)$ data matrix.


### 10.6 Exercises

Exercise 10.1 Prove that $n^{-1} \mathcal{Z}^{\top} \mathcal{Z}$ is the covariance of the centred data matrix, where $\mathcal{Z}$ is the matrix formed by the columns $z_{k}=\mathcal{X} u_{k}$.

Exercise 10.2 Compute the SVD of the French food data (Table 22.6).
Exercise 10.3 Compute $\tau_{3}, \tau_{4}, \ldots$ for the French food data (Table 22.6).
Exercise 10.4 Apply the factorial techniques to the Swiss bank notes (Sect. 22.2).
Exercise 10.5 Apply the factorial techniques to the time budget data (Table 22.14).
Exercise 10.6 Assume that you wish to analyse $p$ independent identically distributed random variables. What is the percentage of the inertia explained by the first factor? What is the percentage of the inertia explained by the first q factors?

Exercise 10.7 Assume that you have p i.i.d. r.v.'s. What does the eigenvector, corresponding to the first factor, look like.

Exercise 10.8 Assume that you have two random variables, $X_{1}$ and $X_{2}=2 X_{1}$. What do the eigenvalues and eigenvectors of their correlation matrix look like? How many eigenvalues are nonzero?

Exercise 10.9 What percentage of inertia is explained by the first factor in the previous exercise?

Exercise 10.10 How do the eigenvalues and eigenvectors in Example 10.1 change if we take the prices in USD instead of in EUR? Does it make a difference if some of the prices are in EUR and others in USD?

## Chapter 11 <br> Principal Components Analysis

Chapter 10 presented the basic geometric tools needed to produce a lower dimensional description of the rows and columns of a multivariate data matrix. Principal components analysis (PCA) has the same objective with the exception that the rows of the data matrix $\mathcal{X}$ will now be considered as observations from a $p$-variate random variable $X$. The principle idea of reducing the dimension of $X$ is achieved through linear combinations. Low dimensional linear combinations are often easier to interpret and serve as an intermediate step in a more complex data analysis. More precisely one looks for linear combinations which create the largest spread among the values of $X$. In other words, one is searching for linear combinations with the largest variances.

Section 11.1 introduces the basic ideas and technical elements behind principal components. No particular assumption will be made on $X$ except that the mean vector and the covariance matrix exist. When reference is made to a data matrix $\mathcal{X}$ in Sect. 11.2, the empirical mean and covariance matrix will be used. Section 11.3 shows how to interpret the principal components by studying their correlations with the original components of $X$. Often analyses are performed in practice by looking at two-dimensional scatterplots. Section 11.4 develops inference techniques on principal components. This is particularly helpful in establishing the appropriate dimension reduction and thus in determining the quality of the resulting lower dimensional representations. Since principal component analysis is performed on covariance matrices, it is not scale invariant. Often, the measurement units of the components of $X$ are quite different, so it is reasonable to standardise the measurement units. The normalised version of principal components is defined in Sect. 11.5. In Sect. 11.6 it is discovered that the empirical principal components are the factors of appropriate transformations of the data matrix. The classical way of defining principal components through linear combinations with respect to the largest variance is described here in geometric terms, i.e. in terms of the optimal fit within subspaces generated by the columns and/or the rows of $\mathcal{X}$ as was discussed in Chap. 10. Section 11.9 concludes with additional examples.

### 11.1 Standardised Linear Combination

The main objective of PCA is to reduce the dimension of the observations. The simplest way of dimension reduction is to take just one element of the observed vector and to discard all others. This is not a very reasonable approach, as we have seen in the earlier chapters, since strength may be lost in interpreting the data. In the bank notes example we have seen that just one variable (e.g. $X_{1}=$ length) had no discriminatory power in distinguishing counterfeit from genuine bank notes. An alternative method is to weight all variables equally, i.e. to consider the simple average $p^{-1} \sum_{j=1}^{p} X_{j}$ of all the elements in the vector $X=\left(X_{1}, \ldots, X_{p}\right)^{\top}$. This again is undesirable, since all of the elements of $X$ are considered with equal importance (weight).

A more flexible approach is to study a weighted average, namely

$$
\begin{equation*}
\delta^{\top} X=\sum_{j=1}^{p} \delta_{j} X_{j}, \quad \text { such that } \quad \sum_{j=1}^{p} \delta_{j}^{2}=1 \tag{11.1}
\end{equation*}
$$

The weighting vector $\delta=\left(\delta_{1}, \ldots, \delta_{p}\right)^{\top}$ can then be optimised to investigate and to detect specific features. We call (11.1) a standardised linear combination (SLC). Which SLC should we choose? One aim is to maximise the variance of the projection $\delta^{\top} X$, i.e. to choose $\delta$ according to

$$
\begin{equation*}
\max _{\{\delta:\|\delta\|=1\}} \operatorname{Var}\left(\delta^{\top} X\right)=\max _{\{\delta:\|\delta\|=1\}} \delta^{\top} \operatorname{Var}(X) \delta . \tag{11.2}
\end{equation*}
$$

The interesting "directions" of $\delta$ are found through the spectral decomposition of the covariance matrix. Indeed, from Theorem 2.5 , the direction $\delta$ is given by the eigenvector $\gamma_{1}$ corresponding to the largest eigenvalue $\lambda_{1}$ of the covariance matrix $\Sigma=\operatorname{Var}(X)$.

Figures 11.1 and 11.2 show two such projections (SLCs) of the same data set with zero mean. In Fig. 11.1 an arbitrary projection is displayed. The upper window shows the data point cloud and the line onto which the data are projected. The middle window shows the projected values in the selected direction. The lower window shows the variance of the actual projection and the percentage of the total variance that is explained.

Figure 11.2 shows the projection that captures the majority of the variance in the data. This direction is of interest and is located along the main direction of the point cloud. The same line of thought can be applied to all data orthogonal to this direction leading to the second eigenvector. The SLC with the highest variance obtained from maximising (11.2) is the first principal component (PC) $y_{1}=\gamma_{1}^{\top} X$. Orthogonal to the direction $\gamma_{1}$ we find the SLC with the second highest variance: $y_{2}=\gamma_{2}^{\top} X$, the second PC.


Fig. 11.1 An arbitrary SLC $Q_{\text {MVApcasimu }}$


Fig. 11.2 The most interesting SLC 0 MVApcasimu

Proceeding in this way and writing in matrix notation, the result for a random variable $X$ with $\mathrm{E}(X)=\mu$ and $\operatorname{Var}(X)=\Sigma=\Gamma \Lambda \Gamma^{\top}$ is the PC transformation which is defined as

$$
\begin{equation*}
Y=\Gamma^{\top}(X-\mu) . \tag{11.3}
\end{equation*}
$$

Here we have centred the variable $X$ in order to obtain a zero mean PC variable $Y$.

Example 11.1 Consider a bivariate normal distribution $N(0, \Sigma)$ with $\Sigma=\left(\begin{array}{ll}1 & \rho \\ \rho & 1\end{array}\right)$ and $\rho>0$ (see Example 3.13). Recall that the 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

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

or

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

So 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 using formulas (4.22)-(4.26):

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

Similarly we find that

$$
\operatorname{Var}\left(Y_{2}\right)=\lambda_{2}
$$

This can be expressed more generally and is given in the next theorem.

Theorem 11.1 For a given $X \sim(\mu, \Sigma)$ let $Y=\Gamma^{\top}(X-\mu)$ be the $P C$ transformation. Then

$$
\begin{align*}
& \mathrm{E} Y_{j}=0, \quad j=1, \ldots, p  \tag{11.4}\\
& \operatorname{Var}\left(Y_{j}\right)=\lambda_{j}, \quad j=1, \ldots, p  \tag{11.5}\\
& \mathrm{Cov}\left(Y_{i}, Y_{j}\right)=0, \quad i \neq j  \tag{11.6}\\
& \operatorname{Var}\left(Y_{1}\right) \geq \operatorname{Var}\left(Y_{2}\right) \geq \cdots \geq \operatorname{Var}\left(Y_{p}\right) \geq 0  \tag{11.7}\\
& \sum_{j=1}^{p} \operatorname{Var}\left(Y_{j}\right)=\operatorname{tr}(\Sigma)  \tag{11.8}\\
& \prod_{j=1}^{p} \operatorname{Var}\left(Y_{j}\right)=|\Sigma| \tag{11.9}
\end{align*}
$$

Proof To prove (11.6), we use $\gamma_{i}$ to denote the $i$ th column of $\Gamma$. Then

$$
\operatorname{Cov}\left(Y_{i}, Y_{j}\right)=\gamma_{i}^{\top} \operatorname{Var}(X-\mu) \gamma_{j}=\gamma_{i}^{\top} \operatorname{Var}(X) \gamma_{j}
$$

As $\operatorname{Var}(X)=\Sigma=\Gamma \Lambda \Gamma^{\top}, \Gamma^{\top} \Gamma=\mathcal{I}$, we obtain via the orthogonality of $\Gamma$ :

$$
\gamma_{i}^{\top} \Gamma \Lambda \Gamma^{\top} \gamma_{j}=\left\{\begin{array}{cc}
0 & i \neq j \\
\lambda_{i} & i=j
\end{array}\right.
$$

In fact, as $Y_{i}=\gamma_{i}^{\top}(X-\mu)$ lies in the eigenvector space corresponding to $\gamma_{i}$, and eigenvector spaces corresponding to different eigenvalues are orthogonal to each other, we can directly see $Y_{i}$ and $Y_{j}$ are orthogonal to each other, so their covariance is 0 .

The connection between the PC transformation and the search for the best SLC is made in the following theorem, which follows directly from (11.2) and Theorem 2.5.

Theorem 11.2 There exists no SLC that has larger variance than $\lambda_{1}=\operatorname{Var}\left(Y_{1}\right)$.
Theorem 11.3 If $Y=a^{\top} X$ is an SLC that is not correlated with the first $k$ PCs of $X$, then the variance of $Y$ is maximised by choosing it to be the $(k+1)$-st PC.

|  | An SLC is a weighted average $\delta^{\top} X=\sum_{j=1}^{p} \delta_{j} X_{j}$ where $\delta$ is a |
| :--- | :--- |
| vector of length 1. |  |


| Summary (continued) |
| :---: |
| $\hookrightarrow$ Maximising the variance of $\delta^{\top} X$ leads to the choice $\delta=\gamma_{1}$, the eigenvector corresponding to the largest eigenvalue $\lambda_{1}$ of $\Sigma=$ $\operatorname{Var}(X)$. <br> This is a projection of $X$ into the one-dimensional space, where the components of $X$ are weighted by the elements of $\gamma_{1} . Y_{1}=$ $\gamma_{1}^{\top}(X-\mu)$ <br> is called the first principal component (PC). |
| $\hookrightarrow$ This projection can be generalised for higher dimensions. The PC transformation is the linear transformation $Y=\Gamma^{\top}(X-\mu)$, where $\Sigma=\operatorname{Var}(X)=\Gamma \Lambda \Gamma^{\top}$ and $\mu=\mathrm{E} X$. <br> $Y_{1}, Y_{2}, \ldots, Y_{p}$ are called the first, second, $\ldots$, and $p$-th PCs. |
| $\hookrightarrow$ The PCs have zero means, variance $\operatorname{Var}\left(Y_{j}\right)=\lambda_{j}$, and zero covariances. From $\lambda_{1} \geq \cdots \geq \lambda_{p}$ it follows that $\operatorname{Var}\left(Y_{1}\right) \geq$ $\cdots \geq \operatorname{Var}\left(Y_{p}\right)$. It holds that $\sum_{j=1}^{p} \operatorname{Var}\left(Y_{j}\right)=\operatorname{tr}(\Sigma)$ and $\prod_{j=1}^{p} \operatorname{Var}\left(Y_{j}\right)=\|\Sigma\|$. |
| $\hookrightarrow$ If $Y=a^{\top} X$ is an SLC which is not correlated with the first $k$ PCs of $X$, then the variance of $Y$ is maximised by choosing it to be the $(k+1)$-st PC. |

### 11.2 Principal Components in Practice

In practice the PC transformation has to be replaced by the respective estimators: $\mu$ becomes $\bar{x}, \Sigma$ is replaced by $\mathcal{S}$, etc. If $g_{1}$ denotes the first eigenvector of $\mathcal{S}$, the first principal component is given by $y_{1}=\left(\mathcal{X}-1_{n} \bar{x}^{\top}\right) g_{1}$. More generally if $\mathcal{S}=\mathcal{G} \mathcal{L} \mathcal{G}^{\top}$ is the spectral decomposition of $\mathcal{S}$, then the PCs are obtained by

$$
\begin{equation*}
\mathcal{Y}=\left(\mathcal{X}-1_{n} \bar{x}^{\top}\right) \mathcal{G} . \tag{11.10}
\end{equation*}
$$

Note that with the centering matrix $\mathcal{H}=\mathcal{I}-\left(n^{-1} 1_{n} 1_{n}^{\top}\right)$ and $\mathcal{H} 1_{n} \bar{x}^{\top}=0$ we can write

$$
\begin{align*}
\mathcal{S}_{\mathcal{Y}} & =n^{-1} \mathcal{Y}^{\top} \mathcal{H Y}=n^{-1} \mathcal{G}^{\top}\left(\mathcal{X}-1_{n} \bar{x}^{\top}\right)^{\top} \mathcal{H}\left(\mathcal{X}-1_{n} \bar{x}^{\top}\right) \mathcal{G} \\
& =n^{-1} \mathcal{G}^{\top} \mathcal{X}^{\top} \mathcal{H} \mathcal{X} \mathcal{G}=\mathcal{G}^{\top} \mathcal{S G}=\mathcal{L} \tag{11.11}
\end{align*}
$$

where $\mathcal{L}=\operatorname{diag}\left(\ell_{1}, \ldots, \ell_{p}\right)$ is the matrix of eigenvalues of $\mathcal{S}$. Hence the variance of $y_{i}$ equals the eigenvalue $\ell_{i}$ !

The PC technique is sensitive to scale changes. If we multiply one variable by a scalar we obtain different eigenvalues and eigenvectors. This is due to the fact that


Fig. 11.3 Principal components of the bank data $Q_{\text {mVApcabank }}$
an eigenvalue decomposition is performed on the covariance matrix and not on the correlation matrix (see Sect. 11.5). The following warning is therefore important:

The PC transformation should be applied to data that have approximately the same scale in each variable.

Example 11.2 Let us apply this technique to the bank data set. In this example we do not standardise the data. Figure 11.3 shows some PC plots of the bank data set. The genuine and counterfeit bank notes are marked by " 0 " and " + ", respectively.

Recall that the mean vector of $\mathcal{X}$ is

$$
\bar{x}=(214.9,130.1,129.9,9.4,10.6,140.5)^{\top}
$$

The vector of eigenvalues of $\mathcal{S}$ is

$$
\ell=(2.985,0.931,0.242,0.194,0.085,0.035)^{\top}
$$

The eigenvectors $g_{j}$ are given by the columns of the matrix

$$
\mathcal{G}=\left(\begin{array}{rrrrrr}
-0.044 & 0.011 & 0.326 & 0.562 & -0.753 & 0.098 \\
0.112 & 0.071 & 0.259 & 0.455 & 0.347 & -0.767 \\
0.139 & 0.066 & 0.345 & 0.415 & 0.535 & 0.632 \\
0.768 & -0.563 & 0.218 & -0.186 & -0.100 & -0.022 \\
0.202 & 0.659 & 0.557 & -0.451 & -0.102 & -0.035 \\
-0.579 & -0.489 & 0.592 & -0.258 & 0.085 & -0.046
\end{array}\right) .
$$

The first column of $\mathcal{G}$ is the first eigenvector and gives the weights used in the linear combination of the original data in the first PC.

Example 11.3 To see how sensitive the PCs are to a change in the scale of the variables, assume that $X_{1}, X_{2}, X_{3}$ and $X_{6}$ are measured in cm and that $X_{4}$ and $X_{5}$ remain in mm in the bank data set. This leads to:

$$
\bar{x}=(21.49,13.01,12.99,9.41,10.65,14.05)^{\top}
$$

The covariance matrix can be obtained from $S$ in (3.4) by dividing rows 1, 2, 3, 6 and columns $1,2,3,6$ by 10 . We obtain:

$$
\ell=(2.101,0.623,0.005,0.002,0.001,0.0004)^{\top}
$$

which clearly differs from Example 11.2. Only the first two eigenvectors are given:

$$
\begin{gathered}
g_{1}=(-0.005,0.011,0.014,0.992,0.113,-0.052)^{\top} \\
g_{2}=(-0.001,0.013,0.016,-0.117,0.991,-0.069)^{\top} .
\end{gathered}
$$

Comparing these results to the first two columns of $\mathcal{G}$ from Example 11.2, a completely different story is revealed. Here the first component is dominated by $X_{4}$ (lower margin) and the second by $X_{5}$ (upper margin), while all of the other variables have much less weight. The results are shown in Fig. 11.4. Section 11.5 will show how to select a reasonable standardisation of the variables when the scales are too different.


| Summary (continued) |  |
| :---: | :---: |
| $\hookrightarrow \quad$For the practical implementation of PCA we replace $\mu$ by the mean <br> $\bar{x}$ and $\Sigma$ by the empirical covariance $\mathcal{S}$. Then we compute the <br> eigenvalues $\ell_{1}, \ldots, \ell_{p}$ and the eigenvectors $g_{1}, \ldots, g_{p}$ of $\mathcal{S}$. The <br> graphical representation of the PCs is obtained by plotting the first <br> PC vs. the second (and eventually vs. the third). |  |
| $\hookrightarrow$ | The components of the eigenvectors $g_{i}$ are the weights of the <br> original variables in the PCs. |

First vs. Second PC


First vs. Third PC


Second vs. Third PC


Eigenvalues of $S$


Fig. 11.4 Principal components of the rescaled bank data $Q$ MVApcabankr

### 11.3 Interpretation of the PCs

Recall that the main idea of PC transformations is to find the most informative projections that maximise variances. The most informative SLC is given by the first eigenvector. In Sect. 11.2 the eigenvectors were calculated for the bank data. In particular, with centred $x$ 's, we had:

$$
\begin{aligned}
& y_{1}=-0.044 x_{1}+0.112 x_{2}+0.139 x_{3}+0.768 x_{4}+0.202 x_{5}-0.579 x_{6} \\
& y_{2}=0.011 x_{1}+0.071 x_{2}+0.066 x_{3}-0.563 x_{4}+0.659 x_{5}-0.489 x_{6}
\end{aligned}
$$

and

$$
\begin{aligned}
& x_{1}=\text { length } \\
& x_{2}=\text { left height } \\
& x_{3}=\text { right height } \\
& x_{4}=\text { bottom frame } \\
& x_{5}=\text { top frame } \\
& x_{6}=\text { diagonal } .
\end{aligned}
$$

Hence, the first PC is essentially the difference between the bottom frame variable and the diagonal. The second PC is best described by the difference between the top frame variable and the sum of bottom frame and diagonal variables.

The weighting of the PCs tells us in which directions, expressed in original coordinates, the best variance explanation is obtained. A measure of how well the first $q$ PCs explain variation is given by the relative proportion:

$$
\begin{equation*}
\psi_{q}=\frac{\sum_{j=1}^{q} \lambda_{j}}{\sum_{j=1}^{p} \lambda_{j}}=\frac{\sum_{j=1}^{q} \operatorname{Var}\left(Y_{j}\right)}{\sum_{j=1}^{p} \operatorname{Var}\left(Y_{j}\right)} . \tag{11.12}
\end{equation*}
$$

Referring to the bank data Example 11.2, the (cumulative) proportions of explained variance are given in Table 11.1. The first PC $(q=1)$ already explains $67 \%$ of the variation. The first three $(q=3)$ PCs explain $93 \%$ of the variation. Once again it should be noted that PCs are not scale invariant, e.g. the PCs derived from the correlation matrix give different results than the PCs derived from the covariance matrix (see Sect. 11.5).

A good graphical representation of the ability of the PCs to explain the variation in the data is given by the scree plot shown in the lower right-hand window of Fig. 11.3. The screeplot can be modified by using the relative proportions on the $y$-axis, as is shown in Fig. 11.5 for the bank data set.

Table 11.1 Proportion of variance of PC's

| Eigenvalue | Proportion of variance | Cumulated proportion |
| :--- | :--- | :--- |
| 2.985 | 0.67 | 0.67 |
| 0.931 | 0.21 | 0.88 |
| 0.242 | 0.05 | 0.93 |
| 0.194 | 0.04 | 0.97 |
| 0.085 | 0.02 | 0.99 |
| 0.035 | 0.01 | 1.00 |

Fig. 11.5 Relative proportion of variance explained by PCs $\mathbf{Q}$ MVApcabanki


The covariance between the PC vector $Y$ and the original vector $X$ is calculated with the help of (11.4) as follows:

$$
\begin{align*}
\operatorname{Cov}(X, Y) & =\mathrm{E}\left(X Y^{\top}\right)-\mathrm{E} X \mathrm{E} Y^{\top}=\mathrm{E}\left(X Y^{\top}\right) \\
& =\mathrm{E}\left(X X^{\top} \Gamma\right)-\mu \mu^{\top} \Gamma=\operatorname{Var}(X) \Gamma \\
& =\Sigma \Gamma  \tag{11.13}\\
& =\Gamma \Lambda \Gamma^{\top} \Gamma \\
& =\Gamma \Lambda
\end{align*}
$$

Hence, the correlation, $\rho_{X_{i} Y_{j}}$, between variable $X_{i}$ and the PC $Y_{j}$ is

$$
\begin{equation*}
\rho_{X_{i} Y_{j}}=\frac{\gamma_{i j} \lambda_{j}}{\left(\sigma_{X_{i} X_{i}} \lambda_{j}\right)^{1 / 2}}=\gamma_{i j}\left(\frac{\lambda_{j}}{\sigma_{X_{i} X_{i}}}\right)^{1 / 2} . \tag{11.14}
\end{equation*}
$$

Using actual data, this of course translates into

$$
\begin{equation*}
r_{X_{i} Y_{j}}=g_{i j}\left(\frac{\ell_{j}}{s_{X_{i} X_{i}}}\right)^{1 / 2} . \tag{11.15}
\end{equation*}
$$

The correlations can be used to evaluate the relations between the PCs $Y_{j}$ where $j=1, \ldots, q$, and the original variables $X_{i}$ where $i=1, \ldots, p$. Note that

$$
\begin{equation*}
\sum_{j=1}^{p} r_{X_{i} Y_{j}}^{2}=\frac{\sum_{j=1}^{p} \ell_{j} g_{i j}^{2}}{s_{X_{i} X_{i}}}=\frac{s_{X_{i} X_{i}}}{s_{X_{i} X_{i}}}=1 \tag{11.16}
\end{equation*}
$$

Fig. 11.6 The correlation of the original variable with the PCs $\mathbf{Q}$ MVApcabanki


Indeed, $\sum_{j=1}^{p} \ell_{j} g_{i j}^{2}=g_{i}^{\top} \mathcal{L} g_{i}$ is the $(i, i)$-element of the matrix $\mathcal{G} \mathcal{L G}{ }^{\top}=\mathcal{S}$, so that $r_{X_{i} Y_{j}}^{2}$ may be seen as the proportion of variance of $X_{i}$ explained by $Y_{j}$.

In the space of the first two PCs we plot these proportions, i.e. $r_{X_{i} Y_{1}}$ versus $r_{X_{i} Y_{2}}$. Figure 11.6 shows this for the bank notes example. This plot shows which of the original variables are most strongly correlated with PC $Y_{1}$ and $Y_{2}$.

From (11.16) it obviously follows that $r_{X_{i} Y_{1}}^{2}+r_{X_{i} Y_{2}}^{2} \leq 1$ so that the points are always inside the circle of radius 1 . In the bank notes example, the variables $X_{4}, X_{5}$ and $X_{6}$ correspond to correlations near the periphery of the circle and are thus well explained by the first two PCs. Recall that we have interpreted the first PC as being essentially the difference between $X_{4}$ and $X_{6}$. This is also reflected in Fig. 11.6 since the points corresponding to these variables lie on different sides of the vertical axis. An analogous remark applies to the second PC. We had seen that the second PC is well described by the difference between $X_{5}$ and the sum of $X_{4}$ and $X_{6}$. Now we are able to see this result again from Fig. 11.6 since the point corresponding to $X_{5}$ lies above the horizontal axis and the points corresponding to $X_{4}$ and $X_{6}$ lie below.

The correlations of the original variables $X_{i}$ and the first two PCs are given in Table 11.2 along with the cumulated percentage of variance of each variable explained by $Y_{1}$ and $Y_{2}$. This table confirms the above results. In particular, it confirms that the percentage of variance of $X_{1}$ (and $X_{2}, X_{3}$ ) explained by the first two PCs is relatively small and so are their weights in the graphical representation of the individual bank notes in the space of the first two PCs (as can be seen in the upper left plot in Fig. 11.3). Looking simultaneously at Fig. 11.6 and the upper left plot of Fig. 11.3 shows that the genuine bank notes are roughly characterised by large values of $X_{6}$ and smaller values of $X_{4}$. The counterfeit bank notes show larger values of $X_{5}$ (see Example 7.15).

Table 11.2 Correlation between the original variables and the PCs

|  | $r_{X_{i} Y_{1}}$ | $r_{X_{i} Y_{2}}$ | $r_{X_{i} Y_{1}}^{2}+r_{X_{i} Y_{2}}^{2}$ |
| :--- | ---: | ---: | :--- |
| $X_{1}$ length | -0.201 | 0.028 | 0.041 |
| $X_{2}$ left h. | 0.538 | 0.191 | 0.326 |
| $X_{3}$ right h. | 0.597 | 0.159 | 0.381 |
| $X_{4}$ lower | 0.921 | -0.377 | 0.991 |
| $X_{5}$ upper | 0.435 | 0.794 | 0.820 |
| $X_{6}$ diagonal | -0.870 | -0.410 | 0.926 |



## Summary

$\hookrightarrow$ The weighting of the PCs tells us in which directions, expressed in original coordinates, the best explanation of the variance is obtained. Note that the PCs are not scale invariant.
$\hookrightarrow$ A measure of how well the first $q$ PCs explain variation is given by the relative proportion $\psi_{q}=\sum_{j=1}^{q} \lambda_{j} / \sum_{j=1}^{p} \lambda_{j}$. A good graphical representation of the ability of the PCs to explain the variation in the data is the scree plot of these proportions.
$\hookrightarrow$ The correlation between PC $Y_{j}$ and an original variable $X_{i}$ is $\rho_{X_{i} Y_{j}}=\gamma_{i j}\left(\frac{\lambda_{j}}{\sigma X_{i} X_{i}}\right)^{1 / 2}$. For a data matrix this translates into $r_{X_{i} Y_{j}}^{2}=\frac{\ell_{j} g_{i j}^{2}}{s_{X_{i} X_{i}}} . r_{X_{i} Y_{j}}^{2}$ can be interpreted as the proportion of variance of $X_{i}$ explained by $Y_{j}$. A plot of $r_{X_{i} Y_{1}}$ vs. $r_{X_{i} Y_{2}}$ shows which of the original variables are most strongly correlated with the PCs, namely those that are close to the periphery of the circle of radius 1.

### 11.4 Asymptotic Properties of the PCs

In practice, PCs are computed from sample data. The following theorem yields results on the asymptotic distribution of the sample PCs.

Theorem 11.4 Let $\Sigma>0$ with distinct eigenvalues, and let $\mathcal{U} \sim m^{-1} W_{p}(\Sigma, m)$ with spectral decompositions $\Sigma=\Gamma \Lambda \Gamma^{\top}$, and $\mathcal{U}=\mathcal{G} \mathcal{L G}{ }^{\top}$. Then
(a) $\sqrt{m}(\ell-\lambda) \xrightarrow{\mathcal{L}} N_{p}\left(0,2 \Lambda^{2}\right)$,
where $\ell=\left(\ell_{1}, \ldots, \ell_{p}\right)^{\top}$ and $\lambda=\left(\lambda_{1}, \ldots, \lambda_{p}\right)^{\top}$ are the diagonals of $\mathcal{L}$ and $\Lambda$,
(b) $\sqrt{m}\left(g_{j}-\gamma_{j}\right) \xrightarrow{\mathcal{L}} N_{p}\left(0, \mathcal{V}_{j}\right)$,
with $\mathcal{V}_{j}=\lambda_{j} \sum_{k \neq j} \frac{\lambda_{k}}{\left(\lambda_{k}-\lambda_{j}\right)^{2}} \gamma_{k} \gamma_{k}^{\top}$,
(c) $\operatorname{Cov}\left(g_{j}, g_{k}\right)=\mathcal{V}_{j k}$,
where the $(r, s)$-element of the matrix $\mathcal{V}_{j k}(p \times p)$ is $-\frac{\lambda_{j} \lambda_{k} \gamma_{r k} \gamma_{s j}}{m\left(\lambda_{j}-\lambda_{k}\right)^{2}}$,
(d) the elements in $\ell$ are asymptotically independent of the elements in $\mathcal{G}$.

Example 11.4 Since $n \mathcal{S} \sim W_{p}(\Sigma, n-1)$ if $X_{1}, \ldots, X_{n}$ are drawn from $N(\mu, \Sigma)$, we have that

$$
\begin{equation*}
\sqrt{n-1}\left(\ell_{j}-\lambda_{j}\right) \xrightarrow{\mathcal{L}} N\left(0,2 \lambda_{j}^{2}\right), \quad j=1, \ldots, p \tag{11.17}
\end{equation*}
$$

Since the variance of (11.17) depends on the true mean $\lambda_{j}$ a $\log$ transformation is useful. Consider $f\left(\ell_{j}\right)=\log \left(\ell_{j}\right)$. Then $\left.\frac{d}{d \ell_{j}} f\right|_{\ell_{j}=\lambda_{j}}=\frac{1}{\lambda_{j}}$ and by the Transformation Theorem 4.11 we have from (11.17) that

$$
\begin{equation*}
\sqrt{n-1}\left(\log \ell_{j}-\log \lambda_{j}\right) \xrightarrow{\mathcal{L}} N(0,2) . \tag{11.18}
\end{equation*}
$$

Hence,

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

and a two-sided confidence interval at the $1-\alpha=0.95$ significance level is given by

$$
\log \left(\ell_{j}\right)-1.96 \sqrt{\frac{2}{n-1}} \leq \log \lambda_{j} \leq \log \left(\ell_{j}\right)+1.96 \sqrt{\frac{2}{n-1}}
$$

In the bank data example we have that

$$
\ell_{1}=2.98
$$

Therefore,

$$
\log (2.98) \pm 1.96 \sqrt{\frac{2}{199}}=\log (2.98) \pm 0.1965
$$

It can be concluded for the true eigenvalue that

$$
\mathrm{P}\left\{\lambda_{1} \in(2.448,3.62)\right\} \approx 0.95
$$

## Variance Explained by the First q PCs

The variance explained by the first $q$ PCs is given by

$$
\psi=\frac{\lambda_{1}+\cdots+\lambda_{q}}{\sum_{j=1}^{p} \lambda_{j}}
$$

In practice this is estimated by

$$
\hat{\psi}=\frac{\ell_{1}+\cdots+\ell_{q}}{\sum_{j=1}^{p} \ell_{j}}
$$

From Theorem 11.4 we know the distribution of $\sqrt{n-1}(\ell-\lambda)$. Since $\psi$ is a nonlinear function of $\lambda$, we can again apply the Transformation Theorem 4.11 to obtain that

$$
\sqrt{n-1}(\hat{\psi}-\psi) \xrightarrow{\mathcal{L}} N\left(0, \mathcal{D}^{\top} \mathcal{V} \mathcal{D}\right)
$$

where $\mathcal{V}=2 \Lambda^{2}\left(\right.$ from Theorem 11.4) and $\mathcal{D}=\left(d_{1}, \ldots, d_{p}\right)^{\top}$ with

$$
d_{j}=\frac{\partial \psi}{\partial \lambda_{j}}= \begin{cases}\frac{1-\psi}{\operatorname{tr}(\Sigma)} & \text { for } 1 \leq j \leq q \\ \frac{-\psi}{\operatorname{tr}(\Sigma)} & \text { for } q+1 \leq j \leq p\end{cases}
$$

Given this result, the following theorem can be derived.

## Theorem 11.5

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

where

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

and

$$
\beta=\frac{\lambda_{1}^{2}+\cdots+\lambda_{q}^{2}}{\lambda_{1}^{2}+\cdots+\lambda_{p}^{2}}
$$

Example 11.5 From Sect. 11.3 it is known that the first PC for the Swiss bank notes resolves $67 \%$ of the variation. It can be tested whether the true proportion is actually $75 \%$. Computing

$$
\begin{aligned}
\hat{\beta} & =\frac{\ell_{1}^{2}}{\ell_{1}^{2}+\cdots+\ell_{p}^{2}}=\frac{(2.985)^{2}}{(2.985)^{2}+(0.931)^{2}+\cdots(0.035)^{2}}=0.902 \\
\operatorname{tr}(\mathcal{S}) & =4.472 \\
\operatorname{tr}\left(\mathcal{S}^{2}\right) & =\sum_{j=1}^{p} \ell_{j}^{2}=9.883 \\
\hat{\omega}^{2} & =\frac{2 \operatorname{tr}\left(\mathcal{S}^{2}\right)}{\{\operatorname{tr}(\mathcal{S})\}^{2}}\left(\hat{\psi}^{2}-2 \hat{\beta} \hat{\psi}+\hat{\beta}\right) \\
& =\frac{2 \cdot 9.883}{(4.472)^{2}}\left\{(0.668)^{2}-2(0.902)(0.668)+0.902\right\}=0.142 .
\end{aligned}
$$

Hence, a confidence interval at a significance of level $1-\alpha=0.95$ is given by

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

Clearly the hypothesis that $\psi=75 \%$ can be rejected!

| $\hookrightarrow$ | The eigenvalues $\ell_{j}$ and eigenvectors $g_{j}$ are asymptotically, nor- |
| :--- | :--- |
|  | mally distributed, in particular $\sqrt{n-1}(\ell-\lambda) \xrightarrow{\mathcal{L}} N_{p}\left(0,2 \Lambda^{2}\right)$. |


| Summary (continued) |  |
| :--- | :--- |
| $\hookrightarrow$ | It holds for $\hat{\psi}$, the estimate of $\psi$ (the proportion of the variance |
|  | explained by the first $q$ PCs) that $\sqrt{n-1}(\hat{\psi}-\psi) \xrightarrow{\mathcal{L}} N\left(0, \omega^{2}\right)$, |
|  | where $\omega$ is given in Theorem 11.5. |

### 11.5 Normalised Principal Components Analysis

In certain situations the original variables can be heterogeneous w.r.t. their variances. This is particularly true when the variables are measured on heterogeneous scales (such as years, kilograms, dollars, ...). In this case a description of the information contained in the data needs to be provided which is robust w.r.t. the choice of scale. This can be achieved through a standardisation of the variables, namely

$$
\begin{equation*}
\mathcal{X}_{S}=\mathcal{H X} \mathcal{D}^{-1 / 2} \tag{11.19}
\end{equation*}
$$

where $\mathcal{D}=\operatorname{diag}\left(s_{X_{1} X_{1}}, \ldots, s_{X_{p} X_{p}}\right)$. Note that $\bar{x}_{S}=0$ and $\mathcal{S}_{\mathcal{X}_{S}}=\mathcal{R}$, the correlation matrix of $\mathcal{X}$. The PC transformations of the matrix $\mathcal{X}_{S}$ are referred to as the Normalised Principal Components (NPCs). The spectral decomposition of $\mathcal{R}$ is

$$
\begin{equation*}
\mathcal{R}=\mathcal{G}_{\mathcal{R}} \mathcal{L}_{\mathcal{R}} \mathcal{G}_{\mathcal{R}}^{\top} \tag{11.20}
\end{equation*}
$$

where $\mathcal{L}_{\mathcal{R}}=\operatorname{diag}\left(\ell_{1}^{\mathcal{R}}, \ldots, \ell_{p}^{\mathcal{R}}\right)$ and $\ell_{1}^{\mathcal{R}} \geq \cdots \geq \ell_{p}^{\mathcal{R}}$ are the eigenvalues of $\mathcal{R}$ with corresponding eigenvectors $g_{1}^{\mathcal{R}}, \ldots, g_{p}^{\mathcal{R}}$ (note that here $\sum_{j=1}^{p} \ell_{j}^{\mathcal{R}}=\operatorname{tr}(\mathcal{R})=p$ ).

The NPCs, $Z_{j}$, provide a representation of each individual, and is given by

$$
\begin{equation*}
\mathcal{Z}=\mathcal{X}_{S} \mathcal{G}_{\mathcal{R}}=\left(z_{1}, \ldots, z_{p}\right) \tag{11.21}
\end{equation*}
$$

After transforming the variables, once again, we have that

$$
\begin{align*}
\bar{z} & =0  \tag{11.22}\\
\mathcal{S}_{\mathcal{Z}} & =\mathcal{G}_{\mathcal{R}}^{\top} \mathcal{S}_{\mathcal{X}_{S}} \mathcal{G}_{\mathcal{R}}=\mathcal{G}_{\mathcal{R}}^{\top} \mathcal{R} \mathcal{G}_{\mathcal{R}}=\mathcal{L}_{\mathcal{R}} \tag{11.23}
\end{align*}
$$



The NPCs provide a perspective similar to that of the PCs, but in terms of the relative position of individuals, NPC gives each variable the same weight (with the PCs the variable with the largest variance received the largest weight).

Computing the covariance and correlation between $X_{i}$ and $Z_{j}$ is straightforward:

$$
\begin{align*}
\mathcal{S}_{X_{S}, Z} & =\frac{1}{n} \mathcal{X}_{S}^{\top} \mathcal{Z}=\mathcal{G}_{\mathcal{R}} \mathcal{L}_{\mathcal{R}}  \tag{11.24}\\
\mathcal{R}_{X_{S}, Z} & =\mathcal{G}_{\mathcal{R}} \mathcal{L}_{\mathcal{R}} \mathcal{L}_{\mathcal{R}}^{-1 / 2}=\mathcal{G}_{\mathcal{R}} \mathcal{L}_{\mathcal{R}}^{1 / 2} \tag{11.25}
\end{align*}
$$

The correlations between the original variables $X_{i}$ and the $\mathrm{NPCs} Z_{j}$ are:

$$
\begin{gather*}
r_{X_{i} Z_{j}}=\sqrt{\ell_{j}} g_{R, i j}  \tag{11.26}\\
\sum_{j=1}^{p} r_{X_{i} Z_{j}}^{2}=1 \tag{11.27}
\end{gather*}
$$

(compare this to (11.15) and (11.16)). The resulting NPCs, the $Z_{j}$, can be interpreted in terms of the original variables and the role of each PC in explaining the variation in variable $X_{i}$ can be evaluated.

### 11.6 Principal Components as a Factorial Method

The empirical PCs (normalised or not) turn out to be equivalent to the factors that one would obtain by decomposing the appropriate data matrix into its factors (see Chap. 10). It will be shown that the PCs are the factors representing the rows of the centred data matrix and that the NPCs correspond to the factors of the standardised data matrix. The representation of the columns of the standardised data matrix provides (at a scale factor) the correlations between the NPCs and the original variables. The derivation of the ( N )PCs presented above will have a nice geometric justification here since they are the best fit in subspaces generated by the columns of the (transformed) data matrix $\mathcal{X}$. This analogy provides complementary interpretations of the graphical representations shown above.

Assume, as in Chap. 10, that we want to obtain representations of the individuals (the rows of $\mathcal{X}$ ) and of the variables (the columns of $\mathcal{X}$ ) in spaces of smaller dimension. To keep the representations simple, some prior transformations are performed. Since the origin has no particular statistical meaning in the space of individuals, we will first shift the origin to the centre of gravity, $\bar{x}$, of the point cloud. This is the same as analysing the centred data matrix $\mathcal{X}_{C}=\mathcal{H} \mathcal{X}$. Now all of the variables have zero means, thus the technique used in Chap. 10 can be applied to the matrix $\mathcal{X}_{C}$. Note that the spectral decomposition of $\mathcal{X}_{C}^{\top} \mathcal{X}_{C}$ is related to that of $\mathcal{S}_{X}$, namely

$$
\begin{equation*}
\mathcal{X}_{C}^{\top} \mathcal{X}_{C}=\mathcal{X}^{\top} \mathcal{H}^{\top} \mathcal{H} \mathcal{X}=n \mathcal{S}_{X}=n \mathcal{G} \mathcal{L} \mathcal{G}^{\top} \tag{11.28}
\end{equation*}
$$

The factorial variables are obtained by projecting $\mathcal{X}_{C}$ on $\mathcal{G}$,

$$
\begin{equation*}
\mathcal{Y}=\mathcal{X}_{C} \mathcal{G}=\left(y_{1}, \ldots, y_{p}\right) \tag{11.29}
\end{equation*}
$$

These are the same principal components obtained above, see formula (11.10). (Note that the $y$ 's here correspond to the $z$ 's in Sect. 10.2.) Since $\mathcal{H} \mathcal{X}_{C}=\mathcal{X}_{C}$, it immediately follows that

$$
\begin{align*}
\bar{y} & =0  \tag{11.30}\\
\mathcal{S}_{Y} & =\mathcal{G}^{\top} \mathcal{S}_{X} \mathcal{G}=\mathcal{L}=\operatorname{diag}\left(\ell_{1}, \ldots, \ell_{p}\right) \tag{11.31}
\end{align*}
$$

The scatterplot of the individuals on the factorial axes are thus centred around the origin and are more spread out in the first direction (first PC has variance $\ell_{1}$ ) than in the second direction (second PC has variance $\ell_{2}$ ).

The representation of the variables can be obtained using the Duality Relations (10.11), and (10.12). The projections of the columns of $\mathcal{X}_{C}$ onto the eigenvectors $v_{k}$ of $\mathcal{X}_{C} \mathcal{X}_{C}^{\top}$ are

$$
\begin{equation*}
\mathcal{X}_{C}^{\top} v_{k}=\frac{1}{\sqrt{n \ell_{k}}} \mathcal{X}_{C}^{\top} \mathcal{X}_{C} g_{k}=\sqrt{n \ell_{k}} g_{k} \tag{11.32}
\end{equation*}
$$

Thus the projections of the variables on the first $p$ axes are the columns of the matrix

$$
\begin{equation*}
\mathcal{X}_{C}^{\top} \mathcal{V}=\sqrt{n} \mathcal{G} \mathcal{L}^{1 / 2} \tag{11.33}
\end{equation*}
$$

Considering the geometric representation, there is a nice statistical interpretation of the angle between two columns of $\mathcal{X}_{C}$. Given that

$$
\begin{array}{r}
x_{C[j]}^{\top} x_{C[k]}=n s_{X_{j} X_{k}}, \\
\left\|x_{C[j]}\right\|^{2}=n s_{X_{j} X_{j}}, \tag{11.35}
\end{array}
$$

where $x_{C[j]}$ and $x_{C[k]}$ denote the $j$-th and $k$-th column of $\mathcal{X}_{C}$, it holds that in the full space of the variables, if $\theta_{j k}$ is the angle between two variables, $x_{C[j]}$ and $x_{C[k]}$, then

$$
\begin{equation*}
\cos \theta_{j k}=\frac{x_{C[j]}^{\top} x_{C[k]}}{\left\|x_{C[j]}\right\|\left\|x_{C[k]}\right\|}=r_{X_{j} X_{k}} . \tag{11.36}
\end{equation*}
$$

(Example 2.11 shows the general connection that exists between the angle and correlation of two variables). As a result, the relative positions of the variables in the scatterplot of the first columns of $\mathcal{X}_{C}^{\top} \mathcal{V}$ may be interpreted in terms of their correlations; the plot provides a picture of the correlation structure of the original data set. Clearly, one should take into account the percentage of variance explained by the chosen axes when evaluating the correlation.

The NPCs can also be viewed as a factorial method for reducing the dimension. The variables are again standardised so that each one has mean zero and unit variance and is independent of the scale of the variables. The factorial analysis of $\mathcal{X}_{S}$ provides the NPCs. The spectral decomposition of $\mathcal{X}_{S}^{\top} \mathcal{X}_{S}$ is related to that of $\mathcal{R}$, namely

$$
\mathcal{X}_{S}^{\top} \mathcal{X}_{S}=\mathcal{D}^{-1 / 2} \mathcal{X}^{\top} \mathcal{H} \mathcal{X} \mathcal{D}^{-1 / 2}=n \mathcal{R}=n \mathcal{G}_{\mathcal{R}} \mathcal{L}_{\mathcal{R}} \mathcal{G}_{\mathcal{R}}^{\top}
$$

The NPCs $Z_{j}$, given by (11.21), may be viewed as the projections of the rows of $\mathcal{X}_{S}$ onto $\mathcal{G}_{R}$.

The representation of the variables are again given by the columns of

$$
\begin{equation*}
\mathcal{X}_{S}^{\top} \mathcal{V}_{\mathcal{R}}=\sqrt{n} \mathcal{G}_{\mathcal{R}} \mathcal{L}_{\mathcal{R}}^{1 / 2} \tag{11.37}
\end{equation*}
$$

Comparing (11.37) and (11.25) we see that the projections of the variables in the factorial analysis provide the correlation between the NPCs $\mathcal{Z}_{k}$ and the original variables $x_{[j]}$ (up to the factor $\sqrt{n}$ which could be the scale of the axes).

This implies that a deeper interpretation of the representation of the individuals can be obtained by looking simultaneously at the graphs plotting the variables. Note that

$$
\begin{align*}
x_{S[j]}^{\top} x_{S[k]} & =n r_{X_{j} X_{k}},  \tag{11.38}\\
\left\|x_{S[j]}\right\|^{2} & =n, \tag{11.39}
\end{align*}
$$

where $x_{S[j]}$ and $x_{S[k]}$ denote the $j$-th and $k$-th column of $\mathcal{X}_{S}$. Hence, in the full space, all the standardised variables (columns of $\mathcal{X}_{S}$ ) are contained within the "sphere" in $\mathbb{R}^{n}$, which is centred at the origin and has radius $\sqrt{n}$ (the scale of the graph). As in (11.36), given the angle $\theta_{j k}$ between two columns $x_{S[j]}$ and $x_{S[k]}$, it holds that

$$
\begin{equation*}
\cos \theta_{j k}=r_{X_{j} X_{k}} \tag{11.40}
\end{equation*}
$$

Therefore, when looking at the representation of the variables in the spaces of reduced dimension (for instance the first two factors), we have a picture of the correlation structure between the original $X_{i}$ 's in terms of their angles. Of course, the quality of the representation in those subspaces has to be taken into account, which is presented in the next section.

## Quality of the Representations

As said before, an overall measure of the quality of the representation is given by

$$
\psi=\frac{\ell_{1}+\ell_{2}+\cdots+\ell_{q}}{\sum_{j=1}^{p} \ell_{j}}
$$

In practice, $q$ is chosen to be equal to 1,2 or 3 . Suppose for instance that $\psi=0.93$ for $q=2$. This means that the graphical representation in two dimensions captures $93 \%$ of the total variance. In other words, there is minimal dispersion in a third direction (no more than $7 \%$ ).

It can be useful to check if each individual is well represented by the PCs. Clearly, the proximity of two individuals on the projected space may not necessarily coincide with the proximity in the full original space $\mathbb{R}^{p}$, which may lead to erroneous interpretations of the graphs. In this respect, it is worth computing the angle $\vartheta_{i k}$ between the representation of an individual $i$ and the $k$-th PC or NPC axis. This can be done using (2.40), i.e.

$$
\cos \vartheta_{i k}=\frac{y_{i}^{\top} e_{k}}{\left\|y_{i}\right\|\left\|e_{k}\right\|}=\frac{y_{i k}}{\left\|x_{C i}\right\|}
$$

for the PCs or analogously

$$
\cos \zeta_{i k}=\frac{z_{i}^{\top} e_{k}}{\left\|z_{i}\right\|\left\|e_{k}\right\|}=\frac{z_{i k}}{\left\|x_{S i}\right\|}
$$

for the NPCs, where $e_{k}$ denotes the $k$-th unit vector $e_{k}=(0, \ldots, 1, \ldots, 0)^{\top}$. An individual $i$ will be represented on the $k$-th PC axis if its corresponding angle is small, i.e. if $\cos ^{2} \vartheta_{i k}$ for $k=1, \ldots, p$ is close to one. Note that for each individual $i$,

$$
\sum_{k=1}^{p} \cos ^{2} \vartheta_{i k}=\frac{y_{i}^{\top} y_{i}}{x_{C i}^{\top} x_{C i}}=\frac{x_{C i}^{\top} \mathcal{G} \mathcal{G}^{\top} x_{C i}}{x_{C i}^{\top} x_{C i}}=1
$$

The values $\cos ^{2} \vartheta_{i k}$ are sometimes called the relative contributions of the $k$-th axis to the representation of the $i$-th individual, e.g. if $\cos ^{2} \vartheta_{i 1}+\cos ^{2} \vartheta_{i 2}$ is large (near one), we know that the individual $i$ is well represented on the plane of the first two principal axes since its corresponding angle with the plane is close to zero.

We already know that the quality of the representation of the variables can be evaluated by the percentage of $X_{i}$ 's variance that is explained by a PC, which is given by $r_{X_{i} Y_{j}}^{2}$ or $r_{X_{i} Z_{j}}^{2}$ according to (11.16) and (11.27) respectively.


Fig. 11.7 Representation of the individuals $\mathbf{Q}$ MVAnpcafood

Example 11.6 Let us return to the French food expenditure example, see Sect. 22.6. This yields a two-dimensional representation of the individuals as shown in Fig. 11.7.

Calculating the matrix $\mathcal{G}_{\mathcal{R}}$ we have

$$
\mathcal{G}_{\mathcal{R}}=\left(\begin{array}{rrrrrr}
-0.240 & 0.622 & -0.011 & -0.544 & 0.036 & 0.508 \\
-0.466 & 0.098 & -0.062 & -0.023 & -0.809 & -0.301 \\
-0.446 & -0.205 & 0.145 & 0.548 & -0.067 & 0.625 \\
-0.462 & -0.141 & 0.207 & -0.053 & 0.411 & -0.093 \\
-0.438 & -0.197 & 0.356 & -0.324 & 0.224 & -0.350 \\
-0.281 & 0.523 & -0.444 & 0.450 & 0.341 & -0.332 \\
0.206 & 0.479 & 0.780 & 0.306 & -0.069 & -0.138
\end{array}\right),
$$

which gives the weights of the variables (milk, vegetables, etc.). The eigenvalues $\ell_{j}$ and the proportions of explained variance are given in Table 11.3.

The interpretation of the principal components are best understood when looking at the correlations between the original $X_{i}$ 's and the PCs. Since the first two PCs explain $88.1 \%$ of the variance, we limit ourselves to the first two PCs. The results are shown in Table 11.4. The two-dimensional graphical representation of the variables in Fig. 11.8 is based on the first two columns of Table 11.4.

The plots are the projections of the variables into $\mathbb{R}^{2}$. Since the quality of the representation is good for all the variables (except maybe $X_{7}$ ), their relative angles give a picture of their original correlation: wine is negatively correlated with the

Table 11.3 Eigenvalues and explained variance

| Eigenvalues | Proportion of variance | Cumulated proportion |
| :--- | :--- | :---: |
| 4.333 | 0.6190 | 61.9 |
| 1.830 | 0.2620 | 88.1 |
| 0.631 | 0.0900 | 97.1 |
| 0.128 | 0.0180 | 98.9 |
| 0.058 | 0.0080 | 99.7 |
| 0.019 | 0.0030 | 99.9 |
| 0.001 | 0.0001 | 100.0 |

Table 11.4 Correlations with PCs

|  | $r_{X_{i} Z_{1}}$ | $r_{X_{i} Z_{2}}$ | $r_{X_{i} Z_{1}}^{2}+r_{X_{i} Z_{2}}^{2}$ |
| :--- | ---: | ---: | :--- |
| $X_{1}:$ bread | -0.499 | 0.842 | 0.957 |
| $X_{2}:$ vegetables | -0.970 | 0.133 | 0.958 |
| $X_{3}:$ fruits | -0.929 | -0.278 | 0.941 |
| $X_{4}:$ meat | -0.962 | -0.191 | 0.962 |
| $X_{5}:$ poultry | -0.911 | -0.266 | 0.901 |
| $X_{6}:$ milk | -0.584 | 0.707 | 0.841 |
| $X_{7}:$ wine | 0.428 | 0.648 | 0.604 |

Fig. 11.8 Representation of the variables $\mathbf{Q}$ MVAnpcafood

vegetables, fruits, meat and poultry groups $\left(\theta>90^{\circ}\right)$, whereas taken individually this latter grouping of variables are highly positively correlated with each other $(\theta \approx$ 0 ). Bread and milk are positively correlated but poorly correlated with meat, fruits and poultry ( $\theta \approx 90^{\circ}$ ).

Now the representation of the individuals in Fig. 11.7 can be interpreted better. From Fig. 11.8 and Table 11.4 we can see that the first factor $Z_{1}$ is a vegetable-meat-poultry-fruit factor (with a negative sign), whereas the second
factor is a milk-bread-wine factor (with a positive sign). Note that this corresponds to the most important weights in the first columns of $\mathcal{G}_{\mathcal{R}}$. In Fig. 11.7 lines were drawn to connect families of the same size and families of the same professional types. A grid can clearly be seen (with a slight deformation by the manager families) that shows the families with higher expenditures (higher number of children) on the left.

Considering both figures together explains what types of expenditures are responsible for similarities in food expenditures. Bread, milk and wine expenditures are similar for manual workers and employees. Families of managers are characterised by higher expenditures on vegetables, fruits, meat and poultry. Very often when analysing NPCs (and PCs), it is illuminating to use such a device to introduce qualitative aspects of individuals in order to enrich the interpretations of the graphs.

| $\hookrightarrow$ | NPCs are PCs applied to the standardised (normalised) data matrix <br> $\mathcal{X}_{S}$. |
| :--- | :--- |
|  | The graphical representation of NPCs provides a similar type of <br> picture as that of PCs, the difference being in the relative position <br> of individuals, i.e. each variable in NPCs has the same weight (in <br> PCs, the variable with the largest variance has the largest weight). |
| $\hookrightarrow$ | The quality of the representation is evaluated by $\psi$ <br> $\left(\sum_{j=1}^{p} \ell_{j}\right)^{-1}\left(\ell_{1}+\ell_{2}+\cdots+\ell_{q}\right)$. |
| $\hookrightarrow$ | The quality of the representation of a variable can be evaluated by <br> the percentage of $X_{i}$ 's variance that is explained by a PC, i.e. $r_{X_{i} Y_{j}}^{2}$. |

### 11.7 Common Principal Components

In many applications a statistical analysis is simultaneously done for groups of data. In this section a technique is presented that allows us to analyse group elements that have common PCs. From a statistical point of view, estimating PCs simultaneously in different groups will result in a joint dimension reducing transformation. This multi-group PCA, the so-called common principle components analysis (CPCA), yields the joint eigenstructure across groups.

In addition to traditional PCA, the basic assumption of CPCA is that the space spanned by the eigenvectors is identical across several groups, whereas variances associated with the components are allowed to vary.

More formally, the hypothesis of common principle components can be stated in the following way (Flury, 1988):

$$
H_{\mathrm{CPC}}: \Sigma_{i}=\Gamma \Lambda_{i} \Gamma^{\top}, \quad i=1, \ldots, k
$$

where $\Sigma_{i}$ is a positive definite $p \times p$ population covariance matrix for every $i, \Gamma=\left(\gamma_{1}, \ldots, \gamma_{p}\right)$ is an orthogonal $p \times p$ transformation matrix and $\Lambda_{i}=$ $\operatorname{diag}\left(\lambda_{i 1}, \ldots, \lambda_{i p}\right)$ is the matrix of eigenvalues. Moreover, assume that all $\lambda_{i}$ are distinct.

Let $\mathcal{S}$ be the (unbiased) sample covariance matrix of an underlying $p$-variate normal distribution $N_{p}(\mu, \Sigma)$ with sample size $n$. Then the distribution of $n S$ has $n-1$ degrees of freedom and is known as the Wishart distribution (Muirhead, 1982, p. 86):

$$
n S \sim \mathcal{W}_{p}(\Sigma, n-1)
$$

The density is given in (5.16). Hence, for a given Wishart matrix $\mathcal{S}_{i}$ with sample size $n_{i}$, the likelihood function can be written as

$$
\begin{equation*}
L\left(\Sigma_{1}, \ldots, \Sigma_{k}\right)=C \prod_{i=1}^{k} \exp \left[\operatorname{tr}\left\{-\frac{1}{2}\left(n_{i}-1\right) \Sigma_{i}^{-1} \mathcal{S}_{i}\right\}\right]\left|\Sigma_{i}\right|^{-\frac{1}{2}\left(n_{i}-1\right)} \tag{11.41}
\end{equation*}
$$

where $C$ is a constant independent of the parameters $\Sigma_{i}$. Maximising the likelihood is equivalent to minimising the function

$$
g\left(\Sigma_{1}, \ldots, \Sigma_{k}\right)=\sum_{i=1}^{k}\left(n_{i}-1\right)\left\{\log \left|\Sigma_{i}\right|+\operatorname{tr}\left(\Sigma_{i}^{-1} \mathcal{S}_{i}\right)\right\} .
$$

Assuming that $H_{\mathrm{CPC}}$ holds, i.e. in replacing $\Sigma_{i}$ by $\Gamma \Lambda_{i} \Gamma^{\top}$, after some manipulations one obtains

$$
g\left(\Gamma, \Lambda_{1}, \ldots, \Lambda_{k}\right)=\sum_{i=1}^{k}\left(n_{i}-1\right) \sum_{j=1}^{p}\left(\log \lambda_{i j}+\frac{\gamma_{j}^{\top} \mathcal{S}_{i} \gamma_{j}}{\lambda_{i j}}\right) .
$$

As we know from Sect.2.2, the vectors $\gamma_{j}$ in $\Gamma$ have to be orthogonal. Orthogonality of the vectors $\gamma_{j}$ is achieved using the Lagrange method, i.e. we impose the $p$ constraints $\gamma_{j}^{\top} \gamma_{j}=1$ using the Lagrange multipliers $\mu_{j}$, and the remaining $p(p-1) / 2$ constraints $\gamma_{h}^{\top} \gamma_{j}=0$ for $h \neq j$ using the multiplier $2 \mu_{h j}$ (Flury, 1988). This yields

$$
g^{*}\left(\Gamma, \Lambda_{1}, \ldots, \Lambda_{k}\right)=g(\cdot)-\sum_{j=1}^{p} \mu_{j}\left(\gamma_{j}^{\top} \gamma_{j}-1\right)-2 \sum_{h=1}^{p} \sum_{j=h+1}^{p} \mu_{h j} \gamma_{h}^{\top} \gamma_{j}
$$

Taking partial derivatives with respect to all $\lambda_{i m}$ and $\gamma_{m}$, it can be shown that the solution of the CPC model is given by the generalised system of characteristic equations

$$
\begin{equation*}
\gamma_{m}^{\top}\left\{\sum_{i=1}^{k}\left(n_{i}-1\right) \frac{\lambda_{i m}-\lambda_{i j}}{\lambda_{i m} \lambda_{i j}} \mathcal{S}_{i}\right\} \gamma_{j}=0, \quad m, j=1, \ldots, p, \quad m \neq j \tag{11.42}
\end{equation*}
$$

This system can be solved using

$$
\lambda_{i m}=\gamma_{m}^{\top} \mathcal{S}_{i} \gamma_{m}, \quad i=1, \ldots, k, \quad m=1, \ldots, p
$$

under the constraints

$$
\gamma_{m}^{\top} \gamma_{j}= \begin{cases}0 & m \neq j \\ 1 & m=j\end{cases}
$$

Flury (1988) proves existence and uniqueness of the maximum of the likelihood function, and Flury and Gautschi (1986) provide a numerical algorithm.

Example 11.7 As an example we provide the data sets XFGvolsurf01, XFGvolsurf02 and XFGvolsurf03 that have been used in Fengler, Härdle, and Villa (2003) to estimate common principle components for the implied volatility surfaces of the DAX 1999. The data has been generated by smoothing an implied volatility surface day by day. Next, the estimated grid points have been grouped into maturities of $\tau=1, \tau=2$ and $\tau=3$ months and transformed into a vector of time series of the "smile", i.e. each element of the vector belongs to a distinct moneyness ranging from 0.85 to 1.10 .

Figure 11.9 shows the first three eigenvectors in a parallel coordinate plot. The basic structure of the first three eigenvectors is not altered. We find a shift, a slope and a twist structure. This structure is common to all maturity groups, i.e. when exploiting PCA as a dimension reducing tool, the same transformation applies to each group! However, by comparing the size of eigenvalues among groups we find that variability is decreasing across groups as we move from the short-term contracts to long-term contracts.

Before drawing conclusions we should convince ourselves that the CPC model is truly a good description of the data. This can be done by using a likelihood ratio test. The likelihood ratio statistic for comparing a restricted (the CPC) model against the unrestricted model (the model where all covariances are treated separately) is given by

$$
T_{\left(n_{1}, n_{2}, \ldots, n_{k}\right)}=-2 \log \frac{L\left(\hat{\Sigma}_{1}, \ldots, \hat{\Sigma}_{k}\right)}{L\left(\mathcal{S}_{1}, \ldots, \mathcal{S}_{k}\right)}
$$

Fig. 11.9 Factor loadings of the first (thick), the second (medium), and the third (thin) PCOMVAcpcaiv

PCP for CPCA, 3 eigenvectors


Inserting the likelihood function, we find that this is equivalent to

$$
T_{\left(n_{1}, n_{2}, \ldots, n_{k}\right)}=\sum_{i=1}^{k}\left(n_{i}-1\right) \frac{\operatorname{det}\left(\hat{\Sigma}_{i}\right)}{\operatorname{det}\left(\mathcal{S}_{i}\right)}
$$

which has a $\chi^{2}$ distribution as $\min \left(n_{i}\right)$ tends to infinity with

$$
k\left\{\frac{1}{2} p(p-1)+1\right\}-\left\{\frac{1}{2} p(p-1)+k p\right\}=\frac{1}{2}(k-1) p(p-1)
$$

degrees of freedom. This test is included in the quantlet Q MVAcpcaiv.
The calculations yield $T_{\left(n_{1}, n_{2}, \ldots, n_{k}\right)}=31.836$, which corresponds to the $p$-value $p=0.37512$ for the $\chi^{2}(30)$ distribution. Hence we cannot reject the CPC model against the unrestricted model, where PCA is applied to each maturity separately.

Using the methods in Sect.11.3, we can estimate the amount of variability, $\zeta_{l}$, explained by the first $l$ principal components: (only a few factors, three at the most, are needed to capture a large amount of the total variability present in the data). Since the model now captures the variability in both the strike and maturity dimensions, this is a suitable starting point for a simplified VaR calculation for delta-gamma neutral option portfolios using Monte Carlo methods, and is hence a valuable insight in risk management.

### 11.8 Boston Housing

A set of transformations were defined in Chap. 1 for the Boston Housing data set that resulted in "regular" marginal distributions. The usefulness of principal component analysis with respect to such high-dimensional data sets will now be shown. The variable $X_{4}$ is dropped because it is a discrete $0-1$ variable. It will be used later, however, in the graphical representations. The scale difference of the remaining 13 variables motivates a NPCA based on the correlation matrix.

The eigenvalues and the percentage of explained variance are given in Table 11.5.

The first principal component explains $56 \%$ of the total variance and the first three components together explain more than $75 \%$. These results imply that it is sufficient to look at 2, maximum 3, principal components.

Table 11.6 provides the correlations between the first three PCs and the original variables. These can be seen in Fig. 11.10.

The correlations with the first PC show a very clear pattern. The variables $X_{2}, X_{6}, X_{8}, X_{12}$, and $X_{14}$ are strongly positively correlated with the first PC, whereas the remaining variables are highly negatively correlated. The minimal correlation in the absolute value is 0.5 . The first PC axis could be interpreted as a quality of life and house indicator. The second axis, given the polarities of $X_{11}$ and $X_{13}$ and of $X_{6}$ and $X_{14}$, can be interpreted as a social factor explaining only $10 \%$ of the total variance. The third axis is dominated by a polarity between $X_{2}$ and $X_{12}$.

The set of individuals from the first two PCs can be graphically interpreted if the plots are colour coded with respect to some particular variable of interest.

Table 11.5 Eigenvalues and percentage of explained variance for Boston Housing data Q MVAnpcahousi

| Eigenvalue | Percentages | Cumulated percentages |
| :--- | :--- | :--- |
| 7.2852 | 0.5604 | 0.5604 |
| 1.3517 | 0.1040 | 0.6644 |
| 1.1266 | 0.0867 | 0.7510 |
| 0.7802 | 0.0600 | 0.8111 |
| 0.6359 | 0.0489 | 0.8600 |
| 0.5290 | 0.0407 | 0.9007 |
| 0.3397 | 0.0261 | 0.9268 |
| 0.2628 | 0.0202 | 0.9470 |
| 0.1936 | 0.0149 | 0.9619 |
| 0.1547 | 0.0119 | 0.9738 |
| 0.1405 | 0.0108 | 0.9846 |
| 0.1100 | 0.0085 | 0.9931 |
| 0.0900 | 0.0069 | 1.0000 |

Table 11.6 Correlations of the first three PC's with the original variables $Q$ MVAnpcahous

|  | $\mathrm{PC}_{1}$ | $\mathrm{PC}_{2}$ | $\mathrm{PC}_{3}$ |
| :--- | ---: | ---: | ---: |
| $X_{1}$ | -0.9076 | 0.2247 | 0.1457 |
| $X_{2}$ | 0.6399 | -0.0292 | 0.5058 |
| $X_{3}$ | -0.8580 | 0.0409 | -0.1845 |
| $X_{5}$ | -0.8737 | 0.2391 | -0.1780 |
| $X_{6}$ | 0.5104 | 0.7037 | 0.0869 |
| $X_{7}$ | -0.7999 | 0.1556 | -0.2949 |
| $X_{8}$ | 0.8259 | -0.2904 | 0.2982 |
| $X_{9}$ | -0.7531 | 0.2857 | 0.3804 |
| $X_{10}$ | -0.8114 | 0.1645 | 0.3672 |
| $X_{11}$ | -0.5674 | -0.2667 | 0.1498 |
| $X_{12}$ | 0.4906 | -0.1041 | -0.5170 |
| $X_{13}$ | -0.7996 | -0.4253 | -0.0251 |
| $X_{14}$ | 0.7366 | 0.5160 | -0.1747 |



Fig. 11.10 NPCA for the Boston housing data, correlations of first three PCs with the original variables Q MVAnpcahousi

Figure 11.11 colour codes $X_{14}>$ median as red points. Clearly the first and second PCs are related to house value. The situation is less clear in Fig. 11.12 where the colour code corresponds to $X_{4}$, the Charles River indicator, i.e. houses near the river are coloured red.

Fig. 11.11 NPC analysis for the Boston housing data, scatterplot of the first two PCs. More expensive houses are marked with red colour $\mathbf{Q}$ MVAnpcahous

Fig. 11.12 NPC analysis for the Boston housing data, scatterplot of the first two PCs. Houses close to the Charles River are indicated with red squares $\mathbf{Q}$ MVAnpcahous

First vs. Second PC


### 11.9 More Examples

Example 11.8 Let us now apply the PCA to the standardised bank data set (Sect. 22.2). Figure 11.13 shows some PC plots of the bank data set. The genuine and counterfeit bank notes are marked by " $о$ " and "+", respectively.

The vector of eigenvalues of $\mathcal{R}$ is

$$
\ell=(2.946,1.278,0.869,0.450,0.269,0.189)^{\top}
$$



Fig. 11.13 Principal components of the standardised bank data $Q$ MVAnpcabank
Table 11.7 Eigenvalues and proportions of explained variance

| $\ell_{j}$ | Proportion of variances | Cumulated proportion |
| :--- | :--- | :--- |
| 2.946 | 0.491 | 49.1 |
| 1.278 | 0.213 | 70.4 |
| 0.869 | 0.145 | 84.9 |
| 0.450 | 0.075 | 92.4 |
| 0.264 | 0.045 | 96.9 |
| 0.189 | 0.032 | 100.0 |

The eigenvectors $g_{j}$ are given by the columns of the matrix

$$
\mathcal{G}=\left(\begin{array}{rrrrrr}
-0.007 & -0.815 & 0.018 & 0.575 & 0.059 & 0.031 \\
0.468 & -0.342 & -0.103 & -0.395 & -0.639 & -0.298 \\
0.487 & -0.252 & -0.123 & -0.430 & 0.614 & 0.349 \\
0.407 & 0.266 & -0.584 & 0.404 & 0.215 & -0.462 \\
0.368 & 0.091 & 0.788 & 0.110 & 0.220 & -0.419 \\
-0.493 & -0.274 & -0.114 & -0.392 & 0.340 & -0.632
\end{array}\right) .
$$

Each original variable has the same weight in the analysis and the results are independent of the scale of each variable.

The proportions of explained variance are given in Table 11.7. It can be concluded that the representation in two dimensions should be sufficient. The

Fig. 11.14 The correlations of the original variable with the PCs $\mathbf{Q}$ MVAnpcabanki


Table 11.8 Correlations with PCs

|  | $r_{X_{i} Z_{1}}$ | $r_{X_{i} Z_{2}}$ | $r_{X_{i} Z_{1}}^{2}+r_{X_{i} Z_{2}}^{2}$ |
| :--- | ---: | ---: | :--- |
| $X_{1}:$ length | -0.012 | -0.922 | 0.85 |
| $X_{2}:$ left height | 0.803 | -0.387 | 0.79 |
| $X_{3}:$ right height | 0.835 | -0.285 | 0.78 |
| $X_{4}:$ lower | 0.698 | 0.301 | 0.58 |
| $X_{5}:$ upper | 0.631 | 0.104 | 0.41 |
| $X_{6}:$ diagonal | -0.847 | -0.310 | 0.81 |

correlations leading to Fig. 11.14 are given in Table 11.8. The picture is different from the one obtained in Sect. 11.3 (see Table 11.2). Here, the first factor is mainly a left-right vs. diagonal factor and the second one is a length factor (with negative weight). Take another look at Fig. 11.13, where the individual bank notes are displayed. In the upper left graph it can be seen that the genuine bank notes are for the most part in the south-eastern portion of the graph featuring a larger diagonal, smaller height $\left(Z_{1}<0\right)$ and also a larger length $\left(Z_{2}<0\right)$. Note also that Fig. 11.14 gives an idea of the correlation structure of the original data matrix.

Example 11.9 Consider the data of 79 US companies given in Table 22.5. The data is first standardised by subtracting the mean and dividing by the standard deviation. Note that the data set contains six variables: assets $\left(X_{1}\right)$, sales $\left(X_{2}\right)$, market value $\left(X_{3}\right)$, profits $\left(X_{4}\right)$, cash flow $\left(X_{5}\right)$, number of employees $\left(X_{6}\right)$.

Calculating the corresponding vector of eigenvalues gives

$$
\ell=(5.039,0.517,0.359,0.050,0.029,0.007)^{\top}
$$



Fig. 11.15 Principal components of the US company data $Q$ MVAnpcausco
and the matrix of eigenvectors is

$$
\mathcal{G}=\left(\begin{array}{rrrrrr}
0.340 & -0.849 & -0.339 & 0.205 & 0.077 & -0.006 \\
0.423 & -0.170 & 0.379 & -0.783 & -0.006 & -0.186 \\
0.434 & 0.190 & -0.192 & 0.071 & -0.844 & 0.149 \\
0.420 & 0.364 & -0.324 & 0.156 & 0.261 & -0.703 \\
0.428 & 0.285 & -0.267 & -0.121 & 0.452 & 0.667 \\
0.397 & 0.010 & 0.726 & 0.548 & 0.098 & 0.065
\end{array}\right) .
$$

Using this information the graphical representations of the first two principal components are given in Fig. 11.15. The different sectors are marked by the following symbols:

| H | $\ldots$ | Hi Tech and Communication |
| ---: | :--- | :--- |
| E | $\ldots$ | Energy |
| F | $\ldots$ | Finance |
| M | $\ldots$ | Manufacturing |
| R | $\ldots$ | Retail |
| $\star$ | $\ldots$ | all other sectors. |

The two outliers in the right-hand side of the graph are IBM and General Electric (GE), which differ from the other companies with their high market values. As can


Fig. 11.16 Principal components of the US company data (without IBM and General Electric) $\mathbf{Q}$ MVAnpcausco2
be seen in the first column of $\mathcal{G}$, market value has the largest weight in the first PC, adding to the isolation of these two companies. If IBM and GE were to be excluded from the data set, a completely different picture would emerge, as shown in Fig. 11.16. In this case the vector of eigenvalues becomes

$$
\ell=(3.191,1.535,0.791,0.292,0.149,0.041)^{\top}
$$

and the corresponding matrix of eigenvectors is

$$
\mathcal{G}=\left(\begin{array}{rrrrrr}
0.263 & -0.408 & -0.800 & -0.067 & 0.333 & 0.099 \\
0.438 & -0.407 & 0.162 & -0.509 & -0.441 & -0.403 \\
0.500 & -0.003 & -0.035 & 0.801 & -0.264 & -0.190 \\
0.331 & 0.623 & -0.080 & -0.192 & 0.426 & -0.526 \\
0.443 & 0.450 & -0.123 & -0.238 & -0.335 & 0.646 \\
0.427 & -0.277 & 0.558 & 0.021 & 0.575 & 0.313
\end{array}\right) .
$$

The percentage of variation explained by each component is given in Table 11.9. The first two components explain almost $79 \%$ of the variance. The interpretation of the factors (the axes of Fig. 11.16) is given in the table of correlations (Table 11.10). The first two columns of this table are plotted in Fig. 11.17.

Table 11.9 Eigenvalues and proportions of explained variance

| $\ell_{j}$ | Proportion of variance | Cumulated proportion |
| :--- | :--- | :--- |
| 3.191 | 0.532 | 0.532 |
| 1.535 | 0.256 | 0.788 |
| 0.791 | 0.132 | 0.920 |
| 0.292 | 0.049 | 0.968 |
| 0.149 | 0.025 | 0.993 |
| 0.041 | 0.007 | 1.000 |

Table 11.10 Correlations with PCs

|  | $r_{X_{i} Z_{1}}$ | $r_{X_{i} Z_{2}}$ | $r_{X_{i} Z_{1}}^{2}+r_{X_{i} Z_{2}}^{2}$ |
| :--- | :--- | :--- | :--- |
| $X_{1}:$ assets | 0.47 | -0.510 | 0.48 |
| $X_{2}:$ sales | 0.78 | -0.500 | 0.87 |
| $X_{3}:$ market value | 0.89 | -0.003 | 0.80 |
| $X_{4}:$ profits | 0.59 | 0.770 | 0.95 |
| $X_{5}:$ cash flow | 0.79 | 0.560 | 0.94 |
| $X_{6}:$ employees | 0.76 | -0.340 | 0.70 |

U.S. Company Data


Fig. 11.17 The correlation of the original variables with the PCs $Q$ MVAnpcausco2i

From Fig. 11.17 (and Table 11.10) it appears that the first factor is a "size effect", it is positively correlated with all the variables describing the size of the activity of the companies. It is also a measure of the economic strength of the firms. The second factor describes the "shape" of the companies ("profit-cash flow" vs. "assets-sales" factor), which is more difficult to interpret from an economic point of view.

Example 11.10 Volle (1985) analyses data on 28 individuals (Table 22.14). For each individual, the time spent (in hours) on 10 different activities has been recorded over 100 days, as well as informative statistics such as the individual's sex, country of residence, professional activity and matrimonial status. The results of a NPCA are given below.

The eigenvalues of the correlation matrix are given in Table 11.11. Note that the last eigenvalue is exactly zero since the correlation matrix is singular (the sum of all the variables is always equal to $2,400=24 \times 100$ ). The results of the 4 first PCs are given in Tables 11.12 and 11.13.

From these tables (and Figs. 11.18 and 11.19), it appears that the professional and household activities are strongly contrasted in the first factor. Indeed on the horizontal axis of Fig. 11.18 it can be seen that all the active men are on the right and all the inactive women are on the left. Active women and/or single women are in between. The second factor contrasts meal/sleeping vs. toilet/shopping (note the high correlation between meal and sleeping). Along the vertical axis of Fig. 11.18 we see near the bottom of the graph the people from Western-European countries, who spend more time on meals and sleeping than people from the US (who can be found close to the top of the graph). The other categories are in between.

Table 11.11 Eigenvalues of correlation matrix for the time budget data

| $\ell_{j}$ | Proportion of variance | Cumulated proportion |
| :--- | :--- | :--- |
| 4.59 | 0.459 | 0.460 |
| 2.12 | 0.212 | 0.670 |
| 1.32 | 0.132 | 0.800 |
| 1.20 | 0.120 | 0.920 |
| 0.47 | 0.047 | 0.970 |
| 0.20 | 0.020 | 0.990 |
| 0.05 | 0.005 | 0.990 |
| 0.04 | 0.004 | 0.999 |
| 0.02 | 0.002 | 1.000 |
| 0.00 | 0.000 | 1.000 |

Table 11.12 Correlation of variables with PCs

|  |  | $r_{X_{i} W_{1}}$ | $r_{X_{i} W_{2}}$ | $r_{X_{i} W_{3}}$ | $r_{X_{i} W_{4}}$ |
| :--- | :--- | ---: | ---: | ---: | ---: |
| $X_{1}:$ | prof | 0.9772 | -0.1210 | -0.0846 | 0.0669 |
| $X_{2}:$ | tran | 0.9798 | 0.0581 | -0.0084 | 0.4555 |
| $X_{3}:$ | hous | -0.8999 | 0.0227 | 0.3624 | 0.2142 |
| $X_{4}:$ | kids | -0.8721 | 0.1786 | 0.0837 | 0.2944 |
| $X_{5}:$ | shop | -0.5636 | 0.7606 | -0.0046 | -0.1210 |
| $X_{6}:$ | pers | -0.0795 | 0.8181 | -0.3022 | -0.0636 |
| $X_{7}:$ | eati | -0.5883 | -0.6694 | -0.4263 | 0.0141 |
| $X_{8}:$ | slee | -0.6442 | -0.5693 | -0.1908 | -0.3125 |
| $X_{9}:$ | tele | -0.0994 | 0.1931 | -0.9300 | 0.1512 |
| $X_{10}:$ | leis | -0.0922 | 0.1103 | 0.0302 | -0.9574 |

Table 11.13 PCs for time budget data

|  | $Z_{1}$ | $Z_{2}$ | $Z_{3}$ | $Z_{4}$ |
| :--- | ---: | ---: | :--- | ---: |
| maus | 0.0633 | 0.0245 | -0.0668 | 0.0205 |
| waus | 0.0061 | 0.0791 | -0.0236 | 0.0156 |
| wnus | -0.1448 | 0.0813 | -0.0379 | -0.0186 |
| mmus | 0.0635 | 0.0105 | -0.0673 | 0.0262 |
| wmus | -0.0934 | 0.0816 | -0.0285 | 0.0038 |
| msus | 0.0537 | 0.0676 | -0.0487 | -0.0279 |
| wsus | 0.0166 | 0.1016 | -0.0463 | -0.0053 |
| mawe | 0.0420 | -0.0846 | -0.0399 | -0.0016 |
| wawe | -0.0111 | -0.0534 | -0.0097 | 0.0337 |
| wnwe | -0.1544 | -0.0583 | -0.0318 | -0.0051 |
| mmwe | 0.0402 | -0.0880 | -0.0459 | 0.0054 |
| wmwe | -0.1118 | -0.0710 | -0.0210 | 0.0262 |
| mswe | 0.0489 | -0.0919 | -0.0188 | -0.0365 |
| wswe | -0.0393 | -0.0591 | -0.0194 | -0.0534 |
| mayo | 0.0772 | -0.0086 | 0.0253 | -0.0085 |
| wayo | 0.0359 | 0.0064 | 0.0577 | 0.0762 |
| wnyo | -0.1263 | -0.0135 | 0.0584 | -0.0189 |
| mmyo | 0.0793 | -0.0076 | 0.0173 | -0.0039 |
| wmyo | -0.0550 | -0.0077 | 0.0579 | 0.0416 |
| msyo | 0.0763 | 0.0207 | 0.0575 | -0.0778 |
| wsyo | 0.0120 | 0.0149 | 0.0532 | -0.0366 |
| maes | 0.0767 | -0.0025 | 0.0047 | 0.0115 |
| waes | 0.0353 | 0.0209 | 0.0488 | 0.0729 |
| wnes | -0.1399 | 0.0016 | 0.0240 | -0.0348 |
| mmes | 0.0742 | -0.0061 | -0.0152 | 0.0283 |
| wmes | -0.0175 | 0.0073 | 0.0429 | 0.0719 |
| mses | 0.0903 | 0.0052 | 0.0379 | -0.0701 |
| fses | 0.0020 | 0.0287 | 0.0358 | -0.0346 |
|  |  |  |  |  |

In Fig. 11.19 the variables television and other leisure activities hardly play any role (look at Table 11.12). The variable television appears in $Z_{3}$ (negatively correlated). Table 11.13 shows that this factor contrasts people from Eastern countries and Yugoslavia with men living in the US The variable other leisure activities is the factor $Z_{4}$. It merely distinguishes between men and women in Eastern countries and in Yugoslavia. These last two factors are orthogonal to the preceding axes and of course their contribution to the total variation is less important.


Fig. 11.18 Representation of the individuals $Q$ MVAnpcatime


Fig. 11.19 Representation of the variables $Q$ MVAnpcatime

### 11.10 Exercises

Exercise 11.1 Prove Theorem 11.1. (Hint: use (4.23).)
Exercise 11.2 Interpret the results of the PCA of the US companies. Use the analysis of the bank notes in Sect. 11.3 as a guide. Compare your results with those in Example 11.9.

Exercise 11.3 Test the hypothesis that the proportion of variance explained by the first two PCs for the US companies is $\psi=0.75$.

Exercise 11.4 Apply the PCA to the car data (Table 22.7). Interpret the first two PCs. Would it be necessary to look at the third PC?

Exercise 11.5 Take the athletic records for 55 countries (Sect. 22.18) and apply the NPCA. Interpret your results.

Exercise 11.6 Apply a PCA to $\Sigma=\left(\begin{array}{ll}1 & \rho \\ \rho & 1\end{array}\right)$, where $\rho>0$. Now change the scale of $X_{1}$, i.e. consider the covariance of $c X_{1}$ and $X_{2}$. How do the PC directions change with the screeplot?

Exercise 11.7 Suppose that we have standardised some data using the Mahalanobis transformation. Would it be reasonable to apply a PCA?

Exercise 11.8 Apply a NPCA to the US CRIME data set (Table 22.10). Interpret the results. Would it be necessary to look at the third PC? Can you see any difference between the four regions? Redo the analysis excluding the variable "area of the state".

Exercise 11.9 Repeat Exercise 11.8 using the US HEALTH data set (Table 22.16).
Exercise 11.10 Do a NPCA on the GEOPOL data set (see Table 22.15) which compares 41 countries w.r.t. different aspects of their development. Why or why not would a PCA be reasonable here?

Exercise 11.11 Let $U$ be an uniform r.v. on $[0,1]$. Let $a \in \mathbb{R}^{3}$ be a vector of constants. Suppose that $X=U a^{\top}=\left(X_{1}, X_{2}, X_{3}\right)$. What do you expect the NPCs of $X$ to be?

Exercise 11.12 Let $U_{1}$ and $U_{2}$ be two independent uniform random variables on $[0,1]$. Suppose that $X=\left(X_{1}, X_{2}, X_{3}, X_{4}\right)^{\top}$ where $X_{1}=U_{1}, X_{2}=U_{2}, X_{3}=$ $U_{1}+U_{2}$ and $X_{4}=U_{1}-U_{2}$. Compute the correlation matrix $P$ of $X$. How many PCs are of interest? Show that $\gamma_{1}=\left(\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}}, 1,0\right)^{\top}$ and $\gamma_{2}=\left(\frac{1}{\sqrt{2}}, \frac{-1}{\sqrt{2}}, 0,1\right)^{\top}$ are eigenvectors of $P$ corresponding to the non trivial $\lambda$ 's. Interpret the first two NPCs obtained.

Exercise 11.13 Simulate a sample of size $n=50$ for the r.v. $X$ in Exercise 11.12 and analyse the results of a NPCA.

Exercise 11.14 Bouroche and Saporta (1980) reported the data on the state expenses of France from the period 1872 to 1971 (24 selected years) by noting the percentage of 11 categories of expenses. Do a NPCA of this data set. Do the three main periods (before WWI, between WWI and WWII, and after WWII) indicate a change in behaviour w.r.t. state expenses?

## Chapter 12 <br> Factor Analysis

A frequently applied paradigm in analysing data from multivariate observations is to model the relevant information (represented in a multivariate variable $X$ ) as coming from a limited number of latent factors. In a survey on household consumption, for example, the consumption levels, $X$, of $p$ different goods during 1 month could be observed. The variations and covariations of the $p$ components of $X$ throughout the survey might in fact be explained by two or three main social behaviour factors of the household. For instance, a basic desire of comfort or the willingness to achieve a certain social level or other social latent concepts might explain most of the consumption behaviour. These unobserved factors are much more interesting to the social scientist than the observed quantitative measures $(X)$ themselves, because they give a better understanding of the behaviour of households. As shown in the examples below, the same kind of factor analysis is of interest in many fields such as psychology, marketing, economics, and politic sciences.

How can we provide a statistical model addressing these issues and how can we interpret the obtained model? This is the aim of factor analysis. As in Chaps. 10 and 11, the driving statistical theme of this chapter is to reduce the dimension of the observed data. The perspective used, however, is different: we assume that there is a model (it will be called the "Factor Model") stating that most of the covariances between the $p$ elements of $X$ can be explained by a limited number of latent factors. Section 12.1 defines the basic concepts and notations of the orthogonal factor model, stressing the non-uniqueness of the solutions. We show how to take advantage of this non-uniqueness to derive techniques which lead to easier interpretations. This will involve (geometric) rotations of the factors. Section 12.2 presents an empirical approach to factor analysis. Various estimation procedures are proposed and an optimal rotation procedure is defined. Many examples are used to illustrate the method.

### 12.1 The Orthogonal Factor Model

The aim of factor analysis is to explain the outcome of $p$ variables in the data matrix $\mathcal{X}$ using fewer variables, the so-called factors. Ideally all the information in $\mathcal{X}$ can be reproduced by a smaller number of factors. These factors are interpreted as latent (unobserved) common characteristics of the observed $x \in \mathbb{R}^{p}$. The case just described occurs when every observed $x=\left(x_{1}, \ldots, x_{p}\right)^{\top}$ can be written as

$$
\begin{equation*}
x_{j}=\sum_{\ell=1}^{k} q_{j \ell} f_{\ell}+\mu_{j}, j=1, \ldots, p \tag{12.1}
\end{equation*}
$$

Here $f_{\ell}$, for $\ell=1, \ldots, k$ denotes the factors. The number of factors, $k$, should always be much smaller than $p$. For instance, in psychology $x$ may represent $p$ results of a test measuring intelligence scores. One common latent factor explaining $x \in \mathbb{R}^{p}$ could be the overall level of "intelligence". In marketing studies, $x$ may consist of $p$ answers to a survey on the levels of satisfaction of the customers. These $p$ measures could be explained by common latent factors like the attraction level of the product or the image of the brand, and so on. Indeed it is possible to create a representation of the observations that is similar to the one in (12.1) by means of principal components, but only if the last $p-k$ eigenvalues corresponding to the covariance matrix are equal to zero. Consider a $p$-dimensional random vector X with mean $\mu$ and covariance matrix $\operatorname{Var}(X)=\Sigma$. A model similar to (12.1) can be written for $X$ in matrix notation, namely

$$
\begin{equation*}
X=\mathcal{Q} F+\mu, \tag{12.2}
\end{equation*}
$$

where $F$ is the $k$-dimensional vector of the $k$ factors. When using the factor model (12.2) it is often assumed that the factors $F$ are centred, uncorrelated and standardised: $\mathrm{E}(F)=0$ and $\operatorname{Var}(F)=\mathcal{I}_{k}$. We will now show that if the last $p-k$ eigenvalues of $\Sigma$ are equal to zero, we can easily express $X$ by the factor model (12.2).

The spectral decomposition of $\Sigma$ is given by $\Gamma \Lambda \Gamma^{\top}$. Suppose that only the first $k$ eigenvalues are positive, i.e. $\lambda_{k+1}=\cdots=\lambda_{p}=0$. Then the (singular) covariance matrix can be written as

$$
\Sigma=\sum_{\ell=1}^{k} \lambda_{\ell} \gamma_{\ell} \gamma_{\ell}^{\top}=\left(\Gamma_{1} \Gamma_{2}\right)\left(\begin{array}{cc}
\Lambda_{1} & 0 \\
0 & 0
\end{array}\right)\binom{\Gamma_{1}^{\top}}{\Gamma_{2}^{\top}}
$$

In order to show the connection to the factor model (12.2), recall that the PCs are given by $Y=\Gamma^{\top}(X-\mu)$. Rearranging we have $X-\mu=\Gamma Y=\Gamma_{1} Y_{1}+\Gamma_{2} Y_{2}$, where the components of $Y$ are partitioned according to the partition of $\Gamma$ above, namely

$$
Y=\binom{Y_{1}}{Y_{2}}=\binom{\Gamma_{1}^{\top}}{\Gamma_{2}^{\top}}(X-\mu), \text { where }\binom{\Gamma_{1}^{\top}}{\Gamma_{2}^{\top}}(X-\mu) \sim\left(0,\left(\begin{array}{cc}
\Lambda_{1} & 0 \\
0 & 0
\end{array}\right)\right) .
$$

In other words, $Y_{2}$ has a singular distribution with mean and covariance matrix equal to zero. Therefore, $X-\mu=\Gamma_{1} Y_{1}+\Gamma_{2} Y_{2}$ implies that $X-\mu$ is equivalent to $\Gamma_{1} Y_{1}$, which can be written as

$$
X=\Gamma_{1} \Lambda_{1}^{1 / 2} \Lambda_{1}^{-1 / 2} Y_{1}+\mu
$$

Defining $\mathcal{Q}=\Gamma_{1} \Lambda_{1}^{1 / 2}$ and $F=\Lambda_{1}^{-1 / 2} Y_{1}$, we obtain the factor model (12.2).
Note that the covariance matrix of model (12.2) can be written as

$$
\begin{equation*}
\Sigma=\mathrm{E}(X-\mu)(X-\mu)^{\top}=\mathcal{Q} \mathrm{E}\left(F F^{\top}\right) \mathcal{Q}^{\top}=\mathcal{Q} \mathcal{Q}^{\top}=\sum_{j=1}^{k} \lambda_{j} \gamma_{j} \gamma_{j}^{\top} \tag{12.3}
\end{equation*}
$$

We have just shown how the variable $X$ can be completely determined by a weighted sum of $k$ (where $k<p$ ) uncorrelated factors. The situation used in the derivation, however, is too idealistic. In practice the covariance matrix is rarely singular.

It is a common praxis in factor analysis to split the influences of the factors into common and specific ones. There are, for example, highly informative factors that are common to all of the components of $X$ and factors that are specific to certain components. The factor analysis model used in praxis is a generalisation of (12.2):

$$
\begin{equation*}
X=\mathcal{Q} F+U+\mu \tag{12.4}
\end{equation*}
$$

where $\mathcal{Q}$ is a $(p \times k)$ matrix of the (non-random) loadings of the common factors $F(k \times 1)$ and $U$ is a $(p \times 1)$ matrix of the (random) specific factors. It is assumed that the factor variables $F$ are uncorrelated random vectors and that the specific factors are uncorrelated and have zero covariance with the common factors. More precisely, it is assumed that:

$$
\begin{align*}
\mathrm{E} F & =0, \\
\operatorname{Var}(F) & =\mathcal{I}_{k}, \\
\mathrm{E} U & =0,  \tag{12.5}\\
\operatorname{Cov}\left(U_{i}, U_{j}\right) & =0, \quad i \neq j \\
\operatorname{Cov}(F, U) & =0 .
\end{align*}
$$

Define

$$
\operatorname{Var}(U)=\Psi=\operatorname{diag}\left(\psi_{11}, \ldots, \psi_{p p}\right)
$$

The generalised factor model (12.4) together with the assumptions given in (12.5) constitute the orthogonal factor model.


Note that (12.4) implies for the components of $X=\left(X_{1}, \ldots, X_{p}\right)^{\top}$ that

$$
\begin{equation*}
X_{j}=\sum_{\ell=1}^{k} q_{j \ell} F_{\ell}+U_{j}+\mu_{j}, \quad j=1, \ldots, p \tag{12.6}
\end{equation*}
$$

Using (12.5) we obtain $\sigma_{X_{j} X_{j}}=\operatorname{Var}\left(X_{j}\right)=\sum_{\ell=1}^{k} q_{j \ell}^{2}+\psi_{j j}$. The quantity $h_{j}^{2}=\sum_{\ell=1}^{k} q_{j \ell}^{2}$ is called the communality and $\psi_{j j}$ the specific variance. Thus the covariance of $X$ can be rewritten as

$$
\begin{align*}
\Sigma & =\mathrm{E}(X-\mu)(X-\mu)^{\top}=\mathrm{E}(\mathcal{Q} F+U)(\mathcal{Q} F+U)^{\top} \\
& =\mathcal{Q} \mathrm{E}\left(F F^{\top}\right) \mathcal{Q}^{\top}+\mathrm{E}\left(U U^{\top}\right)=\mathcal{Q} \operatorname{Var}(F) \mathcal{Q}^{\top}+\operatorname{Var}(U) \\
& =\mathcal{Q} \mathcal{Q}^{\top}+\Psi . \tag{12.7}
\end{align*}
$$

In a sense, the factor model explains the variations of $X$ for the most part by a small number of latent factors $F$ common to its $p$ components and entirely explains all the correlation structure between its components, plus some "noise" $U$ which allows specific variations of each component to enter. The specific factors adjust to capture the individual variance of each component. Factor analysis relies on the assumptions presented above. If the assumptions are not met, the analysis could be spurious. Although principal components analysis and factor analysis might be related (this was hinted at in the derivation of the factor model), they are quite different in nature. PCs are linear transformations of $X$ arranged in decreasing order of variance and used to reduce the dimension of the data set, whereas in factor analysis, we try to model the variations of $X$ using a linear transformation of a fixed, limited number of latent factors. The objective of factor analysis is to find the loadings $\mathcal{Q}$ and the specific variance $\Psi$. Estimates of $\mathcal{Q}$ and $\Psi$ are deduced from the covariance structure (12.7).

## Interpretation of the Factors

Assume that a factor model with $k$ factors was found to be reasonable, i.e. most of the (co)variations of the $p$ measures in $X$ were explained by the $k$ fixed latent factors. The next natural step is to try to understand what these factors represent. To interpret $F_{\ell}$, it makes sense to compute its correlations with the original variables $X_{j}$ first. This is done for $\ell=1, \ldots, k$ and for $j=1, \ldots, p$ to obtain the matrix $P_{X F}$. The sequence of calculations used here are in fact the same that were used to interpret the PCs in the principal components analysis.

The following covariance between $X$ and $F$ is obtained via (12.5),

$$
\Sigma_{X F}=\mathrm{E}\left\{(\mathcal{Q} F+U) F^{\top}\right\}=\mathcal{Q}
$$

The correlation is

$$
\begin{equation*}
P_{X F}=D^{-1 / 2} \mathcal{Q} \tag{12.8}
\end{equation*}
$$

where $D=\operatorname{diag}\left(\sigma_{X_{1} X_{1}}, \ldots, \sigma_{X_{p} X_{p}}\right)$. Using (12.8) it is possible to construct a figure analogous to Fig. 11.6 and thus to consider which of the original variables $X_{1}, \ldots, X_{p}$ play a role in the unobserved common factors $F_{1}, \ldots, F_{k}$.

Returning to the psychology example where $X$ are the observed scores to $p$ different intelligence tests (the WAIS data set in Table 22.12 provides an example), we would expect a model with one factor to produce a factor that is positively correlated with all of the components in $X$. For this example the factor represents the overall level of intelligence of an individual. A model with two factors could produce a refinement in explaining the variations of the $p$ scores. For example, the first factor could be the same as before (overall level of intelligence), whereas the second factor could be positively correlated with some of the tests, $X_{j}$, that are related to the individual's ability to think abstractly and negatively correlated with other tests, $X_{i}$, that are related to the individual's practical ability. The second factor would then concern a particular dimension of the intelligence stressing the distinctions between the "theoretical" and "practical" abilities of the individual. If the model is true, most of the information coming from the $p$ scores can be summarised by these two latent factors. Other practical examples are given below.

## Invariance of Scale

What happens if we change the scale of $X$ to $Y=\mathcal{C} X$ with $\mathcal{C}=\operatorname{diag}\left(c_{1}, \ldots, c_{p}\right)$ ? If the $k$-factor model (12.6) is true for $X$ with $\mathcal{Q}=\mathcal{Q}_{X}, \Psi=\Psi_{X}$, then, since

$$
\operatorname{Var}(Y)=\mathcal{C} \Sigma \mathcal{C}^{\top}=\mathcal{C} \mathcal{Q}_{X} \mathcal{Q}_{X}^{\top} \mathcal{C}^{\top}+\mathcal{C} \Psi_{X} \mathcal{C}^{\top}
$$

the same $k$-factor model is also true for $Y$ with $\mathcal{Q}_{Y}=\mathcal{C} \mathcal{Q}_{X}$ and $\Psi_{Y}=\mathcal{C} \Psi_{X} \mathcal{C}^{\top}$. In many applications, the search for the loadings $\mathcal{Q}$ and for the specific variance $\Psi$ will be done by the decomposition of the correlation matrix of $X$ rather than the covariance matrix $\Sigma$. This corresponds to a factor analysis of a linear transformation of $X$ (i.e. $Y=D^{-1 / 2}(X-\mu)$ ). The goal is to try to find the loadings $\mathcal{Q}_{Y}$ and the specific variance $\Psi_{Y}$ such that

$$
\begin{equation*}
P=\mathcal{Q}_{Y} \mathcal{Q}_{Y}^{\top}+\Psi_{Y} \tag{12.9}
\end{equation*}
$$

In this case the interpretation of the factors $F$ immediately follows from (12.8) given the following correlation matrix:

$$
\begin{equation*}
P_{X F}=P_{Y F}=\mathcal{Q}_{Y} . \tag{12.10}
\end{equation*}
$$

Because of the scale invariance of the factors, the loadings and the specific variance of the model, where $X$ is expressed in its original units of measure, are given by

$$
\begin{aligned}
& \mathcal{Q}_{X}=D^{1 / 2} \mathcal{Q}_{Y} \\
& \Psi_{X}=D^{1 / 2} \Psi_{Y} D^{1 / 2}
\end{aligned}
$$

It should be noted that although the factor analysis model (12.4) enjoys the scale invariance property, the actual estimated factors could be scale dependent. We will come back to this point later when we discuss the method of principal factors.

## Non-uniqueness of Factor Loadings

The factor loadings are not unique! Suppose that $\mathcal{G}$ is an orthogonal matrix. Then $X$ in (12.4) can also be written as

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

This implies that, if a $k$-factor of $X$ with factors $F$ and loadings $\mathcal{Q}$ is true, then the $k$-factor model with factors $\mathcal{G}^{\top} F$ and loadings $\mathcal{Q G}$ is also true. In practice, we will take advantage of this non-uniqueness. Indeed, referring back to Sect. 2.6 we can conclude that premultiplying a vector $F$ by an orthogonal matrix corresponds to a rotation of the system of axis, the direction of the first new axis being given by the first row of the orthogonal matrix. It will be shown that choosing an appropriate rotation will result in a matrix of loadings $\mathcal{Q G}$ that will be easier to interpret. We have seen that the loadings provide the correlations between the factors and the original variables; therefore, it makes sense to search for rotations that give factors that are maximally correlated with various groups of variables.

From a numerical point of view, the non-uniqueness is a drawback. We have to find loadings $\mathcal{Q}$ and specific variances $\Psi$ satisfying the decomposition $\Sigma=$ $\mathcal{Q} \mathcal{Q}^{\top}+\Psi$, but no straightforward numerical algorithm can solve this problem due to the multiplicity of the solutions. An acceptable technique is to impose some chosen constraints in order to get-in the best case-an unique solution to the decomposition. Then, as suggested above, once we have a solution we will take advantage of the rotations in order to obtain a solution that is easier to interpret.

An obvious question is: what kind of constraints should we impose in order to eliminate the non-uniqueness problem? Usually, we impose additional constraints where

$$
\begin{equation*}
\mathcal{Q}^{\top} \Psi^{-1} \mathcal{Q} \quad \text { is diagonal } \tag{12.11}
\end{equation*}
$$

or

$$
\begin{equation*}
\mathcal{Q}^{\top} \mathcal{D}^{-1} \mathcal{Q} \quad \text { is diagonal. } \tag{12.12}
\end{equation*}
$$

How many parameters does the model (12.7) have without constraints?

$$
\begin{aligned}
& \mathcal{Q}(p \times k) \quad \text { has } \quad p \cdot k \quad \text { parameters, and } \\
& \Psi(p \times p) \text { has } \quad p \text { parameters. }
\end{aligned}
$$

Hence we have to determine $p k+p$ parameters! Conditions (12.11) respectively (12.12) introduce $\frac{1}{2}\{k(k-1)\}$ constraints, since we require the matrices to be diagonal. Therefore, the degrees of freedom of a model with $k$ factors is:

$$
\begin{aligned}
d & =(\# \text { parameters for } \Sigma \text { unconstrained })-(\# \text { parameters for } \Sigma \text { constrained }) \\
& =\frac{1}{2} p(p+1)-\left(p k+p-\frac{1}{2} k(k-1)\right) \\
& =\frac{1}{2}(p-k)^{2}-\frac{1}{2}(p+k) .
\end{aligned}
$$

If $d<0$, then the model is undetermined: there are infinitely many solutions to (12.7). This means that the number of parameters of the factorial model is larger than the number of parameters of the original model, or that the number of factors $k$ is "too large" relative to $p$. In some cases $d=0$ : there is a unique solution to the problem (except for rotation). In practice we usually have that $d>0$ : there are more equations than parameters, thus an exact solution does not exist. In this case approximate solutions are used. An approximation of $\Sigma$, for example, is $\mathcal{Q} \mathcal{Q}^{\top}+\Psi$. The last case is the most interesting since the factorial model has less parameters than the original one. Estimation methods are introduced in the next section.

Evaluating the degrees of freedom, $d$, is particularly important, because it already gives an idea of the upper bound on the number of factors we can hope to identify in a factor model. For instance, if $p=4$, we could not identify a factor model with two factors (this results in $d=-1$ which has infinitely many solutions). With $p=4$,
only a one factor model gives an approximate solution $(d=2)$. When $p=6$, models with 1 and 2 factors provide approximate solutions and a model with three factors results in an unique solution (up to the rotations) since $d=0$. A model with four or more factors would not be allowed, but of course, the aim of factor analysis is to find suitable models with a small number of factors, i.e. smaller than $p$. The next two examples give more insights into the notion of degrees of freedom.
Example 12.1 Let $p=3$ and $k=1$, then $d=0$ and

$$
\Sigma=\left(\begin{array}{ccc}
\sigma_{11} & \sigma_{12} & \sigma_{13} \\
\sigma_{21} & \sigma_{22} & \sigma_{23} \\
\sigma_{31} & \sigma_{32} & \sigma_{33}
\end{array}\right)=\left(\begin{array}{ccc}
q_{1}^{2}+\psi_{11} & q_{1} q_{2} & q_{1} q_{3} \\
q_{1} q_{2} & q_{2}^{2}+\psi_{22} & q_{2} q_{3} \\
q_{1} q_{3} & q_{2} q_{3} & q_{3}^{2}+\psi_{33}
\end{array}\right)
$$

with $\mathcal{Q}=\left(\begin{array}{l}q_{1} \\ q_{2} \\ q_{3}\end{array}\right)$ and $\Psi=\left(\begin{array}{ccc}\psi_{11} & 0 & 0 \\ 0 & \psi_{22} & 0 \\ 0 & 0 & \psi_{33}\end{array}\right)$. Note that here the constraint (12.11) is automatically verified since $k=1$. We have

$$
q_{1}^{2}=\frac{\sigma_{12} \sigma_{13}}{\sigma_{23}} ; q_{2}^{2}=\frac{\sigma_{12} \sigma_{23}}{\sigma_{13}} ; q_{3}^{2}=\frac{\sigma_{13} \sigma_{23}}{\sigma_{12}}
$$

and

$$
\psi_{11}=\sigma_{11}-q_{1}^{2} ; \psi_{22}=\sigma_{22}-q_{2}^{2} ; \psi_{33}=\sigma_{33}-q_{3}^{2}
$$

In this particular case ( $k=1$ ), the only rotation is defined by $\mathcal{G}=-1$, so the other solution for the loadings is provided by $-\mathcal{Q}$.

Example 12.2 Suppose now $p=2$ and $k=1$, then $d<0$ and

$$
\Sigma=\left(\begin{array}{ll}
1 & \rho \\
\rho & 1
\end{array}\right)=\left(\begin{array}{cc}
q_{1}^{2}+\psi_{11} & q_{1} q_{2} \\
q_{1} q_{2} & q_{2}^{2}+\psi_{22}
\end{array}\right) .
$$

We have infinitely many solutions: for any $\alpha$ ( $\rho<\alpha<1$ ), a solution is provided by

$$
q_{1}=\alpha ; q_{2}=\rho / \alpha ; \psi_{11}=1-\alpha^{2} ; \psi_{22}=1-(\rho / \alpha)^{2} .
$$

The solution in Example 12.1 may be unique (up to a rotation), but it is not proper in the sense that it cannot be interpreted statistically. Exercise 12.5 gives an example where the specific variance $\psi_{11}$ is negative.


Even in the case of a unique solution $(d=0)$, the solution may be inconsistent with statistical interpretations.


### 12.2 Estimation of the Factor Model

In practice, we have to find estimates $\hat{\mathcal{Q}}$ of the loadings $\mathcal{Q}$ and estimates $\hat{\Psi}$ of the specific variances $\Psi$ such that analogously to (12.7)

$$
\mathcal{S}=\hat{\mathcal{Q}} \hat{\mathcal{Q}}^{\top}+\hat{\Psi},
$$

where $\mathcal{S}$ denotes the empirical covariance of $\mathcal{X}$. Given an estimate $\hat{\mathcal{Q}}$ of $\mathcal{Q}$, it is natural to set

$$
\hat{\psi}_{i j}=s_{X_{j} X_{j}}-\sum_{\ell=1}^{k} \hat{q}_{j \ell}^{2} .
$$

We have that $\hat{h}_{j}^{2}=\sum_{\ell=1}^{k} \hat{q}_{j \ell}^{2}$ is an estimate for the communality $h_{j}^{2}$.
In the ideal case $d=0$, there is an exact solution. However, $d$ is usually greater than zero, therefore we have to find $\hat{\mathcal{Q}}$ and $\hat{\Psi}$ such that $S$ is approximated by $\hat{\mathcal{Q}} \hat{\mathcal{Q}}^{\top}+$ $\hat{\Psi}$. As mentioned above, it is often easier to compute the loadings and the specific variances of the standardised model.

Define $\mathcal{Y}=\mathcal{H} \mathcal{X} \mathcal{D}^{-1 / 2}$, the standardisation of the data matrix $\mathcal{X}$, where $\mathcal{D}=$ $\operatorname{diag}\left(s_{X_{1} X_{1}}, \ldots, s_{X_{p} X_{p}}\right)$ and the centering matrix $\mathcal{H}=\mathcal{I}-n^{-1} 1_{n} 1_{n}^{\top}$ (recall from Chap. 2 that $\mathcal{S}=\frac{1}{n} \mathcal{X}^{\top} \mathcal{H} \mathcal{X}$ ). The estimated factor loading matrix $\hat{\mathcal{Q}}_{Y}$ and the estimated specific variance $\hat{\Psi}_{Y}$ of $\mathcal{Y}$ are

$$
\hat{\mathcal{Q}}_{Y}=\mathcal{D}^{-1 / 2} \hat{\mathcal{Q}}_{X} \quad \text { and } \quad \hat{\Psi}_{Y}=\mathcal{D}^{-1} \hat{\Psi}_{X} .
$$

For the correlation matrix $\mathcal{R}$ of $\mathcal{X}$, we have that

$$
\mathcal{R}=\hat{\mathcal{Q}}_{Y} \hat{\mathcal{Q}}_{Y}^{\top}+\hat{\Psi}_{Y} .
$$

The interpretations of the factors are formulated from the analysis of the loadings $\hat{\mathcal{Q}}_{Y}$.

Example 12.3 Let us calculate the matrices just defined for the car data given in Table 22.7. This data set consists of the averaged marks (from $1=$ low to $6=$ high) for 24 car types. Considering the three variables price, security and easy handling, we get the following correlation matrix:

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

We will first look for one factor, i.e. $k=1$. Note that (\# number of parameters of $\Sigma$ unconstrained - \# parameters of $\Sigma$ constrained) is equal to $\frac{1}{2}(p-k)^{2}-\frac{1}{2}(p+k)=$ $\frac{1}{2}(3-1)^{2}-\frac{1}{2}(3+1)=0$. This implies that there is an exact solution! The equation

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

yields the communalities $\hat{h}_{i}^{2}=\hat{q}_{i}^{2}$, where

$$
\hat{q}_{1}^{2}=\frac{r_{X_{1} X_{2}} r_{X_{1} X_{3}}}{r_{X_{2} X_{3}}}, \quad \hat{q}_{2}^{2}=\frac{r_{X_{1} X_{2}} r_{X_{2} X_{3}}}{r_{X_{1} X_{3}}} \quad \text { and } \quad \hat{q}_{3}^{2}=\frac{r_{X_{1} X_{3}} r_{X_{2} X_{3}}}{r_{X_{1} X_{2}}} .
$$

Combining this with the specific variances $\hat{\psi}_{11}=1-\hat{q}_{1}^{2}, \hat{\psi}_{22}=1-\hat{q}_{2}^{2}$ and $\hat{\psi}_{33}=1-\hat{q}_{3}^{2}$, we obtain the following solution

$$
\begin{aligned}
\hat{q}_{1} & =0.982 & \hat{q}_{2} & =0.993
\end{aligned} r \hat{q}_{3}=0.624, ~ 子 \hat{\psi}_{33}=0.610 .
$$

Since the first two communalities ( $\hat{h}_{i}^{2}=\hat{q}_{i}^{2}$ ) are close to one, we can conclude that the first two variables, namely price and security, are explained by the single factor quite well. This factor can be interpreted as a "price+security" factor.

## The Maximum Likelihood Method

Recall from Chap. 6 the log-likelihood function $\ell$ for a data matrix $\mathcal{X}$ of observations of $X \sim N_{p}(\mu, \Sigma)$ :

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

This can be rewritten as

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

Replacing $\mu$ by $\hat{\mu}=\bar{x}$ and substituting $\Sigma=\mathcal{Q} \mathcal{Q}^{\top}+\Psi$ this becomes

$$
\begin{equation*}
\ell(\mathcal{X} ; \hat{\mu}, \mathcal{Q}, \Psi)=-\frac{n}{2}\left[\log \left\{\left|2 \pi\left(\mathcal{Q} \mathcal{Q}^{\top}+\Psi\right)\right|\right\}+\operatorname{tr}\left\{\left(\mathcal{Q} \mathcal{Q}^{\top}+\Psi\right)^{-1} \mathcal{S}\right\}\right] \tag{12.13}
\end{equation*}
$$

Even in the case of a single factor $(k=1)$, these equations are rather complicated and iterative numerical algorithms have to be used [for more details see Mardia, Kent \& Bibby, 1979, p. 263]. A practical computation scheme is also given in Supplement 9A of Johnson and Wichern (1998).

## Likelihood Ratio Test for the Number of Common Factors

Using the methodology of Chap. 7, it is easy to test the adequacy of the factor analysis model by comparing the likelihood under the null (factor analysis) and alternative (no constraints on covariance matrix) hypotheses.

Assuming that $\hat{\mathcal{Q}}$ and $\hat{\Psi}$ are the maximum likelihood estimates corresponding to (12.13), we obtain the following LR test statistic:

$$
\begin{equation*}
-2 \log \left(\frac{\text { maximised likelihood under } H_{0}}{\text { maximised likelihood }}\right)=n \log \left(\frac{\left|\hat{\mathcal{Q}} \hat{\mathcal{Q}}^{\top}+\hat{\Psi}\right|}{|\mathcal{S}|}\right) \tag{12.14}
\end{equation*}
$$

which asymptotically has the $\chi_{\frac{1}{2}\left\{(p-k)^{2}-p-k\right\}}^{2}$ distribution.

The $\chi^{2}$ approximation can be improved if we replace $n$ by $n-1-(2 p+4 k+5) / 6$ in (12.14) (Bartlett, 1954). Using Bartlett's correction, we reject the factor analysis model at the $\alpha$ level if

$$
\begin{equation*}
\{n-1-(2 p+4 k+5) / 6\} \log \left(\frac{\left|\hat{\mathcal{Q}} \hat{\mathcal{Q}}^{\top}+\hat{\Psi}\right|}{|\mathcal{S}|}\right)>\chi_{1-\alpha ;\left\{(p-k)^{2}-p-k\right\} / 2}^{2} \tag{12.15}
\end{equation*}
$$

and if the number of observations $n$ is large and the number of common factors $k$ is such that the $\chi^{2}$ statistic has a positive number of degrees of freedom.

## The Method of Principal Factors

The method of principal factors concentrates on the decomposition of the correlation matrix $\mathcal{R}$ or the covariance matrix $\mathcal{S}$. For simplicity, only the method for the correlation matrix $\mathcal{R}$ will be discussed. As pointed out in Chap. 11, the spectral decompositions of $\mathcal{R}$ and $\mathcal{S}$ yield different results and therefore, the method of principal factors may result in different estimators. The method can be motivated as follows: Suppose we know the exact $\Psi$, then the constraint (12.12) implies that the columns of $\mathcal{Q}$ are orthogonal since $\mathcal{D}=\mathcal{I}$ and it implies that they are eigenvectors of $\mathcal{Q} \mathcal{Q}^{\top}=\mathcal{R}-\Psi$. Furthermore, assume that the first $k$ eigenvalues are positive. In this case we could calculate $\mathcal{Q}$ by means of a spectral decomposition of $\mathcal{Q} \mathcal{Q}^{\top}$ and $k$ would be the number of factors.

The principal factors algorithm is based on good preliminary estimators $\hat{h}_{j}^{2}$ of the communalities $h_{j}^{2}$, for $j=1, \ldots, p$. There are two traditional proposals:

- $\hat{h}_{j}^{2}$, defined as the square of the multiple correlation coefficient of $X_{j}$ with $\left(X_{l}\right)$, for $l \neq j$, i.e. $\rho^{2}(V, W \hat{\beta})$ with $V=X_{j}, W=\left(X_{\ell}\right)_{\ell \neq j}$ and where $\hat{\beta}$ is the least squares regression parameter of a regression of $V$ on $W$.
- $\hat{h}_{j}^{2}=\max _{\ell \neq j}\left|r_{X_{j} X_{\ell}}\right|$, where $\mathcal{R}=\left(r_{X_{j} X_{\ell}}\right)$ is the correlation matrix of $\mathcal{X}$.

Given $\tilde{\psi}_{j j}=1-\tilde{h}_{j}^{2}$ we can construct the reduced correlation matrix, $\mathcal{R}-\hat{\Psi}$. The Spectral Decomposition Theorem says that

$$
\mathcal{R}-\hat{\Psi}=\sum_{\ell=1}^{p} \lambda_{\ell} \gamma_{\ell} \gamma_{\ell}^{\top}
$$

with eigenvalues $\lambda_{1} \geq \cdots \geq \lambda_{p}$. Assume that the first $k$ eigenvalues $\lambda_{1}, \ldots, \lambda_{k}$ are positive and large compared to the others. Then we can set

$$
\hat{q}_{\ell}=\sqrt{\lambda_{\ell}} \gamma_{\ell}, \quad \ell=1, \ldots, k
$$

or

$$
\hat{\mathcal{Q}}=\Gamma_{1} \Lambda_{1}^{1 / 2}
$$

with

$$
\Gamma_{1}=\left(\gamma_{1}, \ldots, \gamma_{k}\right) \quad \text { and } \quad \Lambda_{1}=\operatorname{diag}\left(\lambda_{1}, \ldots, \lambda_{k}\right)
$$

In the next step set

$$
\hat{\psi}_{i j}=1-\sum_{\ell=1}^{k} \hat{q}_{j \ell}^{2}, \quad j=1, \ldots, p
$$

Note that the procedure can be iterated: from $\hat{\psi}_{j j}$ we can compute a new reduced correlation matrix $\mathcal{R}-\hat{\Psi}$ following the same procedure. The iteration usually stops when the $\hat{\psi}_{j j}$ have converged to a stable value.

Example 12.4 Consider once again the car data given in Table 22.7. From Exercise 11.4 we know that the first PC is mainly influenced by $X_{2}-X_{7}$. Moreover, we know that most of the variance is already captured by the first PC. Thus we can conclude that the data are mainly determined by one factor $(k=1)$.

The eigenvalues of $\mathcal{R}-\hat{\Psi}$ for $\hat{\Psi}=\left(\max _{j \neq i}\left|r_{X_{i} X_{j}}\right|\right)$ are

$$
(4.628,1.340,1.201,1.045,1.007,0.993,0.980,-4.028)^{\top} .
$$

It would suffice to choose only one factor. Nevertheless, we have computed two factors. The result (the factor loadings for two factors) is shown in Fig. 12.1.


Fig. 12.1 Loadings of the evaluated car qualities, factor analysis with $k=2 \mathrm{Q}$ MVAfactcarm

We can clearly see a cluster of points to the right, which contain the factor loadings for the variables $X_{2}-X_{7}$. This shows, as did the PCA, that these variables are highly dependent and are thus more or less equivalent. The factor loadings for $X_{1}$ (economy) and $X_{8}$ (easy handling) are separate, but note the different scales on the horizontal and vertical axes! Although there are two or three sets of variables in the plot, the variance is already explained by the first factor, the "price+security" factor.

## The Principal Component Method

The principal factor method (PFM) involves finding an approximation $\tilde{\Psi}$ of $\Psi$, the matrix of specific variances, and then correcting $\mathcal{R}$, the correlation matrix of $X$, by $\tilde{\Psi}$. The principal component method (PCM) starts with an approximation $\hat{\mathcal{Q}}$ of $\mathcal{Q}$, the factor loadings matrix. The sample covariance matrix is diagonalised, $\mathcal{S}=\Gamma \Lambda \Gamma^{\top}$. Then the first $k$ eigenvectors are retained to build

$$
\begin{equation*}
\hat{\mathcal{Q}}=\left(\sqrt{\lambda_{1}} \gamma_{1}, \ldots, \sqrt{\lambda_{k}} \gamma_{k}\right) \tag{12.16}
\end{equation*}
$$

The estimated specific variances are provided by the diagonal elements of the matrix $\mathcal{S}-\hat{\mathcal{Q}} \hat{\mathcal{Q}}^{\top}$,

$$
\hat{\Psi}=\left(\begin{array}{cccc}
\hat{\psi}_{11} & & & 0  \tag{12.17}\\
& \hat{\psi}_{22} & & \\
& & \ddots & \\
0 & & & \hat{\psi}_{p p}
\end{array}\right) \text { with } \hat{\psi}_{j j}=s_{X j X j}-\sum_{\ell=1}^{k} \hat{q}_{j \ell}^{2} .
$$

By definition, the diagonal elements of $\mathcal{S}$ are equal to the diagonal elements of $\hat{\mathcal{Q}} \hat{\mathcal{Q}}^{\top}+\hat{\Psi}$. The off-diagonal elements are not necessarily estimated. How good then is this approximation? Consider the residual matrix

$$
\mathcal{S}-\left(\hat{\mathcal{Q}} \hat{\mathcal{Q}}^{\top}+\hat{\Psi}\right)
$$

resulting from the principal component solution. Analytically we have that

$$
\sum_{i, j}\left(\mathcal{S}-\hat{\mathcal{Q}} \hat{\mathcal{Q}}^{\top}-\hat{\Psi}\right)_{i j}^{2} \leq \lambda_{k+1}^{2}+\cdots+\lambda_{p}^{2}
$$

This implies that a small value of the neglected eigenvalues can result in a small approximation error. A heuristic device for selecting the number of factors is to
consider the proportion of the total sample variance due to the $j$ th factor. This quantity is in general equal to
(a) $\lambda_{j} / \sum_{j=1}^{p} s_{j j}$ for a factor analysis of $\mathcal{S}$,
(b) $\lambda_{j} / p$ for a factor analysis of $\mathcal{R}$.

Example 12.5 This example uses a consumer-preference study from Johnson and Wichern (1998). Customers were asked to rate several attributes of a new product. The responses were tabulated and the following correlation matrix $\mathcal{R}$ was constructed:

Attribute (Variable)

| Taste | 1 |
| :--- | :--- |
| Good buy for money | 2 |
| Flavor | 3 |
| Suitable for snack | 4 |
| Provides lots of energy | 5 |\(\quad\left(\begin{array}{lllllll}1.00 \& 0.02 \& \mathbf{0 . 9 6} \& 0.42 \& 0.01 <br>

0.02 \& 1.00 \& 0.13 \& 0.71 \& \mathbf{0 . 8 5} <br>
0.96 \& 0.13 \& 1.00 \& 0.50 \& 0.11 <br>
0.42 \& 0.71 \& 0.50 \& 1.00 \& \mathbf{0 . 7 9} <br>
0.01 \& 0.85 \& 0.11 \& 0.79 \& 1.00\end{array}\right)\)

The bold entries of $\mathcal{R}$ show that variables 1 and 3 and variables 2 and 5 are highly correlated. Variable 4 is more correlated with variables 2 and 5 than with variables 1 and 3 . Hence, a model with 2 (or 3 ) factors seems to be reasonable.

The first two eigenvalues $\lambda_{1}=2.85$ and $\lambda_{2}=1.81$ of $\mathcal{R}$ are the only eigenvalues greater than one. Moreover, $k=2$ common factors account for a cumulative proportion

$$
\frac{\lambda_{1}+\lambda_{2}}{p}=\frac{2.85+1.81}{5}=0.93
$$

of the total (standardised) sample variance. Using the PCM, the estimated factor loadings, communalities, and specific variances are calculated from formulas (12.16) and (12.17), and the results are given in Table 12.1.

Table 12.1 Estimated factor loadings, communalities, and specific variances

| Variable |  | Estimated factor loadings |  | Communalities$\hat{h}_{j}^{2}$ | Specific variances$\hat{\psi}_{j j}=1-\hat{h}_{j}^{2}$ |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | $\hat{q}_{1}$ | $\hat{q}_{2}$ |  |  |
| 1. | Taste | 0.56 | 0.82 | 0.98 | 0.02 |
| 2. | Good buy for money | 0.78 | $-0.53$ | 0.88 | 0.12 |
| 3. | Flavor | 0.65 | 0.75 | 0.98 | 0.02 |
| 4. | Suitable for snack | 0.94 | -0.11 | 0.89 | 0.11 |
| 5. | Provides lots of energy | 0.80 | -0.54 | 0.93 | 0.07 |
|  | Prvalues | 2.85 | 1.81 |  |  |
|  | mulative proportion of total ndardised) sample variance | 0.571 | 0.932 |  |  |

Take a look at:

$$
\begin{aligned}
& \hat{Q} \hat{Q}^{\top}+\hat{\Psi}=\left(\begin{array}{rrr}
0.56 & 0.82 \\
0.78 & -0.53 \\
0.65 & 0.75 \\
0.94 & -0.11 \\
0.80 & -0.54
\end{array}\right)\left(\begin{array}{rrrrr}
0.56 & 0.78 & 0.65 & 0.94 & 0.80 \\
0.82 & -0.53 & 0.75 & -0.11 & -0.54
\end{array}\right) \\
& +\left(\begin{array}{lllll}
0.02 & 0 & 0 & 0 & 0 \\
0 & 0.12 & 0 & 0 \\
0 & 0 & 0.02 & 0 & 0 \\
0 & 0 & 0 & 0.11 & 0 \\
0 & 0 & 0 & 0 & 0.07
\end{array}\right)=\left(\begin{array}{rrrrrr}
1.00 & 0.01 & 0.97 & 0.44 & 0.00 \\
0.01 & 1.00 & 0.11 & 0.79 & 0.91 \\
0.97 & 0.11 & 1.00 & 0.53 & 0.11 \\
0.44 & 0.79 & 0.53 & 1.00 & 0.81 \\
0.00 & 0.91 & 0.11 & 0.81 & 1.00
\end{array}\right) .
\end{aligned}
$$

This nearly reproduces the correlation matrix $\mathcal{R}$. We conclude that the twofactor model provides a good fit of the data. The communalities $(0.98,0.88,0.98$, $0.89,0.93$ ) indicate that the two factors account for a large percentage of the sample variance of each variable. Due to the nonuniqueness of factor loadings, the interpretation might be enhanced by rotation. This is the topic of the next subsection.

## Rotation

The constraints (12.11) and (12.12) are given as a matter of mathematical convenience (to create unique solutions) and can therefore complicate the problem of interpretation. The interpretation of the loadings would be very simple if the variables could be split into disjoint sets, each being associated with one factor. A well-known analytical algorithm to rotate the loadings is given by the varimax rotation method proposed by Kaiser (1985). In the simplest case of $k=2$ factors, a rotation matrix $\mathcal{G}$ is given by

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

representing a clockwise rotation of the coordinate axes by the angle $\theta$. The corresponding rotation of loadings is calculated via $\hat{\mathcal{Q}}^{*}=\hat{\mathcal{Q}} \mathcal{G}(\theta)$. The idea of the varimax method is to find the angle $\theta$ that maximises the sum of the variances of the squared loadings $\hat{q}_{i j}^{*}$ within each column of $\hat{\mathcal{Q}}^{*}$. More precisely, defining $\tilde{q}_{j l}^{*}=\hat{q}_{j l}^{*} / \hat{h}_{j}^{*}$, the varimax criterion chooses $\theta$ so that

$$
\mathcal{V}=\frac{1}{p} \sum_{\ell=1}^{k}\left[\sum_{j=1}^{p}\left(\hat{q}_{j l}^{*}\right)^{4}-\left\{\frac{1}{p} \sum_{j=1}^{p}\left(\hat{q}_{j l}^{*}\right)^{2}\right\}^{2}\right]
$$

is maximised.
Example 12.6 Let us return to the marketing example of Johnson and Wichern (1998) (Example 12.5). The basic factor loadings given in Table 12.1 of the first factor and a second factor are almost identical making it difficult to interpret the factors. Applying the varimax rotation we obtain the loadings $\tilde{q}_{1}=$ $(0.02, \boldsymbol{0} .94,0.13, \mathbf{0 . 8 4}, \mathbf{0 . 9 7})^{\top}$ and $\tilde{q}_{2}=(\mathbf{0 . 9 9},-0.01, \mathbf{0 . 9 8}, 0.43,-0.02)^{\top}$. The high loadings, indicated as bold entries, show that variables $2,4,5$ define factor 1 , a nutritional factor. Variables 1 and 3 define factor 2 which might be referred to as a taste factor.

|  |
| :---: |
| $\hookrightarrow$ In practice, $\mathcal{Q}$ and $\Psi$ have to be estimated from $\mathcal{S}=\hat{\mathcal{Q}} \hat{\mathcal{Q}}^{\top}+\hat{\Psi}$. The number of parameters is $d=\frac{1}{2}(p-k)^{2}-\frac{1}{2}(p+k)$. |
| $\hookrightarrow$ If $d=0$, then there exists an exact solution. In practice, $d$ is usually greater than 0 , thus approximations must be considered. |
| $\hookrightarrow$ The maximum-likelihood method assumes a normal distribution for the data. A solution can be found using numerical algorithms. |
| $\hookrightarrow$ The method of principal factors is a two-stage method which calculates $\hat{\mathcal{Q}}$ from the reduced correlation matrix $\mathcal{R}-\hat{\Psi}$, where $\hat{\Psi}$ is a pre-estimate of $\Psi$. The final estimate of $\Psi$ is found by $\hat{\psi}_{i i}=1-\sum_{j=1}^{k} \hat{q}_{i j}^{2}$. |
| $\hookrightarrow$ The PCM is based on an approximation, $\hat{\mathcal{Q}}$, of $\mathcal{Q}$. |
| $\hookrightarrow$ Often a more informative interpretation of the factors can be found by rotating the factors. |
| $\hookrightarrow$ The varimax rotation chooses a rotation $\theta$ that maximises $\mathcal{V}=\frac{1}{p} \sum_{\ell=1}^{k}\left[\sum_{j=1}^{p}\left(\tilde{q}_{j l}^{*}\right)^{4}-\left\{\frac{1}{p} \sum_{j=1}^{p}\left(\tilde{q}_{j l}^{*}\right)^{2}\right\}^{2}\right]$ |

### 12.3 Factor Scores and Strategies

Up to now strategies have been presented for factor analysis that have concentrated on the estimation of loadings and communalities and on their interpretations. This was a logical step since the factors $F$ were considered to be normalised random sources of information and were explicitly addressed as nonspecific (common factors). The estimated values of the factors, called the factor scores, may also be useful in the interpretation as well as in the diagnostic analysis. To be more precise, the factor scores are estimates of the unobserved random vectors $F_{l}, l=1, \ldots, k$, for each individual $x_{i}, i=1, \ldots, n$. Johnson and Wichern (1998) describe three methods which in practice yield very similar results. Here, we present the regression method which has the advantage of being the simplest technique and is easy to implement.

The idea is to consider the joint distribution of $(X-\mu)$ and $F$, and then to proceed with the regression analysis presented in Chap. 5. Under the factor model (12.4), the joint covariance matrix of $(X-\mu)$ and $F$ is:

$$
\operatorname{Var}\binom{X-\mu}{F}=\left(\begin{array}{cc}
\mathcal{Q} \mathcal{Q}^{\top}+\Psi & \mathcal{Q}  \tag{12.18}\\
\mathcal{Q}^{\top} & \mathcal{I}_{k}
\end{array}\right)
$$

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

Assuming joint normality, the conditional distribution of $F \mid X$ is multinormal, see Theorem 5.1, with

$$
\begin{equation*}
\mathrm{E}(F \mid X=x)=\mathcal{Q}^{\top} \Sigma^{-1}(X-\mu) \tag{12.19}
\end{equation*}
$$

and using (5.7) the covariance matrix can be calculated:

$$
\begin{equation*}
\operatorname{Var}(F \mid X=x)=\mathcal{I}_{k}-\mathcal{Q}^{\top} \Sigma^{-1} \mathcal{Q} \tag{12.20}
\end{equation*}
$$

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

$$
\begin{equation*}
\hat{f_{i}}=\hat{\mathcal{Q}}^{\top} \mathcal{S}^{-1}\left(x_{i}-\bar{x}\right) \tag{12.21}
\end{equation*}
$$

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

The same rule can be followed when using $\mathcal{R}$ instead of $\mathcal{S}$. Then (12.18) remains valid when standardised variables, i.e. $Z=\mathcal{D}_{\Sigma}^{-1 / 2}(X-\mu)$, are considered if $\mathcal{D}_{\Sigma}=$ $\operatorname{diag}\left(\sigma_{11}, \ldots, \sigma_{p p}\right)$. In this case the factors are given by

$$
\begin{equation*}
\hat{f_{i}}=\hat{\mathcal{Q}}^{\top} \mathcal{R}^{-1}\left(z_{i}\right) \tag{12.22}
\end{equation*}
$$

where $z_{i}=\mathcal{D}_{S}^{-1 / 2}\left(x_{i}-\bar{x}\right), \hat{\mathcal{Q}}$ is the loading obtained with the matrix $\mathcal{R}$, and $\mathcal{D}_{S}=$ $\operatorname{diag}\left(s_{11}, \ldots, s_{p p}\right)$.

If the factors are rotated by the orthogonal matrix $\mathcal{G}$, the factor scores have to be rotated accordingly, that is

$$
\begin{equation*}
\hat{f}_{i}^{*}=\mathcal{G}^{\top} \hat{f_{i}} \tag{12.23}
\end{equation*}
$$

A practical example is presented in Sect. 12.4 using the Boston Housing data.

## Practical Suggestions

No one method outperforms another in the practical implementation of factor analysis. However, by applying a tâtonnement process, the factor analysis view of the data can be stabilised. This motivates the following procedure.

1. Fix a reasonable number of factors, say $k=2$ or 3 , based on the correlation structure of the data and/or screeplot of eigenvalues.
2. Perform several of the presented methods, including rotation. Compare the loadings, communalities, and factor scores from the respective results.
3. If the results show significant deviations, check for outliers (based on factor scores), and consider changing the number of factors $k$.

For larger data sets, cross-validation methods are recommended. Such methods involve splitting the sample into a training set and a validation data set. On the training sample one estimates the factor model with the desired methodology and uses the obtained parameters to predict the factor scores for the validation data set. The predicted factor scores should be comparable to the factor scores obtained using only the validation data set. This stability criterion may also involve the loadings and communalities.

## Factor Analysis Versus PCA

Factor analysis and principal component analysis use the same set of mathematical tools (spectral decomposition, projections, ...). One could conclude, on first sight, that they share the same view and strategy and therefore yield very similar results. This is not true. There are substantial differences between these two data analysis techniques that we would like to describe here.

The biggest difference between PCA and factor analysis comes from the model philosophy. Factor analysis imposes a strict structure of a fixed number of common (latent) factors whereas the PCA determines $p$ factors in decreasing order of importance. The most important factor in PCA is the one that maximises the
projected variance. The most important factor in factor analysis is the one that (after rotation) gives the maximal interpretation. Often this is different from the direction of the first principal component.

From an implementation point of view, the PCA is based on a well-defined, unique algorithm (spectral decomposition), whereas fitting a factor analysis model involves a variety of numerical procedures. The non-uniqueness of the factor analysis procedure opens the door for subjective interpretation and yields therefore a spectrum of results. This data analysis philosophy makes factor analysis difficult especially if the model specification involves cross-validation and a data-driven selection of the number of factors.

### 12.4 Boston Housing

To illustrate how to implement factor analysis we will use the Boston Housing data set and the by now well-known set of transformations. Once again, the variable $X_{4}$ (Charles River indicator) will be excluded. As before, standardised variables are used and the analysis is based on the correlation matrix.

In Sect. 12.3, we described a practical implementation of factor analysis. Based on principal components, three factors were chosen and factor analysis was applied using the maximum likelihood method (MLM), the PFM, and the PCM. For illustration, the MLM will be presented with and without varimax rotation.

Table 12.2 gives the MLM factor loadings without rotation and Table 12.3 gives the varimax version of this analysis. The corresponding graphical representations of the loadings are displayed in Figs. 12.2 and 12.3. We can see that the varimax

Table 12.2 Estimated factor loadings, communalities, and specific variances, MLM Q MVAfacthous

|  |  | Estimated factor loadings |  |  | Communalities $\hat{h}_{j}^{2}$ | Specific <br> variances $\hat{\psi}_{j j}=1-\hat{h}_{j}^{2}$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | $\hat{q}_{1}$ | $\hat{q}_{2}$ | $\hat{q}_{3}$ |  |  |
| 1 | Crime | 0.9295 | 0.1653 | 0.1107 | 0.9036 | 0.0964 |
| 2 | Large lots | -0.5823 | 0.0379 | 0.2902 | 0.4248 | 0.5752 |
| 3 | Nonretail acres | 0.8192 | -0.0296 | -0.1378 | 0.6909 | 0.3091 |
| 5 | Nitric oxides | 0.8789 | 0.0987 | -0.2719 | 0.8561 | 0.1439 |
| 6 | Rooms | $-0.4447$ | 0.5311 | $-0.0380$ | 0.4812 | 0.5188 |
| 7 | Prior 1940 | 0.7837 | -0.0149 | -0.3554 | 0.7406 | 0.2594 |
| 8 | Empl. centers | -0.8294 | -0.1570 | 0.4110 | 0.8816 | 0.1184 |
| 9 | Accessibility | 0.7955 | 0.3062 | 0.4053 | 0.8908 | 0.1092 |
| 10 | Tax-rate | 0.8262 | 0.1401 | 0.2906 | 0.7867 | 0.2133 |
| 11 | Pupil/teacher | 0.5051 | -0.1850 | 0.1553 | 0.3135 | 0.6865 |
| 12 | African American | 0.4701 | -0.0227 | $-0.1627$ | 0.2480 | 0.7520 |
| 13 | Lower status | 0.7601 | -0.5059 | -0.0070 | 0.8337 | 0.1663 |
| 14 | Value | -0.6942 | 0.5904 | -0.1798 | 0.8628 | 0.1371 |

Table 12.3 Estimated factor loadings, communalities, and specific variances, MLM, varimax rotation Q MVAfacthous

|  |  | Estimated factor loadings |  |  | Communalities $\hat{h}_{j}^{2}$ | Specific variances $\hat{\psi}_{j j}=1-\hat{h}_{j}^{2}$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | $\hat{q}_{1}$ | $\hat{q}_{2}$ | $\hat{q}_{3}$ |  |  |
| 1 | Crime | 0.7247 | -0.2705 | -0.5525 | 0.9036 | 0.0964 |
| 2 | Large lots | -0.1570 | 0.2377 | 0.5858 | 0.4248 | 0.5752 |
| 3 | Nonretail acres | 0.4195 | -0.3566 | -0.6287 | 0.6909 | 0.3091 |
| 5 | Nitric oxides | 0.4141 | -0.2468 | -0.7896 | 0.8561 | 0.1439 |
| 6 | Rooms | -0.0799 | 0.6691 | 0.1644 | 0.4812 | 0.5188 |
| 7 | Prior 1940 | 0.2518 | -0.2934 | -0.7688 | 0.7406 | 0.2594 |
| 8 | Empl. centers | -0.3164 | 0.1515 | 0.8709 | 0.8816 | 0.1184 |
| 9 | Accessibility | 0.8932 | -0.1347 | -0.2736 | 0.8908 | 0.1092 |
| 10 | Tax-rate | 0.7673 | -0.2772 | -0.3480 | 0.7867 | 0.2133 |
| 11 | Pupil/teacher | 0.3405 | -0.4065 | -0.1800 | 0.3135 | 0.6865 |
| 12 | African American | -0.3917 | 0.2483 | 0.1813 | 0.2480 | 0.7520 |
| 13 | Lower status | 0.2586 | -0.7752 | -0.4072 | 0.8337 | 0.1663 |
| 14 | Value | -0.3043 | 0.8520 | 0.2111 | 0.8630 | 0.1370 |



Fig. 12.2 Factor analysis for Boston housing data, MLM $Q$ MVAfacthous


Fig. 12.3 Factor analysis for Boston housing data, MLM after varimax rotation $\mathbf{a}$ MVAfacthous
does not significantly change the interpretation of the factors obtained by the MLM. Factor 1 can be roughly interpreted as a "quality of life factor" because it is positively correlated with variables like $X_{11}$ and negatively correlated with $X_{8}$, both having low specific variances. The second factor may be interpreted as a "residential factor", since it is highly correlated with variables $X_{6}$, and $X_{13}$. The most striking difference between the results with and without varimax rotation can be seen by comparing the lower left corners of Figs. 12.2 and 12.3. There is a clear separation of the variables in the varimax version of the MLM. Given this arrangement of the variables in Fig. 12.3, we can interpret factor 3 as an employment factor, since we observe high correlations with $X_{8}$ and $X_{5}$.

We now turn to the PCM and PFM analyses. The results are presented in Tables 12.4 and 12.5 and in Figs. 12.4 and 12.5. We would like to focus on the PCM, because this three-factor model yields only one specific variance (unexplained variation) above 0.5 . Looking at Fig. 12.4, it turns out that factor 1 remains a "quality of life factor" which is clearly visible from the clustering of $X_{5}, X_{3}, X_{10}$ and $X_{1}$ on the right-hand side of the graph, while the variables $X_{8}, X_{2}, X_{14}, X_{12}$ and $X_{6}$ are on the left-hand side. Again, the second factor is a "residential factor", clearly

Table 12.4 Estimated factor loadings, communalities, and specific variances, PCM, varimax rotation Q MVAfacthous

|  |  | Estimated factor loadings |  |  | Communalities$\hat{h}_{j}^{2}$ | Specific variances$\hat{\psi}_{j j}=1-\hat{h}_{j}^{2}$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | $\hat{q}_{1}$ | $\hat{q}_{2}$ | $\hat{q}_{3}$ |  |  |
| 1 | Crime | 0.6034 | -0.2456 | 0.6864 | 0.8955 | 0.1045 |
| 2 | Large lots | -0.7722 | 0.2631 | 0.0270 | 0.6661 | 0.3339 |
| 3 | Nonretail acres | 0.7183 | -0.3701 | 0.3449 | 0.7719 | 0.2281 |
| 5 | Nitric oxides | 0.7936 | -0.2043 | 0.4250 | 0.8521 | 0.1479 |
| 6 | Rooms | -0.1601 | 0.8585 | 0.0218 | 0.7632 | 0.2368 |
| 7 | Prior 1940 | 0.7895 | -0.2375 | 0.2670 | 0.7510 | 0.2490 |
| 8 | Empl. centers | -0.8562 | 0.1318 | $-0.3240$ | 0.8554 | 0.1446 |
| 9 | Accessibility | 0.3681 | -0.1268 | 0.8012 | 0.7935 | 0.2065 |
| 10 | Tax-rate | 0.3744 | -0.2604 | 0.7825 | 0.8203 | 0.1797 |
| 11 | Pupil/teacher | 0.1982 | -0.5124 | 0.3372 | 0.4155 | 0.5845 |
| 12 | African American | 0.1647 | 0.0368 | -0.7002 | 0.5188 | 0.4812 |
| 13 | Lower status | 0.4141 | -0.7564 | 0.2781 | 0.8209 | 0.1791 |
| 14 | Value | -0.2111 | 0.8131 | -0.3671 | 0.8394 | 0.1606 |

Table 12.5 Estimated factor loadings, communalities, and specific variances, PFM, varimax rotation Q MVAfacthous

|  |  | Estimated factor loadings |  |  | Communalities $\hat{h}_{j}^{2}$ | Specific variances$\hat{\psi}_{j j}=1-\hat{h}_{j}^{2}$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | $\hat{q}_{1}$ | $\hat{q}_{2}$ | $\hat{q}_{3}$ |  |  |
| 1 | Crime | 0.5477 | -0.2558 | -0.7387 | 0.9111 | 0.0889 |
| 2 | Large lots | -0.6148 | 0.2668 | 0.1281 | 0.4655 | 0.5345 |
| 3 | Nonretail acres | 0.6523 | -0.3761 | -0.3996 | 0.7266 | 0.2734 |
| 5 | Nitric oxides | 0.7723 | -0.2291 | -0.4412 | 0.8439 | 0.1561 |
| 6 | Rooms | -0.1732 | 0.6783 | 0.1296 | 0.0699 | 0.5046 |
| 7 | Prior 1940 | 0.7390 | -0.2723 | -0.2909 | 0.7049 | 0.2951 |
| 8 | Empl. centers | -0.8565 | 0.1485 | 0.3395 | 0.8708 | 0.1292 |
| 9 | Accessibility | 0.2855 | -0.1359 | $-0.8460$ | 0.8156 | 0.1844 |
| 10 | Tax-rate | 0.3062 | -0.2656 | -0.8174 | 0.8325 | 0.1675 |
| 11 | Pupil/teacher | 0.2116 | -0.3943 | -0.3297 | 0.3090 | 0.6910 |
| 12 | African American | 0.1994 | 0.0666 | 0.4217 | 0.2433 | 0.7567 |
| 13 | Lower status | 0.4005 | -0.7743 | -0.2706 | 0.8333 | 0.1667 |
| 14 | Value | -0.1885 | 0.8400 | 0.3473 | 0.8611 | 0.1389 |



Fig. 12.4 Factor analysis for Boston housing data, PCM after varimax rotation $Q_{\text {mVAf acthous }}$
demonstrated by the location of variables $X_{6}, X_{14}, X_{11}$, and $X_{13}$. The interpretation of the third factor is more difficult because all of the loadings (except for $X_{12}$ ) are very small.

### 12.5 Exercises

Exercise 12.1 In Example 12.4 we have computed $\hat{\mathcal{Q}}$ and $\hat{\Psi}$ using the method of principal factors. We used a two-step iteration for $\hat{\Psi}$. Perform the third iteration step and compare the results (i.e. use the given $\hat{\mathcal{Q}}$ as a pre-estimate to find the final $\Psi$ ).

Exercise 12.2 Using the bank data set, how many factors can you find with the Method of Principal Factors?

Exercise 12.3 Repeat Exercise 12.2 with the US company data set!
Exercise 12.4 Generalise the two-dimensional rotation matrix in Sect. 12.2 to n-dimensional space.


Fig. 12.5 Factor analysis for Boston housing data, PFM after varimax rotation $Q$ MVAfacthous

Exercise 12.5 Compute the orthogonal factor model for

$$
\Sigma=\left(\begin{array}{ccc}
1 & 0.9 & 0.7 \\
0.9 & 1 & 0.4 \\
0.7 & 0.4 & 1
\end{array}\right)
$$

[Solution: $\left.\psi_{11}=-0.575, q_{11}=1.255\right]$
Exercise 12.6 Perform a factor analysis on the type of families in the French food data set. Rotate the resulting factors in a way which provides the most reasonable interpretation. Compare your result with the varimax method.

Exercise 12.7 Perform a factor analysis on the variables $X_{3}$ to $X_{9}$ in the US crime data set (Table 22.10). Would it make sense to use all of the variables for the analysis?

Exercise 12.8 Analyse the athletic records data set (Table 22.18). Can you recognise any patterns if you sort the countries according to the estimates of the factor scores?

Exercise 12.9 Perform a factor analysis on the US health data set (Table 22.16) and estimate the factor scores.

Exercise 12.10 Redo Exercise 12.9 using the US crime data in Table 22.10. Compare the estimated factor scores of the two data sets.

Exercise 12.11 Analyse the vocabulary data given in Table 22.17.

## Chapter 13 <br> Cluster Analysis

The next two chapters address classification issues from two varying perspectives. When considering groups of objects in a multivariate data set, two situations can arise. Given a data set containing measurements on individuals, in some cases we want to see if some natural groups or classes of individuals exist, and in other cases, we want to classify the individuals according to a set of existing groups. Cluster analysis develops tools and methods concerning the former case, that is, given a data matrix containing multivariate measurements on a large number of individuals (or objects), the objective is to build some natural sub-groups or clusters of individuals. This is done by grouping individuals that are "similar" according to some appropriate criterion. Once the clusters are obtained, it is generally useful to describe each group using some descriptive tool from Chaps. 1, 10 or 11 to create a better understanding of the differences that exist among the formulated groups.

Cluster analysis is applied in many fields such as the natural sciences, the medical sciences, economics, marketing, etc. In marketing, for instance, it is useful to build and describe the different segments of a market from a survey on potential consumers. An insurance company, on the other hand, might be interested in the distinction among classes of potential customers so that it can derive optimal prices for its services. Other examples are provided below.

Discriminant analysis presented in Chap. 14 addresses the other issue of classification. It focuses on situations where the different groups are known a priori. Decision rules are provided in classifying a multivariate observation into one of the known groups.

Section 13.1 states the problem of cluster analysis where the criterion chosen to measure the similarity among objects clearly plays an important role. Section 13.2
shows how to precisely measure the proximity between objects. Finally, Sect. 13.3 provides some algorithms. We will concentrate on hierarchical algorithms only where the number of clusters is not known in advance.

### 13.1 The Problem

Cluster analysis is a set of tools for building groups (clusters) from multivariate data objects. The aim is to construct groups with homogeneous properties out of heterogeneous large samples. The groups or clusters should be as homogeneous as possible and the differences among the various groups as large as possible. Cluster analysis can be divided into two fundamental steps.

1. Choice of a proximity measure:

One checks each pair of observations (objects) for the similarity of their values. A similarity (proximity) measure is defined to measure the "closeness" of the objects. The "closer" they are, the more homogeneous they are.
2. Choice of group-building algorithm:

On the basis of the proximity measures the objects assigned to groups so that differences between groups become large and observations in a group become as close as possible.

In marketing, for example, cluster analysis is used to select test markets. Other applications include the classification of companies according to their organisational structures, technologies and types. In psychology, cluster analysis is used to find types of personalities on the basis of questionnaires. In archaeology, it is applied to classify art objects in different time periods. Other scientific branches that use cluster analysis are medicine, sociology, linguistics and biology. In each case a heterogeneous sample of objects are analysed with the aim to identify homogeneous sub-groups.


## Summary

$\hookrightarrow$ Cluster analysis is a set of tools for building groups (clusters) from multivariate data objects.
$\hookrightarrow$ The methods used are usually divided into two fundamental steps: The choice of a proximity measure and the choice of a groupbuilding algorithm.

### 13.2 The Proximity Between Objects

The starting point of a cluster analysis is a data matrix $\mathcal{X}(n \times p)$ with $n$ measurements (objects) of $p$ variables. The proximity (similarity) among objects is described by a matrix $\mathcal{D}(n \times n)$

$$
\mathcal{D}=\left(\begin{array}{ccccc}
d_{11} & d_{12} & \ldots & \ldots & \ldots  \tag{13.1}\\
\vdots & d_{1 n} \\
\vdots & \vdots & \ddots & & \\
\vdots & \vdots & & \ddots & \\
\vdots \\
\vdots & \vdots & & & \ddots \\
\vdots \\
d_{n 1} & d_{n 2} & \ldots & \ldots & \ldots
\end{array}\right)
$$

The matrix $\mathcal{D}$ contains measures of similarity or dissimilarity among the $n$ objects. If the values $d_{i j}$ are distances, then they measure dissimilarity. The greater the distance, the less similar are the objects. If the values $d_{i j}$ are proximity measures, then the opposite is true, i.e. the greater the proximity value, the more similar are the objects. A distance matrix, for example, could be defined by the $L_{2}$-norm: $d_{i j}=\left\|x_{i}-x_{j}\right\|_{2}$, where $x_{i}$ and $x_{j}$ denote the rows of the data matrix $\mathcal{X}$. Distance and similarity are of course dual. If $d_{i j}$ is a distance, then $d_{i j}^{\prime}=\max _{i, j}\left\{d_{i j}\right\}-d_{i j}$ is a proximity measure.

The nature of the observations plays an important role in the choice of proximity measure. Nominal values (like binary variables) lead in general to proximity values, whereas metric values lead (in general) to distance matrices. We first present possibilities for $\mathcal{D}$ in the binary case and then consider the continuous case.

## Similarity of Objects with Binary Structure

In order to measure the similarity between objects we always compare pairs of observations $\left(x_{i}, x_{j}\right)$ where $x_{i}^{\top}=\left(x_{i 1}, \ldots, x_{i p}\right), x_{j}^{\top}=\left(x_{j 1}, \ldots, x_{j p}\right)$, and $x_{i k}, x_{j k} \in$ $\{0,1\}$. Obviously there are four cases:

$$
\begin{aligned}
x_{i k} & =x_{j k}=1, \\
x_{i k} & =0, x_{j k}=1, \\
x_{i k} & =1, x_{j k}=0, \\
x_{i k} & =x_{j k}=0 .
\end{aligned}
$$

Define

$$
\begin{aligned}
& a_{1}=\sum_{k=1}^{p} \boldsymbol{I}\left(x_{i k}=x_{j k}=1\right), \\
& a_{2}=\sum_{k=1}^{p} \boldsymbol{I}\left(x_{i k}=0, x_{j k}=1\right), \\
& a_{3}=\sum_{k=1}^{p} \boldsymbol{I}\left(x_{i k}=1, x_{j k}=0\right), \\
& a_{4}=\sum_{k=1}^{p} \boldsymbol{I}\left(x_{i k}=x_{j k}=0\right)
\end{aligned}
$$

Note that each $a_{l}, l=1, \ldots, 4$, depends on the pair $\left(x_{i}, x_{j}\right)$.
The following proximity measures are used in practice:

$$
\begin{equation*}
d_{i j}=\frac{a_{1}+\delta a_{4}}{a_{1}+\delta a_{4}+\lambda\left(a_{2}+a_{3}\right)} \tag{13.2}
\end{equation*}
$$

where $\delta$ and $\lambda$ are weighting factors. Table 13.1 shows some similarity measures for given weighting factors.

These measures provide alternative ways of weighting mismatching and positive (presence of a common character) or negative (absence of a common character) matchings. In principle, we could also consider the Euclidean distance. However, the disadvantage of this distance is that it treats the observations 0 and 1 in the same way. If $x_{i k}=1$ denotes, say, knowledge of a certain language, then the contrary, $x_{i k}=0$ (not knowing the language) should eventually be treated differently.
Example 13.1 Let us consider binary variables computed from the car data set (Table 22.7). We define the new binary data by

$$
y_{i k}= \begin{cases}1 & \text { if } x_{i k}>\bar{x}_{k}, \\ 0 & \text { otherwise },\end{cases}
$$

Table 13.1 The common similarity coefficients

| 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 | $\frac{a_{1}+a_{4}}{p}$ |
| Russel and Rao (RR) | - | - | $\frac{a_{1}}{p}$ |
| Dice | 0 | 0.5 | $\frac{2 a_{1}}{2 a_{1}+\left(a_{2}+a_{3}\right)}$ |
| Kulczynski | - | - | $\frac{a_{1}}{a_{2}+a_{3}}$ |

for $i=1, \ldots, n$ and $k=1, \ldots, p$. This means that we transform the observations of the $k$-th variable to 1 if it is larger than the mean value of all observations of the $k$-th variable. Let us only consider the data points 17 to 19 (Renault 19, Rover and Toyota Corolla) which lead to $(3 \times 3)$ distance matrices. The Jaccard measure gives the similarity matrix

$$
\mathcal{D}=\left(\begin{array}{r}
1.0000 .0000 .400 \\
1.0000 .167 \\
1.000
\end{array}\right),
$$

the Tanimoto measure yields

$$
\mathcal{D}=\left(\begin{array}{rr}
1.000 & 0.000 \\
1.000 & 0.2351 \\
& 1.000
\end{array}\right)
$$

whereas the Simple Matching measure gives

$$
\mathcal{D}=\left(\begin{array}{rr}
1.000 & 0.000 \\
& 0.625 \\
& 1.000 \\
& 1.375 \\
& 1.000
\end{array}\right) .
$$

## Distance Measures for Continuous Variables

A wide variety of distance measures can be generated by the $L_{r}$-norms, $r \geq 1$,

$$
\begin{equation*}
d_{i j}=\left\|x_{i}-x_{j}\right\|_{r}=\left\{\sum_{k=1}^{p}\left|x_{i k}-x_{j k}\right|^{r}\right\}^{1 / r} . \tag{13.3}
\end{equation*}
$$

Here $x_{i k}$ denotes the value of the $k$-th variable on object $i$. It is clear that $d_{i i}=0$ for $i=1, \ldots, n$. The class of distances (13.3) for varying $r$ measures the dissimilarity of different weights. The $L_{1}$-metric, for example, gives less weight to outliers than the $L_{2}$-norm (Euclidean norm). It is common to consider the squared $L_{2}$-norm.

Example 13.2 Suppose we have $x_{1}=(0,0), x_{2}=(1,0)$ and $x_{3}=(5,5)$. Then the distance matrix for the $L_{1}$-norm is

$$
\mathcal{D}_{1}=\left(\begin{array}{rrr}
0 & 1 & 10 \\
1 & 0 & 9 \\
10 & 9 & 0
\end{array}\right),
$$

and for the squared $L_{2}$ - or Euclidean norm

$$
\mathcal{D}_{2}=\left(\begin{array}{rrr}
0 & 1 & 50 \\
1 & 0 & 41 \\
50 & 41 & 0
\end{array}\right)
$$

One can see that the third observation $x_{3}$ receives much more weight in the squared $L_{2}$-norm than in the $L_{1}$-norm.

An underlying assumption in applying distances based on $L_{r}$-norms is that the variables are measured on the same scale. If this is not the case, a standardisation should first be applied. This corresponds to using a more general $L_{2}$ - or Euclidean norm with a metric $\mathcal{A}$, where $\mathcal{A}>0$ (see Sect. 2.6):

$$
\begin{equation*}
d_{i j}^{2}=\left\|x_{i}-x_{j}\right\|_{\mathcal{A}}=\left(x_{i}-x_{j}\right)^{\top} \mathcal{A}\left(x_{i}-x_{j}\right) . \tag{13.4}
\end{equation*}
$$

$L_{2}$-norms are given by $\mathcal{A}=\mathcal{I}_{p}$, but if a standardisation is desired, then the weight matrix $\mathcal{A}=\operatorname{diag}\left(s_{X_{1} X_{1}}^{-1}, \ldots, s_{X_{p} X_{p}}^{-1}\right)$ may be suitable. Recall that $s_{X_{k} X_{k}}$ is the variance of the $k$-th component. Hence we have

$$
\begin{equation*}
d_{i j}^{2}=\sum_{k=1}^{p} \frac{\left(x_{i k}-x_{j k}\right)^{2}}{s_{X_{k} X_{k}}} . \tag{13.5}
\end{equation*}
$$

Here each component has the same weight in the computation of the distances and the distances do not depend on a particular choice of the units of measure.

Example 13.3 Consider the French Food expenditures (Table 22.6). The Euclidean distance matrix (squared $L_{2}$-norm) is

$$
\mathcal{D}=10^{4} .\left(\begin{array}{rrrrrrrrrrr}
0.00 & 5.82 & 58.19 & 3.54 & 5.15 & 151.44 & 16.91 & 36.15 & 147.99 & 51.84 & 102.56 \\
& 0.00 & 41.73 & 4.53 & 2.93 & 120.59 & 13.52 & 25.39 & 116.31 & 43.68 & 76.81 \\
& & 0.00 & 44.14 & 40.10 & 24.12 & 29.95 & 8.17 & 25.57 & 20.81 & 20.30 \\
\\
& & 0.00 & 0.76 & 127.85 & 5.62 & 21.70 & 124.98 & 31.21 & 72.97 & 231.57 \\
& & & 0.00 & 121.05 & 5.70 & 19.85 & 118.77 & 30.82 & 67.39 & 220.72 \\
& & & & & 0.00 & 96.57 & 48.16 & 1.80 & 60.52 & 28.90 \\
29.56 \\
& & & & & 0.00 & 9.20 & 94.87 & 11.07 & 42.12 & 179.84 \\
& & & & & & & 0.00 & 46.95 & 6.17 & 18.76 \\
113.03 \\
& & & & & & & & 0.00 & 61.08 & 29.62 \\
31.86 \\
& & & & & & & & & 0.00 & 15.83 \\
116.11 \\
& & & & & & & & & 0.00 & 53.77 \\
& & & & & & & & & & 0.00
\end{array}\right) .
$$

Taking the weight matrix $\mathcal{A}=\operatorname{diag}\left(s_{X_{1} X_{1}}^{-1}, \ldots, s_{X_{7} X_{7}}^{-1}\right)$, we obtain the distance matrix (squared $L_{2}$-norm)

$$
\mathcal{D}=\left(\begin{array}{rrrrrrrrrrr}
0.00 & 6.85 & 10.04 & 1.68 & 2.66 & 24.90 & 8.28 & 8.56 & 24.61 & 21.55 & 30.68  \tag{13.6}\\
& 0.00 & 13.11 & 6.59 & 3.75 & 20.12 & 13.13 & 12.38 & 15.88 & 31.52 & 25.65 \\
& & 0.00 & 8.03 & 7.27 & 4.99 & 9.27 & 3.88 & 7.46 & 14.92 & 15.08 \\
& & 0.00 & 0.64 & 20.06 & 2.76 & 3.82 & 19.63 & 12.81 & 19.28 & 45.01 \\
& & & 0.00 & 17.00 & 3.54 & 3.81 & 15.76 & 14.98 & 16.89 & 39.87 \\
& & & & 0.00 & 17.51 & 9.79 & 1.58 & 21.32 & 11.36 & 13.40 \\
& & & & & 0.00 & 1.80 & 17.92 & 4.39 & 9.93 & 33.61 \\
& & & & & & & 0.00 & 10.50 & 5.70 & 7.97 \\
& & & & & & & 0.00 & 24.75 & 11.02 & 13.07 \\
& & & & & & & & & 0.00 & 9.13 \\
& & & & & & & & 29.78 \\
& & & & & & & & & 0.00 & 9.39 \\
& & & & & & & & & & \\
& & & & 0.00
\end{array}\right) .
$$

When applied to contingency tables, a $\chi^{2}$-metric is suitable to compare (and cluster) rows and columns of a contingency table.

If $\mathcal{X}$ is a contingency table, row $i$ is characterised by the conditional frequency distribution $\frac{x_{i j}}{x_{i} \bullet}$, where $x_{i \bullet}=\sum_{j=1}^{p} x_{i j}$ indicates the marginal distributions over the rows: $\frac{x_{i}}{x_{\bullet \bullet}}, x_{\bullet \bullet}=\sum_{i=1}^{n} x_{i}$. Similarly, column $j$ of $\mathcal{X}$ is characterised by the conditional frequencies $\frac{x_{i j}}{x_{\bullet j}}$, where $x_{\bullet j}=\sum_{i=1}^{n} x_{i j}$. The marginal frequencies of the columns are $\frac{x_{\bullet j}}{x_{\bullet \bullet}}$.

The distance between two rows, $i_{1}$ and $i_{2}$, corresponds to the distance between their respective frequency distributions. It is common to define this distance using the $\chi^{2}$-metric:

$$
\begin{equation*}
d^{2}\left(i_{1}, i_{2}\right)=\sum_{j=1}^{p} \frac{1}{\left(\frac{x_{\bullet} j}{x_{\bullet \bullet}}\right)}\left(\frac{x_{i_{1} j}}{x_{i_{1} \bullet}}-\frac{x_{i_{2} j}}{x_{i_{2} \bullet}}\right)^{2} . \tag{13.7}
\end{equation*}
$$

Note that this can be expressed as a distance between the vectors $x_{1}=\left(\frac{x_{i_{1 j}}}{x_{\bullet \bullet}}\right)$ and $x_{2}=\left(\frac{x_{i_{2} j}}{x_{\bullet \bullet}}\right)$ as in (13.4) with weighting matrix $\mathcal{A}=\left\{\operatorname{diag}\left(\frac{x_{\bullet j}}{x_{\bullet \bullet \bullet}}\right)\right\}^{-1}$. Similarly, if we are interested in clusters among the columns, we can define:

$$
d^{2}\left(j_{1}, j_{2}\right)=\sum_{i=1}^{n} \frac{1}{\left(\frac{x_{i}}{x_{\bullet \bullet}}\right)}\left(\frac{x_{i j_{1}}}{x_{\bullet j_{1}}}-\frac{x_{i j_{2}}}{x_{\bullet j_{2}}}\right)^{2} .
$$

Apart from the Euclidean and the $L_{r}$-norm measures one can use a proximity measure such as the $Q$-correlation coefficient

$$
\begin{equation*}
d_{i j}=\frac{\sum_{k=1}^{p}\left(x_{i k}-\bar{x}_{i}\right)\left(x_{j k}-\bar{x}_{j}\right)}{\left\{\sum_{k=1}^{p}\left(x_{i k}-\bar{x}_{i}\right)^{2} \sum_{k=1}^{p}\left(x_{j k}-\bar{x}_{j}\right)^{2}\right\}^{1 / 2}} . \tag{13.8}
\end{equation*}
$$

Here $\bar{x}_{i}$ denotes the mean over the variables $\left(x_{i 1}, \ldots, x_{i p}\right)$.

| $\hookrightarrow$ | The proximity between data points is measured by a distance <br> or similarity matrix $\mathcal{D}$ whose components $d_{i j}$ give the similarity <br> coefficient or the distance between two points $x_{i}$ and $x_{j}$. |
| :--- | :--- |
| $\hookrightarrow$ | A variety of similarity (distance) measures exist for binary data <br> (e.g. Jaccard, Tanimoto, Simple Matching coefficients) and for <br> continuous data (e.g. $L_{r}$-norms). |
| $\hookrightarrow$ | The nature of the data could impose the choice of a particular <br> metric $\mathcal{A}$ in defining the distances (standardisation, $\chi^{2}$-metric etc.). |

### 13.3 Cluster Algorithms

There are essentially two types of clustering methods: hierarchical algorithms and partitioning algorithms. The hierarchical algorithms can be divided into agglomerative and splitting procedures. The first type of hierarchical clustering starts from the finest partition possible (each observation forms a cluster) and groups them. The second type starts with the coarsest partition possible: one cluster contains all of the observations. It proceeds by splitting the single cluster up into smaller sized clusters.

The partitioning algorithms start from a given group definition and proceed by exchanging elements between groups until a certain score is optimised. The main difference between the two clustering techniques is that in hierarchical clustering once groups are found and elements are assigned to the groups, this assignment cannot be changed. In partitioning techniques, on the other hand, the assignment of objects into groups may change during the algorithm application.

## Hierarchical Algorithms, Agglomerative Techniques

Agglomerative algorithms are used quite frequently in practice. The algorithm consists of the following steps:

```
Algorithm Hierarchical algorithms-agglomerative technique
    Construct the finest partition
    Compute the distance matrix \(\mathcal{D}\).
    repeat
        Find the two clusters with the closest distance
        Put those two clusters into one cluster
        Compute the distance between the new groups and obtain a reduced distance matrix \(\mathcal{D}\)
    until all clusters are agglomerated into \(\mathcal{X}\)
```

If two objects or groups say, $P$ and $Q$, are united, one computes the distance between this new group (object) $P+Q$ and group $R$ using the following distance function:

$$
\begin{equation*}
d(R, P+Q)=\delta_{1} d(R, P)+\delta_{2} d(R, Q)+\delta_{3} d(P, Q)+\delta_{4}|d(R, P)-d(R, Q)| . \tag{13.9}
\end{equation*}
$$

The $\delta_{j}$ 's are weighting factors that lead to different agglomerative algorithms as described in Table 13.2. Here $n_{P}=\sum_{i=1}^{n} \boldsymbol{I}\left(x_{i} \in P\right)$ is the number of objects in group $P$. The values of $n_{Q}$ and $n_{R}$ are defined analogously.

For the most common used Single and Complete linkages, below are the modified agglomerative algorithm steps:

As instead of computing new distance matrixes every step, a linear search in the original distance matrix is enough for clustering in the modified algorithm, it is more efficient in practice.

Table 13.2 Computations of group distances

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

```
Algorithm Modified hierarchical algorithms-agglomerative technique
    Construct the finest partition
    Compute the distance matrix \(\mathcal{D}\).
    repeat
        Find the smallest (Single linkage)/ largest (Complete linkage) value \(d\) (between objects \(m\)
    and \(n\) ) in \(\mathcal{D}\)
            If \(m\) and \(n\) are not in the same cluster, combine the clusters \(m\) and \(n\) belonging to together,
    and delete the smallest value
    until all clusters are agglomerated into \(\mathcal{X}\) or the value \(d\) exceeds the preset level
```

Example 13.4 Let us examine the agglomerative algorithm for the three points in Example 13.2, $x_{1}=(0,0), x_{2}=(1,0)$ and $x_{3}=(5,5)$, and the squared Euclidean distance matrix with single linkage weighting. The algorithm starts with $N=3$ clusters: $P=\left\{x_{1}\right\}, Q=\left\{x_{2}\right\}$ and $R=\left\{x_{3}\right\}$. The distance matrix $\mathcal{D}_{2}$ is given in Example 13.2. The smallest distance in $\mathcal{D}_{2}$ is the one between the clusters $P$ and $Q$. Therefore, applying step 4 in the above algorithm we combine these clusters to form $P+Q=\left\{x_{1}, x_{2}\right\}$. The single linkage distance between the remaining two clusters is from Table 13.2 and (13.9) equal to

$$
\begin{align*}
d(R, P+Q) & =\frac{1}{2} d(R, P)+\frac{1}{2} d(R, Q)-\frac{1}{2}|d(R, P)-d(R, Q)| \\
& =\frac{1}{2} d_{13}+\frac{1}{2} d_{23}-\frac{1}{2} \cdot\left|d_{13}-d_{23}\right| \\
& =\frac{50}{2}+\frac{41}{2}-\frac{1}{2} \cdot|50-41| \\
& =41 . \tag{13.10}
\end{align*}
$$

The reduced distance matrix is then $\left(\begin{array}{cc}0 & 41 \\ 41 & 0\end{array}\right)$. The next and last step is to unite the clusters $R$ and $P+Q$ into a single cluster $\mathcal{X}$, the original data matrix.

When there are more data points than in the example above, a visualisation of the implication of clusters is desirable. A graphical representation of the sequence of clustering is called a dendrogram. It displays the observations, the sequence of clusters and the distances between the clusters. The vertical axis displays the indices of the points, whereas the horizontal axis gives the distance between the clusters. Large distances indicate the clustering of heterogeneous groups. Thus, if we choose to "cut the tree" at a desired level, the branches describe the corresponding clusters.

Fig. 13.1 The 8-point example Q MVAclus8p

8 points


Example 13.5 Here we describe the single linkage algorithm for the eight data points displayed in Fig. 13.1. The distance matrix ( $L_{2}$-norms) is
and the dendrogram is shown in Fig. 13.2.
If we decide to cut the tree at the level 10 , three clusters are defined: $\{1,2\}$, $\{3,4,5\}$ and $\{6,7,8\}$.

The single linkage algorithm defines the distance between two groups as the smallest value of the individual distances. Table 13.2 shows that in this case

$$
\begin{equation*}
d(R, P+Q)=\min \{d(R, P), d(R, Q)\} . \tag{13.11}
\end{equation*}
$$

This algorithm is also called the Nearest Neighbour algorithm. As a consequence of its construction, single linkage tends to build large groups. Groups that differ but are not well separated may thus be classified into one group as long as they have

Fig. 13.2 The dendrogram for the 8 -point example, single linkage algorithm O MVAclus8p

Single Linkage Dendrogram - 8 points

two approximate points. The complete linkage algorithm tries to correct this kind of grouping by considering the largest (individual) distances. Indeed, the complete linkage distance can be written as

$$
\begin{equation*}
d(R, P+Q)=\max \{d(R, P), d(R, Q)\} . \tag{13.12}
\end{equation*}
$$

It is also called the Farthest Neighbour algorithm. This algorithm will cluster groups where all the points are proximate, since it compares the largest distances. The average linkage algorithm (weighted or unweighted) proposes a compromise between the two preceding algorithms, in that it computes an average distance:

$$
\begin{equation*}
d(R, P+Q)=\frac{n_{P}}{n_{P}+n_{Q}} d(R, P)+\frac{n_{Q}}{n_{P}+n_{Q}} d(R, Q) \tag{13.13}
\end{equation*}
$$

The centroid algorithm is quite similar to the average linkage algorithm and uses the natural geometrical distance between $R$ and the weighted centre of gravity of $P$ and $Q$ (see Fig. 13.3):
$d(R, P+Q)=\frac{n_{P}}{n_{P}+n_{Q}} d(R, P)+\frac{n_{Q}}{n_{P}+n_{Q}} d(R, Q)-\frac{n_{P} n_{Q}}{\left(n_{P}+n_{Q}\right)^{2}} d(P, Q)$.

The Ward clustering algorithm computes the distance between groups according to the formula in Table 13.2. The main difference between this algorithm and the linkage procedures is in the unification procedure. The Ward algorithm does not put together groups with smallest distance. Instead, it joins groups that do not increase a given measure of heterogeneity "too much". The aim of the Ward procedure is

Fig. 13.3 The centroid algorithm

weighted centre of gravity of $P+Q$
to unify groups such that the variation inside these groups does not increase too drastically: the resulting groups are as homogeneous as possible.

The heterogeneity of group $R$ is measured by the inertia inside the group. This inertia is defined as follows:

$$
\begin{equation*}
I_{R}=\frac{1}{n_{R}} \sum_{i=1}^{n_{R}} d^{2}\left(x_{i}, \bar{x}_{R}\right) \tag{13.15}
\end{equation*}
$$

where $\bar{x}_{R}$ is the centre of gravity (mean) over the groups. $I_{R}$ clearly provides a scalar measure of the dispersion of the group around its centre of gravity. If the usual Euclidean distance is used, then $I_{R}$ represents the sum of the variances of the $p$ components of $x_{i}$ inside group $R$.

When two objects or groups $P$ and $Q$ are joined, the new group $P+Q$ has a larger inertia $I_{P+Q}$. It can be shown that the corresponding increase of inertia is given by

$$
\begin{equation*}
\Delta(P, Q)=\frac{n_{P} n_{Q}}{n_{P}+n_{Q}} d^{2}(P, Q) \tag{13.16}
\end{equation*}
$$

In this case, the Ward algorithm is defined as an algorithm that "joins the groups that give the smallest increase in $\Delta(P, Q)$ ". It is easy to prove that when $P$ and $Q$ are joined, the new criterion values are given by (13.9) along with the values of $\delta_{i}$ given in Table 13.2, when the centroid formula is used to modify $d^{2}(R, P+Q)$. So, the Ward algorithm is related to the centroid algorithm, but with an "inertial" distance $\Delta$ rather than the "geometric" distance $d^{2}$.

As pointed out in Sect. 13.2, all the algorithms above can be adjusted by the choice of the metric $\mathcal{A}$ defining the geometric distance $d^{2}$. If the results of a clustering algorithm are illustrated as graphical representations of individuals in spaces of low dimension (using principal components (normalised or not) or using a correspondence analysis for contingency tables), it is important to be coherent in the choice of the metric used.

Fig. 13.4 PCA for 20
randomly chosen bank notes Q MVAclusbank

Fig. 13.5 The dendrogram for the 20 bank notes, Ward algorithm O MVAclusbank


Example 13.6 As an example we randomly select 20 observations from the bank notes data and apply the Ward technique using Euclidean distances. Figure 13.4 shows the first two PCs of these data, Fig. 13.5 displays the dendrogram.

Example 13.7 Consider the French food expenditures. As in Chap. 11 we use standardised data which is equivalent to using $\mathcal{A}=\operatorname{diag}\left(s_{X_{1} X_{1}}^{-1}, \ldots, s_{X_{7} X_{7}}^{-1}\right)$ as the weight matrix in the $L_{2}$-norm. The NPCA plot of the individuals was given

Fig. 13.6 The dendrogram for the French food expenditures, Ward algorithm Q MVAclusfood

in Fig. 11.7. The Euclidean distance matrix is of course given by (13.6). The dendrogram obtained by using the Ward algorithm is shown in Fig. 13.6.

If the aim was to have only two groups, as can be seen in Fig. 13.6, they would be \{CA2, CA3, CA4, CA5, EM5\} and \{MA2, MA3, MA4, MA5, EM2, EM3, EM4\}. Clustering three groups is somewhat arbitrary (the levels of the distances are too similar). If we were interested in four groups, we would obtain \{CA2, CA3, CA4\}, \{EM2, MA2, EM3, MA3\}, \{EM4, MA4, MA5\} and \{EM5, CA5\}. This grouping shows a balance between socio-professional levels and size of the families in determining the clusters. The four groups are clearly well represented in the NPCA plot in Fig. 11.7.

| $\hookrightarrow$ | The class of clustering algorithms can be divided into two types: <br> hierarchical and partitioning algorithms. Hierarchical algorithms <br> start with the finest (coarsest) possible partition and put groups <br> together (split groups apart) step by step. Partitioning algorithms <br> start from a preliminary clustering and exchange group elements <br> until a certain score is reached. |
| :--- | :--- |
|  | Hierarchical agglomerative techniques are frequently used in prac- <br> tice. They start from the finest possible structure (each data point <br> forms a cluster), compute the distance matrix for the clusters <br> and join the clusters that have the smallest distance. This step is <br> repeated until all points are united in one cluster. |


| Summary (continued) |  |
| :---: | :---: |
|  | The agglomerative procedure depends on the definition of the <br> distance between two clusters. Single linkage, complete linkage, <br> and Ward distance are frequently used distances. |
| $\qquad$ | The process of the unification of clusters can be graphically <br> represented by a dendrogram. |

### 13.4 Boston Housing

Presented multivariate techniques are now applied to the Boston Housing data. We focus our attention to 14 transformed and standardised variables, see e.g. Fig. 13.7 that provides descriptive statistics via boxplots for two clusters, as discussed in the


Fig. 13.7 Boxplots of the 14 standardised variables of the Boston housing data $\mathbf{Q}$ MVAclusbh

Fig. 13.8 Dendrogram of the Boston housing data using the Ward algorithm 0 MVAclusbh


Table 13.3 Means and standard errors of the 13 standardised variables for Cluster 1 (251 observations) and Cluster 2 (255 observations) Q MVAclusbh

| Variable | Mean C1 | SE C1 | Mean C2 | SE C2 |
| :--- | ---: | :--- | ---: | ---: |
| 1 | -0.7105 | 0.0332 | 0.6994 | 0.0535 |
| 2 | 0.4848 | 0.0786 | -0.4772 | 0.0047 |
| 3 | -0.7665 | 0.0510 | 0.7545 | 0.0279 |
| 5 | -0.7672 | 0.0365 | 0.7552 | 0.0447 |
| 6 | 0.4162 | 0.0571 | -0.4097 | 0.0576 |
| 7 | -0.7730 | 0.0429 | 0.7609 | 0.0378 |
| 8 | 0.7140 | 0.0472 | -0.7028 | 0.0417 |
| 9 | -0.5429 | 0.0358 | 0.5344 | 0.0656 |
| 10 | -0.6932 | 0.0301 | 0.6823 | 0.0569 |
| 11 | -0.5464 | 0.0469 | 0.5378 | 0.0582 |
| 12 | 0.3547 | 0.0080 | -0.3491 | 0.0824 |
| 13 | -0.6899 | 0.0401 | 0.6791 | 0.0509 |
| 14 | 0.5996 | 0.0431 | -0.5902 | 0.0570 |

sequel. A dendrogram for 13 variables(excluding the dummy variable $\tilde{X}_{4}$ - Charles River indicator) using the Ward method is displayed in Fig. 13.8. One observes two dominant clusters. A further refinement of say, four clusters, could be considered at a lower level of distance.

To interpret the two clusters, we present the mean values and their respective standard errors of the $13 \tilde{\mathcal{X}}$ variables by groups in Table 13.3. Comparison of the mean values for both groups shows that all the differences in the means are individually significant. Moreover, cluster one corresponds to housing districts with better living quality and higher house prices, whereas cluster two corresponds to less favored districts in Boston. This can be confirmed, for instance, by a lower crime rate, a higher proportion of residential land, lower proportion of African American,


Fig. 13.9 Scatterplot matrix for variables $\tilde{X}_{1}$ to $\tilde{X}_{7}$ of the Boston housing data $Q$ mVAclusbh
etc. for cluster one. Cluster two is identified by a higher proportion of older houses, a higher pupil/teacher ratio and a higher percentage of the lower status population.

This interpretation is underlined by visual inspection of all the variables via scatterplot matrices, see e.g. Figs. 13.9 and 13.10. For example, the lower right boxplot of Fig. 13.7 and the correspondingly coloured clusters in the last row of Fig. 13.10 confirm the role of each variable in determining the clusters. This interpretation perfectly coincides with the previous PC analysis (Fig. 11.11). The quality of life factor is clearly visible in Fig. 13.11, where cluster membership is


Fig. 13.10 Scatterplot matrix for variables $\tilde{X}_{8}$ to $\tilde{X}_{14}$ of the Boston housing data $Q$ mVAclusbh
distinguished by the shape and colour of the points graphed according to the first two principal components. Clearly, the first PC completely separates the two clusters and corresponds, as we have discussed in Chap. 11, to a quality of life and house indicator.

Fig. 13.11 Scatterplot of the first two PCs displaying the two clusters $\mathbf{Q}$ MVAclusbh


### 13.5 Exercises

Exercise 13.1 Prove formula (13.16).
Exercise 13.2 Prove that $I_{R}=\operatorname{tr}\left(\mathcal{S}_{R}\right)$, where $\mathcal{S}_{R}$ denotes the empirical covariance matrix of the observations contained in $R$.
Exercise 13.3 Prove that

$$
\begin{aligned}
\Delta(R, P+Q)= & \frac{n_{R}+n_{P}}{n_{R}+n_{P}+n_{Q}} \Delta(R, P)+\frac{n_{R}+n_{Q}}{n_{R}+n_{P}+n_{Q}} \Delta(R, Q) \\
& -\frac{n_{R}}{n_{R}+n_{P}+n_{Q}} \Delta(P, Q)
\end{aligned}
$$

when the centroid formula is used to define $d^{2}(R, P+Q)$.
Exercise 13.4 Repeat the 8-point example (Example 13.5) using the complete linkage and the Ward algorithm. Explain the difference to single linkage.

Exercise 13.5 Explain the differences between various proximity measures by means of an example.
Exercise 13.6 Repeat the bank notes example (Example 13.6) with another random sample of 20 notes.
Exercise 13.7 Repeat the bank notes example (Example 13.6) with another clustering algorithm.

Exercise 13.8 Repeat the bank notes example (Example 13.6) or the 8-point example (Example 13.5) with the $L_{1}$-norm.

Exercise 13.9 Analyse the US companies example (Table 22.5) using the Ward algorithm and the $L_{2}$-norm.

Exercise 13.10 Analyse the US crime data set (Table 22.10) with the Ward algorithm and the $L_{2}$-norm on standardised variables (use only the crime variables).

Exercise 13.11 Repeat Exercise 13.10 with the US health data set (use only the number of deaths variables).

Exercise 13.12 Redo Exercise 13.10 with the $\chi^{2}$-metric. Compare the results.
Exercise 13.13 Redo Exercise 13.11 with the $\chi^{2}$-metric and compare the results.

## Chapter 14 <br> Discriminant Analysis

Discriminant analysis is used in situations where the clusters are known a priori. The aim of discriminant analysis is to classify an observation, or several observations, into these known groups. For instance, in credit scoring, a bank knows from past experience that there are good customers (who repay their loan without any problems) and bad customers (who showed difficulties in repaying their loan). When a new customer asks for a loan, the bank has to decide whether or not to give the loan. The past records of the bank provides two data sets: multivariate observations $x_{i}$ on the two categories of customers (including for example age, salary, marital status, the amount of the loan, etc.). The new customer is a new observation $x$ with the same variables. The discrimination rule has to classify the customer into one of the two existing groups and the discriminant analysis should evaluate the risk of a possible "bad decision".

Many other examples are described below, and in most applications, the groups correspond to natural classifications or to groups known from history (like in the credit scoring example). These groups could have been formed by a cluster analysis performed on past data.

Section 14.1 presents the allocation rules when the populations are known, i.e. when we know the distribution of each population. As described in Sect. 14.2 in practice the population characteristics have to be estimated from history. The methods are illustrated in several examples.

### 14.1 Allocation Rules for Known Distributions

Discriminant analysis is a set of methods and tools used to distinguish between groups of populations $\Pi_{j}$ and to determine how to allocate new observations into groups. In one of our running examples we are interested in discriminating between counterfeit and true bank notes on the basis of measurements of these bank notes,
see Sect. 22.2. In this case we have two groups (counterfeit and genuine bank notes) and we would like to establish an algorithm (rule) that can allocate a new observation (a new bank note) into one of the groups.

Another example is the detection of "fast" and "slow" consumers of a newly introduced product. Using a consumer's characteristics like education, income, family size, amount of previous brand switching, we want to classify each consumer into the two groups just identified.

In poetry and literary studies the frequencies of spoken or written words and lengths of sentences indicate profiles of different artists and writers. It can be of interest to attribute unknown literary or artistic works to certain writers with a specific profile. Anthropological measures on ancient sculls help in discriminating between male and female bodies. Good and poor credit risk ratings constitute a discrimination problem that might be tackled using observations on income, age, number of credit cards, family size, etc.

In general we have populations $\Pi_{j}, j=1,2, \ldots, J$ and we have to allocate an observation $x$ to one of these groups. A discriminant rule is a separation of the sample space (in general $\mathbb{R}^{p}$ ) into sets $R_{j}$ such that if $x \in R_{j}$, it is identified as a member of population $\Pi_{j}$.

The main task of discriminant analysis is to find "good" regions $R_{j}$ such that the error of misclassification is small. In the following we describe such rules when the population distributions are known.

## Maximum Likelihood Discriminant Rule

Denote the densities of each population $\Pi_{j}$ by $f_{j}(x)$. The maximum likelihood discriminant rule (ML rule) is given by allocating $x$ to $\Pi_{j}$ maximising the likelihood $L_{j}(x)=f_{j}(x)=\arg \max _{i} f_{i}(x)$.

If several $f_{i}$ give the same maximum then any of them may be selected. Mathematically, the sets $R_{j}$ given by the ML discriminant rule are defined as

$$
\begin{equation*}
R_{j}=\left\{x: L_{j}(x)>L_{i}(x) \text { for } i=1, \ldots, J, i \neq j\right\} . \tag{14.1}
\end{equation*}
$$

By classifying the observation into a certain group we may encounter a misclassification error. For $J=2$ groups the probability of putting $x$ into group 2 although it is from population 1 can be calculated as

$$
\begin{equation*}
p_{21}=\mathrm{P}\left(X \in R_{2} \mid \Pi_{1}\right)=\int_{R_{2}} f_{1}(x) d x \tag{14.2}
\end{equation*}
$$

Similarly the conditional probability of classifying an object as belonging to the first population $\Pi_{1}$ although it actually comes from $\Pi_{2}$ is

$$
\begin{equation*}
p_{12}=\mathrm{P}\left(X \in R_{1} \mid \Pi_{2}\right)=\int_{R_{1}} f_{2}(x) d x \tag{14.3}
\end{equation*}
$$

The misclassified observations create a cost $C(i \mid j)$ when a $\Pi_{j}$ observation is assigned to $R_{i}$. In the credit risk example, this might be the cost of a "sour" credit. The cost structure can be pinned down in a cost matrix:


Let $\pi_{j}$ be the prior probability of population $\Pi_{j}$, where "prior" means the a priori probability that an individual selected at random belongs to $\Pi_{j}$ (i.e. before looking to the value $x$ ). Prior probabilities should be considered if it is clear ahead of time that an observation is more likely to stem from a certain population $\Pi_{j}$. An example is the classification of musical tunes. If it is known that during a certain period of time a majority of tunes were written by a certain composer, then there is a higher probability that a certain tune was composed by this composer. Therefore, he should receive a higher prior probability when tunes are assigned to a specific group.

The expected cost of misclassification (ECM) is given by

$$
\begin{equation*}
\mathrm{ECM}=C(2 \mid 1) p_{21} \pi_{1}+C(1 \mid 2) p_{12} \pi_{2} \tag{14.4}
\end{equation*}
$$

We will be interested in classification rules that keep the ECM small or minimise it over a class of rules. The discriminant rule minimising the ECM (14.4) for two populations is given below.
Theorem 14.1 For two given populations, the rule minimising the ECM is given by

$$
\begin{aligned}
& R_{1}=\left\{x: \frac{f_{1}(x)}{f_{2}(x)} \geq\left(\frac{C(1 \mid 2)}{C(2 \mid 1)}\right)\left(\frac{\pi_{2}}{\pi_{1}}\right)\right\} \\
& R_{2}=\left\{x: \frac{f_{1}(x)}{f_{2}(x)}<\left(\frac{C(1 \mid 2)}{C(2 \mid 1)}\right)\left(\frac{\pi_{2}}{\pi_{1}}\right)\right\}
\end{aligned}
$$

The ML discriminant rule is thus a special case of the ECM rule for equal misclassification costs and equal prior probabilities. For simplicity the unity cost case, $C(1 \mid 2)=C(2 \mid 1)=1$, and equal prior probabilities, $\pi_{2}=\pi_{1}$, are assumed in the following.

Theorem 14.1 will be proven by an example from credit scoring.
Example 14.1 Suppose that $\Pi_{1}$ represents the population of bad clients who create the cost $C(2 \mid 1)$ if they are classified as good clients. Analogously, define $C(1 \mid 2)$ as the cost of loosing a good client classified as a bad one. Let $\gamma$ denote the gain of the bank for the correct classification of a good client. The total gain of the bank is then

$$
\begin{aligned}
G\left(R_{2}\right)= & -C(2 \mid 1) \pi_{1} \int \boldsymbol{I}\left(x \in R_{2}\right) f_{1}(x) d x \\
& -C(1 \mid 2) \pi_{2} \int\left\{1-\boldsymbol{I}\left(x \in R_{2}\right)\right\} f_{2}(x) d x+\gamma \pi_{2} \int \boldsymbol{I}\left(x \in R_{2}\right) f_{2}(x) d x \\
= & -C(1 \mid 2) \pi_{2}+\int \boldsymbol{I}\left(x \in R_{2}\right)\left\{-C(2 \mid 1) \pi_{1} f_{1}(x)\right. \\
& \left.+(C(1 \mid 2)+\gamma) \pi_{2} f_{2}(x)\right\} d x
\end{aligned}
$$

Since the first term in this equation is constant, the maximum is obviously obtained for

$$
R_{2}=\left\{x:-C(2 \mid 1) \pi_{1} f_{1}(x)+\{C(1 \mid 2)+\gamma\} \pi_{2} f_{2}(x) \geq 0\right\} .
$$

This is equivalent to

$$
R_{2}=\left\{x: \frac{f_{2}(x)}{f_{1}(x)} \geq \frac{C(2 \mid 1) \pi_{1}}{\{C(1 \mid 2)+\gamma\} \pi_{2}}\right\}
$$

which corresponds to the set $R_{2}$ in Theorem 14.1 for a gain of $\gamma=0$.
Example 14.2 Suppose $x \in\{0,1\}$ and

$$
\begin{aligned}
& \Pi_{1}: \mathrm{P}(X=0)=\mathrm{P}(X=1)=\frac{1}{2} \\
& \Pi_{2}: \mathrm{P}(X=0)=\frac{1}{4}=1-\mathrm{P}(X=1)
\end{aligned}
$$

The sample space is the set $\{0,1\}$. The ML discriminant rule is to allocate $x=0$ to $\Pi_{1}$ and $x=1$ to $\Pi_{2}$, defining the sets $R_{1}=\{0\}, R_{2}=\{1\}$ and $R_{1} \cup R_{2}=\{0,1\}$.

Example 14.3 Consider two normal populations

$$
\begin{aligned}
& \Pi_{1}: N\left(\mu_{1}, \sigma_{1}^{2}\right), \\
& \Pi_{2}: N\left(\mu_{2}, \sigma_{2}^{2}\right)
\end{aligned}
$$

Then

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

Hence $x$ is allocated to $\Pi_{1}\left(x \in R_{1}\right)$ if $L_{1}(x) \geq L_{2}(x)$. Note that $L_{1}(x) \geq L_{2}(x)$ is equivalent to

$$
\frac{\sigma_{2}}{\sigma_{1}} \exp \left[-\frac{1}{2}\left\{\left(\frac{x-\mu_{1}}{\sigma_{1}}\right)^{2}-\left(\frac{x-\mu_{2}}{\sigma_{2}}\right)^{2}\right\}\right] \geq 1
$$

or

$$
\begin{equation*}
x^{2}\left(\frac{1}{\sigma_{1}^{2}}-\frac{1}{\sigma_{2}^{2}}\right)-2 x\left(\frac{\mu_{1}}{\sigma_{1}^{2}}-\frac{\mu_{2}}{\sigma_{2}^{2}}\right)+\left(\frac{\mu_{1}^{2}}{\sigma_{1}^{2}}-\frac{\mu_{2}^{2}}{\sigma_{2}^{2}}\right) \leq 2 \log \frac{\sigma_{2}}{\sigma_{1}} . \tag{14.5}
\end{equation*}
$$

Suppose that $\mu_{1}=0, \sigma_{1}=1$ and $\mu_{2}=1, \sigma_{2}=\frac{1}{2}$. Formula (14.5) leads to

$$
\begin{aligned}
& R_{1}=\left\{x: x \leq \frac{1}{3}(4-\sqrt{4+6 \log (2)}) \text { or } x \geq \frac{1}{3}(4+\sqrt{4+6 \log (2)})\right\}, \\
& R_{2}=\mathbb{R} \backslash R_{1} .
\end{aligned}
$$

This situation is shown in Fig. 14.1.

Fig. 14.1 Maximum
likelihood rule for normal distributions O MVAdisnorm


The situation simplifies in the case of equal variances $\sigma_{1}=\sigma_{2}$. The discriminant rule (14.5) is then (for $\mu_{1}<\mu_{2}$ )

$$
\begin{array}{ll}
x \rightarrow \Pi_{1}, \text { if } & x \in R_{1}=\left\{x: x \leq \frac{1}{2}\left(\mu_{1}+\mu_{2}\right)\right\}  \tag{14.6}\\
x \rightarrow \Pi_{2}, \text { if } & x \in R_{2}=\left\{x: x>\frac{1}{2}\left(\mu_{1}+\mu_{2}\right)\right\} .
\end{array}
$$

Theorem 14.2 shows that the ML discriminant rule for multinormal observations is intimately connected with the Mahalanobis distance. The discriminant rule is based on linear combinations and belongs to the family of linear discriminant analysis (LDA) methods.

Theorem 14.2 Suppose $\Pi_{i}=N_{p}\left(\mu_{i}, \Sigma\right)$.
(a) The ML rule allocates $x$ to $\Pi_{j}$, where $j \in\{1, \ldots, J\}$ is the value minimising the square Mahalanobis distance between $x$ and $\mu_{i}$ :

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

(b) In the case of $J=2$,

$$
x \in R_{1} \Longleftrightarrow \alpha^{\top}(x-\mu) \geq 0
$$

where $\alpha=\Sigma^{-1}\left(\mu_{1}-\mu_{2}\right)$ and $\mu=\frac{1}{2}\left(\mu_{1}+\mu_{2}\right)$.
Proof Part (a) of the theorem follows directly from comparison of the likelihoods.
For $J=2$, part (a) says that $x$ is allocated to $\Pi_{1}$ if

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

Rearranging terms leads to

$$
-2 \mu_{1}^{\top} \Sigma^{-1} x+2 \mu_{2}^{\top} \Sigma^{-1} x+\mu_{1}^{\top} \Sigma^{-1} \mu_{1}-\mu_{2}^{\top} \Sigma^{-1} \mu_{2} \leq 0
$$

which is equivalent to

$$
\begin{gathered}
2\left(\mu_{2}-\mu_{1}\right)^{\top} \Sigma^{-1} x+\left(\mu_{1}-\mu_{2}\right)^{\top} \Sigma^{-1}\left(\mu_{1}+\mu_{2}\right) \leq 0, \\
\left(\mu_{1}-\mu_{2}\right)^{\top} \Sigma^{-1}\left\{x-\frac{1}{2}\left(\mu_{1}+\mu_{2}\right)\right\} \geq 0, \\
\alpha^{\top}(x-\mu) \geq 0 .
\end{gathered}
$$

## Bayes Discriminant Rule

We have seen an example where prior knowledge on the probability of classification into $\Pi_{j}$ was assumed. Denote the prior probabilities by $\pi_{j}$ and note that $\sum_{j=1}^{J} \pi_{j}=1$. The Bayes rule of discrimination allocates $x$ to the $\Pi_{j}$ that gives the largest value of $\pi_{i} f_{i}(x), \pi_{j} f_{j}(x)=\max _{i} \pi_{i} f_{i}(x)$. Hence, the discriminant rule is defined by $R_{j}=\left\{x: \pi_{j} f_{j}(x) \geq \pi_{i} f_{i}(x)\right.$ for $\left.i=1, \ldots, J\right\}$. Obviously the Bayes rule is identical to the ML discriminant rule for $\pi_{j}=1 / J$.

A further modification is to allocate $x$ to $\Pi_{j}$ with a certain probability $\phi_{j}(x)$, such that $\sum_{j=1}^{J} \phi_{j}(x)=1$ for all $x$. This is called a randomised discriminant rule. A randomised discriminant rule is a generalisation of deterministic discriminant rules since

$$
\phi_{j}(x)= \begin{cases}1 & \text { if } \quad \pi_{j} f_{j}(x)=\max _{i} \pi_{i} f_{i}(x) \\ 0 & \text { otherwise }\end{cases}
$$

reflects the deterministic rules.
Which discriminant rules are good? We need a measure of comparison. Denote

$$
\begin{equation*}
p_{i j}=\int \phi_{i}(x) f_{j}(x) d x \tag{14.7}
\end{equation*}
$$

as the probability of allocating $x$ to $\Pi_{i}$ if it in fact belongs to $\Pi_{j}$. A discriminant rule with probabilities $p_{i j}$ is as good as any other discriminant rule with probabilities $p_{i j}^{\prime}$ if

$$
\begin{equation*}
p_{i i} \geq p_{i i}^{\prime} \quad \text { for all } \quad i=1, \ldots, J \tag{14.8}
\end{equation*}
$$

We call the first rule better if the strict inequality in (14.8) holds for at least one i. A discriminant rule is called admissible if there is no better discriminant rule.

Theorem 14.3 All Bayes discriminant rules (including the ML rule) are admissible.

## Probability of Misclassification for the ML Rule ( $J=2$ )

Suppose that $\Pi_{i}=N_{p}\left(\mu_{i}, \Sigma\right)$. In the case of two groups, it is not difficult to derive the probabilities of misclassification for the ML discriminant rule. Consider for instance $p_{12}=\mathrm{P}\left(x \in R_{1} \mid \Pi_{2}\right)$. By part (b) in Theorem 14.2 we have

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

If $X \in R_{2}, \alpha^{\top}(X-\mu) \sim N\left(-\frac{1}{2} \delta^{2}, \delta^{2}\right)$ 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, the probability of being classified into population 2 although $x$ stems from $\Pi_{1}$ is equal to $p_{21}=\Phi\left(-\frac{1}{2} \delta\right)$.

## Classification with Different Covariance Matrices

The minimum ECM depends on the ratio of the densities $\frac{f_{1}(x)}{f_{2}(x)}$ or equivalently on the difference $\log \left\{f_{1}(x)\right\}-\log \left\{f_{2}(x)\right\}$. When the covariance for both density functions differ, the allocation rule becomes more complicated:

$$
\begin{aligned}
R_{1}=\left\{x:-\frac{1}{2} x^{\top}\left(\Sigma_{1}^{-1}-\Sigma_{2}^{-1}\right) x+\left(\mu_{1}^{\top} \Sigma_{1}^{-1}\right.\right. & \left.-\mu_{2}^{\top} \Sigma_{2}^{-1}\right) x-k \\
& \left.\geq \log \left[\left(\frac{C(1 \mid 2)}{C(2 \mid 1)}\right)\left(\frac{\pi_{2}}{\pi_{1}}\right)\right]\right\} \\
R_{2}=\left\{x:-\frac{1}{2} x^{\top}\left(\Sigma_{1}^{-1}-\Sigma_{2}^{-1}\right) x+\left(\mu_{1}^{\top} \Sigma_{1}^{-1}\right.\right. & \left.-\mu_{2}^{\top} \Sigma_{2}^{-1}\right) x-k \\
& \left.<\log \left[\left(\frac{C(1 \mid 2)}{C(2 \mid 1)}\right)\left(\frac{\pi_{2}}{\pi_{1}}\right)\right]\right\}
\end{aligned}
$$

where $k=\frac{1}{2} \log \left(\frac{\left|\Sigma_{1}\right|}{\left|\Sigma_{2}\right|}\right)+\frac{1}{2}\left(\mu_{1}^{\top} \Sigma_{1}^{-1} \mu_{1}-\mu_{2}^{\top} \Sigma_{2}^{-1} \mu_{2}\right)$. The classification regions are defined by quadratic functions. Therefore they belong to the family of quadratic discriminant analysis (QDA) methods. This quadratic classification rule coincides with the rules used when $\Sigma_{1}=\Sigma_{2}$, since the term $\frac{1}{2} x^{\top}\left(\Sigma_{1}^{-1}-\Sigma_{2}^{-1}\right) x$ disappears.


## Summary

$\hookrightarrow$ Discriminant analysis is a set of methods used to distinguish among groups in data and to allocate new observations into the existing groups.
$\hookrightarrow$ Given that data are from populations $\Pi_{j}$ with densities $f_{j}, j=$ $1, \ldots, J$, the maximum likelihood discriminant rule (ML rule) allocates an observation $x$ to that population $\Pi_{j}$ which has the maximum likelihood $L_{j}(x)=f_{j}(x)=\max _{i} f_{i}(x)$.

| Summary (continued) |  |
| :---: | :--- |
| $\hookrightarrow$ | Given prior probabilities $\pi_{j}$ for populations $\Pi_{j}$, Bayes discrim- <br> inant rule allocates an observation $x$ to the population $\Pi_{j}$ that <br> maximises $\pi_{i} f_{i}(x)$ with respect to $i$. All Bayes discriminant rules <br> (incl. the ML rule) are admissible. |
| $\hookrightarrow$ | For the ML rule and $J=2$ normal populations, the probabilities <br> of misclassification are given by $p_{12}=p_{21}=\Phi\left(-\frac{1}{2} \delta\right)$ where $\delta$ is <br> the Mahalanobis distance between the two populations. |
| $\hookrightarrow$ | Classification of two normal populations with different covariance <br> matrices (ML rule) leads to regions defined by a quadratic function. |
| $\hookrightarrow$ | Desirable discriminant rules have a low ECM. |

### 14.2 Discrimination Rules in Practice

The ML rule is used if the distribution of the data is known up to parameters. Suppose for example that the data come from multivariate normal distributions $N_{p}\left(\mu_{j}, \Sigma\right)$. If we have $J$ groups with $n_{j}$ observations in each group, we use $\bar{x}_{j}$ to estimate $\mu_{j}$, and $\mathcal{S}_{j}$ to estimate $\Sigma$. The common covariance may be estimated by

$$
\begin{equation*}
\mathcal{S}_{u}=\sum_{j=1}^{J} n_{j}\left(\frac{\mathcal{S}_{j}}{n-J}\right), \tag{14.9}
\end{equation*}
$$

with $n=\sum_{j=1}^{J} n_{j}$. Thus the empirical version of the ML rule of Theorem 14.2 is to allocate a new observation $x$ to $\Pi_{j}$ such that $j$ minimises

$$
\left(x-\bar{x}_{i}\right)^{\top} \mathcal{S}_{u}^{-1}\left(x-\bar{x}_{i}\right) \quad \text { for } \quad i \in\{1, \ldots, J\} .
$$

Example 14.4 Let us apply this rule to the Swiss bank notes. The 20 randomly chosen bank notes which we had clustered into two groups in Example 13.6 are used. First the covariance $\Sigma$ is estimated by the average of the covariances of $\Pi_{1}$ (cluster 1) and $\Pi_{2}$ (cluster 2). The hyperplane $\hat{\alpha}^{\top}(x-\bar{x})=0$ which separates the two populations is given by

$$
\begin{aligned}
& \hat{\alpha}=\mathcal{S}_{u}^{-1}\left(\bar{x}_{1}-\bar{x}_{2}\right)=(-12.18,20.54,-19.22,-15.55,-13.06,21.43)^{\top}, \\
& \bar{x}=\frac{1}{2}\left(\bar{x}_{1}+\bar{x}_{2}\right)=(214.79,130.05,129.92,9.23,10.48,140.46)^{\top} .
\end{aligned}
$$

Now let us apply the discriminant rule to the entire bank notes data set. Counting the number of misclassifications by

$$
\sum_{i=1}^{100} \boldsymbol{I}\left\{\hat{\alpha}^{\top}\left(x_{i}-\bar{x}\right)<0\right\}, \sum_{i=101}^{200} \boldsymbol{I}\left\{\hat{\alpha}^{\top}\left(x_{i}-\bar{x}\right)>0\right\}
$$

we obtain 1 misclassified observation for the counterfeit bank notes and 0 misclassification for the genuine bank notes.

When $J=3$ groups, the allocation regions can be calculated using

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

The rule is to allocate $x$ to

## Estimation of the Probabilities of Misclassifications

Misclassification probabilities are given by (14.7) and can be estimated by replacing the unknown parameters by their corresponding estimators.

For the ML rule for two normal populations we obtain

$$
\hat{p}_{12}=\hat{p}_{21}=\Phi\left(-\frac{1}{2} \hat{\delta}\right)
$$

where $\hat{\delta}^{2}=\left(\bar{x}_{1}-\bar{x}_{2}\right)^{\top} \mathcal{S}_{u}^{-1}\left(\bar{x}_{1}-\bar{x}_{2}\right)$ is the estimator for $\delta^{2}$.
The probabilities of misclassification may also be estimated by the re-substitution method. We reclassify each original observation $x_{i}, i=1, \ldots, n$ into $\Pi_{1}, \ldots, \Pi_{J}$ according to the chosen rule. Then denoting the number of individuals coming from $\Pi_{j}$ which have been classified into $\Pi_{i}$ by $n_{i j}$, we have $\hat{p}_{i j}=\frac{n_{i j}}{n_{j}}$, an estimator of $p_{i j}$. Clearly, this method leads to too optimistic estimators of $p_{i j}$, but it provides a
rough measure of the quality of the discriminant rule. The matrix $\left(\hat{p}_{i j}\right)$ is called the confusion matrix in Johnson and Wichern (1998).

Example 14.5 In the above classification problem for the Swiss bank notes (Sect. 22.2), we have the following confusion matrix: Q MVAaper


The apparent error rate (APER) is defined as the fraction of observations that are misclassified. The APER, expressed as a percentage, is

$$
\mathrm{APER}=\left(\frac{1}{200}\right) 100 \%=0.5 \%
$$

For the calculation of the APER we use the observations twice: the first time to construct the classification rule and the second time to evaluate this rule. An APER of $0.5 \%$ might therefore be too optimistic. An approach that corrects for this bias is based on the holdout procedure of Lachenbruch and Mickey (1968). For two populations this procedure is as follows:

1. Start with the first population $\Pi_{1}$. Omit one observation and develop the classification rule based on the remaining $n_{1}-1, n_{2}$ observations.
2. Classify the "holdout" observation using the discrimination rule in Step 1.
3. Repeat steps 1 and 2 until all of the $\Pi_{1}$ observations are classified. Count the number $n_{21}^{\prime}$ of misclassified observations.
4. Repeat steps 1 through 3 for population $\Pi_{2}$. Count the number $n_{12}^{\prime}$ of misclassified observations.

Estimates of the misclassification probabilities are given by

$$
\hat{p}_{12}^{\prime}=\frac{n_{12}^{\prime}}{n_{2}}
$$

and

$$
\hat{p}_{21}^{\prime}=\frac{n_{21}^{\prime}}{n_{1}} .
$$

A more realistic estimator of the actual error rate (AER) is given by

$$
\begin{equation*}
\frac{n_{12}^{\prime}+n_{21}^{\prime}}{n_{2}+n_{1}} \tag{14.10}
\end{equation*}
$$

Statisticians favor the AER (for its unbiasedness) over the APER. In large samples, however, the computational costs might counterbalance the statistical advantage. This is not a real problem since the two misclassification measures are asymptotically equivalent.

## Fisher's Linear Discrimination Function

Another approach stems from R.A. Fisher. His idea was to base the discriminant rule on a projection $a^{\top} x$ such that a good separation was achieved. This LDA projection method is called Fisher's linear discrimination function. If

$$
\mathcal{Y}=\mathcal{X} a
$$

denotes a linear combination of observations, then the total sum of squares of $y$, $\sum_{i=1}^{n}\left(y_{i}-\bar{y}\right)^{2}$, is equal to

$$
\begin{equation*}
\mathcal{Y}^{\top} \mathcal{H} \mathcal{Y}=a^{\top} \mathcal{X}^{\top} \mathcal{H} \mathcal{X} a=a^{\top} \mathcal{T} a \tag{14.11}
\end{equation*}
$$

with the centering matrix $\mathcal{H}=\mathcal{I}-n^{-1} 1_{n} 1_{n}^{\top}$ and $\mathcal{T}=\mathcal{X}^{\top} \mathcal{H} \mathcal{X}$.
Suppose we have samples $\mathcal{X}_{j}, j=1, \ldots, J$, from $J$ populations. Fisher's suggestion was to find the linear combination $a^{\top} x$ which maximises the ratio of the between-group-sum of squares to the within-group-sum of squares.

The within-group-sum of squares is given by

$$
\begin{equation*}
\sum_{j=1}^{J} \mathcal{Y}_{j}^{\top} \mathcal{H}_{j} \mathcal{Y}_{j}=\sum_{j=1}^{J} a^{\top} \mathcal{X}_{j}^{\top} \mathcal{H}_{j} \mathcal{X}_{j} a=a^{\top} \mathcal{W} a \tag{14.12}
\end{equation*}
$$

where $\mathcal{Y}_{j}$ denotes the $j$-th sub-matrix of $\mathcal{Y}$ corresponding to observations of group $j$ and $\mathcal{H}_{j}$ denotes the $\left(n_{j} \times n_{j}\right)$ centering matrix. The within-group-sum of squares measures the sum of variations within each group.

The between-group-sum of squares is

$$
\begin{equation*}
\sum_{j=1}^{J} n_{j}\left(\bar{y}_{j}-\bar{y}\right)^{2}=\sum_{j=1}^{J} n_{j}\left\{a^{\top}\left(\bar{x}_{j}-\bar{x}\right)\right\}^{2}=a^{\top} \mathcal{B} a \tag{14.13}
\end{equation*}
$$

where $\bar{y}_{j}$ and $\bar{x}_{j}$ denote the means of $\mathcal{Y}_{j}$ and $\mathcal{X}_{j}$ and $\bar{y}$ and $\bar{x}$ denote the sample means of $\mathcal{Y}$ and $\mathcal{X}$. The between-group-sum of squares measures the variation of the means across groups.

The total sum of squares (14.11) is the sum of the within-group-sum of squares and the between-group-sum of squares, i.e.

$$
a^{\top} \mathcal{T} a=a^{\top} \mathcal{W} a+a^{\top} \mathcal{B} a .
$$

Fisher's idea was to select a projection vector $a$ that maximises the ratio

$$
\begin{equation*}
\frac{a^{\top} \mathcal{B} a}{a^{\top} \mathcal{W} a} \tag{14.14}
\end{equation*}
$$

The solution is found by applying Theorem 2.5.
Theorem 14.4 The vector a that maximises (14.14) is the eigenvector of $\mathcal{W}^{-1} \mathcal{B}$ that corresponds to the largest eigenvalue.

Now a discrimination rule is easy to obtain: classify $x$ into group $j$ where $a^{\top} \bar{x}_{j}$ is closest to $a^{\top} x$, i.e.

$$
x \rightarrow \Pi_{j} \text { where } j=\arg \min _{i}\left|a^{\top}\left(x-\bar{x}_{i}\right)\right| .
$$

When $J=2$ groups, the discriminant rule is easy to compute. Suppose that group 1 has $n_{1}$ elements and group 2 has $n_{2}$ elements. In this case

$$
\mathcal{B}=\left(\frac{n_{1} n_{2}}{n}\right) d d^{\top}
$$

where $d=\left(\bar{x}_{1}-\bar{x}_{2}\right) . \mathcal{W}^{-1} \mathcal{B}$ has only one eigenvalue which equals

$$
\operatorname{tr}\left(\mathcal{W}^{-1} \mathcal{B}\right)=\left(\frac{n_{1} n_{2}}{n}\right) d^{\top} \mathcal{W}^{-1} d
$$

and the corresponding eigenvector is $a=\mathcal{W}^{-1} d$. The corresponding discriminant rule is

$$
\begin{array}{ll}
x \rightarrow \Pi_{1} & \text { if }  \tag{14.15}\\
a^{\top}\left\{x-\frac{1}{2}\left(\bar{x}_{1}+\bar{x}_{2}\right)\right\}>0 \\
x \rightarrow \Pi_{2} & \text { if }
\end{array} a^{\top}\left\{x-\frac{1}{2}\left(\bar{x}_{1}+\bar{x}_{2}\right)\right\} \leq 0 .
$$

The Fisher LDA is closely related to projection pursuit (Chap. 20) since the statistical technique is based on a one-dimensional index $a^{\top} x$.

Example 14.6 Consider the bank notes data again. Let us use the subscript " $g$ " for the genuine and " $f$ " for the counterfeit bank notes, e.g. $\mathcal{X}_{g}$ denotes the first hundred observations of $\mathcal{X}$ and $\mathcal{X}_{f}$ the second hundred. In the context of the bank data set the "between-group-sum of squares" is defined as

$$
\begin{equation*}
100\left\{\left(\bar{y}_{g}-\bar{y}\right)^{2}+\left(\bar{y}_{f}-\bar{y}\right)^{2}\right\}=a^{\top} \mathcal{B} a \tag{14.16}
\end{equation*}
$$

for some matrix $\mathcal{B}$. Here, $\bar{y}_{g}$ and $\bar{y}_{f}$ denote the means for the genuine and counterfeit bank notes and $\bar{y}=\frac{1}{2}\left(\bar{y}_{g}+\bar{y}_{f}\right)$. The "within-group-sum of squares" is

$$
\begin{equation*}
\sum_{i=1}^{100}\left\{\left(y_{g}\right)_{i}-\bar{y}_{g}\right\}^{2}+\sum_{i=1}^{100}\left\{\left(y_{f}\right)_{i}-\bar{y}_{f}\right\}^{2}=a^{\top} \mathcal{W} a \tag{14.17}
\end{equation*}
$$

with $\left(y_{g}\right)_{i}=a^{\top} x_{i}$ and $\left(y_{f}\right)_{i}=a^{\top} x_{i+100}$ for $i=1, \ldots, 100$.
The resulting discriminant rule consists of allocating an observation $x_{0}$ to the genuine sample space if

$$
a^{\top}\left(x_{0}-\bar{x}\right)>0,
$$

with $a=\mathcal{W}^{-1}\left(\bar{x}_{g}-\bar{x}_{f}\right)$ (see Exercise 14.8) and of allocating $x_{0}$ to the counterfeit sample space when the opposite is true. In our case

$$
a=(0.000,0.029,-0.029,-0.039,-0.041,0.054)^{\top} .
$$

One genuine and no counterfeit bank notes are misclassified. Figure 14.2 shows the estimated densities for $y_{g}=a^{\top} \mathcal{X}_{g}$ and $y_{f}=a^{\top} \mathcal{X}_{f}$. They are separated better than those of the diagonals in Fig. 1.9.

Note that the allocation rule (14.15) is exactly the same as the ML rule for $J=2$ groups and for normal distributions with the same covariance. For $J=3$ groups this rule will be different, except for the special case of collinear sample means.

Fig. 14.2 Densities of projections of genuine and counterfeit bank notes by Fisher's discrimination rule Q MVAdisfbank


| Sum Sumary |
| :---: |
| $\hookrightarrow$ A discriminant rule is a separation of the sample space into sets $R_{j}$. An observation $x$ is classified as coming from population $\Pi_{j}$ if it lies in $R_{j}$. |
| $\hookrightarrow$ The ECM for two populations is given by $\mathrm{ECM}=C(2 \mid 1) p_{21} \pi_{1}+$ $C(1 \mid 2) p_{12} \pi_{2}$. |
| $\hookrightarrow$ The ML rule is applied if the distributions in the populations are known up to parameters, e.g. for normal distributions $N_{p}\left(\mu_{j}, \Sigma\right)$. |
| $\hookrightarrow$ The ML rule allocates $x$ to the population that exhibits the smallest Mahalanobis distance $\delta^{2}\left(x ; \mu_{i}\right)=\left(x-\mu_{i}\right)^{\top} \Sigma^{-1}\left(x-\mu_{i}\right) .$ |
| $\hookrightarrow$ The probability of misclassification is given by $p_{12}=p_{21}=\Phi\left(-\frac{1}{2} \delta\right)$ <br> where $\delta$ is the Mahalanobis distance between $\mu_{1}$ and $\mu_{2}$. |
| $\hookrightarrow$ Classification for different covariance structures in the two populations leads to quadratic discrimination rules. |
| $\hookrightarrow$ A different approach is Fisher's linear discrimination rule which finds a linear combination $a^{\top} x$ that maximises the ratio of the "between-group-sum of squares" and the "within-group-sum of squares". This rule turns out to be identical to the ML rule when $J=2$ for normal populations. |

### 14.3 Boston Housing

One interesting application of discriminant analysis with respect to the Boston housing data is the classification of the districts according to the house values. The rationale behind this is that certain observable must determine the value of a district, as in Sect. 3.7 where the house value was regressed on the other variables. Two groups are defined according to the median value of houses $\tilde{X}_{14}$ : in group $\Pi_{1}$ the value of $\tilde{X}_{14}$ is greater than or equal to the median of $\tilde{X}_{14}$ and in group $\Pi_{2}$ the value of $\tilde{X}_{14}$ is less than the median of $\tilde{X}_{14}$.

Table 14.1 APER for price
of Boston houses $\mathbf{Q}$
MVAdiscbh


Table 14.2 AER for price of Boston houses 0 mVAaerbh

|  |  | True |  |
| :---: | :---: | :---: | :---: |
|  |  | $\Pi_{1}$ | $\Pi_{2}$ |
|  | $\Pi_{1}$ | 211 | 42 |
| Predicted |  |  |  |
|  | $\Pi_{2}$ | 39 | 214 |
|  |  |  |  |

Table 14.3 APER for

|  |  | True |  |
| :---: | :---: | :---: | :---: |
|  |  | $\Pi_{1}$ |  |
| Predicted | $\Pi_{2}$ |  |  |
|  | $\Pi_{1}$ | 244 |  |
|  |  | 13 |  |
|  | $\Pi_{2}$ | 7 |  |
|  |  | 242 |  |
|  |  |  |  |

Table 14.4 AER for clusters of Boston houses $\mathbf{Q}$ MVAaerbh

|  |  | True |  |
| :---: | :---: | :---: | :---: |
|  |  | $\Pi_{1}$ |  |
|  | $\Pi_{2}$ |  |  |
| Predicted | $\Pi_{1}$ | 244 |  |
|  |  | 14 |  |
|  | $\Pi_{2}$ | 7 |  |
|  |  | 241 |  |
|  |  |  |  |

The linear discriminant rule, defined on the remaining 12 variables (excluding $\tilde{X}_{4}$ and $\tilde{X}_{14}$ ) is applied. After reclassifying the 506 observations, we obtain an APER of 0.146 . The details are given in Table 14.1. The more appropriate error rate, given by the AER, is 0.160 (see Table 14.2).

Let us now turn to a group definition suggested by the Cluster Analysis in Sect. 13.4. Group $\Pi_{1}$ was defined by higher quality of life and house. We define the linear discriminant rule using the 13 variables from $\widetilde{\mathcal{X}}$ excluding $\tilde{X}_{4}$. Then we reclassify the 506 observations and we obtain an APER of 0.0395 . Details are summarised in Table 14.3. The AER turns out to be 0.0415 (see Table 14.4).

Figure 14.3 displays the values of the linear discriminant scores (see Theorem 14.2) for all of the 506 observations, coloured by groups. One can clearly see the APER is derived from the seven observations from group $\Pi_{1}$ with a negative score and the 13 observations from group $\Pi_{2}$ with positive score.


Fig. 14.3 Discrimination scores for the two clusters created from the Boston housing data $\mathbf{Q}$ MVAdiscbh

### 14.4 Exercises

Exercise 14.1 Prove Theorem 14.2 (a) and 14.2 (b).
Exercise 14.2 Apply the rule from Theorem 14.2 (b) for $p=1$ and compare the result with that of Example 14.3.

Exercise 14.3 Calculate the ML discrimination rule based on observations of a one-dimensional variable with an exponential distribution.

Exercise 14.4 Calculate the ML discrimination rule based on observations of a two-dimensional random variable, where the first component has an exponential distribution and the other has an alternative distribution. What is the difference between the discrimination rule obtained in this exercise and the Bayes discrimination rule?

Exercise 14.5 Apply the Bayes rule to the car data (Table 22.3) in order to discriminate between Japanese, European and US cars, i.e. $J=3$. Consider only the "miles per gallon" variable and take the relative frequencies as prior probabilities.
Exercise 14.6 Compute Fisher's linear discrimination function for the 20 bank notes from Example 13.6. Apply it to the entire bank data set. How many observations are misclassified?

Exercise 14.7 Use the Fisher's linear discrimination function on the WAIS data set (Table 22.12) and evaluate the results by re-substitution the probabilities of misclassification.

Exercise 14.8 Show that in Example 14.6
(a) $\mathcal{W}=100\left(\mathcal{S}_{g}+\mathcal{S}_{f}\right)$, where $\mathcal{S}_{g}$ and $\mathcal{S}_{f}$ denote the empirical covariances (3.6) and (3.5) w.r.t. the genuine and counterfeit bank notes,
(b) $\mathcal{B}=100\left\{\left(\bar{x}_{g}-\bar{x}\right)\left(\bar{x}_{g}-\bar{x}\right)^{\top}+\left(\bar{x}_{f}-\bar{x}\right)\left(\bar{x}_{f}-\bar{x}\right)^{\top}\right\}$, where $\bar{x}=\frac{1}{2}\left(\bar{x}_{g}+\right.$ $\bar{x}_{f}$ ),
(c) $a=\mathcal{W}^{-1}\left(\bar{x}_{g}-\bar{x}_{f}\right)$.

Exercise 14.9 Recalculate Example 14.3 with the prior probability $\pi_{1}=\frac{1}{3}$ and $C(2 \mid 1)=2 C(1 \mid 2)$.

Exercise 14.10 Explain the effect of changing $\pi_{1}$ or $C(1 \mid 2)$ on the relative location of the region $R_{j}, j=1,2$.

Exercise 14.11 Prove that Fisher's linear discrimination function is identical to the ML rule when the covariance matrices are identical $(J=2)$.

Exercise 14.12 Suppose that $x \in\{0,1,2,3,4,5,6,7,8,9,10\}$ and
$\Pi_{1}: X \sim \operatorname{Bi}(10,0.2) \quad$ with the prior probability $\pi_{1}=0.5$;
$\Pi_{2}: X \sim \operatorname{Bi}(10,0.3) \quad$ with the prior probability $\pi_{2}=0.3$;
$\Pi_{3}: X \sim \operatorname{Bi}(10,0.5) \quad$ with the prior probability $\pi_{3}=0.2$.
Determine the sets $R_{1}, R_{2}$ and $R_{3}$. (Use the Bayes discriminant rule.)

## Chapter 15 <br> Correspondence Analysis

Correspondence analysis provides tools for analysing the associations between rows and columns of contingency tables. A contingency table is a two-entry frequency table where the joint frequencies of two qualitative variables are reported. For instance a $(2 \times 2)$ table could be formed by observing from a sample of $n$ individuals two qualitative variables: the individual's sex and whether the individual smokes. The table reports the observed joint frequencies. In general $(n \times p)$ tables may be considered.

The main idea of correspondence analysis is to develop simple indices that will show the relations between the row and the columns categories. These indices will tell us simultaneously which column categories have more weight in a row category and vice versa. Correspondence analysis is also related to the issue of reducing the dimension of the table, similar to principal component analysis in Chap. 11, and to the issue of decomposing the table into its factors as discussed in Chap. 10. The idea is to extract the indices in decreasing order of importance so that the main information of the table can be summarised in spaces with smaller dimensions. For instance, if only two factors (indices) are used, the results can be shown in twodimensional graphs, showing the relationship between the rows and the columns of the table.

Section 15.1 defines the basic notation and motivates the approach and Sect. 15.2 gives the basic theory. The indices will be used to describe the $\chi^{2}$ statistic measuring the associations in the table. Several examples in Sect. 15.3 show how to provide and interpret, in practice, the two-dimensional graphs displaying the relationship between the rows and the columns of a contingency table.

### 15.1 Motivation

The aim of correspondence analysis is to develop simple indices that show relations between the row and columns of a contingency tables. Contingency tables are very useful to describe the association between two variables in very general situations. The two variables can be qualitative (nominal), in which case they are also referred to as categorical variables. Each row and each column in the table represents one category of the corresponding variable. The entry $x_{i j}$ in the table $\mathcal{X}$ (with dimension $(n \times p)$ ) is the number of observations in a sample which simultaneously fall in the $i$ th row category and the $j$ th column category, for $i=1, \ldots, n$ and $j=1, \ldots, p$. Sometimes a "category" of a nominal variable is also called a "modality" of the variable.

The variables of interest can also be discrete quantitative variables, such as the number of family members or the number of accidents an insurance company had to cover during 1 year, etc. Here, each possible value that the variable can have defines a row or a column category. Continuous variables may be taken into account by defining the categories in terms of intervals or classes of values which the variable can take on. Thus contingency tables can be used in many situations, implying that correspondence analysis is a very useful tool in many applications.

The graphical relationships between the rows and the columns of the table $\mathcal{X}$ that result from correspondence analysis are based on the idea of representing all the row and column categories and interpreting the relative positions of the points in terms of the weights corresponding to the column and the row. This is achieved by deriving a system of simple indices providing the coordinates of each row and each column. These row and column coordinates are simultaneously represented in the same graph. It is then clear to see which column categories are more important in the row categories of the table (and the other way around).

As was already eluded to, the construction of the indices is based on an idea similar to that of PCA. Using PCA the total variance was partitioned into independent contributions stemming from the principal components. Correspondence analysis, on the other hand, decomposes a measure of association, typically the total $\chi^{2}$ value used in testing independence, rather than decomposing the total variance.

Example 15.1 The French "baccalauréat" frequencies have been classified into regions and different baccalauréat categories, see Chap. 22, Table 22.8. Altogether $n=202,100$ baccalauréats were observed. The joint frequency of the region Ile-de-France and the modality Philosophy, for example, is 9,724. That is, 9,724 baccalauréats were in Ile-de-France and the category Philosophy.

The question is whether certain regions prefer certain baccalauréat types. If we consider, for instance, the region Lorraine, we have the following percentages:

| A | B | C | D | E | F | G | H |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 20.5 | 7.6 | 15.3 | 19.6 | 3.4 | 14.5 | 18.9 | 0.2 |

The total percentages of the different modalities of the variable baccalauréat are as follows:

| A | B | C | D | E | F | G | H |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 22.6 | 10.7 | 16.2 | 22.8 | 2.6 | 9.7 | 15.2 | 0.2 |

One might argue that the region Lorraine seems to prefer the modalities E, F, G and dislike the specialisations $\mathrm{A}, \mathrm{B}, \mathrm{C}, \mathrm{D}$ relative to the overall frequency of baccalauréat type.

In correspondence analysis we try to develop an index for the regions so that this over- or underrepresentation can be measured in just one single number. Simultaneously we try to weight the regions so that we can see in which region certain baccalauréat types are preferred.

Example 15.2 Consider $n$ types of companies and $p$ locations of these companies. Is there a certain type of company that prefers a certain location? Or is there a location index that corresponds to a certain type of company?

Assume that $n=3, p=3$, and that the frequencies are as follows:

$$
\mathcal{X}=\left(\begin{array}{lll}
4 & 0 & 2 \\
0 & 1 & 1 \\
1 & 1 & 4
\end{array}\right) \quad \begin{aligned}
& \leftarrow \text { Finance } \\
& \leftarrow \text { Energy } \\
& \leftarrow \text { HiTech }
\end{aligned}
$$

$$
\begin{aligned}
& \uparrow \text { Frankfurt } \\
& \qquad \begin{array}{l}
\text { Berlin } \\
\\
\quad \uparrow \text { Munich }
\end{array}
\end{aligned}
$$

The frequencies imply that four type three companies (HiTech) are in location 3 (Munich), and so on. Suppose there is a (company) weight vector $r=\left(r_{1}, \ldots, r_{n}\right)^{\top}$ such that a location index $s_{j}$ could be defined as

$$
\begin{equation*}
s_{j}=c \sum_{i=1}^{n} r_{i} \frac{x_{i j}}{x_{\bullet j}}, \tag{15.1}
\end{equation*}
$$

where $x_{\bullet j}=\sum_{i=1}^{n} x_{i j}$ is the number of companies in location $j$ and $c$ is a constant. $s_{1}$, for example, would give the average weighted frequency (by $r$ ) of companies in location 1 (Frankfurt).

Given a location weight vector $s^{*}=\left(s_{1}^{*}, \ldots, s_{p}^{*}\right)^{\top}$, we can define a company index in the same way as

$$
\begin{equation*}
r_{i}^{*}=c^{*} \sum_{j=1}^{p} s_{j}^{*} \frac{x_{i j}}{x_{i} \bullet} \tag{15.2}
\end{equation*}
$$

where $c^{*}$ is a constant and $x_{i} \bullet=\sum_{j=1}^{p} x_{i j}$ is the sum of the $i$ th row of $\mathcal{X}$, i.e. the number of type $i$ companies. Thus $r_{2}^{*}$, for example, would give the average weighted frequency (by $s^{*}$ ) of energy companies.

If (15.1) and (15.2) can be solved simultaneously for a "row weight" vector $r=\left(r_{1}, \ldots, r_{n}\right)^{\top}$ and a "column weight" vector $s=\left(s_{1}, \ldots, s_{p}\right)^{\top}$, we may represent each row category by $r_{i}, i=1, \ldots, n$ and each column category by $s_{j}, j=1, \ldots, p$ in a one-dimensional graph. If in this graph $r_{i}$ and $s_{j}$ are in close proximity (far from the origin), this would indicate that the $i$ th row category has an important conditional frequency $x_{i j} / x_{\bullet j}$ in (15.1) and that the $j$ th column category has an important conditional frequency $x_{i j} / x_{i} \bullet$ in (15.2). This would indicate a positive association between the $i$ th row and the $j$ th column. A similar line of argument could be used if $r_{i}$ was very far away from $s_{j}$ (and far from the origin). This would indicate a small conditional frequency contribution, or a negative association between the $i$ th row and the $j$ th column.

| $\hookrightarrow$ | The aim of correspondence analysis is to develop simple indices <br> that show relations among qualitative variables in a contingency <br> table. |
| :--- | :--- |
|  | The joint representation of the indices reveals relations among the <br> variables. |

### 15.2 Chi-Square Decomposition

An alternative way of measuring the association between the row and column categories is a decomposition of the value of the $\chi^{2}$-test statistic. The wellknown $\chi^{2}$-test for independence in a two-dimensional contingency table consists of two steps. First the expected value of each cell of the table is estimated under the hypothesis of independence. Second, the corresponding observed values are compared to the expected values using the statistic

$$
\begin{equation*}
t=\sum_{i=1}^{n} \sum_{j=1}^{p}\left(x_{i j}-E_{i j}\right)^{2} / E_{i j}, \tag{15.3}
\end{equation*}
$$

where $x_{i j}$ is the observed frequency in cell $(i, j)$ and $E_{i j}$ is the corresponding estimated expected value under the assumption of independence, i.e.

$$
\begin{equation*}
E_{i j}=\frac{x_{i} x_{\bullet j}}{x_{\bullet \bullet}} . \tag{15.4}
\end{equation*}
$$

Here $x_{\bullet \bullet}=\sum_{i=1}^{n} x_{i \bullet \bullet}$ Under the hypothesis of independence, $t$ has a $\chi_{(n-1)(p-1)}^{2}$ distribution. In the industrial location example introduced above the value of $t=$ 6.26 is almost significant at the $5 \%$ level. It is therefore worth investigating the special reasons for departure from independence.

The method of $\chi^{2}$ decomposition consists of finding the SVD of the matrix $\mathcal{C}$ ( $n \times$ $p)$ with elements

$$
\begin{equation*}
c_{i j}=\left(x_{i j}-E_{i j}\right) / E_{i j}^{1 / 2} \tag{15.5}
\end{equation*}
$$

The elements $c_{i j}$ may be viewed as measuring the (weighted) departure between the observed $x_{i j}$ and the theoretical values $E_{i j}$ under independence. This leads to the factorial tools of Chap. 10 which describe the rows and the columns of $\mathcal{C}$.

For simplification define the matrices $\mathcal{A}(n \times n)$ and $\mathcal{B}(p \times p)$ as

$$
\begin{equation*}
\mathcal{A}=\operatorname{diag}\left(x_{i \bullet}\right) \text { and } \mathcal{B}=\operatorname{diag}\left(x_{\bullet j}\right) \tag{15.6}
\end{equation*}
$$

These matrices provide the marginal row frequencies $a(n \times 1)$ and the marginal column frequencies $b(p \times 1)$ :

$$
\begin{equation*}
a=\mathcal{A} 1_{n} \text { and } b=\mathcal{B} 1_{p} \tag{15.7}
\end{equation*}
$$

It is easy to verify that

$$
\begin{equation*}
\mathcal{C} \sqrt{b}=0 \text { and } \mathcal{C}^{\top} \sqrt{a}=0 \tag{15.8}
\end{equation*}
$$

where the square root of the vector is taken element by element and $R=\operatorname{rank}(\mathcal{C}) \leq$ $\min \{(n-1),(p-1)\}$. From (10.14) of Chap. 10, the SVD of $\mathcal{C}$ yields

$$
\begin{equation*}
\mathcal{C}=\Gamma \Lambda \Delta^{\top}, \tag{15.9}
\end{equation*}
$$

where $\Gamma$ contains the eigenvectors of $\mathcal{C C}^{\top}$, $\Delta$ the eigenvectors of $\mathcal{C}^{\top} \mathcal{C}$ and $\Lambda=$ $\operatorname{diag}\left(\lambda_{1}^{1 / 2}, \ldots, \lambda_{R}^{1 / 2}\right)$ with $\lambda_{1} \geq \lambda_{2} \geq \cdots \geq \lambda_{R}$ (the eigenvalues of $\mathcal{C C}^{\top}$ ). Equation (15.9) implies that

$$
\begin{equation*}
c_{i j}=\sum_{k=1}^{R} \lambda_{k}^{1 / 2} \gamma_{i k} \delta_{j k} . \tag{15.10}
\end{equation*}
$$

Note that (15.3) can be rewritten as

$$
\begin{equation*}
\operatorname{tr}\left(\mathcal{C C}^{\top}\right)=\sum_{k=1}^{R} \lambda_{k}=\sum_{i=1}^{n} \sum_{j=1}^{p} c_{i j}^{2}=t \tag{15.11}
\end{equation*}
$$

This relation shows that the SVD of $\mathcal{C}$ decomposes the total $\chi^{2}$ value rather than, as in Chap. 10, the total variance.

The duality relations between the row and the column space (10.11) are now for $k=1, \ldots, R$ given by

$$
\begin{align*}
\delta_{k} & =\frac{1}{\sqrt{\lambda_{k}}} \mathcal{C}^{\top} \gamma_{k},  \tag{15.12}\\
\gamma_{k} & =\frac{1}{\sqrt{\lambda_{k}}} \mathcal{C} \delta_{k} .
\end{align*}
$$

The projections of the rows and the columns of $\mathcal{C}$ are given by

$$
\begin{align*}
& \mathcal{C} \delta_{k}=\sqrt{\lambda_{k}} \gamma_{k},  \tag{15.13}\\
& \mathcal{C}^{\top} \gamma_{k}=\sqrt{\lambda_{k}} \delta_{k} .
\end{align*}
$$

Note that the eigenvectors satisfy

$$
\begin{equation*}
\delta_{k}^{\top} \sqrt{b}=0, \quad \gamma_{k}^{\top} \quad \sqrt{a}=0 \tag{15.14}
\end{equation*}
$$

From (15.10) we see that the eigenvectors $\delta_{k}$ and $\gamma_{k}$ are the objects of interest when analysing the correspondence between the rows and the columns. Suppose that the first eigenvalue in (15.10) is dominant so that

$$
\begin{equation*}
c_{i j} \approx \lambda_{1}^{1 / 2} \gamma_{i 1} \delta_{j 1} \tag{15.15}
\end{equation*}
$$

In this case when the coordinates $\gamma_{i 1}$ and $\delta_{j 1}$ are both large (with the same sign) relative to the other coordinates, then $c_{i j}$ will be large as well, indicating a positive association between the $i$ th row and the $j$ th column category of the contingency table. If $\gamma_{i 1}$ and $\delta_{j 1}$ were both large with opposite signs, then there would be a negative association between the $i$ th row and $j$ th column.

In many applications, the first two eigenvalues, $\lambda_{1}$ and $\lambda_{2}$, dominate and the percentage of the total $\chi^{2}$ explained by the eigenvectors $\gamma_{1}$ and $\gamma_{2}$ and $\delta_{1}$ and $\delta_{2}$ is large. In this case (15.13) and ( $\gamma_{1}, \gamma_{2}$ ) can be used to obtain a graphical display of the $n$ rows of the table $\left(\left(\delta_{1}, \delta_{2}\right)\right.$ play a similar role for the $p$ columns of the table). The interpretation of the proximity between row and column points will be interpreted as above with respect to (15.10).

In correspondence analysis, we use the projections of weighted rows of $\mathcal{C}$ and the projections of weighted columns of $\mathcal{C}$ for graphical displays. Let $r_{k}(n \times 1)$ be the projections of $\mathcal{A}^{-1 / 2} \mathcal{C}$ on $\delta_{k}$ and $s_{k}(p \times 1)$ be the projections of $\mathcal{B}^{-1 / 2} \mathcal{C}^{\top}$ on $\gamma_{k}$ $(k=1, \ldots, R)$ :

$$
\begin{align*}
& r_{k}=\mathcal{A}^{-1 / 2} \mathcal{C} \delta_{k}=\sqrt{\lambda_{k}} \mathcal{A}^{-1 / 2} \gamma_{k}  \tag{15.16}\\
& s_{k}=\mathcal{B}^{-1 / 2} \mathcal{C}^{\top} \gamma_{k}=\sqrt{\lambda_{k}} \mathcal{B}^{-1 / 2} \delta_{k} .
\end{align*}
$$

These vectors have the property that

$$
\begin{align*}
r_{k}^{\top} a & =0,  \tag{15.17}\\
s_{k}^{\top} b & =0 .
\end{align*}
$$

The obtained projections on each axis $k=1, \ldots, R$ are centred at zero with the natural weights given by $a$ (the marginal frequencies of the rows of $\mathcal{X}$ ) for the row coordinates $r_{k}$ and by $b$ (the marginal frequencies of the columns of $\mathcal{X}$ ) for the column coordinates $s_{k}$ (compare this to expression (15.14)). As a result, the origin is the centre of gravity for all of the representations. We also know from (15.16) and the SVD of $\mathcal{C}$ that

$$
\begin{align*}
& r_{k}^{\top} \mathcal{A} r_{k}=\lambda_{k}, \\
& s_{k}^{\top} \mathcal{B} s_{k}=\lambda_{k} . \tag{15.18}
\end{align*}
$$

From the duality relation between $\delta_{k}$ and $\gamma_{k}$ (see (15.12)) we obtain

$$
\begin{align*}
& r_{k}=\frac{1}{\sqrt{\lambda_{k}}} \mathcal{A}^{-1 / 2} \mathcal{C} \mathcal{B}^{1 / 2} s_{k},  \tag{15.19}\\
& s_{k}=\frac{1}{\sqrt{\lambda_{k}}} \mathcal{B}^{-1 / 2} \mathcal{C}^{\top} \mathcal{A}^{1 / 2} r_{k},
\end{align*}
$$

which can be simplified to

$$
\begin{align*}
r_{k} & =\sqrt{\frac{x_{\bullet \bullet}}{\lambda_{k}}} \mathcal{A}^{-1} \mathcal{X} s_{k}  \tag{15.20}\\
s_{k} & =\sqrt{\frac{x_{\bullet \bullet}}{\lambda_{k}}} \mathcal{B}^{-1} \mathcal{X}^{\top} r_{k} .
\end{align*}
$$

These vectors satisfy the relations (15.1) and (15.2) for each $k=1, \ldots, R$ simultaneously.

As in Chap. 10, the vectors $r_{k}$ and $s_{k}$ are referred to as factors (row factor and column factor respectively). They have the following means and variances:

$$
\begin{gather*}
\bar{r}_{k}=\frac{1}{x_{\bullet \bullet}} r_{k}^{\top} a=0, \\
\bar{s}_{k}=\frac{1}{x_{\bullet \bullet}} s_{k}^{\top} b=0,  \tag{15.21}\\
\operatorname{Var}\left(r_{k}\right)=\frac{1}{x_{\bullet \bullet}} \sum_{i=1}^{n} x_{i \bullet} \bullet_{k i}^{2}=\frac{r_{k}^{\top} \mathcal{A} r_{k}}{x_{\bullet \bullet}}=\frac{\lambda_{k}}{x_{\bullet \bullet}},  \tag{15.22}\\
\operatorname{Var}\left(s_{k}\right)=\frac{1}{x_{\bullet \bullet}} \sum_{j=1}^{p} x_{\bullet j} s_{k j}^{2}=\frac{s_{k} \mathcal{B} s_{k}}{x_{\bullet \bullet}}=\frac{\lambda_{k}}{x_{\bullet \bullet}} .
\end{gather*}
$$

Hence, $\lambda_{k} / \sum_{k=1}^{j} \lambda_{j}$, which is the part of the $k$ th factor in the decomposition of the $\chi^{2}$ statistic $t$, may also be interpreted as the proportion of the variance explained by the factor $k$. The proportions

$$
\begin{equation*}
C_{a}\left(i, r_{k}\right)=\frac{x_{i} \bullet r_{k i}^{2}}{\lambda_{k}}, \text { for } i=1, \ldots, n, k=1, \ldots, R \tag{15.23}
\end{equation*}
$$

are called the absolute contributions of row $i$ to the variance of the factor $r_{k}$. They show which row categories are most important in the dispersion of the $k$ th row factors. Similarly, the proportions

$$
\begin{equation*}
C_{a}\left(j, s_{k}\right)=\frac{x_{\bullet j} s_{k j}^{2}}{\lambda_{k}}, \text { for } j=1, \ldots, p, k=1, \ldots, R \tag{15.24}
\end{equation*}
$$

are called the absolute contributions of column $j$ to the variance of the column factor $s_{k}$. These absolute contributions may help to interpret the graph obtained by correspondence analysis.

### 15.3 Correspondence Analysis in Practice

The graphical representations on the axes $k=1,2, \ldots, R$ of the $n$ rows and of the $p$ columns of $\mathcal{X}$ are provided by the elements of $r_{k}$ and $s_{k}$. Typically, two-dimensional displays are often satisfactory if the cumulated percentage of variance explained by the first two factors, $\Psi_{2}=\frac{\lambda_{1}+\lambda_{2}}{\sum_{k=1}^{R} \lambda_{k}}$, is sufficiently large.

The interpretation of the graphs may be summarised as follows:

- The proximity of two rows (two columns) indicates a similar profile in these two rows (two columns), where "profile" refers to the conditional frequency distribution of a row (column); those two rows (columns) are almost proportional. The opposite interpretation applies when the two rows (two columns) are far apart.
- The proximity of a particular row to a particular column indicates that this row (column) has a particularly important weight in this column (row). In contrast to this, a row that is quite distant from a particular column indicates that there are almost no observations in this column for this row (and vice versa). Of course, as mentioned above, these conclusions are particularly true when the points are far away from 0 .
- The origin is the average of the factors $r_{k}$ and $s_{k}$. Hence, a particular point (row or column) projected close to the origin indicates an average profile.
- The absolute contributions are used to evaluate the weight of each row (column) in the variances of the factors.
- All the interpretations outlined above must be carried out in view of the quality of the graphical representation which is evaluated, as in PCA, using the cumulated percentage of variance.

Remark 15.1 Note that correspondence analysis can also be applied to more general $(n \times p)$ tables $\mathcal{X}$ which in a "strict sense" are not contingency tables.

As long as statistical (or natural) meaning can be given to sums over rows and columns, Remark 15.1 holds. This implies, in particular, that all of the variables are measured in the same units. In that case, $x_{\bullet \bullet}$ constitutes the total frequency
of the observed phenomenon, and is shared between individuals ( $n$ rows) and between variables ( $p$ columns). Representations of the rows and columns of $\mathcal{X}, r_{k}$ and $s_{k}$, have the basic property (15.19) and show which variables have important weights for each individual and vice versa. This type of analysis is used as an alternative to PCA. PCA is mainly concerned with covariances and correlations, whereas correspondence analysis analyses a more general kind of association. (See Exercises 15.3 and 15.11.)

Example 15.3 A survey of Belgium citizens who regularly read a newspaper was conducted in the 1980s. They were asked where they lived. The possible answers were ten regions: seven provinces (Antwerp, Western Flanders, Eastern Flanders, Hainant, Liège, Limbourg, Luxembourg) and three regions around Brussels (Flemish-Brabant, Wallon-Brabant and the city of Brussels). They were also asked what kind of newspapers they read on a regular basis. There were 15 possible answers split up into three classes: Flemish newspapers (label begins with the letter $v$ ), French newspapers (label begins with $f$ ) and both languages together (label begins with $b$ ). The data set is given in Table 22.9. The eigenvalues of the factorial correspondence analysis are given in Table 15.1.

Two-dimensional representations will be quite satisfactory since the first two eigenvalues account for $81 \%$ of the variance. Figure 15.1 shows the projections of the rows (the 15 newspapers) and of the columns (the ten regions).

As expected, there is a high association between the regions and the type of newspapers which is read. In particular, $v_{b}$ (Gazet van Antwerp) is almost exclusively read in the province of Antwerp (this is an extreme point in the graph). The points on the left all belong to Flanders, whereas those on the right all belong to Wallonia. Notice that the Wallon-Brabant and the Flemish-Brabant are not far from Brussels. Brussels is close to the centre (average) and also close to the bilingual newspapers. It is shifted a little to the right of the origin due to the majority of French speaking people in the area.

The absolute contributions of the first three factors are listed in Tables 15.2 and 15.3. The row factors $r_{k}$ are in Table 15.2 and the column factors $s_{k}$ are in Table 15.3.

Table 15.1 Eigenvalues and percentages of the variance (Example 15.3)

| $\lambda_{j}$ | Percentage of variance | Cumulated percentage |
| :--- | :--- | :--- |
| 183.40 | 0.653 | 0.653 |
| 43.75 | 0.156 | 0.809 |
| 25.21 | 0.090 | 0.898 |
| 11.74 | 0.042 | 0.940 |
| 8.04 | 0.029 | 0.969 |
| 4.68 | 0.017 | 0.985 |
| 2.13 | 0.008 | 0.993 |
| 1.20 | 0.004 | 0.997 |
| 0.82 | 0.003 | 1.000 |
| 0.00 | 0.000 | 1.000 |



Fig. 15.1 Projection of rows (the 15 newspapers) and columns (the ten regions) a MVAcorrjourn

Table 15.2 Absolute contributions of row factors $r_{k}$

|  | $C_{a}\left(i, r_{1}\right)$ | $C_{a}\left(i, r_{2}\right)$ | $C_{a}\left(i, r_{3}\right)$ |
| :--- | :--- | :--- | :--- |
| $v_{a}$ | 0.0563 | 0.0008 | 0.0036 |
| $v_{b}$ | 0.1555 | 0.5567 | 0.0067 |
| $v_{c}$ | 0.0244 | 0.1179 | 0.0266 |
| $v_{d}$ | 0.1352 | 0.0952 | 0.0164 |
| $v_{e}$ | 0.0253 | 0.1193 | 0.0013 |
| $f_{f}$ | 0.0314 | 0.0183 | 0.0597 |
| $f_{g}$ | 0.0585 | 0.0162 | 0.0122 |
| $f_{h}$ | 0.1086 | 0.0024 | 0.0656 |
| $f_{i}$ | 0.1001 | 0.0024 | 0.6376 |
| $b_{j}$ | 0.0029 | 0.0055 | 0.0187 |
| $b_{k}$ | 0.0236 | 0.0278 | 0.0237 |
| $b_{l}$ | 0.0006 | 0.0090 | 0.0064 |
| $v_{m}$ | 0.1000 | 0.0038 | 0.0047 |
| $f_{n}$ | 0.0966 | 0.0059 | 0.0269 |
| $f_{0}$ | 0.0810 | 0.0188 | 0.0899 |
| Total | 1.0000 | 1.0000 | 1.0000 |

They show, for instance, the important role of Antwerp and the newspaper $v_{b}$ in determining the variance of both factors. Clearly, the first axis expresses linguistic differences between the three parts of Belgium. The second axis shows a larger dispersion between the Flemish region than the French speaking regions.

Table 15.3 Absolute contributions of column factors $s_{k}$

|  | $C_{a}\left(j, s_{1}\right)$ | $C_{a}\left(j, s_{2}\right)$ | $C_{a}\left(j, s_{3}\right)$ |
| :--- | :--- | :--- | :--- |
| brw | 0.0887 | 0.0210 | 0.2860 |
| bxl | 0.1259 | 0.0010 | 0.0960 |
| anv | 0.2999 | 0.4349 | 0.0029 |
| brf | 0.0064 | 0.2370 | 0.0090 |
| foc | 0.0729 | 0.1409 | 0.0033 |
| for | 0.0998 | 0.0023 | 0.0079 |
| hai | 0.1046 | 0.0012 | 0.3141 |
| lig | 0.1168 | 0.0355 | 0.1025 |
| lim | 0.0562 | 0.1162 | 0.0027 |
| lux | 0.0288 | 0.0101 | 0.1761 |
| Total | 1.0000 | 1.0000 | 1.0000 |



Fig. 15.2 Correspondence analysis including Corsica $Q$ MVAcorrbac

Note also that the third axis shows an important role of the category " $f_{i}$ " (other French newspapers) with the Wallon-Brabant "brw" and the Hainant "hai" showing the most important contributions. The coordinate of " $f_{i}$ " on this axis is negative (not shown here) so are the coordinates of "brw" and "hai". Apparently, these two regions also seem to feature a greater proportion of readers of more local newspapers.

Example 15.4 Applying correspondence analysis to the French baccalauréat data (Table 22.8) leads to Fig. 15.2. Excluding Corsica we obtain Fig. 15.3. The different


Fig. 15.3 Correspondence analysis excluding Corsica $Q$ mVAcorrbac
Table 15.4 Eigenvalues and percentages of explained variance (including Corsica)

| Eigenvalues $\lambda$ | Percentage of variances | Cumulated percentage |
| :--- | :--- | :--- |
| $2,436.2$ | 0.5605 | 0.561 |
| $1,052.4$ | 0.2421 | 0.803 |
| 341.8 | 0.0786 | 0.881 |
| 229.5 | 0.0528 | 0.934 |
| 152.2 | 0.0350 | 0.969 |
| 109.1 | 0.0251 | 0.994 |
| 25.0 | 0.0058 | 1.000 |
| 0.0 | 0.0000 | 1.000 |

modalities are labeled $\mathrm{A}, \ldots, \mathrm{H}$ and the regions are labeled ILDF, $\ldots$, CORS. The results of the correspondence analysis are given in Table 15.4 and Fig. 15.2.

The first two factors explain $80 \%$ of the total variance. It is clear from Fig. 15.2 that Corsica (in the upper left) is an outlier. The analysis is therefore redone without Corsica and the results are given in Table 15.5 and Fig. 15.3. Since Corsica has such a small weight in the analysis, the results have not changed much.

The projections on the first three axes, along with their absolute contribution to the variance of the axis, are summarised in Table 15.6 for the regions and in Table 15.7 for baccalauréats.

The interpretation of the results may be summarised as follows. Table 15.7 shows that the baccalauréats B on one side and F on the other side are most strongly

Table 15.5 Eigenvalues and percentages of explained variance (excluding Corsica)

| Eigenvalues $\lambda$ | Percentage of variances | Cumulated percentage |
| :--- | :--- | :--- |
| $2,408.6$ | 0.5874 | 0.587 |
| 909.5 | 0.2218 | 0.809 |
| 318.5 | 0.0766 | 0.887 |
| 195.9 | 0.0478 | 0.935 |
| 149.3 | 0.0304 | 0.971 |
| 96.1 | 0.0234 | 0.994 |
| 22.8 | 0.0056 | 1.000 |
| 0.0 | 0.0000 | 1.000 |

Table 15.6 Coefficients and absolute contributions for regions, Example 15.4

| Region | $r_{1}$ | $r_{2}$ | $r_{3}$ | $C_{a}\left(i, r_{1}\right)$ | $C_{a}\left(i, r_{2}\right)$ | $C_{a}\left(i, r_{3}\right)$ |
| :--- | ---: | ---: | ---: | :--- | :--- | :--- |
| ILDF | 0.1464 | 0.0677 | 0.0157 | 0.3839 | 0.2175 | 0.0333 |
| CHAM | -0.0603 | -0.0410 | -0.0187 | 0.0064 | 0.0078 | 0.0047 |
| PICA | 0.0323 | -0.0258 | -0.0318 | 0.0021 | 0.0036 | 0.0155 |
| HNOR | -0.0692 | 0.0287 | 0.1156 | 0.0096 | 0.0044 | 0.2035 |
| CENT | -0.0068 | -0.0205 | -0.0145 | 0.0001 | 0.0030 | 0.0043 |
| BNOR | -0.0271 | -0.0762 | 0.0061 | 0.0014 | 0.0284 | 0.0005 |
| BOUR | -0.1921 | 0.0188 | 0.0578 | 0.0920 | 0.0023 | 0.0630 |
| NOPC | -0.1278 | 0.0863 | -0.0570 | 0.0871 | 0.1052 | 0.1311 |
| LORR | -0.2084 | 0.0511 | 0.0467 | 0.1606 | 0.0256 | 0.0608 |
| ALSA | -0.2331 | 0.0838 | 0.0655 | 0.1283 | 0.0439 | 0.0767 |
| FRAC | -0.1304 | -0.0368 | -0.0444 | 0.0265 | 0.0056 | 0.0232 |
| PAYL | -0.0743 | -0.0816 | -0.0341 | 0.0232 | 0.0743 | 0.0370 |
| BRET | 0.0158 | 0.0249 | -0.0469 | 0.0011 | 0.0070 | 0.0708 |
| PCHA | -0.0610 | -0.1391 | -0.0178 | 0.0085 | 0.1171 | 0.0054 |
| AQUI | 0.0368 | -0.1183 | 0.0455 | 0.0055 | 0.1519 | 0.0643 |
| MIDI | 0.0208 | -0.0567 | 0.0138 | 0.0018 | 0.0359 | 0.0061 |
| LIMO | -0.0540 | 0.0221 | -0.0427 | 0.0033 | 0.0014 | 0.0154 |
| RHOA | -0.0225 | 0.0273 | -0.0385 | 0.0042 | 0.0161 | 0.0918 |
| AUVE | 0.0290 | -0.0139 | -0.0554 | 0.0017 | 0.0010 | 0.0469 |
| LARO | 0.0290 | -0.0862 | -0.0177 | 0.0383 | 0.0595 | 0.0072 |
| PROV | 0.0469 | -0.0717 | 0.0279 | 0.0142 | 0.0884 | 0.0383 |

responsible for the variation on the first axis. The second axis mostly characterises an opposition between baccalauréats A and C. Regarding the regions, Ile de France plays an important role on each axis. On the first axis, it is opposed to Lorraine and Alsace, whereas on the second axis, it is opposed to Poitou-Charentes and Aquitaine. All of this is confirmed in Fig. 15.3.

On the right side are the more classical baccalauréats and on the left, more technical ones. The regions on the left side have thus larger weights in the technical

Table 15.7 Coefficients and absolute contributions for baccalauréats, Example 15.4

| Baccal | $s_{1}$ | $s_{2}$ | $s_{3}$ | $C_{a}\left(j, s_{1}\right)$ | $C_{a}\left(j, s_{2}\right)$ | $C_{a}\left(j, s_{3}\right)$ |
| :--- | ---: | ---: | ---: | :--- | :--- | :--- |
| A | 0.0447 | -0.0679 | 0.0367 | 0.0376 | 0.2292 | 0.1916 |
| B | 0.1389 | 0.0557 | 0.0011 | 0.1724 | 0.0735 | 0.0001 |
| C | 0.0940 | 0.0995 | 0.0079 | 0.1198 | 0.3556 | 0.0064 |
| D | 0.0227 | -0.0495 | -0.0530 | 0.0098 | 0.1237 | 0.4040 |
| E | -0.1932 | 0.0492 | -0.1317 | 0.0825 | 0.0141 | 0.2900 |
| F | -0.2156 | 0.0862 | 0.0188 | 0.3793 | 0.1608 | 0.0219 |
| G | -0.1244 | -0.0353 | 0.0279 | 0.1969 | 0.0421 | 0.0749 |
| H | -0.0945 | 0.0438 | -0.0888 | 0.0017 | 0.0010 | 0.0112 |

Table 15.8 Eigenvalues and explained proportion of variance, Example 15.5

| $\lambda_{j}$ | Percentage of variance | Cumulated percentage |
| :--- | :--- | :--- |
| $4,399.0$ | 0.4914 | 0.4914 |
| $2,213.6$ | 0.2473 | 0.7387 |
| $1,382.4$ | 0.1544 | 0.8932 |
| 870.7 | 0.0973 | 0.9904 |
| 51.0 | 0.0057 | 0.9961 |
| 34.8 | 0.0039 | 1.0000 |
| 0.0 | 0.0000 | 0.0000 |

baccalauréats. Note also that most of the southern regions of France are concentrated in the lower part of the graph near the baccalauréat $A$.

Finally, looking at the third axis, we see that it is dominated by the baccalauréat E (negative sign) and to a lesser degree by H (negative) (as opposed to A (positive sign)). The dominating regions are HNOR (positive sign), opposed to NOPC and AUVE (negative sign). For instance, HNOR is particularly poor in baccalauréat D.

Example 15.5 The US crime data set (Table 22.10) gives the number of crimes in the 50 states of the US classified in 1985 for each of the following seven categories: murder, rape, robbery, assault, burglary, larceny and auto-theft. The analysis of the contingency table, limited to the first two factors, provides the following results (see Table 15.8).

Looking at the absolute contributions (not reproduced here, see Exercise 15.6), it appears that the first axis is robbery $(+)$ versus larceny $(-)$ and auto-theft $(-)$ axis and that the second factor contrasts assault $(-)$ to auto-theft $(+)$. The dominating states for the first axis are the North-Eastern States MA $(+)$ and NY $(+)$ contrasting the Western States WY ( - ) and ID ( - . For the second axis, the differences are seen between the Northern States (MA $(+)$ and RI $(+))$ and the Southern States AL (-), MS (-) and AR (-). These results can be clearly seen in Fig. 15.4 where all the states and crimes are reported. The figure also shows in which states the proportion of a particular crime category is higher or lower than the national average (the origin).


Fig. 15.4 Projection of rows (the 50 states) and columns (the seven crime categories) Q MVAcorrcrime

## Biplots

The biplot is a low-dimensional display of a data matrix $\mathcal{X}$ where the rows and columns are represented by points. The interpretation of a biplot is specifically directed towards the scalar products of lower dimensional factorial variables and is designed to approximately recover the individual elements of the data matrix in these scalar products. Suppose that we have a $(10 \times 5)$ data matrix with elements $x_{i j}$. The idea of the biplot is to find 10 row points $q_{i} \in \mathbb{R}^{k}(k<p, i=1, \ldots, 10)$ and 5 column points $t_{j} \in \mathbb{R}^{k}(j=1, \ldots, 5)$ such that the 50 scalar products between the row and the column vectors closely approximate the 50 corresponding elements of the data matrix $\mathcal{X}$. Usually we choose $k=2$. For example, the scalar product between $q_{7}$ and $t_{4}$ should approximate the data value $x_{74}$ in the seventh row and the fourth column. In general, the biplot models the data $x_{i j}$ as the sum of a scalar product in some low-dimensional subspace and a residual "error" term:

$$
\begin{align*}
x_{i j} & =q_{i}^{\top} t_{j}+e_{i j} \\
& =\sum_{k} q_{i k} t_{j k}+e_{i j} \tag{15.25}
\end{align*}
$$

To understand the link between correspondence analysis and the biplot, we need to introduce a formula which expresses $x_{i j}$ from the original data matrix (see (15.3)) in terms of row and column frequencies. One such formula, known as the "reconstitution formula", is (15.10):

$$
\begin{equation*}
x_{i j}=E_{i j}\left(1+\frac{\sum_{k=1}^{R} \lambda_{k}^{\frac{1}{2}} \gamma_{i k} \delta_{j k}}{\sqrt{\frac{x_{i \bullet \bullet} x_{\bullet j}}{x_{\bullet \bullet}}}}\right) \tag{15.26}
\end{equation*}
$$

Consider now the row profiles $x_{i j} / x_{i \bullet}$ (the conditional frequencies) and the average row profile $x_{i \bullet} / x_{\bullet \bullet}$. From (15.26) we obtain the difference between each row profile and this average:

$$
\begin{equation*}
\left(\frac{x_{i j}}{x_{i \bullet}}-\frac{x_{i \bullet}}{x_{\bullet \bullet}}\right)=\sum_{k=1}^{R} \lambda_{k}^{\frac{1}{2}} \gamma_{i k}\left(\sqrt{\frac{x_{\bullet} j}{x_{i \bullet \bullet} x_{\bullet \bullet}}}\right) \delta_{j k} . \tag{15.27}
\end{equation*}
$$

By the same argument we can also obtain the difference between each column profile and the average column profile:

$$
\begin{equation*}
\left(\frac{x_{i j}}{x_{\bullet j}}-\frac{x_{\bullet j}}{x_{\bullet \bullet}}\right)=\sum_{k=1}^{R} \lambda_{k}^{\frac{1}{2}} \gamma_{i k}\left(\sqrt{\frac{x_{i \bullet}}{x_{\bullet j} x_{\bullet \bullet}}}\right) \delta_{j k} . \tag{15.28}
\end{equation*}
$$

Now, if $\lambda_{1} \gg \lambda_{2} \gg \lambda_{3} \ldots$, we can approximate these sums by a finite number of $K$ terms (usually $K=2$ ) using (15.16) to obtain

$$
\begin{align*}
& \left(\frac{x_{i j}}{x_{\bullet j}}-\frac{x_{i \bullet}}{x_{\bullet \bullet}}\right)=\sum_{k=1}^{K}\left(\frac{x_{\bullet i}}{\sqrt{\lambda_{k} x_{\bullet \bullet}}} r_{k i}\right) s_{k j}+e_{i j},  \tag{15.29}\\
& \left(\frac{x_{i j}}{x_{i \bullet}}-\frac{x_{\bullet j}}{x_{\bullet \bullet}}\right)=\sum_{k=1}^{K}\left(\frac{x_{\bullet j}}{\sqrt{\lambda_{k} x_{\bullet \bullet}}} s_{k j}\right) r_{k i}+e_{i j}^{\prime}, \tag{15.30}
\end{align*}
$$

where $e_{i j}$ and $e_{i j}^{\prime}$ are error terms. Equation (15.30) shows that if we consider displaying the differences between the row profiles and the average profile, then the projection of the row profile $r_{k}$ and a rescaled version of the projections of the column profile $s_{k}$ constitute a biplot of these differences. Equation (15.29) implies the same for the differences between the column profiles and this average.

| $\hookrightarrow$ | Correspondence analysis is a factorial decomposition of contin- <br> gency tables. The $p$-dimensional individuals and the $n$-dimensional <br> variables can be graphically represented by projecting onto spaces <br> of smaller dimension. |
| :--- | :--- |
| $\hookrightarrow$ | The practical computation consists of first computing a spectral <br> decomposition of $\mathcal{A}^{-1} \mathcal{X} \mathcal{B}^{-1} \mathcal{X}^{\top}$ and $\mathcal{B}^{-1} \mathcal{X}^{\top} \mathcal{A}^{-1} \mathcal{X}$ which have the <br> same first $p$ eigenvalues. The graphical representation is obtained <br> by plotting $\sqrt{\lambda_{1}} r_{1}$ vs. $\sqrt{\lambda_{2}} r_{2}$ and $\sqrt{\lambda_{1}} s_{1}$ vs. $\sqrt{\lambda_{2}} s_{2}$. Both plots <br> maybe displayed in the same graph taking into account the appro- <br> priate orientation of the eigenvectors $r_{i}, s_{j}$. |
| $\hookrightarrow$ | Correspondence analysis provides a graphical display of the asso- <br> ciation measure $c_{i j}=\left(x_{i j}-E_{i j}\right)^{2} / E_{i j}$. |
| $\hookrightarrow$ | Biplot is a low-dimensional display of a data matrix where the rows <br> and columns are represented by points |

### 15.4 Exercises

Exercise 15.1 Show that the matrices $\mathcal{A}^{-1} \mathcal{X} \mathcal{B}^{-1} \mathcal{X}^{\top}$ and $\mathcal{B}^{-1} \mathcal{X}^{\top} \mathcal{A}^{-1} \mathcal{X}$ have an eigenvalue equal to 1 and that the corresponding eigenvectors are proportional to $(1, \ldots, 1)^{\top}$.

Exercise 15.2 Verify the relations in (15.8), (15.14) and (15.17).
Exercise 15.3 Do a correspondence analysis for the car marks data (Table 22.7)! Explain how this table can be considered as a contingency table.
Exercise 15.4 Compute the $\chi^{2}$-statistic of independence for the French baccalauréat data.
Exercise 15.5 Prove that $\mathcal{C}=\mathcal{A}^{-1 / 2}(\mathcal{X}-E) \mathcal{B}^{-1 / 2} \sqrt{x_{\bullet \bullet}}$ and $E=\frac{a b^{\top}}{x_{\bullet \bullet}}$ and verify (15.20).

Exercise 15.6 Do the full correspondence analysis of the US crime data (Table 22.10), and determine the absolute contributions for the first three axes. How can you interpret the third axis? Try to identify the states with one of the four regions to which it belongs. Do you think the four regions have a different behaviour with respect to crime?

Exercise 15.7 Repeat Exercise 15.6 with the US health data (Table 22.16). Only analyse the columns indicating the number of deaths per state.

Exercise 15.8 Consider a $(n \times n)$ contingency table being a diagonal matrix $\mathcal{X}$. What do you expect the factors $r_{k}, s_{k}$ to be like?

Exercise 15.9 Assume that after some reordering of the rows and the columns, the contingency table has the following structure:

$$
\mathcal{X}=\begin{array}{|c|c|c|}
\hline & J_{1} & J_{2} \\
\hline I_{1} & * & 0 \\
\hline I_{2} & 0 & * \\
\hline
\end{array}
$$

That is, the rows $I_{i}$ only have weights in the columns $J_{i}$, for $i=1,2$. What do you expect the graph of the first two factors to look like?

Exercise 15.10 Redo Exercise 15.9 using the following contingency table:

$$
\mathcal{X}=\begin{array}{|c|c|c|c|}
\hline & J_{1} & J_{2} & J_{3} \\
\hline I_{1} & * & 0 & 0 \\
\hline I_{2} & 0 & * & 0 \\
\hline I_{3} & 0 & 0 & * \\
\hline
\end{array}
$$

Exercise 15.11 Consider the French food data (Table 22.6). Given that all of the variables are measured in the same units (Francs), explain how this table can be considered as a contingency table. Perform a correspondence analysis and compare the results to those obtained in the NPCA analysis in Chap. 11.

## Chapter 16 <br> Canonical Correlation Analysis

Complex multivariate data structures are better understood by studying lowdimensional projections. For a joint study of two data sets, we may ask what type of low-dimensional projection helps in finding possible joint structures for the two samples. The canonical correlation analysis (CCA) is a standard tool of multivariate statistical analysis for discovery and quantification of associations between two sets of variables.

The basic technique is based on projections. One defines an index (projected multivariate variable) that maximally correlates with the index of the other variable for each sample separately. The aim of CCA is to maximise the association (measured by correlation) between the low-dimensional projections of the two data sets. The canonical correlation vectors are found by a joint covariance analysis of the two variables. The technique is applied to a marketing example where the association of a price factor and other variables (like design, sportiness etc.) is analysed. Tests are given on how to evaluate the significance of the discovered association.

### 16.1 Most Interesting Linear Combination

The associations between two sets of variables may be identified and quantified by CCA. The technique was originally developed by Hotelling (1935) who analysed how arithmetic speed and arithmetic power are related to reading speed and reading power. Other examples are the relation between governmental policy variables and economic performance variables and the relation between job and company characteristics.

Suppose we are given two random variables $X \in \mathbb{R}^{q}$ and $Y \in \mathbb{R}^{p}$. The idea is to find an index describing a (possible) link between $X$ and $Y$. CCA is based on linear indices, i.e. linear combinations

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

of the random variables. CCA searches for vectors $a$ and $b$ such that the relation of the two indices $a^{\top} x$ and $b^{\top} y$ is quantified in some interpretable way. More precisely, one is looking for the "most interesting" projections $a$ and $b$ in the sense that they maximise the correlation

$$
\begin{equation*}
\rho(a, b)=\rho_{a}^{\top} X b^{\top} Y \tag{16.1}
\end{equation*}
$$

between the two indices.
Let us consider the correlation $\rho(a, b)$ between the two projections in more detail. Suppose that

$$
\binom{X}{Y} \sim\left(\binom{\mu}{v},\left(\begin{array}{ll}
\Sigma_{X X} & \Sigma_{X Y} \\
\Sigma_{Y X} & \Sigma_{Y Y}
\end{array}\right)\right)
$$

where the sub-matrices of this covariance structure are given by

$$
\begin{aligned}
\operatorname{Var}(X) & =\Sigma_{X X}(q \times q) \\
\operatorname{Var}(Y) & =\Sigma_{Y Y}(p \times p) \\
\operatorname{Cov}(X, Y) & =\mathrm{E}(X-\mu)(Y-v)^{\top}=\Sigma_{X Y}=\Sigma_{Y X}^{\top} \quad(q \times p)
\end{aligned}
$$

Using (3.7) and (4.26),

$$
\begin{equation*}
\rho(a, b)=\frac{a^{\top} \Sigma_{X Y} b}{\left(a^{\top} \Sigma_{X X} a\right)^{1 / 2}\left(b^{\top} \Sigma_{Y Y} b\right)^{1 / 2}} . \tag{16.2}
\end{equation*}
$$

Therefore, $\rho(c a, b)=\rho(a, b)$ for any $c \in \mathbb{R}^{+}$. Given the invariance of scale we may rescale projections $a$ and $b$ and thus we can equally solve

$$
\max _{a, b} \quad a^{\top} \Sigma_{X Y} b
$$

under the constraints

$$
\begin{aligned}
a^{\top} \Sigma_{X X} a & =1 \\
b^{\top} \Sigma_{Y Y} b & =1
\end{aligned}
$$

For this problem, define

$$
\begin{equation*}
\mathcal{K}=\Sigma_{X X}^{-1 / 2} \Sigma_{X Y} \Sigma_{Y Y}^{-1 / 2} \tag{16.3}
\end{equation*}
$$

Recall the singular value decomposition of $\mathcal{K}(q \times p)$ from Theorem 2.2. The matrix $\mathcal{K}$ may be decomposed as

$$
\mathcal{K}=\Gamma \Lambda \Delta^{\top}
$$

with

$$
\begin{align*}
\Gamma & =\left(\gamma_{1}, \ldots, \gamma_{k}\right) \\
\Delta & =\left(\delta_{1}, \ldots, \delta_{k}\right)  \tag{16.4}\\
\Lambda & =\operatorname{diag}\left(\lambda_{1}^{1 / 2}, \ldots, \lambda_{k}^{1 / 2}\right)
\end{align*}
$$

where by (16.3) and (2.15),

$$
k=\operatorname{rank}(\mathcal{K})=\operatorname{rank}\left(\Sigma_{X Y}\right)=\operatorname{rank}\left(\Sigma_{Y X}\right),
$$

and $\lambda_{1} \geq \lambda_{2} \geq \cdots \lambda_{k}$ are the nonzero eigenvalues of $\mathcal{N}_{1}=\mathcal{K} \mathcal{K}^{\top}$ and $\mathcal{N}_{2}=\mathcal{K}^{\top} \mathcal{K}$ and $\gamma_{i}$ and $\delta_{j}$ are the standardised eigenvectors of $\mathcal{N}_{1}$ and $\mathcal{N}_{2}$ respectively.

Define now for $i=1, \ldots, k$ the vectors

$$
\begin{align*}
a_{i} & =\Sigma_{X X}^{-1 / 2} \gamma_{i},  \tag{16.5}\\
b_{i} & =\Sigma_{Y Y}^{-1 / 2} \delta_{i}, \tag{16.6}
\end{align*}
$$

which are called the canonical correlation vectors. Using these canonical correlation vectors we define the canonical correlation variables

$$
\begin{align*}
\eta_{i} & =a_{i}^{\top} X  \tag{16.7}\\
\varphi_{i} & =b_{i}^{\top} Y . \tag{16.8}
\end{align*}
$$

The quantities $\rho_{i}=\lambda_{i}^{1 / 2}$ for $i=1, \ldots, k$ are called the canonical correlation coefficients.

From the properties of the singular value decomposition given in (16.4) we have

$$
\operatorname{Cov}\left(\eta_{i}, \eta_{j}\right)=a_{i}^{\top} \Sigma_{X X} a_{j}=\gamma_{i}^{\top} \gamma_{j}= \begin{cases}1 & i=j  \tag{16.9}\\ 0 & i \neq j\end{cases}
$$

The same is true for $\operatorname{Cov}\left(\varphi_{i}, \varphi_{j}\right)$. The following theorem tells us that the canonical correlation vectors are the solution to the maximisation problem of (16.1).

Theorem 16.1 For any given $r, 1 \leq r \leq k$, the maximum

$$
\begin{equation*}
C(r)=\max _{a, b} a^{\top} \Sigma_{X Y} b \tag{16.10}
\end{equation*}
$$

subject to

$$
a^{\top} \Sigma_{X X} a=1, \quad b^{\top} \Sigma_{Y Y} b=1
$$

and

$$
a_{i}^{\top} \Sigma_{X X} a=0 \text { for } i=1, \ldots, r-1
$$

is given by

$$
C(r)=\rho_{r}=\lambda_{r}^{1 / 2}
$$

and is attained when $a=a_{r}$ and $b=b_{r}$.
Proof The proof is given in three steps.
(i) Fix $a$ and maximise over $b$, i.e. solve:

$$
\max _{b}\left(a^{\top} \Sigma_{X Y} b\right)^{2}=\max _{b}\left(b^{\top} \Sigma_{Y X} a\right)\left(a^{\top} \Sigma_{X Y} b\right)
$$

subject to $b^{\top} \Sigma_{Y Y} b=1$. By Theorem 2.5 the maximum is given by the largest eigenvalue of the matrix

$$
\Sigma_{Y Y}^{-1} \Sigma_{Y X} a a^{\top} \Sigma_{X Y} .
$$

By Corollary 2.2, the only nonzero eigenvalue equals

$$
\begin{equation*}
a^{\top} \Sigma_{X Y} \Sigma_{Y Y}^{-1} \Sigma_{Y X} a . \tag{16.11}
\end{equation*}
$$

(ii) Maximise (16.11) over $a$ subject to the constraints of the theorem. Put $\gamma=$ $\Sigma_{X X}^{1 / 2} a$ and observe that (16.11) equals

$$
\gamma^{\top} \Sigma_{X X}^{-1 / 2} \Sigma_{X Y} \Sigma_{Y Y}^{-1} \Sigma_{Y X} \Sigma_{X X}^{-1 / 2} \gamma=\gamma^{\top} \mathcal{K}^{\top} \mathcal{K} \gamma .
$$

Thus, solve the equivalent problem

$$
\begin{equation*}
\max _{\gamma} \gamma^{\top} \mathcal{N}_{1} \gamma \tag{16.12}
\end{equation*}
$$

subject to $\gamma^{\top} \gamma=1, \gamma_{i}^{\top} \gamma=0$ for $i=1, \ldots, r-1$.

Note that the $\gamma_{i}$ 's are the eigenvectors of $\mathcal{N}_{1}$ corresponding to its first $r-1$ largest eigenvalues. Thus, as in Theorem 11.3, the maximum in (16.12) is obtained by setting $\gamma$ equal to the eigenvector corresponding to the $r$ th largest eigenvalue, i.e. $\gamma=\gamma_{r}$ or equivalently $a=a_{r}$. This yields

$$
C^{2}(r)=\gamma_{r}^{\top} \mathcal{N}_{1} \gamma_{r}=\lambda_{r} \gamma_{r}^{\top} \gamma=\lambda_{r}
$$

(iii) Show that the maximum is attained for $a=a_{r}$ and $b=b_{r}$. From the SVD of $\mathcal{K}$ we conclude that $\mathcal{K} \delta_{r}=\rho_{r} \gamma_{r}$ and hence

$$
a_{r}^{\top} \Sigma_{X Y} b_{r}=\gamma_{r}^{\top} \mathcal{K} \delta_{r}=\rho_{r} \gamma_{r}^{\top} \gamma_{r}=\rho_{r}
$$

Let

$$
\binom{X}{Y} \sim\left(\binom{\mu}{v},\left(\begin{array}{cc}
\Sigma_{X X} & \Sigma_{X Y} \\
\Sigma_{Y X} & \Sigma_{Y Y}
\end{array}\right)\right)
$$

The canonical correlation vectors

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

maximise the correlation between the canonical variables

$$
\begin{aligned}
\eta_{1} & =a_{1}^{\top} X \\
\varphi_{1} & =b_{1}^{\top} Y
\end{aligned}
$$

The covariance of the canonical variables $\eta$ and $\varphi$ is given in the next theorem.
Theorem 16.2 Let $\eta_{i}$ and $\varphi_{i}$ be the $i$ th canonical correlation variables $(i=$ $1, \ldots, k)$. Define $\eta=\left(\eta_{1}, \ldots, \eta_{k}\right)$ and $\varphi=\left(\varphi_{1}, \ldots, \varphi_{k}\right)$. Then

$$
\operatorname{Var}\binom{\eta}{\varphi}=\left(\begin{array}{cc}
\mathcal{I}_{k} & \Lambda \\
\Lambda & \mathcal{I}_{k}
\end{array}\right)
$$

with $\Lambda$ given in (16.4).
This theorem shows that the canonical correlation coefficients, $\rho_{i}=\lambda_{i}^{1 / 2}$, are the covariances between the canonical variables $\eta_{i}$ and $\varphi_{i}$ and that the indices $\eta_{1}=$ $a_{1}^{\top} X$ and $\varphi_{1}=b_{1}^{\top} Y$ have the maximum covariance $\sqrt{\lambda_{1}}=\rho_{1}$.

The following theorem shows that canonical correlations are invariant w.r.t. linear transformations of the original variables.

Theorem 16.3 Let $X^{*}=\mathcal{U}^{\top} X+u$ and $Y^{*}=\mathcal{V}^{\top} Y+v$ where $\mathcal{U}$ and $\mathcal{V}$ are nonsingular matrices. Then the canonical correlations between $X^{*}$ and $Y^{*}$ are the same as those between $X$ and $Y$. The canonical correlation vectors of $X^{*}$ and $Y^{*}$ are given by

$$
\begin{align*}
a_{i}^{*} & =\mathcal{U}^{-1} a_{i}, \\
b_{i}^{*} & =\mathcal{V}^{-1} b_{i} . \tag{16.13}
\end{align*}
$$

| Sumpary |
| :--- | :--- |

### 16.2 Canonical Correlation in Practice

In practice we have to estimate the covariance matrices $\Sigma_{X X}, \Sigma_{X Y}$ and $\Sigma_{Y Y}$. Let us apply the CCA to the car marks data (see Table 22.7). In the context of this data set one is interested in relating price variables with variables such as sportiness and safety. In particular, we would like to investigate the relation between the two variables non-depreciation of value and price of the car and all other variables.

Example 16.1 We perform the CCA on the data matrices $\mathcal{X}$ and $\mathcal{Y}$ that correspond to the set of values $\{$ Price, Value Stability $\}$ and $\{$ Economy, Service, Design, Sporty car, Safety, Easy handling\}, respectively. The estimated covariance matrix $\mathcal{S}$ is given by

$$
\begin{aligned}
& \text { Price Value Econ. Serv. Design Sport. Safety Easy h. } \\
& \mathcal{S}=\left(\begin{array}{rr|rrrrrr}
1.41 & -1.11 & 0.78 & -0.71 & -0.90 & -1.04 & -0.95 & 0.18 \\
-1.11 & 1.19 & -0.42 & 0.82 & 0.77 & 0.90 & 1.12 & 0.11 \\
--- & ---\mid & --- & --- & --- & --- & ---- & --- \\
0.78 & -0.42 & 0.75 & -0.23 & -0.45 & -0.42 & -0.28 & 0.28 \\
-0.71 & 0.82 & -0.23 & 0.66 & 0.52 & 0.57 & 0.85 & 0.14 \\
-0.90 & 0.77 & -0.45 & 0.52 & 0.72 & 0.77 & 0.68 & -0.10 \\
-1.04 & 0.90 & -0.42 & 0.57 & 0.77 & 1.05 & 0.76 & -0.15 \\
-0.95 & 1.12 & -0.28 & 0.85 & 0.68 & 0.76 & 1.26 & 0.22 \\
0.18 & 0.11 & 0.28 & 0.14 & -0.10 & -0.15 & 0.22 & 0.32
\end{array}\right)
\end{aligned}
$$

Hence,

$$
\begin{aligned}
& \mathcal{S}_{X X}=\left(\begin{array}{rr}
1.41 & -1.11 \\
-1.11 & 1.19
\end{array}\right), \quad \mathcal{S}_{X Y}=\left(\begin{array}{rrrrrr}
0.78 & -0.71 & -0.90 & -1.04 & -0.95 & 0.18 \\
-0.42 & 0.82 & 0.77 & 0.90 & 1.12 & 0.11
\end{array}\right), \\
& \mathcal{S}_{Y Y}=\left(\begin{array}{rrrrrr}
0.75 & -0.23 & -0.45 & -0.42 & -0.28 & 0.28 \\
-0.23 & 0.66 & 0.52 & 0.57 & 0.85 & 0.14 \\
-0.45 & 0.52 & 0.72 & 0.77 & 0.68 & -0.10 \\
-0.42 & 0.57 & 0.77 & 1.05 & 0.76 & -0.15 \\
-0.28 & 0.85 & 0.68 & 0.76 & 1.26 & 0.22 \\
0.28 & 0.14 & -0.10 & -0.15 & 0.22 & 0.32
\end{array}\right) .
\end{aligned}
$$

It is interesting to see that value stability and price have a negative covariance. This makes sense since highly priced vehicles tend to loose their market value at a faster pace than medium priced vehicles.

Now we estimate $\mathcal{K}=\Sigma_{X X}^{-1 / 2} \Sigma_{X Y} \Sigma_{Y Y}^{-1 / 2}$ by

$$
\hat{\mathcal{K}}=\mathcal{S}_{X X}^{-1 / 2} \mathcal{S}_{X Y} \mathcal{S}_{Y Y}^{-1 / 2}
$$

and perform a singular value decomposition of $\hat{\mathcal{K}}$ :

$$
\hat{\mathcal{K}}=\mathcal{G} \mathcal{L D} \mathcal{D}^{\top}=\left(g_{1}, g_{2}\right) \operatorname{diag}\left(\ell_{1}^{1 / 2}, \ell_{2}^{1 / 2}\right)\left(d_{1}, d_{2}\right)^{\top}
$$

## Car Marks Data



Fig. 16.1 The second canonical variables for the car marks data $\mathbf{Q}$ texttMMVAcancarm
where the $\ell_{i}$ 's are the eigenvalues of $\hat{\mathcal{K}} \hat{\mathcal{K}}^{\top}$ and $\hat{\mathcal{K}}^{\top} \hat{\mathcal{K}}$ with $\operatorname{rank}(\hat{\mathcal{K}})=2$, and $g_{i}$ and $d_{i}$ are the eigenvectors of $\hat{\mathcal{K}} \hat{\mathcal{K}}^{\top}$ and $\hat{\mathcal{K}}^{\top} \hat{\mathcal{K}}$, respectively. The canonical correlation coefficients are

$$
r_{1}=\ell_{1}^{1 / 2}=0.98, \quad r_{2}=\ell_{2}^{1 / 2}=0.89
$$

The high correlation of the second two canonical variables can be seen in Fig. 16.1. The second canonical variables are
$\hat{\eta}_{1}=\hat{a}_{1}^{\top} x=1.602 x_{1}+1.686 x_{2}$
$\hat{\varphi}_{1}=\hat{b}_{1}^{\top} y=0.568 y_{1}+0.544 y_{2}-0.012 y_{3}-0.096 y_{4}-0.014 y_{5}+0.915 y_{6}$.
Note that the variables $y_{1}$ (economy), $y_{2}$ (service) and $y_{6}$ (easy handling) have positive coefficients on $\hat{\varphi}_{1}$. The variables $y_{3}$ (design), $y_{4}$ (sporty car) and $y_{5}$ (safety) have a negative influence on $\hat{\varphi}_{1}$.

The canonical variable $\eta_{1}$ may be interpreted as a price and value index. The canonical variable $\varphi_{1}$ is mainly formed from the qualitative variables economy, service and handling with negative weights on design, safety and sportiness. These variables may therefore be interpreted as an appreciation of the value of the car. The sportiness has a negative effect on the price and value index, as do the design and the safety features.

## Testing the Canonical Correlation Coefficients

The hypothesis that the two sets of variables $\mathcal{X}$ and $\mathcal{Y}$ are uncorrelated may be tested (under normality assumptions) with Wilks likelihood ratio statistic (Gibbins, 1985):

$$
T^{2 / n}=\left|\mathcal{I}-S_{Y Y}^{-1} S_{Y X} S_{X X}^{-1} S_{X Y}\right|=\prod_{i=1}^{k}\left(1-\ell_{i}\right)
$$

This statistic unfortunately has a rather complicated distribution. Bartlett (1939) provides an approximation for large $n$ :

$$
\begin{equation*}
-\{n-(p+q+3) / 2\} \log \prod_{i=1}^{k}\left(1-\ell_{i}\right) \sim \chi_{p q}^{2} \tag{16.14}
\end{equation*}
$$

A test of the hypothesis that only $s$ of the canonical correlation coefficients are nonzero may be based (asymptotically) on the statistic

$$
\begin{equation*}
-\{n-(p+q+3) / 2\} \log \prod_{i=s+1}^{k}\left(1-\ell_{i}\right) \sim \chi_{(p-s)(q-s)}^{2} \tag{16.15}
\end{equation*}
$$

Example 16.2 Consider Example 16.1 again. There are $n=40$ persons that have rated the cars according to different categories with $p=2$ and $q=6$. The canonical correlation coefficients were found to be $r_{1}=0.98$ and $r_{2}=0.89$. Bartlett's statistic (16.14) is therefore

$$
-\{40-(2+6+3) / 2\} \log \left\{\left(1-0.98^{2}\right)\left(1-0.89^{2}\right)\right\}=165.59 \sim \chi_{12}^{2}
$$

which is highly significant (the $99 \%$ quantile of the $\chi_{12}^{2}$ is 26.23 ). The hypothesis of no correlation between the variables $\mathcal{X}$ and $\mathcal{Y}$ is therefore rejected.

Let us now test whether the second canonical correlation coefficient is different from zero. We use Bartlett's statistic (16.15) with $s=1$ and obtain

$$
-\{40-(2+6+3) / 2\} \log \left\{\left(1-0.89^{2}\right)\right\}=54.19 \sim \chi_{5}^{2}
$$

which is again highly significant with the $\chi_{5}^{2}$ distribution.

## CCA with Qualitative Data

The canonical correlation technique may also be applied to qualitative data. Consider for example the contingency table $\mathcal{N}$ of the French baccalauréat data. The dataset is given in Table 22.8 in Chap. 22. The CCA cannot be applied directly to
this contingency table since the table does not correspond to the usual data matrix structure. We may wish, however, to explain the relationship between the row $r$ and column $c$ categories. It is possible to represent the data in a $(n \times(r+c))$ data matrix $\mathcal{Z}=(\mathcal{X}, \mathcal{Y})$ where $n$ is the total number of frequencies in the contingency table $\mathcal{N}$ and $\mathcal{X}$ and $\mathcal{Y}$ are matrices of zero-one dummy variables. More precisely, let

$$
x_{k i}= \begin{cases}1 & \text { if the } k \text { th individual belongs to the } i \text { th row category } \\ 0 & \text { otherwise }\end{cases}
$$

and

$$
y_{k j}= \begin{cases}1 & \text { if the } k \text { th individual belongs to the } j \text { th column category } \\ 0 & \text { otherwise }\end{cases}
$$

where the indices range from $k=1, \ldots, n, i=1, \ldots, r$ and $j=1, \ldots, c$. Denote the cell frequencies by $n_{i j}$ so that $\mathcal{N}=\left(n_{i j}\right)$ and note that

$$
x_{(i)}^{\top} y_{(j)}=n_{i j}
$$

where $x_{(i)}\left(y_{(j)}\right)$ denotes the $i$ th $(j$ th $)$ column of $\mathcal{X}(\mathcal{Y})$.

$$
\mathcal{X}=\left(\begin{array}{lll}
1 & 0 \\
1 & 0 \\
1 & 0 \\
1 & 0 \\
1 & 0 \\
0 & 1 \\
0 & 1 \\
0 & 1 \\
0 & 1 \\
0 & 1
\end{array}\right), \mathcal{Y}=\left(\begin{array}{ll}
1 & 0 \\
1 & 0 \\
1 & 0 \\
0 & 1 \\
0 & 1 \\
1 & 0 \\
0 & 1 \\
0 & 1 \\
0 & 1 \\
0 & 1
\end{array}\right), \mathcal{Z}=(\mathcal{X}, \mathcal{Y})=\left(\begin{array}{llll}
1 & 0 & 1 & 0 \\
1 & 0 & 1 & 0 \\
1 & 0 & 1 & 0 \\
1 & 0 & 0 & 1 \\
1 & 0 & 0 & 1 \\
0 & 1 & 1 & 0 \\
0 & 1 & 0 & 1 \\
0 & 1 & 0 & 1 \\
0 & 1 & 0 & 1 \\
0 & 1 & 0 & 1
\end{array}\right) .
$$

Example 16.3 Consider the following example where

$$
\mathcal{N}=\left(\begin{array}{ll}
3 & 2 \\
1 & 4
\end{array}\right)
$$

The matrices $\mathcal{X}, \mathcal{Y}$ and $\mathcal{Z}$ are therefore
The element $n_{12}$ of $\mathcal{N}$ may be obtained by multiplying the first column of $\mathcal{X}$ with the second column of $\mathcal{Y}$ to yield

$$
x_{(1)}^{\top} y_{(2)}=2 .
$$

The purpose is to find the canonical variables $\eta=a^{\top} x$ and $\varphi=b^{\top} y$ that are maximally correlated. Note, however, that $x$ has only one nonzero component and therefore an "individual" may be directly associated with its canonical variables or score $\left(a_{i}, b_{j}\right)$. There will be $n_{i j}$ points at each $\left(a_{i}, b_{j}\right)$ and the correlation represented by these points may serve as a measure of dependence between the rows and columns of $\mathcal{N}$.

Let $\mathcal{Z}=(\mathcal{X}, \mathcal{Y})$ denote a data matrix constructed from a contingency table $\mathcal{N}$. Similar to Chap. 14 define

$$
\begin{aligned}
& c=x_{i \bullet}=\sum_{j=1}^{c} n_{i j}, \\
& d=x_{\bullet j}=\sum_{i=1}^{r} n_{i j},
\end{aligned}
$$

and define $\mathcal{C}=\operatorname{diag}(c)$ and $\mathcal{D}=\operatorname{diag}(d)$. Suppose that $x_{i \bullet}>0$ and $x_{\bullet j}>0$ for all $i$ and $j$. It is not hard to see that

$$
\begin{aligned}
n S & =\mathcal{Z}^{\top} \mathcal{H Z}=\mathcal{Z}^{\top} \mathcal{Z}-n \bar{z}^{\top}=\left(\begin{array}{cc}
n S_{X X} & n S_{X Y} \\
n S_{Y X} & n S_{Y Y}
\end{array}\right) \\
& =\left(\frac{n}{n-1}\right)\left(\begin{array}{cc}
\mathcal{C}-n^{-1} c c^{\top} & \mathcal{N}-\hat{\mathcal{N}} \\
\mathcal{N}^{\top} \hat{\mathcal{N}}^{\top} & \mathcal{D}-n^{-1} d d^{\top}
\end{array}\right)
\end{aligned}
$$

where $\hat{\mathcal{N}}=c d^{\top} / n$ is the estimated value of $\mathcal{N}$ under the assumption of independence of the row and column categories.

Note that

$$
(n-1) S_{X X} 1_{r}=\mathcal{C} 1_{r}-n^{-1} c c^{\top} 1_{r}=c-c\left(n^{-1} c^{\top} 1_{r}\right)=c-c\left(n^{-1} n\right)=0
$$

and therefore $S_{X X}^{-1}$ does not exist. The same is true for $S_{Y Y}^{-1}$. One way out of this difficulty is to drop one column from both $\mathcal{X}$ and $\mathcal{Y}$, say the first column. Let $\bar{c}$ and $\bar{d}$ denote the vectors obtained by deleting the first component of $c$ and $d$.

Define $\overline{\mathcal{C}}, \overline{\mathcal{D}}$ and $\bar{S}_{X X}, \bar{S}_{Y Y}, \bar{S}_{X Y}$ accordingly and obtain

$$
\begin{aligned}
& \left(n \bar{S}_{X X}\right)^{-1}=\overline{\mathcal{C}}^{-1}+n_{i \bullet}^{-1} 1_{r} 1_{r}^{\top} \\
& \left(n \bar{S}_{Y Y}\right)^{-1}=\overline{\mathcal{D}}^{-1}+n_{\bullet j}^{-1} 1_{c} 1_{c}^{\top}
\end{aligned}
$$

so that (16.3) exists. The score associated with an individual contained in the first row (column) category of $\mathcal{N}$ is 0 .

The technique described here for purely qualitative data may also be used when the data is a mixture of qualitative and quantitative characteristics. One has to "blow up" the data matrix by dummy zero-one values for the qualitative data variables.

| $\hookrightarrow$ | In practice we estimate $\Sigma_{X X}, \Sigma_{X Y}, \Sigma_{Y Y}$ by the empirical covari- <br> ances and use them to compute estimates $\ell_{i}, g_{i}, d_{i}$ for $\lambda_{i}, \gamma_{i}, \delta_{i}$ <br> from the SVD of $\hat{\mathcal{K}}=\mathcal{S}_{X X}^{-1 / 2} \mathcal{S}_{X Y} \mathcal{S}_{Y Y}^{-1 / 2}$. |
| :--- | :--- |
| $\hookrightarrow$ | The signs of the coefficients of the canonical variables tell us the <br> direction of the influence of these variables. |

### 16.3 Exercises

Exercise 16.1 Show that the eigenvalues of $\mathcal{K} \mathcal{K}^{\top}$ and $\mathcal{K}^{\top} \mathcal{K}$ are identical. (Hint: Use Theorem 2.6.)

Exercise 16.2 Perform the CCA for the following subsets of variables: $\mathcal{X}$ corresponding to $\{$ price $\}$ and $\mathcal{Y}$ corresponding to \{economy, easy handling\} from the car marks data (Table 22.7).

Exercise 16.3 Calculate the first canonical variables for Example 16.1. Interpret the coefficients.

Exercise 16.4 Use the SVD of matrix $\mathcal{K}$ to show that the canonical variables $\eta_{1}$ and $\eta_{2}$ are not correlated.

Exercise 16.5 Verify that the number of nonzero eigenvalues of matrix $\mathcal{K}$ is equal to $\operatorname{rank}\left(\Sigma_{X Y}\right)$.

Exercise 16.6 Express the singular value decomposition of matrices $\mathcal{K}$ and $\mathcal{K}^{\top}$ using eigenvalues and eigenvectors of matrices $\mathcal{K}^{\top} \mathcal{K}$ and $\mathcal{K} \mathcal{K}^{\top}$.
Exercise 16.7 What will be the result of CCA for $Y=X$ ?
Exercise 16.8 What will be the results of CCA for $Y=2 X$ and for $Y=-X$ ?
Exercise 16.9 What results do you expect if you perform CCA for $X$ and $Y$ such that $\Sigma_{X Y}=0$ ? What if $\Sigma_{X Y}=\mathcal{I}_{p}$ ?

## Chapter 17 <br> Multidimensional Scaling

One major aim of multivariate data analysis is dimension reduction. For data measured in Euclidean coordinates, Factor Analysis and Principal Component Analysis are dominantly used tools. In many applied sciences data is recorded as ranked information. For example, in marketing, one may record "product A is better than product B". High-dimensional observations therefore often have mixed data characteristics and contain relative information (w.r.t. a defined standard) rather than absolute coordinates that would enable us to employ one of the multivariate techniques presented so far.

Multidimensional scaling (MDS) is a method based on proximities between objects, subjects, or stimuli used to produce a spatial representation of these items. Proximities express the similarity or dissimilarity between data objects. It is a dimension reduction technique since the aim is to find a set of points in low dimension (typically two dimensions) that reflect the relative configuration of the high-dimensional data objects. The metric MDS is concerned with such a representation in Euclidean coordinates. The desired projections are found via an appropriate spectral decomposition of a distance matrix.

The metric MDS solution may result in projections of data objects that conflict with the ranking of the original observations. The nonmetric MDS solves this problem by iterating between a monotising algorithmic step and a least squares projection step. The examples presented in this chapter are based on reconstructing a map from a distance matrix and on marketing concerns such as ranking of the outfit of cars.

### 17.1 The Problem

MDS is a mathematical tool that uses proximities between objects, subjects or stimuli to produce a spatial representation of these items. The proximities are defined as any set of numbers that express the amount of similarity or dissimilarity
between pairs of objects, subjects or stimuli. 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 distance matrix, $\mathcal{D}$, with the elements $\delta_{i j}$ and $d_{i j}$ respectively. Hence, the underlying dimensionality of the data under investigation is in general not known .

MDS is a data reduction technique because it is concerned with the problem of finding a set of points in low dimension that represents the "configuration" of data in high dimension. The "configuration" in high dimension is represented by the distance or dissimilarity matrix $\mathcal{D}$.

MDS-techniques are often used to understand how people perceive and evaluate certain signals and information. For instance, political scientists use MDS techniques to understand why political candidates are perceived by voters as being similar or dissimilar. Psychologists use MDS to understand the perceptions and evaluations of speech, colours and personality traits, among other things. Last but not least, in marketing researchers use MDS techniques to shed light on the way consumers evaluate brands and to assess the relationship between product attributes.

In short, the primary purpose of all MDS-techniques is to uncover structural relations or patterns in the data and to represent it in a simple geometrical model or picture. One of the aims is to determine the dimension of the model (the goal is a low-dimensional, easily interpretable model) by finding the $d$-dimensional space in which there is maximum correspondence between the observed proximities and the distances between points measured on a metric scale.

MDS based on proximities is usually referred to as metric MDS, whereas the more popular nonmetric MDS is used when the proximities are measured on an ordinal scale.

Example 17.1 A good example of how MDS works is given by Dillon and Goldstein (1984) (page 108). Suppose one is confronted with a map of Germany and asked to measure, with the use of a ruler and the scale of the map, some intercity distances. Admittedly this is quite an easy exercise. However, let us now reverse the problem: One is given a set of distances, as in Table 17.1, and is asked to recreate the map itself. This is a far more difficult exercise, though it can be solved with a

Table 17.1 Inter-city distances

|  | Berlin | Dresden | Hamburg | Koblenz | Munich | Rostock |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| Berlin | 0 | 214 | 279 | 610 | 596 | 237 |
| Dresden |  | 0 | 492 | 533 | 496 | 444 |
| Hamburg |  |  | 0 | 520 | 772 | 140 |
| Koblenz |  |  |  | 0 | 521 | 687 |
| Munich |  |  |  |  | 0 | 771 |
| Rostock |  |  |  |  |  | 0 |

Fig. 17.1 Metric MDS solution for the inter-city road distances $\mathbf{Q}$ MVAMDScity1


Fig. 17.2 Metric MDS solution for the inter-city road distances after reflection and $90^{\circ}$ rotation Q MVAMDScity2

ruler and a compass in two dimensions. MDS is a method for solving this reverse problem in arbitrary dimensions. In Figure 17.1 and 17.2 you can see the graphical representation of the metric MDS solution to Table 17.1 after rotating and reflecting the points representing the cities. Note that the distances given in Table 17.1 are road distances that in general do not correspond to Euclidean distances. In real-life applications, the problems are exceedingly more complex: there are usually errors in the data and the dimensionality is rarely known in advance.

Table 17.2 Dissimilarities for 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 | $\ldots$ |
| $\vdots$ | $\vdots$ | $\vdots$ | $\vdots$ | $\vdots$ | $\ddots$ |

Fig. 17.3 MDS solution on the car data $\mathbf{Q}$ MVAmdscarm


Example 17.2 A further example is given in Table 17.2 where consumers noted their impressions of the dissimilarity of certain cars. The dissimilarities in this table were in fact computed from Table 22.7 as Euclidean distances

$$
d_{i j}=\sqrt{\sum_{l=1}^{8}\left(x_{i l}-x_{j l}\right)^{2}} .
$$

MDS produces Fig. 17.3 which shows a non-linear relationship for all the cars in the projection. This enables us to build a non-linear (quadratic) index with the Wartburg and the Trabant on the left and the Ferrari and the Jaguar on the right. We can construct an order or ranking of the cars based on the subjective impression of the consumers.

Fig. 17.4 Correlation between the MDS direction and the variables $\mathbf{Q}$ MVAmdscarm


What does the ranking describe? The answer is given by Fig. 17.4 which shows the correlation between the MDS projection and the variables. Apparently, the first MDS direction is highly correlated with service(-), value(-), design(-), sportiness(-), safety(-) and price(+). We can interpret the first direction as the price direction since a bad mark in price ("high price") obviously corresponds with a good mark, say, in sportiness ("very sportive"). The second MDS direction is highly positively correlated with practicability. We observe from this data an almost orthogonal relationship between price and practicability.

In MDS a map is constructed in Euclidean space that corresponds to given distances. Which solution can we expect? The solution is determined only up to rotation, reflection and shifts. In general, if $P_{1}, \ldots, P_{n}$ with coordinates $x_{i}=$ $\left(x_{i 1}, \ldots, x_{i p}\right)^{\top}$ for $i=1, \ldots, n$ represents a MDS solution in $p$ dimensions, then $y_{i}=\mathcal{A} x_{i}+b$ with an orthogonal matrix $\mathcal{A}$ and a shift vector $b$ also represents a MDS solution. A comparison of Figs. 17.1 and 17.2 illustrates this fact.

Solution methods that use only the rank order of the distances are termed nonmetric methods of MDS. Methods aimed at finding the points $P_{i}$ directly from a distance matrix like the one in the Table 17.2 are called metric methods.

|  | MDS is a set of techniques which use distances or dissimilarities <br> to project high-dimensional data into a low-dimensional space <br> essential in understanding respondents perceptions and evaluations <br> for all sorts of items. |
| :--- | :--- |
|  | MDS starts with a $(n \times n)$ proximity matrix $\mathcal{D}$ consisting of <br> dissimilarities $\delta_{i, j}$ or distances $d_{i j}$. |
| $\hookrightarrow$ | MDS is an explorative technique and focuses on data reduction. |
| $\hookrightarrow$ | The MDS-solution is indeterminate with respect to rotation, reflec- <br> tion and shifts. |
| $\hookrightarrow$ | The MDS-techniques are divided into metric MDS and nonmetric <br> MDS. |

### 17.2 Metric MDS

Metric MDS begins with a $(n \times n)$ distance matrix $\mathcal{D}$ with elements $d_{i j}$ where $i, j=$ $1, \ldots, n$. The objective of metric MDS is to find a configuration of points in $p$ dimensional space from the distances between the points such that the coordinates of the $n$ points along the $p$ dimensions yield a Euclidean distance matrix whose elements are as close as possible to the elements of the given distance matrix $\mathcal{D}$.

## The Classical Solution

The classical solution is based on a distance matrix that is computed from a Euclidean geometry.

Definition 17.1 A $(n \times n)$ distance matrix $\mathcal{D}=\left(d_{i j}\right)$ is Euclidean if for some points $x_{1}, \ldots, x_{n} \in \mathbb{R}^{p} ; d_{i j}^{2}=\left(x_{i}-x_{j}\right)^{\top}\left(x_{i}-x_{j}\right)$.

The following result tells us whether a distance matrix is Euclidean or not.
Theorem 17.1 Define $\mathcal{A}=\left(a_{i j}\right), a_{i j}=-\frac{1}{2} d_{i j}^{2}$ and $\mathcal{B}=\mathcal{H} \mathcal{A H}$ where $\mathcal{H}$ is the centering matrix. $\mathcal{D}$ is Euclidean if and only if $\mathcal{B}$ is positive semidefinite. If $\mathcal{D}$ is the distance matrix of a data matrix $\mathcal{X}$, then $\mathcal{B}=\mathcal{H} \mathcal{X} \mathcal{X}^{\top} \mathcal{H}$. $\mathcal{B}$ is called the inner product matrix.

## Recovery of Coordinates

The task of MDS is to find the original Euclidean coordinates from a given distance matrix. Let the coordinates of $n$ points in a $p$ dimensional Euclidean space be given by $x_{i}(i=1, \ldots, n)$ where $x_{i}=\left(x_{i 1}, \ldots, x_{i p}\right)^{\top}$. Call $\mathcal{X}=\left(x_{1}, \ldots, x_{n}\right)^{\top}$ the coordinate matrix and assume $\bar{x}=0$. The Euclidean distance between the $i$ th and $j$ th points is given by:

$$
\begin{equation*}
d_{i j}^{2}=\sum_{k=1}^{p}\left(x_{i k}-x_{j k}\right)^{2} . \tag{17.1}
\end{equation*}
$$

The general $b_{i j}$ term of $\mathcal{B}$ is given by:

$$
\begin{equation*}
b_{i j}=\sum_{k=1}^{p} x_{i k} x_{j k}=x_{i}^{\top} x_{j} \tag{17.2}
\end{equation*}
$$

It is possible to derive $\mathcal{B}$ from the known squared distances $d_{i j}$, and then from $\mathcal{B}$ the unknown coordinates.

$$
\begin{align*}
d_{i j}^{2} & =x_{i}^{\top} x_{i}+x_{j}^{\top} x_{j}-2 x_{i}^{\top} x_{j} \\
& =b_{i i}+b_{j j}-2 b_{i j} . \tag{17.3}
\end{align*}
$$

Centering of the coordinate matrix $\mathcal{X}$ implies that $\sum_{i=1}^{n} b_{i j}=0$. Summing (17.3) over $i$ and $j$, we find:

$$
\begin{align*}
\frac{1}{n} \sum_{i=1}^{n} d_{i j}^{2} & =\frac{1}{n} \sum_{i=1}^{n} b_{i i}+b_{j j} \\
\frac{1}{n} \sum_{j=1}^{n} d_{i j}^{2} & =b_{i i}+\frac{1}{n} \sum_{j=1}^{n} b_{j j} \\
\frac{1}{n^{2}} \sum_{i=1}^{n} \sum_{j=1}^{n} d_{i j}^{2} & =\frac{2}{n} \sum_{i=1}^{n} b_{i i} \tag{17.4}
\end{align*}
$$

Solving (17.3) and (17.4) gives:

$$
\begin{equation*}
b_{i j}=-\frac{1}{2}\left(d_{i j}^{2}-d_{i \bullet}^{2}-d_{\bullet j}^{2}+d_{\bullet \bullet}^{2}\right) . \tag{17.5}
\end{equation*}
$$

With $a_{i j}=-\frac{1}{2} d_{i j}^{2}$, and

$$
\begin{align*}
& a_{i \bullet}=\frac{1}{n} \sum_{j=1}^{n} a_{i j} \\
& a_{\bullet j}=\frac{1}{n} \sum_{i=1}^{n} a_{i j} \\
& a_{\bullet \bullet}=\frac{1}{n^{2}} \sum_{i=1}^{n} \sum_{j=1}^{n} a_{i j} \tag{17.6}
\end{align*}
$$

we get:

$$
\begin{equation*}
b_{i j}=a_{i j}-a_{i \bullet}-a_{\bullet j}+a_{\bullet \bullet} \tag{17.7}
\end{equation*}
$$

Define the matrix $\mathcal{A}$ as $\left(a_{i j}\right)$, and observe that:

$$
\begin{equation*}
\mathcal{B}=\mathcal{H} \mathcal{A H} \tag{17.8}
\end{equation*}
$$

The inner product matrix $\mathcal{B}$ can be expressed as:

$$
\begin{equation*}
\mathcal{B}=\mathcal{X} \mathcal{X}^{\top}, \tag{17.9}
\end{equation*}
$$

where $\mathcal{X}=\left(x_{1}, \ldots, x_{n}\right)^{\top}$ is the $(n \times p)$ matrix of coordinates. The rank of $\mathcal{B}$ is then

$$
\begin{equation*}
\operatorname{rank}(\mathcal{B})=\operatorname{rank}\left(\mathcal{X} \mathcal{X}^{\top}\right)=\operatorname{rank}(\mathcal{X})=p \tag{17.10}
\end{equation*}
$$

As required in Theorem 17.1 the matrix $\mathcal{B}$ is symmetric, positive semidefinite and of rank $p$, and hence it has $p$ non-negative eigenvalues and $n-p$ zero eigenvalues. $\mathcal{B}$ can now be written as:

$$
\begin{equation*}
\mathcal{B}=\Gamma \Lambda \Gamma^{\top} \tag{17.11}
\end{equation*}
$$

where $\Lambda=\operatorname{diag}\left(\lambda_{1}, \ldots, \lambda_{p}\right)$, the diagonal matrix of the eigenvalues of $\mathcal{B}$, and $\Gamma=\left(\gamma_{1}, \ldots, \gamma_{p}\right)$, the matrix of corresponding eigenvectors. Hence the coordinate matrix $\mathcal{X}$ containing the point configuration in $\mathbb{R}^{p}$ is given by:

$$
\begin{equation*}
\mathcal{X}=\Gamma \Lambda^{\frac{1}{2}} \tag{17.12}
\end{equation*}
$$

## How Many Dimensions?

The number of desired dimensions is small in order to provide practical interpretations, and is given by the rank of $\mathcal{B}$ or the number of nonzero eigenvalues $\lambda_{i}$. If $\mathcal{B}$ is positive semidefinite, then the number of nonzero eigenvalues gives the number of eigenvalues required for representing the distances $d_{i j}$.

The proportion of variation explained by $p$ dimensions is given by

$$
\begin{equation*}
\frac{\sum_{i=1}^{p} \lambda_{i}}{\sum_{i=1}^{n-1} \lambda_{i}} . \tag{17.13}
\end{equation*}
$$

It can be used for the choice of $p$. If $\mathcal{B}$ is not positive semidefinite we can modify (17.13) to

$$
\frac{\sum_{i=1}^{p} \lambda_{i}}{\sum(\text { ("positive eigenvalues") }}
$$

In practice the eigenvalues $\lambda_{i}$ are almost always unequal to zero. To be able to represent the objects in a space with dimensions as small as possible we may modify the distance matrix to:

$$
\begin{equation*}
\mathcal{D}^{*}=d_{i j}^{*} \tag{17.15}
\end{equation*}
$$

with

$$
d_{i j}^{*}=\left\{\begin{array}{lr}
0 & ; i=j  \tag{17.16}\\
d_{i j}+e \geq 0 ; i \neq j
\end{array}\right.
$$

where $e$ is determined such that the inner product matrix $\mathcal{B}$ becomes positive semidefinite with a small rank.

## Similarities

In some situations we do not start with distances but with similarities. The standard transformation (see Chap. 13) from a similarity matrix $\mathcal{C}$ to a distance matrix $\mathcal{D}$ is:

$$
\begin{equation*}
d_{i j}=\left(c_{i i}-2 c_{i j}+c_{j j}\right)^{\frac{1}{2}} \tag{17.17}
\end{equation*}
$$

Theorem 17.2 If $\leq 0$, then the distance matrix $\mathcal{D}$ defined by(17.17) is Euclidean with centred inner product matrix $\mathcal{B}=\mathcal{H C H}$.

## Relation to Factorial Analysis

Suppose that the $(n \times p)$ data matrix $\mathcal{X}$ is centred so that $\mathcal{X}^{\top} \mathcal{X}$ equals a multiple of the covariance matrix $n \mathcal{S}$. Suppose that the $p$ eigenvalues $\lambda_{1}, \ldots, \lambda_{p}$ of $n \mathcal{S}$ are distinct and non zero. Using the duality Theorem 10.4 of factorial analysis we see that $\lambda_{1}, \ldots, \lambda_{p}$ are also eigenvalues of $\mathcal{X} \mathcal{X}^{\top}=\mathcal{B}$ when $\mathcal{D}$ is the Euclidean distance matrix between the rows of $\mathcal{X}$. The $k$-dimensional solution to the metric MDS problem is thus given by the $k$ first principal components of $\mathcal{X}$.

## Optimality Properties of the Classical MDS Solution

Let $\mathcal{X}$ be a $(n \times p)$ data matrix with some inter-point distance matrix $\mathcal{D}$. The objective of MDS is thus to find $\mathcal{X}_{1}$, a representation of $\mathcal{X}$ in a lower dimensional Euclidean space $\mathbb{R}^{k}$ whose inter-point distance matrix $\mathcal{D}_{1}$ is not far from $\mathcal{D}$. Let $\mathcal{L}=\left(\mathcal{L}_{1}, \mathcal{L}_{2}\right)$ be a $(p \times p)$ orthogonal matrix where $\mathcal{L}_{1}$ is $(p \times k)$. $\mathcal{X}_{1}=\mathcal{X} \mathcal{L}_{1}$ represents a projection of $\mathcal{X}$ on the column space of $\mathcal{L}_{1}$; in other words, $\mathcal{X}_{1}$ may be viewed as a fitted configuration of $\mathcal{X}$ in $\mathbb{R}^{k}$. A measure of discrepancy between $\mathcal{D}$ and $\mathcal{D}_{1}=\left(d_{i j}^{(1)}\right)$ is given by

$$
\begin{equation*}
\phi=\sum_{i, j=1}^{n}\left(d_{i j}-d_{i j}^{(1)}\right)^{2} \tag{17.18}
\end{equation*}
$$

Theorem 17.3 Among all projections $\mathcal{X} \mathcal{L}_{1}$ of $\mathcal{X}$ onto $k$-dimensional subspaces of $\mathbb{R}^{p}$ the quantity $\phi$ in (17.18) is minimised when $\mathcal{X}$ is projected onto its first $k$ principal factors.

We see therefore that the metric MDS is identical to principal factor analysis as we have defined it in Chap. 10.

| $\hookrightarrow$ | Metric MDS starts with a distance matrix $\mathcal{D}$. |
| :--- | :--- |
|  | The aim of metric MDS is to construct a map in Euclidean space |
| that corresponds to the given distances. |  |


| Summary (continued) |
| :---: |
| $\hookrightarrow$ A practical algorithm is given as: <br> 1. start with distances $d_{i j}$ <br> 2. define $\mathcal{A}=-\frac{1}{2} d_{i j}^{2}$ <br> 3. put $\mathcal{B}=\left(a_{i j}-a_{i \bullet}-a_{\bullet j}+a_{\bullet \bullet}\right)$ <br> 4. find the eigenvalues $\lambda_{1}, \ldots, \lambda_{p}$ and the associated eigenvectors $\gamma_{1}, \ldots, \gamma_{p}$ where the eigenvectors are normalised so that $\gamma_{i}^{\top} \gamma_{i}=1$. <br> 5. Choose an appropriate number of dimensions $p$ (ideally $p=2$ ) <br> 6. The coordinates of the $n$ points in the Euclidean space are given by $x_{i j}=\gamma_{i j} \lambda_{j}^{1 / 2}$ for $i=1, \ldots, n$ and $j=1, \ldots, p$. |
| $\hookrightarrow$ Metric MDS is identical to principal components analysis. |

### 17.3 Nonmetric MDS

The object of nonmetric MDS, as well as of metric MDS, is to find the coordinates of the points in $p$-dimensional space, so that there is a good agreement between the observed proximities and the inter-point distances. The development of nonmetric MDS was motivated by two main weaknesses in the metric MDS (Fahrmeir \& Hamerle, 1984, p. 679):

1. the definition of an explicit functional connection between dissimilarities and distances in order to derive distances out of given dissimilarities, and
2. the restriction to Euclidean geometry in order to determine the object configurations.

The idea of a nonmetric MDS is to demand a less rigid relationship between the dissimilarities and the distances. Suppose that an unknown monotonic increasing function $f$,

$$
\begin{equation*}
d_{i j}=f\left(\delta_{i j}\right), \tag{17.19}
\end{equation*}
$$

is used to generate a set of distances $d_{i j}$ as a function of given dissimilarities $\delta_{i j}$. Here $f$ has the property that if $\delta_{i j}<\delta_{r s}$, then $f\left(\delta_{i j}\right)<f\left(\delta_{r s}\right)$. The scaling is based on the rank order of the dissimilarities. Nonmetric MDS is therefore ordinal in character.

The most common approach used to determine the elements $d_{i j}$ and to obtain the coordinates of the objects $x_{1}, x_{2}, \ldots, x_{n}$ given only rank order information is an iterative process commonly referred to as the Shepard-Kruskal algorithm.

## Shepard-Kruskal Algorithm

In a first step, called the initial phase, we calculate Euclidean distances $d_{i j}^{(0)}$ from an arbitrarily chosen initial configuration $\mathcal{X}_{0}$ in dimension $p^{*}$, provided that all objects have different coordinates. One might use metric MDS to obtain these initial coordinates. The second step or nonmetric phase determines disparities $\hat{d}_{i j}^{(0)}$ from the distances $d_{i j}^{(0)}$ by constructing a monotone regression relationship between the $d_{i j}^{(0)}$ 's and $\delta_{i j}$ 's, under the requirement that if $\delta_{i j}<\delta_{r s}$, then $\hat{d}_{i j}^{(0)} \leq \hat{d}_{r s}^{(0)}$. This is called the weak monotonicity requirement. To obtain the disparities $\hat{d}_{i j}^{(0)}$, a useful approximation method is the pool-adjacent violators (PAV) algorithm (see Fig. 17.5). Let

$$
\begin{equation*}
\left(i_{1}, j_{1}\right)>\left(i_{2}, j_{2}\right)>\cdots>\left(i_{k}, j_{k}\right) \tag{17.20}
\end{equation*}
$$

be the rank order of dissimilarities of the $k=n(n-1) / 2$ pairs of objects. This corresponds to the points in Fig. 17.6. The PAV algorithm is described as follows: "beginning with the lowest ranked value of $\delta_{i j}$, the adjacent $d_{i j}^{(0)}$ values are compared for each $\delta_{i j}$ to determine if they are monotonically related to the $\delta_{i j}$ 's. Whenever a block of consecutive values of $d_{i j}^{(0)}$ are encountered that violate the required monotonicity property the $d_{i j}^{(0)}$ values are averaged together with the most recent non-violator $d_{i j}^{(0)}$ value to obtain an estimator. Eventually this value is assigned to all points in the particular block".

Fig. 17.5 Pool-adjacent violators algorithm O . MVAMDSpooladj

Pool-Adjacent-Violator-Algorithm


Fig. 17.6 Ranks and distances 0
MVAMDSnonmstart


Table 17.3 Dissimilarities $\delta_{i j}$ for car marks

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

In a third step, called the metric phase, the spatial configuration of $\mathcal{X}_{0}$ is altered to obtain $\mathcal{X}_{1}$. From $\mathcal{X}_{1}$ the new distances $d_{i j}^{(1)}$ can be obtained which are more closely related to the disparities $\hat{d}_{i j}^{(0)}$ from step two.
Example 17.3 Consider a small example with 4 objects based on the car marks data set, see (Table 17.3). Our aim is to find a representation with $p^{*}=2$ via MDS. Suppose that we choose as an initial configuration (Fig. 17.7) of $\mathcal{X}_{0}$ the coordinates given in Table 17.4. The corresponding distances $d_{i j}=\sqrt{\left(x_{i}-x_{j}\right)^{\top}\left(x_{i}-x_{j}\right)}$ are calculated in Table 17.5

A plot of the dissimilarities of Table 17.5 against the distance yields Fig. 17.8. This relation 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.

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

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

Fig. 17.7 Initial configuration of the MDS of the car data Q MVAnmdscar1


Table 17.4 Initial coordinates for MDS

Table 17.5 Ranks and distances

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


| $i, j$ | $d_{i j}$ | $\operatorname{rank}\left(d_{i j}\right)$ | $\delta_{i j}$ |
| :--- | :--- | :--- | :--- |
| 1,2 | 5.1 | 3 | 3 |
| 1,3 | 2.2 | 1 | 2 |
| 1,4 | 7.3 | 4 | 5 |
| 2,3 | 4.1 | 2 | 1 |
| 2,4 | 8.5 | 5 | 4 |
| 3,4 | 9.1 | 6 | 6 |

Applying the same procedure to $(2,4)$ and $(1,4)$ we obtain $\hat{d}_{24}=\hat{d}_{14}=7.9$. The plot of $\delta_{i j}$ versus the disparities $\hat{d}_{i j}$ represents a monotone regression relationship.

In the initial configuration (Fig. 17.7), the third point (Ferrari) could be moved so that the distance to object 2 (Jaguar) is reduced. This procedure however also alters the distance between objects 3 and 4. Care should be given when establishing a monotone relation between $\delta_{i j}$ and $d_{i j}$.

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

$$
\begin{equation*}
\operatorname{STRESS} 1=\left(\frac{\sum_{i<j}\left(d_{i j}-\hat{d}_{i j}\right)^{2}}{\sum_{i<j} d_{i j}^{2}}\right)^{\frac{1}{2}} \tag{17.21}
\end{equation*}
$$



Fig. 17.8 Scatterplot of dissimilarities against distances $Q$ MVAnmdscar2
Table 17.6 STRESS calculations for car marks example

| $(i, j)$ | $\delta_{i j}$ | $d_{i j}$ | $\hat{d}_{i j}$ | $\left(d_{i j}-\hat{d}_{i j}\right)^{2}$ | $d_{i j}^{2}$ | $\left(d_{i j}-\bar{d}\right)^{2}$ |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| $(2,3)$ | 1 | 4.1 | 3.15 | 0.9 | 16.8 | 3.8 |
| $(1,3)$ | 2 | 2.2 | 3.15 | 0.9 | 4.8 | 14.8 |
| $(1,2)$ | 3 | 5.1 | 5.1 | 0 | 26.0 | 0.9 |
| $(2,4)$ | 4 | 8.5 | 7.9 | 0.4 | 72.3 | 6.0 |
| $(1,4)$ | 5 | 7.3 | 7.9 | 0.4 | 53.3 | 1.6 |
| $(3,4)$ | 6 | 9.1 | 9.1 | 0 | 82.8 | 9.3 |
| $\Sigma$ |  | 36.3 |  | 2.6 | 256.0 | 36.4 |

An alternative stress measure is given by

$$
\begin{equation*}
\operatorname{STRESS} 2=\left(\frac{\sum_{i<j}\left(d_{i j}-\hat{d}_{i j}\right)^{2}}{\sum_{i<j}\left(d_{i j}-\bar{d}\right)^{2}}\right)^{\frac{1}{2}} \tag{17.22}
\end{equation*}
$$

where $\bar{d}$ denotes the average distance.
Example 17.4 Table 17.6 presents the STRESS calculations for the car example.
The average distance is $\bar{d}=36.4 / 6=6.1$. The corresponding STRESS measures are:

$$
\begin{gathered}
\text { STRESS } 1=\sqrt{2.6 / 256}=0.1 \\
\text { STRESS } 2=\sqrt{2.6 / 36.4}=0.27
\end{gathered}
$$

The goal is to find a point configuration that balances the effects STRESS and non monotonicity. This is achieved by an iterative procedure. More precisely, one defines a new position of object $i$ relative to object $j$ by

$$
\begin{equation*}
x_{i l}^{\mathrm{NEW}}=x_{i l}+\alpha\left(1-\frac{\hat{d}_{i j}}{d_{i j}}\right)\left(x_{j l}-x_{i l}\right), \quad l=1, \ldots, p^{*} . \tag{17.23}
\end{equation*}
$$

Here $\alpha$ denotes the step width of the iteration.
By (17.23) 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:

$$
\begin{equation*}
x_{i l}^{\mathrm{NEW}}=x_{i l}+\frac{\alpha}{n-1} \sum_{j=1, j \neq i}^{n}\left(1-\frac{\hat{d}_{i j}}{d_{i j}}\right)\left(x_{j l}-x_{i l}\right), \quad l=1, \ldots, p^{*} . \tag{17.24}
\end{equation*}
$$

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, such as steepest descent or the Newton-Raphson procedure.

In a fourth step, the evaluation phase, the STRESS measure is used to evaluate whether or not its change as a result of the last iteration is sufficiently small that the procedure is terminated. At this stage the optimal fit has been obtained for a given dimension. Hence, the whole procedure needs to be carried out for several dimensions.

Example 17.5 Let us compute the new point configuration for $i=3$ (Ferrari) (Fig. 17.9). The initial coordinates from Table 17.4 are

$$
x_{31}=1 \text { and } x_{32}=3
$$

Applying (17.24) yields (for $\alpha=3$ ):

$$
\begin{aligned}
x_{31}^{\mathrm{NEW}} & =1+\frac{3}{4-1} \sum_{j=1, j \neq 3}^{4}\left(1-\frac{\hat{d}_{3 j}}{d_{3 j}}\right)\left(x_{j 1}-1\right) \\
& =1+\left(1-\frac{3.15}{2.2}\right)(3-1)+\left(1-\frac{3.15}{4.1}\right)(2-1)+\left(1-\frac{9.1}{9.1}\right)(10-1) \\
& =1-0.86+0.23+0 \\
& =0.37
\end{aligned}
$$

Similarly we obtain $x_{32}^{\mathrm{NEW}}=4.36$.
To find the appropriate number of dimensions, $p^{*}$, a plot of the minimum STRESS value as a function of the dimensionality is made. One possible criterion in selecting the appropriate dimensionality is to look for an elbow in the plot. A rule


Fig. 17.9 First iteration for Ferrari using Shepard-Kruskal algorithm $Q_{\text {MVAnmdscar } 3}$
of thumb that can be used to decide if a STRESS value is sufficiently small or not is provided by Kruskal:
$S>20 \%$, poor; $S=10 \%$, fair; $S<5 \%$, good; $S=0$, perfect.

| $\hookrightarrow$ | Nonmetric MDS is only based on the rank order of dissimilarities. |
| :--- | :--- |
| $\hookrightarrow$ | The object of nonmetric MDS is to create a spatial representation |
| of the objects with low dimensionality. |  |$\quad$| A practical algorithm is given as: |
| :--- |
| 1. Choose an initial configuration. <br> 2. Find $d_{i j}$ from the configuration. <br> 3. Fit $\hat{d}_{i j}$, the disparities, by the PAV algorithm. <br> 4. Find a new configuration $\mathcal{X}_{n+1}$ by using the steepest descent. <br> 5. Go to 2 . |

### 17.4 Exercises

Exercise 17.1 Apply the MDS method to the Swiss bank note data. What do you expect to see?

Exercise 17.2 Using (17.6), show that (17.7) can be written in the form (17.2).
Exercise 17.3 Show that

1. $b_{i i}=a_{\bullet \bullet}-2 a_{i \bullet} ; b_{i j}=a_{i j}-a_{i \bullet}-a_{\bullet j}+a_{\bullet \bullet} ; i \neq j$
2. $\mathcal{B}=\sum_{i=1}^{p} x_{i} x_{i}^{\top}$
3. $\sum_{i=1}^{n} \lambda_{i}=\sum_{i=1}^{n} b_{i i}=\frac{1}{2 n \sum_{i, j=1}^{n} d_{i j}}$.

Exercise 17.4 Redo a careful analysis of the car marks data based on the following dissimilarity matrix:

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

Exercise 17.5 Apply the MDS method to the US health data. Is the result in accordance with the geographic location of the US states?

Exercise 17.6 Redo Exercise 17.5 with the US crime data.
Exercise 17.7 Perform the MDS analysis on the Athletic Records data in Sect. 22.18. Can you see which countries are "close to each other"?

## Chapter 18 <br> Conjoint Measurement Analysis

Conjoint Measurement Analysis plays an important role in marketing. In the design of new products it is valuable to know which components carry what kind of utility for the customer. Marketing and advertisement strategies are based on the perception of the new product's overall utility. It can be valuable information for a car producer to know whether a change in sportiness or a change in safety or comfort equipment is perceived as a higher increase in overall utility. The Conjoint Measurement Analysis is a method for attributing utilities to the components (part worths) on the basis of ranks given to different outcomes (stimuli) of the product. An important assumption is that the overall utility is decomposed as a sum of the utilities of the components.

In Sect. 18.1 we introduce the idea of Conjoint Measurement Analysis. We give two examples from the food and car industries. In Sect. 18.2 we shed light on the problem of designing questionnaires for ranking different product outcomes. In Sect. 18.3 we see that the metric solution of estimating the part-worths is given by solving a least squares problem. The estimated preference ordering may be nonmonotone. The nonmetric solution strategy takes care of this inconsistency by iterating between a least squares solution and the pool adjacent violators algorithm.

### 18.1 Introduction

In the design and perception of new products it is important to specify the contributions made by different facets or elements. The overall utility and acceptance of such a new product can then be estimated and understood as a possibly additive function of the elementary utilities. Examples are the design of cars, a food article or the program of a political party. For a new type of margarine one may ask
whether a change in taste or presentation will enhance the overall perception of the product. The elementary utilities are here the presentation style and the taste (e.g. calory content). For a party program one may want to investigate whether a stronger ecological or a stronger social orientation gives a better overall profile of the party. For the marketing of a new car one may be interested in whether this new car should have a stronger active safety or comfort equipment or a more sporty note or combinations of both.

In Conjoint Measurement Analysis one assumes that the overall utility can be explained as an additive decomposition of the utilities of different elements. In a sample of questionnaires people ranked the product types and thus revealed their preference orderings. The aim is to find the decomposition of the overall utility on the basis of observed data and to interpret the elementary or marginal utilities.

Example 18.1 A car producer plans to introduce a new car with features that appeal to the customer and that may help in promoting future sales. The new elements that are considered are comfort/safety components (e.g. active steering or GPS) and a sporty look (leather steering wheel and additional kW of the engine). The car producer has thus four lines of cars.

| car 1: | basic safety equipment | and | low sportiness |
| :--- | :--- | :--- | :--- |
| car 2: | basic safety equipment | and | high sportiness |
| car 3: | high safety equipment | and | low sportiness |
| car 4: | high safety equipment | and | high sportiness |

For the car producer it is important to rank these cars and to find out customers' attitudes toward a certain product line in order to develop a suitable marketing scheme. A tester may rank the cars as described in Table 18.1.

The elementary utilities here are the comfort equipment and the level of sportiness. Conjoint Measurement Analysis aims at explaining the rank order given by the test person as a function of these elementary utilities.

Example 18.2 A food producer plans to create a new margarine and varies the product characteristics "calories" (low vs. high) and "presentation" (a plastic pot vs. paper package) (Backhaus, Erichson, Plinke, \& Weiber, 1996). We can view this in fact as ranking four products.

| product 1: | low calories | and | plastic pot |
| :--- | :--- | :--- | :--- |
| product 2: | low calories | and | paper package |
| product 3: | high calories | and | plastic pot |
| product 4: | high calories | and | paper package |

Table 18.1 Tester's ranking of cars

| Car | 1 | 2 | 3 | 4 |
| :--- | :--- | :--- | :--- | :--- |
| Ranking | 1 | 2 | 4 | 3 |

Table 18.2 Tester's ranking of margarine

| Product | 1 | 2 | 3 | 4 |
| :--- | :--- | :--- | :--- | :--- |
| Ranking | 3 | 4 | 1 | 2 |

These four fictive products may now be ordered by a set of sample testers as described in Table 18.2.

The Conjoint Measurement Analysis aims to explain such a preference ranking by attributing part-worths to the different elements of the product. The part-worths are the utilities of the elementary components of the product.

In interpreting the part-worths one may find that for a test person one of the elements has a higher value or utility. This may lead to a new design or to the decision that this utility should be emphasised in advertisement schemes.

| $\hookrightarrow$ | Conjoint Measurement Analysis is used in the design of new <br> products. |
| :--- | :--- |
| $\hookrightarrow$ | Conjoint Measurement Analysis tries to identify part-worth utilities <br> that contribute to an overall utility. |
|  | The part-worths enter additively into an overall utility. |
|  | The interpretation of the part-worths gives insight into the percep- |
|  | Tion and acceptance of the product. |

### 18.2 Design of Data Generation

The product is defined through the properties of the components. A stimulus is defined as a combination of the different components. Examples 18.1 and 18.2 had four stimuli each. In the margarine example they were the possible combinations of the factors $X_{1}$ (calories) and $X_{2}$ (presentation). If a product property such as

$$
X_{3}(\text { usage })=\left\{\begin{array}{l}
1 \text { bread } \\
2 \text { cooking } \\
3 \text { universal }
\end{array}\right.
$$

is added, then there are $3 \cdot 2 \cdot 2=12$ stimuli.

For the automobile Example 18.1 additional characteristics may be engine power and the number of doors. Suppose that the engines offered for the new car have $50,70,90 \mathrm{~kW}$ and that the car may be produced in $2-, 4$-, or 5 -door versions. These categories may be coded as

$$
X_{3}(\text { power of engine })=\left\{\begin{array}{l}
150 \mathrm{~kW} \\
270 \mathrm{~kW} \\
390 \mathrm{~kW}
\end{array}\right.
$$

and

$$
X_{4}(\text { doors })=\left\{\begin{array}{l}
12 \text { doors } \\
24 \text { doors } \\
35 \text { doors }
\end{array}\right.
$$

Both $X_{3}$ and $X_{4}$ have three factor levels each, whereas the first two factors $X_{1}$ (safety) and $X_{2}$ (sportiness) have only two levels. Altogether $2 \cdot 2 \cdot 3 \cdot 3=36$ stimuli are possible. In a questionnaire a tester would have to rank all 36 different products.

The profile method asks for the utility of each stimulus. This may be time consuming and tiring for a test person if there are too many factors and factor levels. Suppose that there are six properties of components with three levels each. This results in $3^{6}=729$ stimuli (i.e. 729 different products) that a tester would have to rank.

The two factor method is a simplification and considers only two factors simultaneously. It is also called trade-off analysis. The idea is to present just two stimuli at a time and then to recombine the information. Trade-off analysis is performed by defining the trade-off matrices corresponding to stimuli of two factors only.

The trade-off matrices for the levels $X_{1}, X_{2}$ and $X_{3}$ from the margarine Example 18.2 are given in Table 18.3. The trade-off matrices for the new car outfit are described in Tabel 18.4.

The choice between the profile method and the trade-off analysis should be guided by consideration of the following aspects:

1. requirements on the test person,
2. time consumption,
3. product perception.

Table 18.3 Trade-off matrices for margarine

| $X_{3}$ | $X_{1}$ |  | $X_{3}$ | $X_{2}$ |  |  | $X_{2}$ |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 1 | 1 | 2 | 1 | 1 | 2 | $X_{1}$ |  |  |
| 2 | 1 | 2 | 2 | 1 | 2 | 1 | 1 | 2 |
| 3 | 1 | 2 | 3 | 1 | 2 | 2 | 1 | 2 |

Table 18.4 Trade-off matrices for car design

| $X_{4}$ | $X_{3}$ |  |  | $X_{4}$ | $X_{2}$ |  | $X_{4}$ | $X_{1}$ |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 1 | 1 | 2 | 3 | 1 | 1 | 2 | 1 | 1 | 2 |
| 2 | 1 | 2 | 3 | 2 | 1 | 2 | 2 | 1 | 2 |
| 3 | 1 | 2 | 3 | 3 | 1 | 2 | 3 | 1 | 2 |
| $X_{3}{ }^{\prime} X_{2}$ |  |  |  | $X_{3}$ | $X_{1}$ |  |  |  |  |
|  | 1 | 1 | 2 | 1 | 1 | 2 | $X_{2}$ | $X$ |  |
|  | 2 | 1 | 2 | 2 | 1 | 2 | 1 | 1 | 2 |
|  | 3 | 1 | 2 | 3 | 1 | 2 | 2 | 1 | 2 |

The first aspect relates to the ability of the test person to judge the different stimuli. It is certainly an advantage of the trade-off analysis that one only has to consider two factors simultaneously. The two factor method can be carried out more easily in a questionnaire without an interview.

The profile method incorporates the possibility of a complete product perception since the test person is not confronted with an isolated aspect ( 2 factors) of the product. The stimuli may be presented visually in its final form (e.g. as a picture). With the number of levels and properties the number of stimuli rise exponentially with the profile method. The time to complete a questionnaire is therefore a factor in the choice of method.

In general the product perception is the most important aspect and is therefore the profile method that is used the most. The time consumption aspect speaks for the trade-off analysis. There exist, however, clever strategies on selecting representation subsets of all profiles that bound the time investment. We therefore concentrate on the profile method in the following.

|  | A stimulus is a combination of different properties of a product. |
| :--- | :--- |
|  | Summary |
|  | Conjoint measurement analysis is based either on a list of all factors <br> (profile method) or on trade-off matrices (two factor method). |
|  | Trade-off matrices are used if there are too many factor levels. |

### 18.3 Estimation of Preference Orderings

On the basis of the reported preference values for each stimulus conjoint analysis determines the part-worths. Conjoint analysis uses an additive model of the form

$$
\begin{equation*}
Y_{k}=\sum_{j=1}^{J} \sum_{l=1}^{L_{j}} \beta_{j l} \boldsymbol{I}\left(X_{j}=x_{j l}\right)+\mu, \text { for } k=1, \ldots, K \text { and } \forall j \sum_{l=1}^{L_{j}} \beta_{j l}=0 \tag{18.1}
\end{equation*}
$$

$X_{j}(j=1, \ldots, J)$ denote the factors, $x_{j l}\left(l=1, \ldots, L_{j}\right)$ are the levels of each factor $X_{j}$ and the coefficients $\beta_{j l}$ are the part-worths. The constant $\mu$ denotes an overall level and $Y_{k}$ is the observed preference for each stimulus and the total number of stimuli are:

$$
K=\prod_{j=1}^{J} L_{j} .
$$

Equation (18.1) is without an error term for the moment. In order to explain how (18.1) may be written in the standard linear model form we first concentrate on $J=2$ factors. Suppose that the factors engine power and airbag safety equipment have been ranked as follows:

|  |  | Airbag |  |  |
| :--- | :--- | :--- | :--- | :--- |
|  |  | 1 | 2 |  |
| Engine | 50 kW | 1 | 1 | 3 |
|  | 70 kW | 2 | 2 | 6 |
|  | 90 kW | 3 | 4 | 5 |

There are $K=6$ preferences altogether. Suppose that the stimuli have been sorted so that $Y_{1}$ corresponds to engine level 1 and airbag level $1, Y_{2}$ corresponds to engine level 1 and airbag level 2, and so on. Then model (18.1) reads:

$$
\begin{aligned}
& Y_{1}=\beta_{11}+\beta_{21}+\mu \\
& Y_{2}=\beta_{11}+\beta_{22}+\mu \\
& Y_{3}=\beta_{12}+\beta_{21}+\mu \\
& Y_{4}=\beta_{12}+\beta_{22}+\mu \\
& Y_{5}=\beta_{13}+\beta_{21}+\mu \\
& Y_{6}=\beta_{13}+\beta_{22}+\mu .
\end{aligned}
$$

Now we would like to estimate the part-worths $\beta_{j l}$.

Table 18.5 Ranked products

|  |  | $X_{2}$ (calories) |  |
| :--- | :--- | :--- | :--- |
|  |  | Low | High |
|  |  | 1 | 2 |
| $X_{1}$ (usage) | Bread | 1 | 2 |
|  |  |  |  |
|  | Cooking | 2 | 3 |
| 4 |  |  |  |
|  | Universal | 3 | 6 |

Example 18.3 In the margarine example let us consider the part-worths of $X_{1}=$ usage and $X_{2}=$ calories. We have $x_{11}=1, x_{12}=2, x_{13}=3, x_{21}=1$ and $x_{22}=2$. (We momentarily re-labeled the factors: $X_{3}$ became $X_{1}$.) Hence $L_{1}=3$ and $L_{2}=2$. Suppose that a person has ranked the six different products as in Table 18.5.

If we order the stimuli as follows:

$$
\begin{aligned}
& Y_{1}=\operatorname{Utility}\left(X_{1}=1 \wedge X_{2}=1\right) \\
& Y_{2}=\operatorname{Utility}\left(X_{1}=1 \wedge X_{2}=2\right) \\
& Y_{3}=\operatorname{Utility}\left(X_{1}=2 \wedge X_{2}=1\right) \\
& Y_{4}=\operatorname{Utility}\left(X_{1}=2 \wedge X_{2}=2\right) \\
& Y_{5}=\operatorname{Utility}\left(X_{1}=3 \wedge X_{2}=1\right) \\
& Y_{6}=\operatorname{Utility}\left(X_{1}=3 \wedge X_{2}=2\right),
\end{aligned}
$$

we obtain from Eq. (18.1) the same decomposition as above:

$$
\begin{aligned}
& Y_{1}=\beta_{11}+\beta_{21}+\mu \\
& Y_{2}=\beta_{11}+\beta_{22}+\mu \\
& Y_{3}=\beta_{12}+\beta_{21}+\mu \\
& Y_{4}=\beta_{12}+\beta_{22}+\mu \\
& Y_{5}=\beta_{13}+\beta_{21}+\mu \\
& Y_{6}=\beta_{13}+\beta_{22}+\mu .
\end{aligned}
$$

Our aim is to estimate the part-worths $\beta_{j l}$ as well as possible from a collection of tables like Table 18.5 that have been generated by a sample of test persons. First, the so-called metric solution to this problem is discussed and then a non-metric solution.

## Metric Solution

The problem of conjoint measurement analysis can be solved by the technique of Analysis of Variance (ANOVA). An important assumption underlying this technique is that the "distance" between any two adjacent preference orderings corresponds to the same difference in utility. That is, the difference in utility between the products ranked 1st and 2 nd is the same as the difference in utility between the products

Table 18.6 Metric solution for car example

|  |  | $X_{2}$ (airbags) |  |  |  |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
|  |  | 1 | 2 | $\bar{p}_{x_{1} \bullet}$ | $\beta_{1 l}$ |  |
| $X_{1}$ (engine) | 50 kW | 1 | 1 | 3 | 2 | -1.5 |
|  | 70 kW | 2 | 2 | 6 | 4 | 0.5 |
|  | 90 kW | 3 | 4 | 5 | 4.5 | 1 |
| $\bar{p}_{x_{2} \bullet}$ | 2.33 | 4.66 | 3.5 |  |  |  |
| $\beta_{2 l}$ |  | -1.16 | 1.16 |  |  |  |

ranked 4th and 5th. Put differently, we treat the ranking of the products-which is a cardinal variable-as if it were a metric variable.

Introducing a mean utility $\mu$ Eq. (18.1) can be rewritten. The mean utility in the above Example 18.3 is $\mu=(1+2+3+4+5+6) / 6=21 / 6=3.5$. In order to check the deviations of the utilities from this mean, we enlarge Table 18.5 by the mean utility $\bar{p}_{x_{j} \bullet}$, given a certain level of the other factor. The metric solution for the car example is given in Table 18.6.

Example 18.4 In the margarine example the resulting part-worths for $\mu=3.5$ are

$$
\begin{aligned}
& \beta_{11}=-2 \beta_{21}=0.16 \\
& \beta_{12}=0 \beta_{22}=-0.16 \\
& \beta_{13}=2
\end{aligned}
$$

Note that $\sum_{l=1}^{L_{j}} \beta_{j l}=0(j=1, \ldots, J)$. The estimated utility $\hat{Y}_{1}$ for the product with low calories and usage of bread, for example, is:

$$
\hat{Y}_{1}=\beta_{11}+\beta_{21}+\mu=-2+0.16+3.5=1.66
$$

The estimated utility $\hat{Y}_{4}$ for product 4 (cooking $\left(X_{1}=2\right)$ and high calories $\left(X_{2}=\right.$ $2)$ ) is:

$$
\hat{Y}_{4}=\beta_{12}+\beta_{22}+\mu=0-0.16+3.5=3.33
$$

The coefficients $\beta_{j l}$ are computed as $\bar{p}_{x_{j l}}-\mu$, where $\bar{p}_{x_{j l}}$ is the average preference ordering for each factor level. For instance, $\bar{p}_{x_{11}}=1 / 2 *(2+1)=1.5$.

The fit can be evaluated by calculating the deviations of the fitted values to the observed preference orderings. In the rightmost column of Table 18.8 the quadratic deviations between the observed rankings (utilities) $Y_{k}$ and the estimated utilities $\hat{Y}_{k}$ are listed.

Table 18.7 Metric solution for Table 18.5

|  |  |  | $X_{2}$ (calories) |  | $\bar{p}_{x_{1} \text { 。 }}$ | $\beta_{1}$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  | $\begin{array}{\|l\|} \hline \text { Low } \\ \hline 1 \\ \hline \end{array}$ | $\begin{aligned} & \hline \text { High } \\ & \hline 2 \end{aligned}$ |  |  |
|  |  |  |  |  |  |  |
| $X_{1}$ (usage) | Bread | 1 | 2 | 1 | 1.5 | -2 |
|  | Cooking | 2 | 3 | 4 | 3.5 | 0 |
|  | Universal | 3 | 6 | 5 | 5.5 | 2 |
| $\bar{p}_{x_{2}}$ |  |  | 3.66 | 3.33 | 3.5 |  |
| $\beta_{2 l}$ |  |  | 0.16 | -0.16 |  |  |

Table 18.8 Deviations between model and data

| Stimulus | $Y_{k}$ | $\hat{Y}_{k}$ | $Y_{k}-\hat{Y}_{k}$ | $\left(Y_{k}-\hat{Y}_{k}\right)^{2}$ |
| :--- | ---: | :--- | :--- | :--- |
| 1 | 2 | 1.66 | 0.33 | 0.11 |
| 2 | 1 | 1.33 | -0.33 | 0.11 |
| 3 | 3 | 3.66 | -0.66 | 0.44 |
| 4 | 4 | 3.33 | 0.66 | 0.44 |
| 5 | 6 | 5.66 | 0.33 | 0.11 |
| 6 | 5 | 5.33 | -0.33 | 0.11 |
| $\sum$ | 21 | 21 | 0 | 1.33 |

The technique described that generated Table 18.7 is in fact the solution to a least squares problem. The conjoint measurement problem (18.1) may be rewritten as a linear regression model (with error $\varepsilon=0$ ):

$$
\begin{equation*}
Y=\mathcal{X} \beta+\varepsilon \tag{18.2}
\end{equation*}
$$

with $\mathcal{X}$ being a design matrix with dummy variables. $\mathcal{X}$ has the row dimension $K=\prod_{j=1}^{J} L_{j}$ (the number of stimuli) and the column dimension $D=\sum_{j=1}^{J} L_{j}-J$. The reason for the reduced column number is that per factor only $\left(L_{j}-1\right)$ vectors are linearly independent. Without loss of generality we may standardise the problem so that the last coefficient of each factor is omitted. The error term $\varepsilon$ is introduced since even for one person the preference orderings may not fit the model (18.1).

Example 18.5 If we rewrite the $\beta$ coefficients in the form

$$
\left(\begin{array}{c}
\beta_{1}  \tag{18.3}\\
\beta_{2} \\
\beta_{3} \\
\beta_{4}
\end{array}\right)=\left(\begin{array}{c}
\mu+\beta_{13}+\beta_{22} \\
\beta_{11}-\beta_{13} \\
\beta_{12}-\beta_{13} \\
\beta_{21}-\beta_{22}
\end{array}\right)
$$

and define the design matrix $\mathcal{X}$ as

$$
\mathcal{X}=\left(\begin{array}{c|cc|c}
1 & 1 & 0 & 1  \tag{18.4}\\
1 & 1 & 0 & 0 \\
\hline & 0 & 1 & 1 \\
1 & 0 & 1 & 0 \\
\hline & 0 & 0 & 1 \\
1 & 0 & 0 & 0
\end{array}\right),
$$

then Eq. (18.1) leads to the linear model (with error $\varepsilon=0$ ):

$$
\begin{equation*}
Y=\mathcal{X} \beta+\varepsilon \tag{18.5}
\end{equation*}
$$

The least squares solution to this problem is the technique used for Table 18.7.
In practice we have more than one person to answer the utility rank question for the different factor levels. The design matrix is then obtained by stacking the above design matrix $n$ times. Hence, for $n$ persons we have as a final design matrix:

$$
\left.\mathcal{X}^{*}=1_{n} \otimes \mathcal{X}=\left(\begin{array}{c}
\mathcal{X} \\
\vdots \\
\vdots \\
\mathcal{X}
\end{array}\right)\right\} n-\text { times }
$$

which has dimension $(n K)(L-J)$ (where $\left.L=\sum_{j=1}^{J} L_{j}\right)$ and $Y^{*}=\left(Y_{1}^{\top}, \ldots, Y_{n}^{\top}\right)^{\top}$. The linear model (18.5) can now be written as:

$$
\begin{equation*}
Y^{*}=\mathcal{X}^{*} \beta+\varepsilon^{*} \tag{18.6}
\end{equation*}
$$

Given that the test people assign different rankings, the error term $\varepsilon^{*}$ is a necessary part of the model.

Example 18.6 If we take the $\beta$ vector as defined in (18.3) and the design matrix $\mathcal{X}$ from (18.4), we obtain the coefficients:

$$
\begin{align*}
& \hat{\beta}_{1}=5.33=\hat{\mu}+\hat{\beta}_{13}+\hat{\beta}_{22} \\
& \hat{\beta}_{2}=-4=\hat{\beta}_{11}-\hat{\beta}_{13} \\
& \hat{\beta}_{3}=-2=\hat{\beta}_{12}-\hat{\beta}_{13}  \tag{18.7}\\
& \hat{\beta}_{4}=0.33=\hat{\beta}_{21}-\hat{\beta}_{22} \\
& \sum_{l=1}^{L_{j}} \hat{\beta}_{j l}=0 .
\end{align*}
$$

Solving (18.7) we have:

$$
\begin{array}{ll}
\hat{\beta}_{11}=\hat{\beta}_{2}-\frac{1}{3}\left(\hat{\beta}_{2}+\hat{\beta}_{3}\right) & =-2 \\
\hat{\beta}_{12}=\hat{\beta}_{3}-\frac{1}{3}\left(\hat{\beta}_{2}+\hat{\beta}_{3}\right) & =0 \\
\hat{\beta}_{13}=-\frac{1}{3}\left(\hat{\beta}_{2}+\hat{\beta}_{3}\right) & =2  \tag{18.8}\\
\hat{\beta}_{21}=\hat{\beta}_{4}-\frac{1}{2} \hat{\beta}_{4}=\frac{1}{2} \hat{\beta}_{4} & =0.16 \\
\hat{\beta}_{31}=-\frac{1}{2} \hat{\beta}_{4} & =-0.16 \\
\hat{\mu}=\hat{\beta}_{1}+\frac{1}{3}\left(\hat{\beta}_{2}+\hat{\beta}_{3}\right)+\frac{1}{2}\left(\hat{\beta}_{4}\right) & =3.5 .
\end{array}
$$

In fact, we obtain the same estimated part-worths as in Table 18.7. The stimulus $k=2$ corresponds to adding up $\beta_{11}, \beta_{22}$, and $\mu$ (see (18.3)). Adding $\hat{\beta}_{1}$ and $\hat{\beta}_{2}$ gives:

$$
\hat{Y}_{2}=5.33-4=1.33 .
$$

## Nonmetric Solution

If we drop the assumption that utilities are measured on a metric scale, we have to use (18.1) to estimate the coefficients from an adjusted set of estimated utilities. More precisely, we may use the monotone ANOVA as developed by Kruskal (1965). The procedure works as follows. First, one estimates model (18.1) with the ANOVA technique described above. Then one applies a monotone transformation $\hat{Z}=f(\hat{Y})$ to the estimated stimulus utilities. The monotone transformation $f$ is used because the fitted values $\hat{Y}_{k}$ from (18.2) of the reported preference orderings $Y_{k}$ may not be monotone. The transformation $\hat{Z}_{k}=f\left(\hat{Y}_{k}\right)$ is introduced to guarantee monotonicity of preference orderings. For the car example the reported $Y_{k}$ values were $Y=$ $(1,3,2,6,4,5)^{\top}$. The estimated values are computed as:

$$
\begin{aligned}
& \hat{Y}_{1}=-1.5-1.16+3.5=0.84 \\
& \hat{Y}_{2}=-1.5+1.16+3.5=3.16 \\
& \hat{Y}_{3}=-0.5-1.16+3.5=2.84 \\
& \hat{Y}_{4}=-0.5+1.16+3.5=5.16 \\
& \hat{Y}_{5}=1.5-1.16+3.5=3.34 \\
& \hat{Y}_{6}=1.5+1.16+3.5=5.66 .
\end{aligned}
$$

If we make a plot of the estimated preference orderings against the revealed ones, we obtain Fig. 18.1.

Fig. 18.1 Plot of estimated preference orderings vs. revealed rankings and PAV fit 0
MVAcarrankings

## Car rankings



We see that the estimated $\hat{Y}_{6}=5.16$ is below the estimated $\hat{Y}_{5}=5.66$ and thus an inconsistency in ranking the utilities occurs. The monotone transformation $\hat{Z}_{k}=f\left(\hat{Y}_{k}\right)$ is introduced to make the relationship in Fig. 18.1 monotone. A very simple procedure consists of averaging the "violators" $\hat{Y}_{6}$ and $\hat{Y}_{5}$ to obtain 5.41. The relationship is then monotone but the model (18.1) may now be violated. The idea is therefore to iterate these two steps. This procedure is iterated until the stress measure (see Chap. 17)

$$
\begin{equation*}
\text { STRESS }=\frac{\sum_{k=1}^{K}\left(\hat{Z}_{k}-\hat{Y}_{k}\right)^{2}}{\sum_{k=1}^{K}\left(\hat{Y}_{k}-\overline{\hat{Y}}\right)^{2}} \tag{18.9}
\end{equation*}
$$

is minimised over $\beta$ and the monotone transformation $f$. The monotone transformation can be computed by the so-called pool-adjacent-violators (PAV) algorithm.

| $\hookrightarrow$ | The part-worths are estimated via the least squares method. |
| :--- | :--- |
| $\hookrightarrow$ | The metric solution corresponds to analysis of variance in a linear <br> model. |
|  | The non-metric solution iterates between a monotone regression <br> curve fitting and determining the part-worths by ANOVA method- <br> ology. |
|  | The fitting of data to a monotone function is done via the PAV <br> algorithm. |

### 18.4 Exercises

Exercise 18.1 Compute the part-worths for the following table of rankings

|  |  | $X_{2}$ |  |
| :---: | :--- | :--- | :--- |
|  |  | 1 | 2 |
| $X_{1}$ | 1 | 1 | 2 |
|  | 2 | 4 | 3 |
|  | 3 | 6 | 5 |

Exercise 18.2 Consider again Example 18.5. Rewrite the design matrix $\mathcal{X}$ and the parameter vector $\beta$ so that the overall mean effect $\mu$ is part of $\mathcal{X}$ and $\beta$, i.e. find the matrix $\mathcal{X}^{\prime}$ and $\beta^{\prime}$ such that $Y=\mathcal{X}^{\prime} \beta^{\prime}$.

Exercise 18.3 Compute the design matrix for Example 18.5 for $n=3$ persons ranking the margarine with $X_{1}$ and $X_{2}$.

Exercise 18.4 Construct an analog for Table 18.8 for the car example.
Exercise 18.5 Compute the part-worths on the basis of the following tables of rankings observed on $n=3$ persons.

|  |  | $X_{2}$ |  |  | $X_{2}$ |  |  | $X_{2}$ |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| $X_{1}$ | 1 | 1 | 2 | $X_{1}$ | 1 | 3 | $X_{1}$ | 3 | 1 |
|  | 2 | 4 | 3 |  | 4 | 2 |  | 5 | 2 |
|  | 3 | 6 | 5 |  | 5 | 6 |  | 6 | 4 |

Exercise 18.6 Suppose that in the car example a person has ranked cars by the profile method on the following characteristics:

$$
\begin{aligned}
& X_{1}=\text { motor } \\
& X_{2}=\text { safety } \\
& X_{3}=\text { doors }
\end{aligned}
$$

| $X_{1}$ | $X_{2}$ | $X_{3}$ | Preference |
| :--- | :--- | :--- | :--- |
| 1 | 1 | 1 | 1 |
| 1 | 1 | 2 | 3 |
| 1 | 1 | 3 | 2 |
| 1 | 2 | 1 | 5 |
| 1 | 2 | 2 | 4 |
| 1 | 2 | 3 | 6 |


| $X_{1}$ | $X_{2}$ | $X_{3}$ | Preference |
| :--- | :--- | :--- | :---: |
| 2 | 1 | 1 | 7 |
| 2 | 1 | 2 | 8 |
| 2 | 1 | 3 | 9 |
| 2 | 2 | 1 | 10 |
| 2 | 2 | 2 | 12 |
| 2 | 2 | 3 | 11 |


| $X_{1}$ | $X_{2}$ | $X_{3}$ | Preference |
| :--- | :--- | :--- | :--- |
| 3 | 1 | 1 | 13 |
| 3 | 1 | 2 | 15 |
| 3 | 1 | 3 | 14 |
| 3 | 2 | 1 | 16 |
| 3 | 2 | 2 | 17 |
| 3 | 2 | 3 | 18 |

There are $k=18$ stimuli.
Estimate and analyse the part-worths.

## Chapter 19 <br> Applications in Finance

A portfolio is a linear combination of assets. Each asset contributes with a weight $c_{j}$ to the portfolio. The performance of such a portfolio is a function of the various returns of the assets and of the weights $c=\left(c_{1}, \ldots, c_{p}\right)^{\top}$. In this chapter we investigate the "optimal choice" of the portfolio weights $c$. The optimality criterion is the mean-variance efficiency of the portfolio. Usually investors are risk-averse, therefore, we can define a mean-variance efficient portfolio to be a portfolio that has a minimal variance for a given desired mean return. Equivalently, we could try to optimise the weights for the portfolios with maximal mean return for a given variance (risk structure). We develop this methodology in the situations of (non)existence of riskless assets and discuss relations with the capital asset pricing model (CAPM).

### 19.1 Portfolio Choice

Suppose that one has a portfolio of $p$ assets. The price of asset $j$ at time $i$ is denoted as $p_{i j}$. The return from asset $j$ in a single time period (day, month, year etc.) is:

$$
x_{i j}=\frac{p_{i j}-p_{i-1, j}}{p_{i-1, j}} .
$$

We observe the vectors $x_{i}=\left(x_{i l}, \ldots, x_{i p}\right)^{\top}$ (i.e. the returns of the assets which are contained in the portfolio) over several time periods. We stack these observations into a data matrix $\mathcal{X}=\left(x_{i j}\right)$ consisting of observations of a random variable

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

The return of the portfolio is the weighted sum of the returns of the $p$ assets:

$$
\begin{equation*}
Q=c^{\top} X \tag{19.1}
\end{equation*}
$$

where $c=\left(c_{1}, \ldots, c_{p}\right)^{\top}$ (with $\sum_{j=1}^{p} c_{j}=1$ ) denotes the proportions of the assets in the portfolio. The mean return of the portfolio is given by the expected value of $Q$, which is $c^{\top} \mu$. The risk or variance (squared volatility) of the portfolio is given by the variance of $Q$ (Theorem 4.6), which is equal to two times

$$
\begin{equation*}
\frac{1}{2} c^{\top} \Sigma c \tag{19.2}
\end{equation*}
$$

The reason for taking half of the variance of $Q$ is merely technical. The optimisation of (19.2) with respect to $c$ is of course equivalent to minimising $c^{\top} \Sigma c$. Our aim is to maximise the portfolio returns (19.1) given a bound on the volatility (19.2) or vice versa to minimise risk given a (desired) mean return of the portfolio.

| $\hookrightarrow$ | Given a matrix of returns $\mathcal{X}$ from $p$ assets in $n$ time periods, and |
| :--- | :--- |
| that the underlying distribution is stationary, i.e. $X \sim(\mu, \Sigma)$, then |  |
| the (theoretical) return of the portfolio is a weighted sum of the |  |
| returns of the $p$ assets, namely $Q=c^{\top} X$. |  |

### 19.2 Efficient Portfolio

A variance efficient portfolio is one that keeps the risk (19.2) minimal under the constraint that the weights sum to 1 , i.e. $c^{\top} 1_{p}=1$. For a variance efficient portfolio, we therefore try to find the value of $c$ that minimises the Lagrangian

$$
\begin{equation*}
\mathcal{L}=\frac{1}{2} c^{\top} \Sigma c-\lambda\left(c^{\top} 1_{p}-1\right) \tag{19.3}
\end{equation*}
$$

A mean-variance efficient portfolio is defined as one that has minimal variance among all portfolios with the same mean. More formally, we have to find a vector
of weights $c$ such that the variance of the portfolio is minimal subject to two constraints:

1. a certain, pre-specified mean return $\bar{\mu}$ has to be achieved,
2. the weights have to sum to one.

Mathematically speaking, we are dealing with an optimisation problem under two constraints.

The Lagrangian function for this problem is given by

$$
\mathcal{L}=c^{\top} \Sigma c+\lambda_{1}\left(\bar{\mu}-c^{\top} \mu\right)+\lambda_{2}\left(1-c^{\top} 1_{p}\right)
$$

With tools presented in Sect. 2.4 we can calculate the first order condition for a minimum:

$$
\begin{equation*}
\frac{\partial \mathcal{L}}{\partial c}=2 \Sigma c-\lambda_{1} \mu-\lambda_{2} 1_{p}=0 \tag{19.4}
\end{equation*}
$$

Example 19.1 Figure 19.1 shows the monthly returns from January 2000 to December 2009 of six stocks. The data is from Yahoo Finance. For each stock


Fig. 19.1 Returns of six firms from January 2000 to December 2009 Q MVAreturns
we have chosen the same scale on the vertical axis (which gives the return of the stock). Note how the return of some stocks, such as Forward Industries and Apple, are much more volatile than the returns of other stocks, such as IBM or Consolidated Edison (Electric utilities).

As a very simple example consider two differently weighted portfolios containing only two assets, IBM and Forward Industries. Figure 19.2 displays the monthly returns of the two portfolios. The portfolio in the upper panel consists of approximately $10 \%$ Forward Industries assets and $90 \%$ IBM assets. The portfolio in the lower panel contains an equal proportion of each of the assets. The text windows on the right of Fig. 19.2 show the exact weights which were used. We can clearly see that the returns of the portfolio with a higher share of the IBM assets (which have a low variance) are much less volatile.

For an exact analysis of the optimisation problem (19.4) we distinguish between two cases: the existence and nonexistence of a riskless asset. A riskless asset is an asset such as a zero bond, i.e. a financial instrument with a fixed nonrandom return (Franke, Härdle \& Hafner, 2011).

Fig. 19.2 Portfolio of IBM and Forward Industries assets, equal and efficient weights Q MVAportfolIBMFord


## Nonexistence of a Riskless Asset

Assume that the covariance matrix $\Sigma$ is invertible (which implies positive definiteness). This is equivalent to the nonexistence of a portfolio $c$ with variance $c^{\top} \Sigma c=0$. If all assets are uncorrelated, $\Sigma$ is invertible if all of the asset returns have positive variances. A riskless asset (uncorrelated with all other assets) would have zero variance since it has fixed, nonrandom returns. In this case $\Sigma$ would not be positive definite.

The optimal weights can be derived from the first order condition (19.4) as

$$
\begin{equation*}
c=\frac{1}{2} \Sigma^{-1}\left(\lambda_{1} \mu+\lambda_{2} 1_{p}\right) . \tag{19.5}
\end{equation*}
$$

Multiplying this by a $(p \times 1)$ vector $1_{p}$ of ones, we obtain

$$
1=1_{p}^{\top} c=\frac{1}{2} 1_{p}^{\top} \Sigma^{-1}\left(\lambda_{1} \mu+\lambda_{2} 1_{p}\right),
$$

which can be solved for $\lambda_{2}$ to get:

$$
\lambda_{2}=\frac{2-\lambda_{1} 1_{p}^{\top} \Sigma^{-1} \mu}{1_{p}^{\top} \Sigma^{-1} 1_{p}} .
$$

Plugging this expression into (19.5) yields

$$
\begin{equation*}
c=\frac{1}{2} \lambda_{1}\left(\Sigma^{-1} \mu-\frac{1_{p}^{\top} \Sigma^{-1} \mu}{1_{p}^{\top} \Sigma^{-1} 1_{p}} \Sigma^{-1} 1_{p}\right)+\frac{\Sigma^{-1} 1_{p}}{1_{p}^{\top} \Sigma^{-1} 1_{p}} . \tag{19.6}
\end{equation*}
$$

For the case of a variance efficient portfolio there is no restriction on the mean of the portfolio $\left(\lambda_{1}=0\right)$. The optimal weights are therefore:

$$
\begin{equation*}
c=\frac{\Sigma^{-1} 1_{p}}{1_{p}^{\top} \Sigma^{-1} 1_{p}} . \tag{19.7}
\end{equation*}
$$

This formula is identical to the solution of (19.3). Indeed, differentiation with respect to $c$ gives

$$
\begin{aligned}
\Sigma c & =\lambda 1_{p} \\
c & =\lambda \Sigma^{-1} 1_{p}
\end{aligned}
$$

If we plug this into (19.3), we obtain

$$
\begin{aligned}
\mathcal{L} & =\frac{1}{2} \lambda^{2} 1_{p}^{\top} \Sigma^{-1} 1_{p}-\lambda\left(\lambda 1_{p}^{\top} \Sigma^{-1} 1_{p}-1\right) \\
& =\lambda-\frac{1}{2} \lambda^{2} 1_{p}^{\top} \Sigma^{-1} 1_{p} .
\end{aligned}
$$

This quantity is a function of $\lambda$ and is minimal for

$$
\lambda=\left(1_{p}^{\top} \Sigma^{-1} 1_{p}\right)^{-1}
$$

since

$$
\frac{\partial^{2} \mathcal{L}}{\partial c^{\top} \partial c}=\Sigma>0
$$

Theorem 19.1 The variance efficient portfolio weights for returns $X \sim(\mu, \Sigma)$ are

$$
\begin{equation*}
c_{\mathrm{opt}}=\frac{\Sigma^{-1} 1_{p}}{1_{p}^{\top} \Sigma^{-1} 1_{p}} . \tag{19.8}
\end{equation*}
$$

## Existence of a Riskless Asset

If an asset exists with variance equal to zero, then the covariance matrix $\Sigma$ is not invertible. The notation can be adjusted for this case as follows: denote the return of the riskless asset by $r$ (under the absence of arbitrage this is the interest rate), and partition the vector and the covariance matrix of returns such that the last component is the riskless asset. Thus, the last equation of the system (19.4) becomes

$$
2 \operatorname{Cov}(r, X)-\lambda_{1} r-\lambda_{2}=0
$$

and, because the covariance of the riskless asset with any portfolio is zero, we have

$$
\begin{equation*}
\lambda_{2}=-r \lambda_{1} . \tag{19.9}
\end{equation*}
$$

Let us for a moment modify the notation in such a way that in each vector and matrix the components corresponding to the riskless asset are excluded. For example, $c$ is the weight vector of the risky assets (i.e. assets with positive variance), and $c_{0}$ denotes the proportion invested in the riskless asset. Obviously, $c_{0}=1-1_{p}^{\top} c$, and $\Sigma$ the covariance matrix of the risky assets, is assumed to be invertible. Solving (19.4) using (19.9) gives

$$
\begin{equation*}
c=\frac{\lambda_{1}}{2} \Sigma^{-1}\left(\mu-r 1_{p}\right) . \tag{19.10}
\end{equation*}
$$

This equation may be solved for $\lambda_{1}$ by plugging it into the condition $\mu^{\top} c=\bar{\mu}$. This is the mean-variance efficient weight vector of the risky assets if a riskless asset exists. The final solution is:

$$
\begin{equation*}
c=\frac{\bar{\mu} \Sigma^{-1}\left(\mu-r 1_{p}\right)}{\mu^{\top} \Sigma^{-1}\left(\mu-r 1_{p}\right)} . \tag{19.11}
\end{equation*}
$$

The variance optimal weighting of the assets in the portfolio depends on the structure of the covariance matrix as the following corollaries show.

Corollary 19.1 A portfolio of uncorrelated assets whose returns have equal variances $\left(\Sigma=\sigma^{2} \mathcal{I}_{p}\right)$ needs to be weighted equally:

$$
c_{\mathrm{opt}}=p^{-1} 1_{p}
$$

Proof Here we obtain $1_{p}^{\top} \Sigma^{-1} 1_{p}=\sigma^{-2} 1_{p}^{\top} 1_{p}=\sigma^{-2} p$ and therefore $c=\frac{\sigma^{-2} 1_{p}}{\sigma^{-2} p}=$ $p^{-1} 1_{p}$.

Corollary 19.2 A portfolio of correlated assets whose returns have equal variances, i.e.

$$
\Sigma=\sigma^{2}\left(\begin{array}{cccc}
1 & \rho & \cdots & \rho \\
\rho & 1 & \cdots & \rho \\
\vdots & \vdots & \ddots & \vdots \\
\rho & \rho & \cdots & 1
\end{array}\right), \quad-\frac{1}{p-1}<\rho<1
$$

needs to be weighted equally:

$$
c_{\mathrm{opt}}=p^{-1} 1_{p}
$$

Proof $\Sigma$ can be rewritten as $\Sigma=\sigma^{2}\left\{(1-\rho) \mathcal{I}_{p}+\rho 1_{p} 1_{p}^{\top}\right\}$. The inverse is

$$
\Sigma^{-1}=\frac{\mathcal{I}_{p}}{\sigma^{2}(1-\rho)}-\frac{\rho 1_{p} 1_{p}^{\top}}{\sigma^{2}(1-\rho)\{1+(p-1) \rho\}}
$$

since for a $(p \times p)$ matrix $\mathcal{A}$ of the form $\mathcal{A}=(a-b) \mathcal{I}_{p}+b 1_{p} 1_{p}^{\top}$ the inverse is generally given by

$$
\mathcal{A}^{-1}=\frac{\mathcal{I}_{p}}{(a-b)}-\frac{b 1_{p} 1_{p}^{\top}}{(a-b)\{a+(p-1) b\}} .
$$

Hence

$$
\begin{aligned}
\Sigma^{-1} 1_{p} & =\frac{1_{p}}{\sigma^{2}(1-\rho)}-\frac{\rho 1_{p} 1_{p}^{\top} 1_{p}}{\sigma^{2}(1-\rho)\{1+(p-1) \rho\}} \\
& =\frac{[\{1+(p-1) \rho\}-\rho p] 1_{p}}{\sigma^{2}(1-\rho)\{1+(p-1) \rho\}}=\frac{\{1-\rho\} 1_{p}}{\sigma^{2}(1-\rho)\{1+(p-1) \rho\}} \\
& =\frac{1_{p}}{\sigma^{2}\{1+(p-1) \rho\}}
\end{aligned}
$$

which yields

$$
1_{p}^{\top} \Sigma^{-1} 1_{p}=\frac{p}{\sigma^{2}\{1+(p-1) \rho\}}
$$

and thus $c=p^{-1} 1_{p}$.
Let us now consider assets with different variances. We will see that in this case the weights are adjusted to the risk.

Corollary 19.3 A portfolio of uncorrelated assets with returns of different variances, i.e. $\Sigma=\operatorname{diag}\left(\sigma_{1}^{2}, \ldots, \sigma_{p}^{2}\right)$, has the following optimal weights

$$
c_{j, \mathrm{opt}}=\frac{\sigma_{j}^{-2}}{\sum_{l=1}^{p} \sigma_{l}^{-2}}, \quad j=1, \ldots, p
$$

Proof From $\Sigma^{-1}=\operatorname{diag}\left(\sigma_{1}^{-2}, \ldots, \sigma_{p}^{-2}\right)$ we have $1_{p}^{\top} \Sigma^{-1} 1_{p}=\sum_{l=1}^{p} \sigma_{l}^{-2}$ and therefore the optimal weights are $c_{j}=\sigma_{j}^{-2} / \sum_{l=1}^{p} \sigma_{l}^{-2}$.

This result can be generalised for covariance matrices with block structures.
Corollary 19.4 A portfolio of assets with returns $X \sim(\mu, \Sigma)$, where the covariance matrix has the form:

$$
\Sigma=\left(\begin{array}{cccc}
\Sigma_{1} & 0 & \ldots & 0 \\
0 & \Sigma_{2} & \ddots & \vdots \\
\vdots & \ddots & \ddots & \vdots \\
0 & \ldots & 0 & \Sigma_{r}
\end{array}\right)
$$

has optimal weights $c=\left(c_{1}, \ldots, c_{r}\right)^{\top}$ given by

$$
c_{j, \mathrm{opt}}=\frac{\Sigma_{j}^{-1} 1}{1^{\top} \Sigma_{j}^{-1} 1}, \quad j=1, \ldots, r
$$

## med

$\hookrightarrow ~ A n ~ e f f i c i e n t ~ p o r t f o l i o ~ i s ~ o n e ~ t h a t ~ k e e p s ~ t h e ~ r i s k ~ m i n i m a l ~ u n d e r ~ t h e ~$ constraint that a given mean return is achieved and that the weights sum to 1 , i.e. that minimises $\mathcal{L}=c^{\top} \Sigma c+\lambda_{1}\left(\bar{\mu}-c^{\top} \mu\right)+\lambda_{2}(1-$ $c^{\top} 1_{p}$ ).
$\hookrightarrow$ If a riskless asset does not exist, the variance efficient portfolio weights are given by

$$
c=\frac{\Sigma^{-1} 1_{p}}{1_{p}^{\top} \Sigma^{-1} 1_{p}} .
$$

$\hookrightarrow$ If a riskless asset exists, the mean-variance efficient portfolio weights are given by

$$
c=\frac{\bar{\mu} \Sigma^{-1}\left(\mu-r 1_{p}\right)}{\mu^{\top} \Sigma^{-1}\left(\mu-r 1_{p}\right)} .
$$

$\hookrightarrow$ The efficient weighting depends on the structure of the covariance matrix $\Sigma$. Equal variances of the assets in the portfolio lead to equal weights, different variances lead to weightings proportional to these variances:

$$
c_{j, \mathrm{opt}}=\frac{\sigma_{j}^{-2}}{\sum_{l=1}^{p} \sigma_{l}^{-2}}, \quad j=1, \ldots, p
$$

### 19.3 Efficient Portfolios in Practice

We can now demonstrate the usefulness of this technique by applying our method to the monthly market returns computed on the basis of transactions at the New York stock market and the NASDAQ stock market between January 2000 and December 2009.

Example 19.2 Recall that we had shown the portfolio returns with uniform and optimal weights in Fig. 19.2. The covariance matrix of the returns of IBM and Forward Industries is

$$
\mathcal{S}=\left(\begin{array}{ll}
0.0073 & 0.0023 \\
0.0023 & 0.0454
\end{array}\right)
$$

Hence by (19.7) the optimal weighting is

$$
\hat{c}=\frac{\mathcal{S}^{-1} 1_{2}}{1_{2}^{\top} \mathcal{S}^{-1} 1_{2}}=(0.8952,0.1048)^{\top} .
$$

The effect of efficient weighting becomes even clearer when we expand the portfolio to six assets. The covariance matrix for the returns of all six firms introduced in Example 19.1 is

$$
\mathcal{S}=10^{-3}\left(\begin{array}{rrrrrr}
7.3 & 6.2 & 3.1 & 2.3 & -0.1 & 5.2 \\
6.2 & 23.9 & 4.3 & 2.1 & 0.4 & 6.4 \\
3.1 & 4.3 & 19.5 & -0.9 & 1.1 & 3.7 \\
2.3 & 2.1 & -0.9 & 45.4 & -2.1 & 0.8 \\
-0.1 & 0.4 & 1.1 & -2.1 & 2.4 & -0.1 \\
5.2 & 6.4 & 3.7 & 0.8 & -0.1 & 14.7
\end{array}\right) \text {. }
$$

Hence the optimal weighting is

$$
\hat{c}=\frac{\mathcal{S}^{-1} 1_{6}}{1_{6}^{\top} \mathcal{S}^{-1} 1_{6}}=(0.1894,-0.0139,0.0094,0.0580,0.7112,0.0458)^{\top}
$$

As we can clearly see, the optimal weights are quite different from the equal weights $\left(c_{j}=1 / 6\right)$. The weights which were used are shown in text windows on the right hand side of Fig. 19.3.

This efficient weighting assumes stable covariances between the assets over time. Changing covariance structure over time implies weights that depend on time as well. This is part of a large body of literature on multivariate volatility models. For a review refer to Franke et al. (2011).

Fig. 19.3 Portfolio of all six
Weights assets, equal and efficient weights Q MVAportfol


Weights


## ${ }^{19+4} 0$

## Summary

| $\hookrightarrow$ | Efficient portfolio weighting in practice consists of estimating the <br> covariances of the assets in the portfolio and then computing <br> efficient weights from this empirical covariance matrix. |
| :---: | :--- |
| $\hookrightarrow$ | Note that this efficient weighting assumes stable covariances <br> between the assets over time. |

### 19.4 The Capital Asset Pricing Model

The CAPM considers the relation between a mean-variance efficient portfolio and an asset uncorrelated with this portfolio. Let us denote this specific asset return by $y_{0}$. The riskless asset with constant return $y_{0}=r$ may be such an asset. Recall from (19.4) the condition for a mean-variance efficient portfolio:

$$
2 \Sigma c-\lambda_{1} \mu-\lambda_{2} 1_{p}=0 .
$$

In order to eliminate $\lambda_{2}$, we can multiply (19.4) by $c^{\top}$ to get:

$$
2 c^{\top} \Sigma c-\lambda_{1} \bar{\mu}=\lambda_{2} .
$$

Plugging this into (19.4), we obtain:

$$
\begin{align*}
2 \Sigma c-\lambda_{1} \mu & =2 c^{\top} \Sigma c 1_{p}-\lambda_{1} \bar{\mu} 1_{p} \\
\mu & =\bar{\mu} 1_{p}+\frac{2}{\lambda_{1}}\left(\Sigma c-c^{\top} \Sigma c 1_{p}\right) \tag{19.12}
\end{align*}
$$

For the asset that is uncorrelated with the portfolio, Eq. (19.12) can be written as:

$$
y_{0}=\bar{\mu}-\frac{2}{\lambda_{1}} c^{\top} \Sigma c
$$

since $y_{0}=r$ is the mean return of this asset and is otherwise uncorrelated with the risky assets. This yields:

$$
\begin{equation*}
\lambda_{1}=2 \frac{c^{\top} \Sigma c}{\bar{\mu}-y_{0}} \tag{19.13}
\end{equation*}
$$

and if (19.13) is plugged into (19.12):

$$
\begin{align*}
\boldsymbol{\mu} & =\bar{\mu} 1_{p}+\frac{\bar{\mu}-y_{0}}{c^{\top} \Sigma c}\left(\Sigma c-c^{\top} \Sigma c 1_{p}\right) \\
\boldsymbol{\mu} & =y_{0} 1_{p}+\frac{\Sigma c}{c^{\top} \Sigma c}\left(\bar{\mu}-y_{0}\right) \\
\boldsymbol{\mu} & =y_{0} 1_{p}+\boldsymbol{\beta}\left(\bar{\mu}-y_{0}\right) \tag{19.14}
\end{align*}
$$

with

$$
\beta \stackrel{\text { def }}{=} \frac{\Sigma c}{c^{\top} \Sigma c} .
$$

The relation (19.14) holds if there exists any asset that is uncorrelated with the mean-variance efficient portfolio $c$. The existence of a riskless asset is not a necessary condition for deriving (19.14). However, for this special case we arrive at the well-known expression

$$
\begin{equation*}
\boldsymbol{\mu}=r 1_{p}+\boldsymbol{\beta}(\bar{\mu}-r), \tag{19.15}
\end{equation*}
$$

which is known as the CAPM, see Franke et al. (2011). The beta factor $\beta$ measures the relative performance with respect to riskless assets or an index. It reflects the sensitivity of an asset with respect to the whole market. The beta factor is close to

1 for most assets. A factor of 1.16, for example, means that the asset reacts in relation to movements of the whole market (expressed through an index like DAX or DOW JONES) $16 \%$ stronger than the index. This is of course true for both positive and negative fluctuations of the whole market.


## Summary

$\hookrightarrow$ The weights of the mean-variance efficient portfolio satisfy $2 \Sigma c-$ $\lambda_{1} \mu-\lambda_{2} 1_{p}=0$.
$\hookrightarrow$ In the CAPM the mean of $X$ depends on the riskless asset and the pre-specified mean $\bar{\mu}$ as follows $\mu=r 1_{p}+\beta(\bar{\mu}-r)$.
$\hookrightarrow$ The beta factor $\beta$ measures the relative performance with respect to riskless assets or an index and reflects the sensitivity of an asset with respect to the whole market.

### 19.5 Exercises

Exercise 19.1 Prove that the inverse of $\mathcal{A}=(a-b) \mathcal{I}_{p}+b 1_{p} 1_{p}^{\top}$ is given by

$$
\mathcal{A}^{-1}=\frac{\mathcal{I}_{p}}{(a-b)}-\frac{b 1_{p} 1_{p}^{\top}}{(a-b)\{a+(p-1) b\}} .
$$

Exercise 19.2 The empirical covariance between the 120 returns of IBM and Forward Industries is 0.0023 (see Example 19.2). Test if the true covariance is zero. Hint: Use Fisher's Z-transform.

Exercise 19.3 Explain why in both Figs. 19.2 and 19.3 the portfolios have negative returns just before the end of the series, regardless of whether they are optimally weighted or not! (What happened in the mid 2007?)

Exercise 19.4 Apply the method used in Example 19.2 on the same data (Table 22.5) including also the Digital Equipment company. Obviously one of the weights is negative. Is this an efficient weighting?

Exercise 19.5 In the CAPM the $\beta$ value tells us about the performance of the portfolio relative to the riskless asset. Calculate the $\beta$ value for each single stock price series relative to the "riskless" asset IBM.

## Chapter 20 <br> Computationally Intensive Techniques

It is generally accepted that training in statistics must include some exposure to the mechanics of computational statistics. This exposure to computational methods is of an essential nature when we consider extremely high-dimensional data. Computeraided techniques can help us to discover dependencies in high dimensions without complicated mathematical tools. A draftman's plot (i.e. a matrix of pairwise scatterplots like in Fig. 1.14) may lead us immediately to a theoretical hypothesis (on a lower dimensional space) on the relationship of the variables. Computer-aided techniques are therefore at the heart of multivariate statistical analysis.

With the rapidly increasing amount of data statistics faces a new challenge. While in the twentieth century the focus was on the mathematical precision of statistical modelling, the twenty-first century relies more and more on data analytic procedures that provide information (even for extremely large data bases) on the fingertip. This demand on fast availability of condensed statistical information has changed the statistical paradigm and has shifted energy from mathematical analysis to computational analysis of course without loosing sight of the statistical core questions.

In this chapter we first present the concept of Simplicial Depth-a multivariate extension of the data depth concept of Sect.1.1. We then present Projection Pursuit-a semiparametric technique which is based on a one-dimensional, flexible regression or on the idea of density smoothing applied to principal component analysis (PCA) type projections. A similar model is underlying the Sliced Inverse Regression (SIR) technique which we discuss in Sect. 20.3.

The next technique is called support vector machines (SVMs) and is motivated by non-linear classification (discrimination) problems. SVMs are classification methods based on statistical learning theory. A quadratic optimisation problem determines so-called support vectors with high margin that guarantee maximal separability. Non-linear classification is achieved by mapping the data into a feature space and finding a linear separating hyperplane in this feature space. Another
advanced technique is CART-Classification and Regression Trees, a decision tree procedure developed by Breiman, Friedman, Olshen, and Stone (1984).

### 20.1 Simplicial Depth

Simplicial depth generalises the notion of data depth as introduced in Sect.1.1. This general definition allows us to define a multivariate median and to visually present high-dimensional data in low dimension. For univariate data we have well known parameters of location which describe the centre of a distribution of a random variable $X$. These parameters are for example the mean

$$
\begin{equation*}
\bar{x}=\frac{1}{n} \sum_{i=1}^{n} x_{i}, \tag{20.1}
\end{equation*}
$$

or the mode

$$
x_{\mathrm{mod}}=\arg \max _{x} \hat{f}(x),
$$

where $\hat{f}$ is the estimated density function of $X$ (see Sect. 1.3). The median

$$
x_{\mathrm{med}}= \begin{cases}x_{\{(n+1) / 2\}} & \text { if } n \text { odd } \\ \frac{x_{(n / 2)}+x_{(n / 2+1)}}{2} & \text { otherwise }\end{cases}
$$

where $x_{(i)}$ is the order statistics of the $n$ observations $x_{i}$, is yet another measure of location.

The first two parameters can be easily extended to multivariate random variables. The mean in higher dimensions is defined as in (20.1) and the mode accordingly,

$$
x_{\mathrm{mod}}=\arg \max _{x} \hat{f}(x)
$$

with $\hat{f}$ the estimated multidimensional density function of $X$ (see Sect. 1.3). The median poses a problem though since in a multivariate sense we cannot interpret the element-wise median

$$
x_{\mathrm{med}, j}= \begin{cases}x_{\{(n+1) / 2\}, j} & \text { if } n \text { odd }  \tag{20.2}\\ \frac{x_{(n / 2), j}+x_{(n / 2+1), j}}{2} & \text { otherwise }\end{cases}
$$

as a point that is "most central". The same argument applies to other observations of a sample that have a certain "depth" as defined in Sect. 1.1. The "fourths" or the
"extremes" are not defined in a straightforward way in higher (not even for two) dimensions.

An equivalent definition of the median in one dimension is given by the simplicial depth. It is defined as follows: For each pair of datapoints $x_{i}$ and $x_{j}$ we generate a closed interval, a one-dimensional simplex, which contains $x_{i}$ and $x_{j}$ as border points. Redefine the median as the datapoint $x_{\text {med }}$, which is enclosed in the maximum number of intervals:

$$
\begin{equation*}
x_{\mathrm{med}}=\arg \max _{i} \#\left\{k, l ; x_{i} \in\left[x_{k}, x_{l}\right]\right\} . \tag{20.3}
\end{equation*}
$$

With this definition of the median, the median is the "deepest" and "most central" point in a data set as discussed in Sect. 1.1. This definition involves a computationally intensive operation since we generate $n(n-1) / 2$ intervals for $n$ observations.

In two dimensions, the computation is even more intensive since the interval $\left[x_{k}, x_{l}\right]$ is replaced by a triangle constructed from three different datapoints. The median as the deepest point is then defined by that datapoint that is covered by the maximum number of triangles. In three dimensions triangles become pyramids formed from 4 points and the median is that datapoint that lies in the maximum number of pyramids.

An example for the depth in two dimensions is given by the constellation of points given in Fig. 20.1. If we build for example the triangle of the points 1, 3, 5 (denoted as $\triangle 135$ in Table 20.1), it contains the point 4 . From Table 20.1 we count the number of coverages to obtain the simplicial depth values of Table 20.2.


Fig. 20.1 Construction of simplicial depth $Q_{\text {MVAs imdep1 }}$

Table 20.1 Coverages for artificial configuration of points

|  | Triangle | Coverages |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 1 | $\triangle 123$ | 1 | 2 | 3 |  |  |  |
| 2 | $\triangle 124$ | 1 | 2 |  | 4 |  |  |
| 3 | $\triangle 125$ | 1 | 2 |  |  | 5 |  |
| 4 | $\triangle 126$ | 1 | 2 | 3 | 4 |  | 6 |
| 5 | $\triangle 134$ | 1 |  | 3 | 4 |  |  |
| 6 | $\triangle 135$ | 1 |  | 3 | 4 | 5 |  |
| 7 | $\triangle 136$ | 1 |  | 3 |  |  | 6 |
| 8 | $\triangle 145$ | 1 |  |  | 4 | 5 |  |
| 9 | $\triangle 146$ | 1 |  | 3 | 4 |  | 6 |
| 10 | $\triangle 156$ | 1 |  | 3 | 4 | 5 | 6 |
| 11 | $\triangle 234$ |  | 2 | 3 | 4 |  |  |
| 12 | $\triangle 235$ |  | 2 | 3 | 4 | 5 |  |
| 13 | $\triangle 236$ |  | 2 | 3 | 4 |  | 6 |
| 14 | $\triangle 245$ |  | 2 |  | 4 | 5 |  |
| 15 | $\triangle 246$ |  | 2 |  | 4 |  | 6 |
| 16 | $\triangle 256$ |  | 2 |  |  | 5 | 6 |
| 17 | $\triangle 345$ |  |  | 3 | 4 | 5 |  |
| 18 | $\triangle 346$ |  |  | 3 | 4 |  | 6 |
| 19 | $\triangle 356$ |  |  | 3 |  | 5 | 6 |
| 20 | $\triangle 456$ |  |  |  | 4 | 5 | 6 |

Table 20.2 Simplicial depths for artificial configuration of points

| Point | 1 | 2 | 3 | 4 | 5 | 6 |
| :--- | ---: | ---: | ---: | ---: | ---: | ---: |
| Depth | 10 | 10 | 12 | 14 | 8 | 8 |

In arbitrary dimension $p$, we look for datapoints that lie inside a simplex (or convex hull) formed from $p+1$ points. We therefore extend the definition of the median to the multivariate case as follows

$$
\begin{equation*}
x_{\mathrm{med}}=\arg \max _{i} \#\left\{k_{0}, \ldots, k_{p} ; x_{i} \in \operatorname{hull}\left(x_{k_{0}}, \ldots, x_{k_{p}}\right)\right\} \tag{20.4}
\end{equation*}
$$

Here $k_{0}, \ldots, k_{p}$ denote the indices of $p+1$ datapoints. Thus for each datapoint we have a multivariate data depth. If we compute all the necessary simplices $\operatorname{hull}\left(x_{k_{0}}, \ldots, x_{k_{p}}\right)$, the computing time will unfortunately be exponential as the dimension increases.

In Fig. 20.2 we calculate the simplicial depth for a two-dimensional, 10 point distribution according to depth. It contains 100 data points with corresponding parameters controlling its spread. The deepest point, the two-dimensional median, is indicated as a big star in the centre. The points with less depth are indicated via grey shades.

Fig. 20.210 point distribution according to depth with the median shown as a big star in the centre $\mathbf{O}$ MVAsimdepex

Simplicial depth


### 20.2 Projection Pursuit

"Projection Pursuit" stands for a class of exploratory projection techniques. This class contains statistical methods designed for analysing high-dimensional data using low-dimensional projections. The aim of projection pursuit is to
reveal possible non-linear and therefore interesting structures hidden in the highdimensional data. To what extent these structures are "interesting" is measured by an index. Exploratory Projection Pursuit (EPP) goes back to Kruskal $(1969,1972)$. The approach was successfully implemented for exploratory purposes by various other authors. The idea has been applied to regression analysis, density estimation, classification and discriminant analysis.

## Exploratory Projection Pursuit

In EPP, we try to find "interesting" low-dimensional projections of the data. For this purpose, a suitable index function $I(\alpha)$, depending on a normalised projection vector $\alpha$, is used. This function will be defined such that "interesting" views correspond to local and global maxima of the function. This approach naturally accompanies the technique of PCA of the covariance structure of a random vector $X$. In PCA we are interested in finding the axes of the covariance ellipsoid. The index function $I(\alpha)$ is in this case the variance of a linear combination $\alpha^{\top} X$ subject to the normalising constraint $\alpha^{\top} \alpha=1$ (see Theorem 11.2). If we analyse a sample with a $p$-dimensional normal distribution, the "interesting" high-dimensional structure we find by maximising this index is of course linear.

There are many possible projection indices, for simplicity the kernel based and polynomial based indices are reported. Assume that the $p$-dimensional random variable $X$ is sphered and centred, that is, $\mathrm{E}(X)=0$ and $\operatorname{Var}(X)=\mathcal{I}_{p}$. This will remove the effect of location, scale, and correlation structure. This covariance structure can be achieved easily by the Mahalanobis transformation (3.26).

Friedman and Tukey (1974) proposed to investigate the high-dimensional distribution of $X$ by considering the index

$$
\begin{equation*}
I_{\mathrm{FT}, h}(\alpha)=n^{-1} \sum_{i=1}^{n} \hat{f}_{h, \alpha}\left(\alpha^{\top} X_{i}\right) \tag{20.5}
\end{equation*}
$$

where $\hat{f}_{h, \alpha}$ denotes the kernel estimator (see Sect. 1.3)

$$
\begin{equation*}
\hat{f}_{h, \alpha}(z)=n^{-1} \sum_{j=1}^{n} K_{h}\left(z-\alpha^{\top} X_{j}\right) \tag{20.6}
\end{equation*}
$$

of the projected data. Note that (20.5) is an estimate of $\int f^{2}(z) d z$ where $z=\alpha^{\top} X$ is a one-dimensional random variable with mean zero and unit variance. If the highdimensional distribution of $X$ is normal, then each projection $z=\alpha^{\top} X$ is standard normal since $\|\alpha\|=1$ and since $X$ has been centred and sphered by, e.g. the Mahalanobis transformation.

The index should therefore be stable as a function of $\alpha$ if the high-dimensional data is in fact normal. Changes in $I_{\mathrm{FT}, h}(\alpha)$ with respect to $\alpha$ therefore indicate deviations from normality. Hodges and Lehman (1956) showed that, given a mean of zero and unit variance, the (compact support) density which minimises $\int f^{2}$ is uniquely given by

$$
f(z)=\max \left\{0, c\left(b^{2}-z^{2}\right)\right\},
$$

where $c=3 /(20 \sqrt{5})$ and $b=\sqrt{5}$. This is a parabolic density function, which is equal to zero outside the interval $(-\sqrt{5}, \sqrt{5})$. A high value of the Friedman-Tukey index indicates a larger departure from the parabolic form.

An alternative index is based on the negative of the entropy measure, i.e. $\int-f \log f$. The density for zero mean and unit variance which minimises the index

$$
\int f \log f
$$

is the standard normal density, a far more plausible candidate than the parabolic density as a norm from which departure is to be regarded as "interesting". Thus in using $\int f \log f$ as a projection index we are really implementing the viewpoint of seeing "interesting" projections as departures from normality. Yet another index could be based on the Fisher information (see Sect. 6.2)

$$
\int\left(f^{\prime}\right)^{2} / f
$$

To optimise the entropy index, it is necessary to recalculate it at each step of the numerical procedure. There is no method of obtaining the index via summary statistics of the multivariate data set, so the workload of the calculation at each iteration is determined by the number of observations. It is therefore interesting to look for approximations to the entropy index. Jones and Sibson (1987) suggested that deviations from the normal density should be considered as

$$
\begin{equation*}
f(x)=\varphi(x)\{1+\varepsilon(x)\} \tag{20.7}
\end{equation*}
$$

where the function $\varepsilon$ satisfies

$$
\begin{equation*}
\int \varphi(u) \varepsilon(u) u^{-r} d u=0, \text { for } r=0,1,2 \tag{20.8}
\end{equation*}
$$

In order to develop the Jones and Sibson (1987) index it is convenient to think in terms of cumulants $\kappa_{3}=\mu_{3}=\mathrm{E}\left(X^{3}\right), \kappa_{4}=\mu_{4}=\mathrm{E}\left(X^{4}\right)-3$ (see Sect. 1.3). The standard normal density satisfies $\kappa_{3}=\kappa_{4}=0$, an index with any hope of tracking the entropy index must at least incorporate information up to the level of symmetric departures ( $\kappa_{3}$ or $\kappa_{4}$ not zero) from normality. The simplest of such indices is a
positive definite quadratic form in $\kappa_{3}$ and $\kappa_{4}$. It must be invariant under sign-reversal of the data since both $\alpha^{\top} X$ and $-\alpha^{\top} X$ should show the same kind of departure from normality. Note that $\kappa_{3}$ is odd under sign-reversal, i.e. $\kappa_{3}\left(\alpha^{\top} X\right)=-\kappa_{3}\left(-\alpha^{\top} X\right)$. The cumulant $\kappa_{4}$ is even under sign-reversal, i.e. $\kappa_{4}\left(\alpha^{\top} X\right)=\kappa_{4}\left(-\alpha^{\top} X\right)$. The quadratic form in $\kappa_{3}$ and $\kappa_{4}$ measuring departure from normality cannot include a mixed $\kappa_{3} \kappa_{4}$ term.

For the density (20.7) one may conclude with (20.8) that

$$
\int f(u) \log (u) d u \approx \frac{1}{2} \int \varphi(u) \varepsilon(u) d u
$$

Now if $f$ is expressed as a Gram-Charliér expansion

$$
\begin{equation*}
f(x) \varphi(x)=\left\{1+\kappa_{3} H_{3}(x) / 6+\kappa_{4} H_{4}(x) / 24+\cdots\right\} \tag{20.9}
\end{equation*}
$$

(Kendall \& Stuart, 1977, p. 169) where $H_{r}$ is the $r$-th Hermite polynomial, then the truncation of (20.9) and use of orthogonality and normalisation properties of Hermite polynomials with respect to $\varphi$ yields

$$
\frac{1}{2} \int \varphi(x) \varepsilon^{2}(x) d x=\left(\kappa_{3}^{2}+\kappa_{4}^{2} / 4\right) / 12
$$

The index proposed by Jones and Sibson (1987) is therefore

$$
I_{\mathrm{JS}}(\alpha)=\left\{\kappa_{3}^{2}\left(\alpha^{\top} X\right)+\kappa_{4}^{2}\left(\alpha^{\top} X\right) / 4\right\} / 12
$$

This index measures in fact the negative entropy difference $\int f \log f-\int \varphi \log \varphi$.
Example 20.1 The EPP is used on the Swiss bank note data. For 50 randomly chosen one-dimensional projections of this six-dimensional dataset we calculate the Friedman-Tukey index to evaluate how "interesting" their structures are.

Figure 20.3 shows the density for the standard, normally distributed data (green) and the estimated densities for the best (red) and the worst (blue) projections found. A dotplot of the projections is also presented. In the lower part of the figure we see the estimated value of the Friedman-Tukey index for each computed projection. From this information we can judge the non normality of the bank note data set since there is a lot of variation across the 50 random projections.

## Projection Pursuit Regression

The problem in projection pursuit regression is to estimate a response surface

$$
f(x)=\mathrm{E}(Y \mid x)
$$



Fig. 20.3 Exploratory Projection Pursuit for the Swiss bank notes data (green $=$ standard normal, red $=$ best, blue $=$ worst) Q MVAppexample
via approximating functions of the form:

$$
\hat{f}(x)=\sum_{k=1}^{M} g_{k}\left(\Lambda_{k}^{\top} x\right)
$$

with non-parametric regression functions $g_{k}$ and projection indices $\Lambda_{k}$. Given observations $\left\{\left(x_{1}, y_{1}\right), \ldots,\left(x_{n}, y_{n}\right)\right\}$ with $x_{i} \in \mathbb{R}^{p}$ and $y_{i} \in \mathbb{R}$ the basic algorithm works as follows.

1. Set $r_{i}^{(0)}=y_{i}$ and $k=1$.
2. Minimise

$$
E_{k}=\sum_{i=1}^{n}\left\{r_{i}^{(k-1)}-g_{k}\left(\Lambda_{k}^{\top} x_{i}\right)\right\}^{2}
$$

where $\Lambda_{k}$ is an orthogonal projection matrix and $g_{k}$ is a non-parametric regression estimator.
3. Compute new residuals

$$
r_{i}^{(k)}=r_{i}^{(k-1)}-g_{k}\left(\Lambda_{k}^{\top} x_{i}\right)
$$

4. Increase $k$ and repeat the last two steps until $E_{k}$ becomes small.

Although this approach seems to be simple, we encounter some problems. One of the most serious is that the decomposition of a function into sums of functions of projections may not be unique. An example is

$$
z_{1} z_{2}=\frac{1}{4 a b}\left\{\left(a z_{1}+b z_{2}\right)^{2}-\left(a z_{1}-b z_{2}\right)^{2}\right\}
$$

Numerical improvements of this algorithm were suggested by Friedman and Stuetzle (1981).

| Summary |
| :---: |
| $\hookrightarrow$ Exploratory Projection Pursuit is a technique used to find interesting structures in high-dimensional data via low-dimensional projections. Since the Gaussian distribution represents a standard situation, we define the Gaussian distribution as the most uninteresting |
| $\hookrightarrow$ The search for interesting structures is done via a projection score like the Friedman-Tukey index $I_{\mathrm{FT}}(\alpha)=\int f^{2}$. The parabolic distribution has the minimal score. We maximise this score over all projections |
| $\hookrightarrow$ The Jones-Sibson index maximises $I_{\mathrm{JS}}(\alpha)=\left\{\kappa_{3}\left(\alpha^{\top} X\right)+\kappa_{4}^{2}\left(\alpha^{\top} X\right) / 4\right\} / 12$ <br> as a function of $\alpha$ |
| $\hookrightarrow$ The entropy index maximises $I_{\mathrm{E}}(\alpha)=\int f\left(\alpha^{\top} X\right) \log f\left(\alpha^{\top} X\right)$ <br> where $f$ is the density of $\alpha^{\top} X$ |
| $\hookrightarrow$ In Projection Pursuit Regression the idea is to represent the unknown function by a sum of non-parametric regression functions on projections. The key problem is in choosing the number of terms and often the interpretability |

### 20.3 Sliced Inverse Regression

SIR is a dimension reduction method proposed by Duan and Li (1991). The idea is to find a smooth regression function that operates on a variable set of projections. Given a response variable $Y$ and a (random) vector $X \in \mathbb{R}^{p}$ of explanatory variables, SIR is based on the model:

$$
\begin{equation*}
Y=m\left(\beta_{1}^{\top} X, \ldots, \beta_{k}^{\top} X, \varepsilon\right) \tag{20.10}
\end{equation*}
$$

where $\beta_{1}, \ldots, \beta_{k}$ are unknown projection vectors, $k$ is unknown and assumed to be less than $p, m: \mathbb{R}^{k+1} \rightarrow \mathbb{R}$ is an unknown function, and $\varepsilon$ is the noise random variable with $\mathrm{E}(\varepsilon \mid X)=0$.

Model (20.10) describes the situation where the response variable $Y$ depends on the $p$-dimensional variable $X$ only through a $k$-dimensional subspace. The unknown $\beta_{i}$ 's, which span this space, are called effective dimension reduction directions (EDR-directions). The span is denoted as effective dimension reduction space (EDR-space). The aim is to estimate the base vectors of this space, for which neither the length nor the direction can be identified. Only the space in which they lie is identifiable.

SIR tries to find this $k$-dimensional subspace of $\mathbb{R}^{p}$ which under the model (20.10) carries the essential information of the regression between $X$ and $Y$. SIR also focuses on small $k$, so that nonparametric methods can be applied for the estimation of $m$. A direct application of nonparametric smoothing to $X$ is for high dimension $p$ generally not possible due to the sparseness of the observations. This fact is well known as the curse of dimensionality, see Huber (1985).

The name of SIR comes from computing the inverse regression (IR) curve. That means instead of looking for $\mathrm{E}(Y \mid X=x)$, we investigate $\mathrm{E}(X \mid Y=y)$, a curve in $\mathbb{R}^{p}$ consisting of $p$ one-dimensional regressions. What is the connection between the IR and the SIR model (20.10)? The answer is given in the following theorem from $\operatorname{Li}$ (1991).

Theorem 20.1 Given the model (20.10) and the assumption

$$
\begin{equation*}
\forall b \in \mathbb{R}^{p}: \mathrm{E}\left(b^{\top} X \mid \beta_{1}^{\top} X=\beta_{1}^{\top} x, \ldots, \beta_{k}^{\top} X=\beta_{k}^{\top} x\right)=c_{0}+\sum_{i=1}^{k} c_{i} \beta_{i}^{\top} x \tag{20.11}
\end{equation*}
$$

the centred IR curve $\mathrm{E}(X \mid Y=y)-\mathrm{E}(X)$ lies in the linear subspace spanned by the vectors $\Sigma \beta_{i}, i=1, \ldots, k$, where $\Sigma=\operatorname{Cov}(X)$.

Assumption (20.11) is equivalent to the fact that $X$ has an elliptically symmetric distribution, see Cook and Weisberg (1991). Hall and Li (1993) have shown that assumption (20.11) only needs to hold for the EDR-directions.

It is easy to see that for the standardised variable $Z=\Sigma^{-1 / 2}\{X-\mathrm{E}(X)\}$ the IR curve $m_{1}(y)=\mathrm{E}(Z \mid Y=y)$ lies in $\operatorname{span}\left(\eta_{1}, \ldots, \eta_{k}\right)$, where $\eta_{i}=\Sigma^{1 / 2} \beta_{i}$. This means that the conditional expectation $m_{1}(y)$ is moving in $\operatorname{span}\left(\eta_{1}, \ldots, \eta_{k}\right)$ depending on $y$. With $b$ orthogonal to $\operatorname{span}\left(\eta_{1}, \ldots, \eta_{k}\right)$, it follows that

$$
b^{\top} m_{1}(y)=0
$$

and further that

$$
m_{1}(y) m_{1}(y)^{\top} b=\operatorname{Cov}\left\{m_{1}(y)\right\} b=0
$$

As a consequence $\operatorname{Cov}\{\mathrm{E}(Z \mid y)\}$ is degenerated in each direction orthogonal to all EDR-directions $\eta_{i}$ of $Z$. This suggests the following algorithm.

First, estimate $\operatorname{Cov}\left\{m_{1}(y)\right\}$ and then calculate the orthogonal directions of this matrix (for example, with eigenvalue/eigenvector decomposition). In general, the estimated covariance matrix will have full rank because of random variability, estimation errors and numerical imprecision. Therefore, we investigate the eigenvalues of the estimate and ignore eigenvectors having small eigenvalues. These eigenvectors $\hat{\eta}_{i}$ are estimates for the EDR-direction $\eta_{i}$ of $Z$. We can easily rescale them to estimates $\hat{\beta}_{i}$ for the EDR-directions of $X$ by multiplying by $\hat{\Sigma}^{-1 / 2}$, but then they are not necessarily orthogonal. SIR is strongly related to PCA. If all of the data falls into a single interval, which means that $\widehat{\operatorname{Cov}}\left\{m_{1}(y)\right\}$ is equal to $\widehat{\operatorname{Cov}}(Z)$, SIR coincides with PCA. Obviously, in this case any information about $y$ is ignored.

## The SIR Algorithm

The algorithm to estimate the EDR-directions via SIR is as follows:

1. Standardise $x$ :

$$
z_{i}=\hat{\Sigma}^{-1 / 2}\left(x_{i}-\bar{x}\right)
$$

2. Divide the range of $y_{i}$ into $S$ nonoverlapping intervals (slices) $H_{s}, s=1, \ldots, S$. $n_{s}$ denotes the number of observations within slice $H_{s}$, and $\boldsymbol{I}_{H_{s}}$ the indicator function for this slice:

$$
n_{s}=\sum_{i=1}^{n} \boldsymbol{I}_{H_{s}}\left(y_{i}\right)
$$

3. Compute the mean of $z_{i}$ over all slices. This is a crude estimate $\hat{m}_{1}$ for the inverse regression curve $m_{1}$ :

$$
\bar{z}_{s}=n_{s}^{-1} \sum_{i=1}^{n} z_{i} \boldsymbol{I}_{H_{s}}\left(y_{i}\right)
$$

4. Calculate the estimate for $\operatorname{Cov}\left\{m_{1}(y)\right\}$ :

$$
\hat{V}=n^{-1} \sum_{s=1}^{S} n_{s} \bar{z}_{s} \bar{z}_{s}^{\top}
$$

5. Identify the eigenvalues $\hat{\lambda}_{i}$ and eigenvectors $\hat{\eta}_{i}$ of $\hat{V}$.
6. Transform the standardised EDR-directions $\hat{\eta}_{i}$ back to the original scale. Now the estimates for the EDR-directions are given by

$$
\hat{\beta}_{i}=\hat{\Sigma}^{-1 / 2} \hat{\eta}_{i} .
$$

Remark 20.1 The number of different eigenvalues unequal to zero depends on the number of slices. The rank of $\hat{V}$ cannot be greater than the number of slices -1 (the $z_{i}$ sum up to zero). This is a problem for categorical response variables, especially for a binary response-where only one direction can be found.

## SIR II

In the previous section we learned that it is interesting to consider the IR curve, that is, $\mathrm{E}(X \mid y)$. In some situations however SIR does not find the EDR-direction. We overcome this difficulty by considering the conditional covariance $\operatorname{Cov}(X \mid y)$ instead of the IR curve. An example where the EDR directions are not found via the SIR curve is given below.
Example 20.2 Suppose that $\left(X_{1}, X_{2}\right)^{\top} \sim N\left(0, \mathcal{I}_{2}\right)$ and $Y=X_{1}^{2}$. Then $\mathrm{E}\left(X_{2} \mid\right.$ $y)=0$ because of independence and $\mathrm{E}\left(X_{1} \mid y\right)=0$ because of symmetry. Hence, the EDR-direction $\beta=(1,0)^{\top}$ is not found when the $\operatorname{IR}$ curve $\mathrm{E}(X \mid y)=0$ is considered.

The conditional variance

$$
\operatorname{Var}\left(X_{1} \mid Y=y\right)=\mathrm{E}\left(X_{1}^{2} \mid Y=y\right)=y
$$

offers an alternative way to find $\beta$. It is a function of $y$ while $\operatorname{Var}\left(X_{2} \mid y\right)$ is a constant.

The idea of SIR II is to consider the conditional covariances. The principle of SIR II is the same as before: investigation of the IR curve (here the conditional covariance instead of the conditional expectation). Unfortunately, the theory of SIR II is more complicated. The assumption of the elliptical symmetrical distribution of $X$ has to be more restrictive, i.e. assuming the normality of $X$.

Given this assumption, one can show that the vectors with the largest distance to $\operatorname{Cov}(Z \mid Y=y)-\mathrm{E}\{\operatorname{Cov}(Z \mid Y=y)\}$ for all $y$ are the most interesting for the

EDR-space. An appropriate measure for the overall mean distance is, according to Li (1992),

$$
\begin{align*}
& \mathrm{E}\left(\|[\operatorname{Cov}(Z \mid Y=y)-\mathrm{E}\{\operatorname{Cov}(Z \mid Y=y)\}] b\|^{2}\right) \\
& \quad=b^{\top} \mathrm{E}\left(\|\operatorname{Cov}(Z \mid y)-\mathrm{E}\{\operatorname{Cov}(Z \mid y)\}\|^{2}\right) b \tag{20.12}
\end{align*}
$$

Equipped with this distance, we conduct again an eigensystem decomposition, this time for the above expectation $\mathrm{E}\left(\|\operatorname{Cov}(Z \mid y)-\mathrm{E}\{\operatorname{Cov}(Z \mid y)\}\|^{2}\right)$. Then we take the rescaled eigenvectors with the largest eigenvalues as estimates for the unknown EDR-directions.

## The SIR II Algorithm

The algorithm of SIR II is very similar to the one for SIR, it differs in only two steps. Instead of merely computing the mean, the covariance of each slice has to be computed. The estimate for the above expectation (20.12) is calculated after computing all slice covariances. Finally, decomposition and rescaling are conducted, as before.

1. Do steps $1-3$ of the SIR algorithm.
2. Compute the slice covariance matrix $\hat{V}_{s}$ :

$$
\hat{V}_{s}=\left(n_{s}-1\right)^{-1} \sum_{i=1}^{n} I_{H_{s}}\left(y_{i}\right) z_{i} z_{i}^{\top}-n_{s} \bar{z}_{s} \bar{z}_{s}^{\top} .
$$

3. Calculate the mean over all slice covariances:

$$
\bar{V}=n^{-1} \sum_{s=1}^{S} n_{s} \hat{V}_{s} .
$$

4. Compute an estimate for (20.12):

$$
\hat{V}=n^{-1} \sum_{s=1}^{S} n_{s}\left(\hat{V}_{s}-\bar{V}\right)^{2}=n^{-1} \sum_{s=1}^{S} n_{s} \hat{V}_{s}^{2}-\bar{V}^{2}
$$

5. Identify the eigenvectors and eigenvalues of $\hat{V}$ and scale back the eigenvectors. This gives estimates for the SIR II EDR-directions:

$$
\hat{\beta}_{i}=\hat{\Sigma}^{-1 / 2} \hat{\eta}_{i} .
$$



Fig. 20.4 SIR: The left plots show the response versus the estimated EDR-directions. The upper right plot is a three-dimensional plot of the first two directions and the response. The lower right plot shows the eigenvalues $\hat{\lambda}_{i}$ (asterisk) and the cumulative sum (open circle) Q MVAsirdata

Example 20.3 The result of SIR is visualised in four plots in Fig. 20.4: the left two show the response variable versus the first respectively second direction. The upper right plot consists of a three-dimensional plot of the first two directions and the response. The last picture shows $\hat{\Psi}_{k}$, the ratio of the sum of the first $k$ eigenvalues and the sum of all eigenvalues, similar to PCA.

The data are generated according to the following model:

$$
y_{i}=\beta_{1}^{\top} x_{i}+\left(\beta_{1}^{\top} x_{i}\right)^{3}+4\left(\beta_{2}^{\top} x_{i}\right)^{2}+\varepsilon_{i},
$$

where the $x_{i}$ 's follow a three-dimensional normal distribution with zero mean, the covariance equal to the identity matrix, $\beta_{2}=(1,-1,-1)^{\top}$, and $\beta_{1}=(1,1,1)^{\top}$.

Fig. 20.5 Plot of the true response versus the true first index. The monotonic and the convex shapes can be clearly seen Q MVAsirdata


Fig. 20.6 Plot of the true response versus the true second index. The monotonic and the convex shapes can be clearly seen $\mathbf{Q}$ MVAsirdata

$\varepsilon_{i}$ is standard, normally distributed and $n=300$. Corresponding to model (20.10), $m(u, v, \varepsilon)=u+u^{3}+v^{2}+\varepsilon$. The situation is depicted in Figs. 20.5 and 20.6.

Both algorithms were conducted using the slicing method with 20 elements in each slice. The goal was to find $\beta_{1}$ and $\beta_{2}$ with SIR. The data are designed such that SIR can detect $\beta_{1}$ because of the monotonic shape of $\left\{\beta_{1}^{\top} x_{i}+\left(\beta_{1}^{\top} x_{i}\right)^{3}\right\}$, while SIR II will search for $\beta_{2}$, as in this direction the conditional variance on $y$ is varying.

Table 20.3 SIR:
EDR-directions for simulated data

| $\hat{\beta}_{1}$ | $\hat{\beta}_{2}$ | $\hat{\beta}_{3}$ |
| :--- | ---: | ---: |
| 0.452 | 0.881 | 0.040 |
| 0.571 | -0.349 | -0.787 |
| 0.684 | -0.320 | 0.615 |

Table 20.4 SIR II:
EDR-directions for simulated data

| $\hat{\beta}_{1}$ | $\hat{\beta}_{2}$ | $\hat{\beta}_{3}$ |
| :---: | :--- | ---: |
| -0.272 | 0.964 | -0.001 |
| 0.670 | 0.100 | 0.777 |
| 0.690 | 0.244 | -0.630 |

If we normalise the eigenvalues for the EDR-directions in Table 20.3 such that they sum up to one, the resulting vector is $(0.852,0.086,0.062)$. As can be seen in the upper left plot of Fig. 20.4, there is a functional relationship found between the first index $\hat{\beta}_{1}^{\top} x$ and the response. Actually, $\beta_{1}$ and $\hat{\beta}_{1}$ are nearly parallel, that is, the normalised inner product $\hat{\beta}_{1}^{\top} \beta_{1} /\left\{\left|\left|\hat{\beta}_{1}\right|\left\|| | \beta_{1}\right\|\right\}=0.9894\right.$ is very close to one.

The second direction along $\beta_{2}$ is probably found due to the good approximation, but SIR does not provide it clearly, because it is "blind" with respect to the change of variance, as the second eigenvalue indicates.

For SIR II, the normalised eigenvalues are $(0.706,0.185,0.108)$, that is, about $69 \%$ of the variance is explained by the first EDR-direction (Table 20.4). Here, the normalised inner product of $\beta_{2}$ and $\hat{\beta}_{1}$ is 0.9992 . The estimator $\hat{\beta}_{1}$ estimates in fact $\beta_{2}$ of the simulated model. In this case, SIR II found the direction where the second moment varies with respect to $\beta_{2}^{\top} x$ (Fig. 20.7).

In summary, SIR has found the direction which shows a strong relation regarding the conditional expectation between $\beta_{1}^{\top} x$ and $y$, and SIR II has found the direction where the conditional variance is varying, namely, $\beta_{2}^{\top} x$.

The behaviour of the two SIR algorithms is as expected. In addition, we have seen that it is worthwhile to apply both versions of SIR. It is possible to combine SIR and SIR II (Cook \& Weisberg, 1991; Li, 1991; Schott, 1994) directly, or to investigate higher conditional moments. For the latter it seems to be difficult to obtain theoretical results.

|  | Summary |
| :--- | :--- |
|  | SIR serves as a dimension reduction tool for regression problems |
|  | Inverse regression avoids the curse of dimensionality |
|  | The dimension reduction can be conducted without estimation of <br> the regression function $y=m(x)$ |


| Summary (continued) |  |
| :---: | :--- |
| $\hookrightarrow$ | SIR searches for the effective dimension reduction (EDR) by |
| computing the inverse regression IR |  |$]$



Fig. 20.7 SIR II mainly sees the direction $\beta_{2}$. The left plots show the response versus the estimated EDR-directions. The upper right plot is a three-dimensional plot of the first two directions and the response. The lower right plot shows the eigenvalues $\hat{\lambda}_{i}$ (asterisk) and the cumulative sum (open circle) Q MVAsir2data

### 20.4 Support Vector Machines

The purpose of this section is to introduce one of the most promising among recently developed multivariate non-linear statistical techniques: the SVM. The SVM is a classification method that is based on statistical learning theory. It has been successfully applied to optical character recognition, early medical diagnostics, and text classification. One application where SVMs outperformed other methods is electric load prediction (EUNITE, 2001), another one is optical character recognition (Vapnik, 1995). In a variety of applications SVMs produce better classification results than parametric methods (e.g. logit analysis) and are outperforming widely used nonparametric techniques, such as neural networks. Here we apply SVMs to corporate bankruptcy analysis.

## Classification Methodology

In order to illustrate the classification methodology we focus for the moment on a company rating example that we will treat further in more detail. Investment risks are evaluated via the default probability (PD) for a company. Each company is described by a set of variables (predictors) $x$, such as financial ratios, and its class $y$ that can be either $y=-1$ ("successful") or $y=1$ ("bankrupt"). Financial ratios are constructed from the variables like net income, total assets, interest payments, etc. A training set represents a sample of data for companies which are known to have survived or gone bankrupt. From the training set one estimates a classifier function $f$ that is then applied to computing PDs. These PDs can be uniquely translated into a company rating.

Classical discriminant analysis is based on the assumption that each group of observations is normally distributed with the same variance-covariance matrix but different means. Under such a formulation the discriminating function will be linear, see Theorem 14.2. Figure 20.8 displays this situation: if some linear combination of predictors (called $Z$-score in the context of bankruptcy analysis) is greater than a particular threshold value $z_{0}$ the observation under consideration is regarded as belonging to $y=1$; if $Z<z_{0}$ the observation would belong to $y=-1$ (successful). One can change the labels " $-1,+1$ " to the more standard notation " 0,1 ". The current labeling is done only for mathematical convenience.

The $Z$-score is:

$$
Z_{i}=a_{1} x_{i 1}+a_{2} x_{i 2}+\ldots+a_{p} x_{i p}=a^{\top} x_{i}
$$

where $x_{i}=\left(x_{i 1}, \ldots, x_{i p}\right)^{\top} \in \mathbb{R}^{p}$ are predictors for the $i$-th company. The classification based on the $Z$-score are necessarily linear and, therefore, may not handle more complex situations as in Fig. 20.9 when non-linear classifiers, such as those generated by SVMs, can produce better results.

Fig. 20.8 A linear classification function in the case of linearly separable data


Fig. 20.9 Different linear classification functions (1) and (2) and a non-linear one (3) in the linearly non-separable case


## Expected vs. Empirical Risk Minimisation

A non-linear classifier function $f$ may be described by a function class $\mathcal{F} . \mathcal{F}$ is fixed a priori, e.g. it can be the class of linear classifiers (hyperplanes). A good classifier optimises some criterion that tells us how well $f$ separates the classes. As in (14.4) one considers the minimisation of the expected risk:

$$
\begin{equation*}
R(f)=\int \frac{1}{2}|f(x)-y| d F(x, y) \tag{20.13}
\end{equation*}
$$

The joint distribution $F(x, y)$, however, is never known in practical applications and must be estimated from the training set $\left\{x_{i}, y_{i}\right\}_{i=1}^{n}$. By replacing $F(x, y)$ with the empirical cdf $F_{n}(x, y)$ one obtains the empirical risk:

$$
\begin{equation*}
\hat{R}(f)=\frac{1}{n} \sum_{i=1}^{n} \frac{1}{2}\left|f\left(x_{i}\right)-y_{i}\right| \tag{20.14}
\end{equation*}
$$

The empirical risk is an average value of loss over the training set, while the expected risk is the expected value of loss under the true probability measure. The loss is given by:

$$
L(x, y)=\frac{1}{2}|f(x)-y|= \begin{cases}0, & \text { if classification is correct }, \\ 1, & \text { if classification is wrong. }\end{cases}
$$

One sees here that it is convenient to work with the labels " $-1,1$ " for $y$. The solutions to the problems of expected and empirical risk minimisation:

$$
\begin{align*}
f_{\mathrm{opt}} & =\arg \min _{f \in \mathcal{F}} R(f),  \tag{20.15}\\
\hat{f_{n}} & =\arg \min _{f \in \mathcal{F}} \hat{R}(f), \tag{20.16}
\end{align*}
$$

generally do not coincide (Fig. 20.10), although converge as $n \rightarrow \infty$ if $\mathcal{F}$ is not too large. According to statistical learning theory (Vapnik, 1995), it is possible to get a uniform upper bound on the difference between $R(f)$ and $\hat{R}(f)$ via the Vapnik-Chervonenkis (VC) theory. The VC bound states that there is a function $\phi$ (monotone increasing in $h$ ) so that for all $f \in \mathcal{F}$ with a probability $1-\eta$ :

$$
\begin{equation*}
R(f) \leq \hat{R}(f)+\phi\left\{\frac{h}{n}, \frac{\log (\eta)}{n}\right\} . \tag{20.17}
\end{equation*}
$$

Here $h$ denotes the VC dimension, a measure of complexity of the involved function class $\mathcal{F}$. For a linear classification rule $g(x)=\operatorname{sign}\left(x^{\top} w+b\right)$ :

$$
\begin{equation*}
\phi\left\{\frac{h}{n}, \frac{\log (\eta)}{n}\right\}=\sqrt{\frac{h\left(\log \frac{2 n}{h}\right)-\log \frac{\eta}{4}}{n}}, \tag{20.18}
\end{equation*}
$$

Fig. 20.10 The minima $f_{\text {opt }}$ and $\hat{f}_{n}$ of the expected $(R)$ and empirical $(\hat{R})$ risk functions generally do not coincide

where $h$ is the VC dimension. By plotting the function $\phi(u, v)=\{-u \cdot \log 2 u+$ $\log 4-v\}^{-1 / 2}$ for small $u$ one sees the monotonicity of $\phi(u, v)$. In fact one can show that

$$
\frac{\partial \phi\left(\frac{h}{n}, \frac{\log (\eta)}{n}\right)}{\partial h} \geqslant 0
$$

if and only if $2 n \geqslant h$. For a linear classifier with $h=p+1$ this is an easy condition to meet.

The VC dimension of a set $\mathcal{F}$ of functions in a $d$-dimensional space is $h$ if some function $f \in \mathcal{F}$ can shatter $h$ objects $\left\{x_{i} \in \mathbb{R}^{d}, i=1, \ldots, h\right\}$, in all $2^{h}$ possible configurations and no set $\left\{x_{j} \in \mathbb{R}^{d}, j=1, \ldots, q\right\}$ with $q>h$, exists that satisfies this property. For example, three points on a plane $(d=2)$ can be shattered by linear indicator functions in $2^{h}=2^{3}=8$ ways, whereas 4 points can not be shattered in $2^{q}=2^{4}=16$ ways. Thus, the VC dimension of the set of linear indicator functions in a two-dimensional space is $h=3$, see Fig. 20.11. The expression for the VC bound (20.17) involves the VC dimension $h$, a parameter controlling complexity of $\mathcal{F}$. The term $\phi\left\{\frac{h}{n}, \frac{\log (\eta)}{n}\right\}$ introduces a penalty for excessive complexity of a classifier function. The higher is the complexity of $f \in \mathcal{F}$ the higher are $h$ and therefore $\phi$. There is a trade-off between the number of classification errors on the training set and the complexity of the classifier function. If the complexity were not controlled for, it would be possible to construct a classifier function with no classification errors on the training set notwithstanding how low its generalisation ability would be.


Fig. 20.11 Eight possible ways of shattering 3 points on the plane with a linear indicator function

## The SVM in the Linearly Separable Case

First we will describe the SVM in the linearly separable case. The family $\mathcal{F}$ of classification functions in the data space is given by:

$$
\begin{equation*}
\mathcal{F}=\left\{x^{\top} w+b, w \in \mathbb{R}^{p}, b \in \mathbb{R}\right\} \tag{20.19}
\end{equation*}
$$

In order to determine the support vectors we choose $f \in \mathcal{F}$ (or equivalently $(w, b)$ ) such that the so-called margin-the corridor between the separating hyperplanesis maximal. This situation is illustrated in Fig. 20.12. The margin is equal to $d_{-}+d_{+}$. The classification function is a hyperplane plus the margin zone, where, in the separable case, no observations can lie. It separates the points from both classes with the highest "safest" distance (margin) between them. It can be shown that margin maximisation corresponds to the reduction of complexity as given by the VCdimension of the SVM classifier. Apparently, the separating hyperplane is defined only by the support vectors that hold the hyperplanes parallel to the separating one. In Fig. 20.12 there are three support vectors that are marked with bold style: two crosses and one circle. We come now to the description of the SVM selection.

Let $x^{\top} w+b=0$ be a separating hyperplane. Then $d_{+}\left(d_{-}\right)$will be the shortest distance to the closest objects from the classes $+1(-1)$. Since the separation can be done without errors, all observations $i=1,2, \ldots, n$ must satisfy:

$$
\begin{array}{lll}
x_{i}^{\top} w+b \geq+1 & \text { for } & y_{i}=+1 \\
x_{i}^{\top} w+b \leq-1 & \text { for } & y_{i}=-1
\end{array}
$$

Fig. 20.12 The separating hyperplane $x^{\top} w+b=0$ and the margin in the linearly separable case


We can combine both constraints into one:

$$
\begin{equation*}
y_{i}\left(x_{i}^{\top} w+b\right)-1 \geq 0 \quad i=1,2, \ldots, n \tag{20.20}
\end{equation*}
$$

The canonical hyperplanes $x_{i}^{\top} w+b= \pm 1$ are parallel and the distance between each of them and the separating hyperplane is $d_{+}=d_{-}=1 /\|w\|$. To maximise the margin $d_{+}+d_{-}=2 /\|w\|$ one therefore minimises the Euclidean norm $\|w\|$ or its square $\|w\|^{2}$.

The Lagrangian for the primal problem that corresponds to margin maximisation subject to constraint (20.20) is:

$$
\begin{equation*}
L_{P}(w, b)=\frac{1}{2}\|w\|^{2}-\sum_{i=1}^{n} \alpha_{i}\left\{y_{i}\left(x_{i}^{\top} w+b\right)-1\right\} \tag{20.21}
\end{equation*}
$$

The Karush-Kuhn-Tucker (KKT) (Gale et al., 1951) first order optimality conditions are:

$$
\begin{array}{ll}
\frac{\partial L_{P}}{\partial w}=0: & w-\sum_{i=1}^{n} \alpha_{i} y_{i} x_{i}=0 \\
\frac{\partial L_{P}}{\partial b}=0: & \sum_{i=1}^{n} \alpha_{i} y_{i}=0 \\
& y_{i}\left(x_{i}^{\top} w+b\right)-1 \geq 0, \quad i=1, \ldots, n \\
& \alpha_{i} \geq 0 \\
& \alpha_{i}\left\{y_{i}\left(x_{i}^{\top} w+b\right)-1\right\}=0
\end{array}
$$

From these first order condition, we can derive $w=\sum_{i=1}^{n} \alpha_{i} y_{i} x_{i}$ and therefore the summands in (20.21) read:

$$
\begin{aligned}
\frac{1}{2}\|w\|^{2} & =\frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n} \alpha_{i} \alpha_{j} y_{i} y_{j} x_{i}^{\top} x_{j} \\
-\sum_{i=1}^{n} \alpha_{i}\left\{y_{i}\left(x_{i}^{\top} w+b\right)-1\right\} & =-\sum_{i=1}^{n} \alpha_{i} y_{i} x_{i}^{\top} \sum_{j=1}^{n} \alpha_{j} y_{j} x_{j}+\sum_{i=1}^{n} \alpha_{i} \\
& =-\sum_{i=1}^{n} \sum_{j=1}^{n} \alpha_{i} \alpha_{j} y_{i} y_{j} x_{i}^{\top} x_{j}+\sum_{i=1}^{n} \alpha_{i}
\end{aligned}
$$

Substituting this into (20.21) we obtain the Lagrangian for the dual problem:

$$
\begin{equation*}
L_{D}(\alpha)=\sum_{i=1}^{n} \alpha_{i}-\frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n} \alpha_{i} \alpha_{j} y_{i} y_{j} x_{i}^{\top} x_{j} . \tag{20.22}
\end{equation*}
$$

The primal and dual problems are:

$$
\begin{gathered}
\min _{w, b} L_{P}(w, b) \\
\max _{\alpha} L_{D}(\alpha) \quad \text { s.t. } \quad \alpha_{i} \geq 0, \quad \sum_{i=1}^{n} \alpha_{i} y_{i}=0 .
\end{gathered}
$$

Since the optimisation problem is convex the dual and primal formulations give the same solution.

Those points $i$ for which the equation $y_{i}\left(x_{i}^{\top} w+b\right)=1$ holds are called support vectors. After "training the SVM" i.e. solving the dual problem above and deriving Lagrange multipliers (they are equal to 0 for non-support vectors) one can classify a company. One uses the classification rule:

$$
\begin{equation*}
g(x)=\operatorname{sign}\left(x^{\top} w+b\right), \tag{20.23}
\end{equation*}
$$

where $w=\sum_{i=1}^{n} \alpha_{i} y_{i} x_{i}$ and $b=\frac{1}{2}\left(x_{+1}+x_{-1}\right) w . x_{+1}$ and $x_{-1}$ are two support vectors belonging to different classes for which $y\left(x^{\top} w+b\right)=1$. The value of the classification function (the score of a company) can be computed as

$$
\begin{equation*}
f(x)=x^{\top} w+b \tag{20.24}
\end{equation*}
$$

Each score $f(x)$ uniquely corresponds to a default probability (PD). The higher $f(x)$ the higher the PD.

## SVMs in the Linearly Non-separable Case

In the linearly non-separable case the situation is like in Fig. 20.13. The slack variables $\xi_{i}$ represent the violation from strict separation. In this case the following inequalities can be induced from Fig. 20.13:

$$
\begin{gathered}
x_{i}^{\top} w+b \geq 1-\xi_{i} \text { for } y_{i}=1, \\
x_{i}^{\top} w+b \leq-1+\xi_{i} \text { for } y_{i}=-1, \\
\xi_{i} \geq 0 .
\end{gathered}
$$

Fig. 20.13 The separating hyperplane $x^{\top} w+b=0$ and the margin in the linearly non-separable case


They can be combined into two constraints:

$$
\begin{gather*}
y_{i}\left(x_{i}^{\top} w+b\right) \geq 1-\xi_{i}  \tag{20.25}\\
\xi_{i} \geq 0 . \tag{20.26}
\end{gather*}
$$

SVM classification again maximises the margin given a family of classification functions $\mathcal{F}$.

The penalty for misclassification, the classification error $\xi_{i} \geq 0$, is related to the distance from a misclassified point $x_{i}$ to the canonical hyperplane bounding its class. If $\xi_{i}>0$, an error in separating the two sets occurs. The objective function corresponding to penalised margin maximisation is then formulated as:

$$
\begin{equation*}
\frac{1}{2}\|w\|^{2}+C \sum_{i=1}^{n} \xi_{i} \tag{20.27}
\end{equation*}
$$

where the parameter $C$ characterises the weight given to the classification errors. The minimisation of the objective function with constraint (20.25) and (20.26) provides the highest possible margin in the case when classification errors are inevitable due to the linearity of the separating hyperplane. Under such a formulation the problem is convex.

The Lagrange function for the primal problem is:

$$
\begin{equation*}
L_{P}(w, b, \xi)=\frac{1}{2}\|w\|^{2}+C \sum_{i=1}^{n} \xi_{i}-\sum_{i=1}^{n} \alpha_{i}\left\{y_{i}\left(x_{i}^{\top} w+b\right)-1+\xi_{i}\right\}-\sum_{i=1}^{n} \mu_{i} \xi_{i} \tag{20.28}
\end{equation*}
$$

where $\alpha_{i} \geq 0$ and $\mu_{i} \geq 0$ are Lagrange multipliers. The primal problem is formulated as:

$$
\min _{w, b, \xi} L_{P}(w, b, \xi) .
$$

The first order conditions in this case are:

$$
\begin{array}{ll}
\frac{\partial L_{P}}{\partial w}=0: & w-\sum_{i=1}^{n} \alpha_{i} y_{i} x_{i}=0 \\
\frac{\partial L_{P}}{\partial b}=0: & \sum_{i=1}^{n} \alpha_{i} y_{i}=0 \\
\frac{\partial L_{P}}{\partial \xi_{i}}=0: & C-\alpha_{i}-\mu_{i}=0
\end{array}
$$

With the conditions for the Lagrange multipliers:

$$
\begin{aligned}
& \alpha_{i} \geq 0 \\
& \mu_{i} \geq 0 \\
& \alpha_{i}\left\{y_{i}\left(x_{i}^{\top} w+b\right)-1+\xi_{i}\right\}=0 \\
& \mu_{i} \xi_{i}=0
\end{aligned}
$$

Note that $\sum_{i=1}^{n} \alpha_{i} y_{i} b=0$ therefore similar to the linear separable case the primal problem translates into:

$$
\begin{aligned}
L_{D}(\alpha)= & \frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n} \alpha_{i} \alpha_{j} y_{i} y_{j} x_{i}^{\top} x_{j}-\sum_{i=1}^{n} \alpha_{i} y_{i} x_{i}^{\top} \sum_{j=1}^{n} \alpha_{j} y_{j} x_{j} \\
& +C \sum_{i=1}^{n} \xi_{i}+\sum_{i=1}^{n} \alpha_{i}-\sum_{i=1}^{n} \alpha_{i} \xi_{i}-\sum_{i=1}^{n} \mu_{i} \xi_{i} \\
= & \sum_{i=1}^{n} \alpha_{i}-\frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n} \alpha_{i} \alpha_{j} y_{i} y_{j} x_{i}^{\top} x_{j}+\sum_{i=1}^{n} \xi_{i}\left(C-\alpha_{i}-\mu_{i}\right)
\end{aligned}
$$

Since the last term is 0 we derive the dual problem as:

$$
\begin{equation*}
L_{D}(\alpha)=\sum_{i=1}^{n} \alpha_{i}-\frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n} \alpha_{i} \alpha_{j} y_{i} y_{j} x_{i}^{\top} x_{j} \tag{20.29}
\end{equation*}
$$

and the dual problem is posed as:

$$
\max _{\alpha} L_{D}(\alpha),
$$

subject to:

$$
\begin{gathered}
0 \leq \alpha_{i} \leq C \\
\sum_{i=1}^{n} \alpha_{i} y_{i}=0
\end{gathered}
$$

## Non-linear Classification

The SVMs can also be generalised to the non-linear case. In order to obtain nonlinear classifiers as in Fig. 20.14 one maps the data with a non-linear structure via a function $\Psi: \mathbb{R}^{p} \mapsto \mathbb{H}$ into a very large dimensional space $\mathbb{H}$ where the classification rule is (almost) linear. Note that all the training vectors $x_{i}$ appear in $L_{D}$ (20.29) only as scalar products of the form $x_{i}^{\top} x_{j}$. In the non-linear SVM situations this transforms to $\psi\left(x_{i}\right)^{\top} \psi\left(x_{j}\right)$.

The so-called kernel trick is to compute this scalar product via a kernel function. These kernel functions are actually related to those we presented in Sect.1.3. If a kernel function $K$ exists such that $K\left(x_{i}, x_{j}\right)=\Psi\left(x_{i}\right)^{\top} \Psi\left(x_{j}\right)$, then it can be used without knowing the transformation $\Psi$ explicitly. A necessary and sufficient condition for a symmetric function $K\left(x_{i}, x_{j}\right)$ to be a kernel is given by Mercer's theorem (Mercer, 1909). It requires positive definiteness, i.e. for


Fig. 20.14 Mapping into a three-dimensional feature space from a two-dimensional data space $\mathbb{R}^{2} \mapsto \mathbb{R}^{3}$. The transformation $\Psi\left(x_{1}, x_{2}\right)=\left(x_{1}^{2}, \sqrt{2} x_{1} x_{2}, x_{2}^{2}\right)^{\top}$ corresponds to the kernel function $K\left(x_{i}, x_{j}\right)=\left(x_{i}^{\top} x_{j}\right)^{2}$
any data set $x_{1}, \ldots, x_{n}$ and any real numbers $\lambda_{1}, \ldots, \lambda_{n}$ the function $K$ must satisfy

$$
\begin{equation*}
\sum_{i=1}^{n} \sum_{j=1}^{n} \lambda_{i} \lambda_{j} K\left(x_{i}, x_{j}\right) \geq 0 \tag{20.30}
\end{equation*}
$$

Some examples of kernel functions are:

- $K\left(x_{i}, x_{j}\right)=e^{-\left\|x_{i}-x_{j}\right\| / 2 \sigma^{2}}$-the isotropic Gaussian kernel with constant $\sigma$
- $K\left(x_{i}, x_{j}\right)=e^{-\left(x_{i}-x_{j}\right)^{\top} r^{-2} \Sigma^{-1}\left(x_{i}-x_{j}\right) / 2}$-the stationary Gaussian kernel with an anisotropic radial basis with constant $r$ and variance-covariance matrix $\Sigma$ from training set
- $K\left(x_{i}, x_{j}\right)=\left(x_{i}^{\top} x_{j}+1\right)^{p}$-the polynomial kernel of degree $p$
- $K\left(x_{i}, x_{j}\right)=\tanh \left(k x_{i}^{\top} x_{j}-\delta\right)$-the hyperbolic tangent kernel with constant $k$ and $\delta$.


## SVMs for Simulated Data

The basic parameters of SVMs are on the scaling $r$ of the anisotropic radial basis functions (in the stationary Gaussian kernel) and the capacity $C$. The parameter $r$ controls the local resolution of the SVM in the sense that smaller $r$ create smaller curvature of the margin. The capacity $C$ controls the amount of slack to allow for unclassified observations. A large $C$ would create a very rough and curved margin where $C$ close to zero makes the margin more smooth.

One of the guinea pig tests for a classification algorithm is the data described as "orange peel", i.e. when two groups of observations have similar means, their variance, however, being different. The classification results in this case are presented in Fig. 20.15. An SVM with a radial basis kernel is highly suitable for such a kind of data.

Another popular non-linear test is the classification of "spiral data". We generated two spirals with the distance between them equal 1.0 that span over $3 \pi$ radian. The SVM was chosen with $r=0.1$ and $C=10 / n$. The SVM was able to separate the classes without an error if noise with parameters $\varepsilon_{i} \sim N\left(0,0.1^{2} \mathcal{I}\right)$ was injected into the pure spiral data (Fig. 20.16). Obviously, both the "orange peel" and the "spiral data" are not linearly separable.

## Solution of the SVM Classification Problem

The standard SVM optimisation problem (20.29), which is a quadratic optimisation problem, is usually solved by means of quadratic programming (QP). This

Fig. 20.15 SVM
classification results for the "orange peel" data, $n=200$, $d=2, n_{-1}=n_{+1}=100$, $x_{+1, i} \sim N\left((0,0)^{\top}, 2^{2} \mathcal{I}\right)$, $x_{-1, i} \sim N\left((0,0)^{\top}, 0.5^{2} \mathcal{I}\right)$ with SVM parameters $r=0.5$ and $C=20 / 200$
Q MVAsvmOrangePeel

Fig. 20.16 SVM
classification results for the noisy spiral data. The spirals spread over $3 \pi$ radian; the distance between the spirals equals 1.0. $d=2$, $n_{-1}=n_{+1}=100, n=200$. The noise was injected with the parameters $\varepsilon_{i} \sim N\left(0,0.1^{2} \mathcal{I}\right)$. The separation is perfect with SVM parameters $r=0.1$ and $C=10 / 200$
Q MVAsvmSpiral


technique, however, is notorious for (i) its bad scaling properties (the time required to solve the problem is proportional to $n^{3}$, where $n$ is the number of observations), (ii) implementation difficulty and (iii) enormous memory requirements. With the QP technique the whole kernel matrix of the size $n \times n$ has to be fit in the memory, which, assuming that each variable takes up 10 bytes of memory, will require $10 \times n \times n$ bytes. This means that 1 million observation (which is not unusual for practical applications such as credit scoring) will require 12,000 TBytes (terabytes)
or $10,000,000$ MBytes of operating memory to store. With a typical size of the computer memory of 512 MBytes no more than around 5,000 observations can be processed. Thus, the main emphasis in designing new algorithms was made on using special properties of SVMs to speed up the solution and reduce memory requirements.

## Scoring Companies

For our illustration we selected the largest bankrupt companies with the capitalisation of no less than 1 billion USD. The dataset used in this work is from the Credit reform database provided by the Research Data Center (RDC) of the Humboldt Universität zu Berlin. It contains financial information from about 20,000 solvent and 1,000 insolvent German companies. The period spans from 1996 to 2002 and in the case of the insolvent companies the information is gathered 2 years before the insolvency took place. The last annual report of a company before it goes bankrupt receives the indicator $y=1$ and for the rest (solvent) companies $y=-1$.

We are given 28 variables, i.e. cash, inventories, equity, EBIT, number of employees, and branch code. From the original data, we create common financial indicators which are denoted as $x 1, \ldots, x 25$. These ratios can be grouped into four categories such as profitability, leverage, liquidity, and activity.

Obviously, data for the year of 1996 are missing and we will exclude them for further calculations. In order to reduce the effect of the outliers on the results, all observations that exceeded the upper limit of IQ (Inter-quartile range) or the lower limit of IQ were replaced with these values. To demonstrate how performance changes, we will use the Accounts Payable (AP) turnover (named X24) and ratio of Operating Income (OI) and Total Asset (TA) (named $X 3$ ). We choose randomnly 50 solvent and 50 insolvent companies. The statistical description of financial ratios is summarized in Table 20.5.

Keep in mind that different kernels will influence performance. We will use one of the most common ones, the isotropic Gaussian kernel. Triangles and circles in Fig. 20.17 represent successful and failing companies from the training set, respectively. The coloured background corresponds to different score values $f$. The more blue the area, the higher the score and the greater the probability of default. Most successful companies lying in the red area have positive profitability and a reasonable activity.

Figure 20.17 presents the classification results for an SVM using isotropic Gaussian kernel with $\sigma=100$ and the fixed capacity $C=1$. With given priors, the SVM has trouble classifying between solvent and insolvent company. The radial base $\sigma$, which determines the minimum radius of a group, is too large. Notice that SVM do a poor job of distinguishing between groups even though most observations are used as support vector.

The applied SVMs differed in two aspects: (i) their capacity that is controlled by the coefficient $C$ in (20.28) and (ii) the complexity of classifier functions controlled in our case by the isotropic radial basis in the Gaussian kernel. In

Fig. 20.17 Ratings of companies in two dimensions. Low complexity of classifier functions with $\sigma=100$ and $C=1$. Percentage of misclassification is 0.43 Q MVAsvmSig100C1


Table 20.5 Descriptive statistics for financial ratios

| Ratio | $q_{0.05}$ | Med. | $q_{0.95}$ | IQR |
| :--- | ---: | :--- | :--- | :--- |
| OI/TA | -0.22 | 0.00 | 0.10 | 0.06 |
| AP/sales | 0.03 | 0.14 | 0.36 | 0.10 |

Fig. 20.18 Ratings of companies in two dimensions. The case of an average complexity of classifier functions with $\sigma=2$ and capacity is fixed at $C=1$. Percentage of misclassification is reduced to 0.27 Q MVAsvmSig2C1


Fig. 20.18 the value $\sigma$ is reduced to 2 while $C$ remains the same. SVM start recognising the difference between solvent and insolvent companies resulting in sharper cluster. Figure 20.19 demonstrate the effect of the changing capacity to the classification result. The optimisation of SVM parameters ( $C$ and $\sigma$ ) can be done


Fig. 20.19 Ratings of companies in two dimensions. High capacity $(C=200)$ with radial basis is fixed at $\sigma=0.5$. Percentage of misclassification is 0.10 Q MVAsvmSig 05 C 200


Fig. 20.20 Cumulative accuracy profile (CAP) curve
by using grid search method or an other advance algorithm the so-called Genetic Algorithm.

Figure 20.20 shows a Cumulative Accuracy Profile (CAP) curve which is particularly useful in that it simultaneously measures Type I and Type II errors. In statistical terms, the CAP curve represents the cumulative probability of default events for different percentiles of the risk score scale. Now, we introduce Accuracy Ratio (AR) derived from CAP curve for measuring and comparing the performance of credit risk model. Therefore, AR is defined as the ratio of the area between a model CAP curve and the random curve to the area between the perfect CAP curve
and the random CAP curve (see Fig. 20.20). Perfect classification is attained if the value of $A R$ is equal to one.

|  | SVM classification is done by mapping the data into feature space <br> and finding a separating hyperplane there |
| :--- | :--- |
|  | The support vectors are determined via a quadratic optimisation <br> problem |
|  | SVM produces highly non-linear classification boundaries |

### 20.5 Classification and Regression Trees

Classification and Regression Trees (CART) is a method of data analysis developed by a group of American statisticians (Breiman et al. , 1984). The aim of CART is to classify observations into a subset of known classes or to predict levels of regression functions. CART is a non-parametric tool which is designed to represent decision rules in a form of the so-called binary trees. Binary trees split a learning sample parallel to the coordinate axis and represent the resulting data clusters hierarchically starting from a root node for the whole learning sample itself and ending with relatively homogenous buckets of observations.

Regression trees are constructed in a similar way but the final buckets do not represent classes but rather approximations to an unknown regression functions at a particular point of the independent variable. In this sense regression trees are estimates via a non-parametric regression model. Here we provide an outlook of how decision trees are created, what challenges arise during practical applications and, of course, a number of examples will illustrate the power of CART.

## How Does CART Work?

Consider the example of how high risk patients (those who will not survive at least 30 days after a heart attack is admitted) were identified at San Diego Medical Center, University of California on the basis of initial 24-h data. A classification rule using at most three decisions (questions) is presented in Fig. 20.21. Left branches of the tree represent cases of positive answers, right branches-negative ones so that e.g. if minimum systolic blood pressure over the last 24 h is less or equal 91 , then the

Fig. 20.21 Decision tree for low/high patients

patient belongs to the high risk group. In this example the dependant variable is binary: low risk (0) and high risk (1).

A different situation occurs when we are interested in the expected amount of days the patient will be able to survive. The decision tree will probably change and the terminal nodes will now indicate a mean expected number of days the patient will survive. This situation describes a regression tree rather than a classification tree.

In a more formal setup let $Y$ be a dependent variable-binary or continuous and $X \in \mathbb{R}^{d}$. We are interested in approximating

$$
f(x)=\mathrm{E}(Y \mid X=x)
$$

For the definition of conditional expectations we refer to Sect. 4.2. CART estimates this function $f$ by a step function that is constructed via splits along the coordinate axis. An illustration is given in Fig. 20.22. The regression function $f(x)$ is approximated by the values of the step function. The splits along the coordinate axes are to be determined from the data.

The following simple one-dimensional example shows that the choice of splits points involves some decisions. Suppose that $f(x)=\boldsymbol{I}(x \in[0,1])+$ $2 \boldsymbol{I}(x \in[1,2])$ is a simple step function with a step at $x=1$. Assume now that one observes $Y_{i}=f\left(x_{i}\right)+\varepsilon_{i}, X_{i} \sim U[0,2], \varepsilon_{i} \sim N(0,1)$. By going through the $X$ data points as possible split points one sees that in the neighbourhood of $x=1$ one has two possibilities: one simply takes the $X_{i}$ left to 1 or the observation right to 1 . In order to make such splits unique one averages these neighbouring points.


Fig. 20.22 CART orthogonal splitting example where each colour corresponds to one cluster

## Impurity Measures

A more formal framework on how to split and where to split needs to be developed. Suppose there are $n$ observations in the learning sample and $n_{j}$ is the overall number of observations belonging to class $j, j=1, \ldots, J$. The class probabilities are:

$$
\begin{equation*}
\pi(j)=\frac{n_{j}}{n}, j=1, \ldots, J \tag{20.31}
\end{equation*}
$$

$\pi(j)$ is the proportion of observations belonging to a particular class. Let $n(t)$ be the number of observations at node $t$ and $n_{j}(t)$-the number of observations belonging to the $j$-th class at $t$. The frequency of the event that an observation of the $j$-th class falls into node $t$ is:

$$
\begin{equation*}
p(j, t)=\pi(j) \frac{n_{j}(t)}{n_{j}} \tag{20.32}
\end{equation*}
$$

The proportion of observations at $t$ are $p(t)=\sum_{j=1}^{J} p(j, t)$ the conditional probability of an observation to belong to class $j$ given that it is at node $t$ is:

$$
\begin{equation*}
p(j \mid t)=\frac{p(j, t)}{p(t)}=\frac{n_{j}(t)}{n(t)} \tag{20.33}
\end{equation*}
$$

Define now a degree of class homogeneity in a given node. This characteristican impurity measure $i(t)$-will represent a class homogeneity indicator for a given tree node and hence will help to find optimal splits. Define an impurity function $\iota(t)$ which is determined on $\left(p_{1}, \ldots, p_{J}\right) \in[0,1]^{J}$ with $\sum_{j=1}^{J} p_{j}=1$ so that:

1. $\iota$ has a unique maximum at point $\left(\frac{1}{J}, \frac{1}{J}, \ldots, \frac{1}{J}\right)$;
2. $\iota$ has a unique minimum at points $(1,0,0, \ldots, 0),(0,1,0, \ldots, 0), \ldots$, $(0,0,0, \ldots, 1)$;
3. $\iota$ is a symmetric function of $p_{1}, \ldots, p_{J}$

Each function satisfying these conditions is called an impurity function. Given $\iota$, define the impurity measure $i(t)$ for a node $t$ as:

$$
\begin{equation*}
i(t)=\iota\{p(1 \mid t), p(2 \mid t), \ldots, p(J \mid t)\} \tag{20.34}
\end{equation*}
$$

Denote an arbitrary data split by $s$, then for a given node $t$ which we will call a parent node two child nodes described in Fig. 20.23 arise: $t_{L}$ and $t_{R}$ representing observations meeting and not meeting the split criterion $s$. A fraction $p_{L}$ of data from $t$ falls to the left child node and $p_{R}=1-p_{L}$ is the share of data in $t_{R}$.

A quality measure of how well split $s$ works is:

$$
\begin{equation*}
\Delta i(s, t)=i(t)-p_{L} i\left(t_{L}\right)-p_{R} i\left(t_{R}\right) \tag{20.35}
\end{equation*}
$$

The higher the value of $\Delta i(s, t)$ the better split we have since data impurity is reduced. In order to find an optimal split $s$ it is natural to maximise $\Delta i(s, t)$. Note that in (20.35) for different splits $s$, the value of $i(t)$ remains constant, hence it is equivalent to find

$$
\begin{aligned}
s^{*} & =\underset{s}{\operatorname{argmax}} \Delta i(s, t) \\
& =\underset{s}{\operatorname{argmax}}\left\{-p_{L} i\left(t_{L}\right)-p_{R} i\left(t_{R}\right)\right\} \\
& =\underset{s}{\operatorname{argmax}}\left\{p_{L} i\left(t_{L}\right)+p_{R} i\left(t_{R}\right)\right\}
\end{aligned}
$$

Fig. 20.23 Parent and child nodes hierarchy

where $t_{L}$ and $t_{R}$ are implicit functions of $s$. This splitting procedure is repeated until one arrives at a minimal bucket size. Classes are then assigned to terminal nodes using the following rule:

$$
\begin{equation*}
\text { If } p(j \mid t)=\max _{i} p(i \mid t) \text {, then } j^{*}(t)=j \tag{20.36}
\end{equation*}
$$

If the maximum is not unique, then $j^{*}(t)$ is assigned randomly to those classes for which $p(i \mid t)$ takes its maximum value. The crucial question is of course to define an impurity function $i(t)$. A natural definition of impurity is via a variance measure: Assign 1 to all observations at node $t$ belonging to class $j$ and 0 to others. A sample variance estimate for node $t$ observations is $p(j \mid t)\{1-p(j \mid t)\}$.

Summing over all $J$ classes we obtain the Gini index:

$$
\begin{equation*}
i(t)=\sum_{j=1}^{J} p(j \mid t)\{1-p(j \mid t)\}=1-\sum_{j=1}^{J} p^{2}(j \mid t) \tag{20.37}
\end{equation*}
$$

The Gini index is an impurity function $\iota\left(p_{1}, \ldots, p_{J}\right), p_{j}=p(j \mid t)$. It is not hard see that the Gini index is a convex function. Since $p_{L}+p_{R}=1$, we get:

$$
\begin{aligned}
i\left(t_{L}\right) p_{L}+i\left(t_{R}\right) p_{R} & =\iota\left\{p\left(1 \mid t_{L}\right), \ldots, p\left(J \mid t_{L}\right)\right\} p_{L}+\iota\left\{p\left(1 \mid t_{R}\right), \ldots, p\left(J \mid t_{R}\right)\right\} p_{R} \\
& \leq \iota\left\{p_{L} p\left(1 \mid t_{L}\right)+p_{R} p\left(1 \mid t_{R}\right), \ldots, p_{L} p\left(J \mid t_{L}\right)+p_{R} p\left(J \mid t_{R}\right)\right\}
\end{aligned}
$$

where inequality becomes an equality in case $p\left(j \mid t_{L}\right)=p\left(j \mid t_{R}\right), j=1, \ldots, J$.
Recall that

$$
\frac{p\left(j, t_{L}\right)}{p(t)}=\frac{p\left(t_{L}\right)}{p(t)} \cdot \frac{p\left(j, t_{L}\right)}{p\left(t_{L}\right)}=p_{L} p\left(j \mid t_{L}\right)
$$

and since

$$
p(j \mid t)=\frac{p\left(j, t_{L}\right)+p\left(j, t_{R}\right)}{p(t)}=p_{L} p\left(j \mid t_{L}\right)+p_{R} p\left(j \mid t_{R}\right)
$$

we can conclude that

$$
\begin{equation*}
i\left(t_{L}\right) p_{L}+i\left(t_{R}\right) p_{R} \leq i(t) \tag{20.38}
\end{equation*}
$$

Hence each variant of data split leads to $\Delta i(s, t)>0$ unless $p\left(j \mid t_{R}\right)=$ $p\left(j \mid t_{L}\right)=p(j \mid t)$, i.e. when no split decreases class heterogeneity.

Impurity measures can be defined in a number of different ways, for practical applications the so-called twoing rule can be considered. Instead of maximising impurity change at a particular node, the twoing rule tries to balance as if the learning sample had only two classes. The reason for such an algorithm is that such
a decision rule is able to distinguish observations between general factors on top levels of the tree and take into account specific data characteristics at lower levels.

If $S=\{1, \ldots, J\}$ is the set of learning sample classes, divide it into two subsets

$$
S_{1}=\left\{j_{1}, \ldots, j_{n}\right\}, \text { and } S_{2}=S \backslash S_{1}
$$

All observations belonging to $S_{1}$ get dummy class 1 , and the rest dummy class 2 . The next step is to calculate $\Delta i(s, t)$ for different $s$ as if there were only two (dummy) classes. Since actually $\Delta i(s, t)$ depends on $S_{1}$, the value $\Delta i\left(s, t, S_{1}\right)$ is maximised. Now apply a two-step procedure: first, find $s^{*}\left(S_{1}\right)$ maximising $\Delta i\left(s, t, S_{1}\right)$ and second, find a superclass $S_{1}^{*}$ maximising $\Delta i\left\{s^{*}\left(S_{1}\right), t, S_{1}\right\}$. In other words the idea of twoing is to find a combination of superclasses at each node that maximises the impurity increment for two classes.

This method provides one big advantage: it finds the so-called strategic nodes, i.e. nodes filtering observations in the way that they are different to the maximum feasible extent. Although applying the twoing rule may seem to be desirable especially for data with a big number of classes, another challenge arises: computational speed. Let's assume that the learning sample has $J$ classes, then a set $S$ can be split into $S_{1}$ and $S_{2}$ by $2^{J-1}$ ways. For 11 classes data this will create more than 1,000 combinations. Fortunately the following result helps to reduce drastically the amount of computations.

It can be proven (Breiman et al., 1984) that in a classification task with two classes and impurity measure $p(1 \mid t) p(2 \mid t)$ for an arbitrary split $s$ a superclass $S_{1}(s)$ is determined by:

$$
\begin{gather*}
S_{1}(s)=\left\{j: p\left(j \mid t_{L}\right) \geq p\left(j \mid t_{R}\right)\right\}, \\
\max _{S_{1}} \Delta i\left(s, t, S_{1}\right)=\frac{p_{L} p_{R}}{4}\left\{\sum_{j=1}^{J}\left|p\left(j \mid t_{L}\right)-p\left(j \mid t_{R}\right)\right|\right\}^{2} \tag{20.39}
\end{gather*}
$$

Hence the twoing rule can be applied in practice as well as Gini index, although the first criterion works a bit slower.

## Gini Index and Twoing Rule in Practice

In this section we look at practical issues of using these two rules. Consider a learning dataset from Salford Systems with 400 observations characterising automobiles: their make, type, colour, technical parameters, age etc. The aim is to build a decision tree splitting different cars by their characteristics based on feasible relevant parameters. The classification tree constructed using the Gini index is given in Fig. 20.24.


Fig. 20.24 Classification tree constructed by Gini index


Fig. 20.25 Classification tree constructed by twoing index

A particular feature here is that at each node observations belonging to one make are filtered out, i.e. observations with most striking characteristics are separated. As a result a decision tree is able to pick out automobile makes quite easily.

The twoing rule based tree Fig. 20.25 for the same data is different. Instead of specifying particular car makes at each node, application of the twoing rule results in strategic nodes, i.e. questions which distinguish between different car classes to the maximum extent. This feature can be vital when high-dimensional datasets with a big number of classes are processed.

## Optimal Size of a Decision Tree

Up to now we were interested in determining the best split $s^{*}$ at a particular node. The next and perhaps more important question is how to determine the optimal tree size, i.e. when to stop splitting. If each terminal node has only class homogenous dataset, then every point of the learning sample can be flawlessly classified using this maximum tree. But can be such an approach fruitful?

The maximum tree is a case of overspecification. Some criterion is required to stop data splitting. Since tree building is dependent on $\Delta i(s, t)$, a criterion is to stop data splitting if

$$
\begin{equation*}
\Delta i(s, t)<\bar{\beta} \tag{20.40}
\end{equation*}
$$

where $\bar{\beta}$ is some threshold value.
The value of $\bar{\beta}$ is to be chosen in a subjective way and this is unfortunately a drawback. Empirical simulations show that the impurity increment is frequently non-monotone, that is why even for small $\bar{\beta}$ the tree may be underparametrised. Setting even smaller values for $\bar{\beta}$ will probably remedy the situation but at the cost of tree overparametrisation.

Another way to determine the adequate shape of a decision tree is to demand a minimum number of observations $\bar{N}$ (bucked size) at each terminal node. A disadvantage is that if at terminal node $t$ the number of observations is higher

$$
\begin{equation*}
N(t)>\bar{N} \tag{20.41}
\end{equation*}
$$

then this node is also being split as data are still not supposed to be clustered well enough.

## Cross-Validation for Tree Pruning

Cross-validation is a procedure which uses the bigger data part as a training set and the rest as a test set. Then the process is looped so that different parts of the data become learning and training set, so that at the end each datapoint was employed both as a member of test and learning sets. The aim of this procedure is to extract maximum information from the learning sample especially in the situations of data scarceness.

The procedure is implemented in the following way. First, the learning sample is randomly divided into $V$ parts. Using the training set from the union of $(V-1)$ subsets a decision tree is constructed while the test set is used to verify the tree quality. This procedure is looped over all possible subsets.

Unfortunately for small values of $V$ cross-validation estimates can be unstable since each iteration a cluster of data is selected randomly and the number of
iterations itself is relatively small, thus the overall estimation result is somewhat random. Nowadays cross-validation with $V=10$ is an industry standard and for many applications a good balance between computational complexity and statistical precision.

## Cost-Complexity Function and Cross-Validation

Another method taken into account is tree complexity, i.e. the number of terminal nodes. The maximum tree will get a penalty for its big size, on the other hand it will be able to make perfect in-sample predictions. Small trees will, of course, get lower penalty for their size but their prediction abilities are limited. Optimisation procedure based on such a trade-off criterion could determine a good decision tree.

Define the internal misclassification error of an arbitrary observation at node $t$ as $e(t)=1-\max _{j} p(j \mid t)$, define also $E(t)=e(t) p(t)$. Then internal misclassification tree error is $E(T)=\sum_{t \in \tilde{T}} E(t)$ where $\tilde{T}$ is a set of terminal nodes. The estimates are called internal because they are based solely on the learning sample. It may seem that $E(T)$ as a tree quality measure is sufficient but unfortunately it is not so. Consider the case of the maximum tree, here $E\left(T_{\mathrm{MAX}}\right)=$ 0 , i.e. the tree is of best configuration.

For any subtree $T\left(\leq T_{\mathrm{MAX}}\right)$ define the number of terminal nodes $|\tilde{T}|$ as a measure of its complexity. The following cost-complexity function can be used:

$$
\begin{equation*}
E_{\alpha}(T)=E(T)+\alpha|\tilde{T}| \tag{20.42}
\end{equation*}
$$

where $\alpha \geq 0$ is a complexity parameter and $\alpha|\tilde{T}|$ is a cost component. The more complex the tree (high number of terminal nodes) the lower is $E(T)$ but at the same time the higher is the penalty $\alpha|\tilde{T}|$ and vice versa.

The number of subtrees of $T_{\text {MAX }}$ is finite. Hence pruning of $T_{\text {MAX }}$ leads to creation of a subtree sequence $T_{1}, T_{2}, T_{3}, \ldots$ with a decreasing number of terminal nodes.

An important question is if a subtree $T \leq T_{\mathrm{MAX}}$ for a given $\alpha$ minimising $E_{\alpha}(T)$ always exists and whether it is unique?

In Breiman et al. (1984) it is shown that for $\forall \alpha \geq 0$ there exists an optimal tree $T(\alpha)$ in the sense that

1. $E_{\alpha}\{T(\alpha)\}=\min _{T \leq T_{\mathrm{MAX}}} E_{\alpha}(T)=\min _{T \leq T_{\mathrm{MAX}}}\{E(T)+\alpha|\tilde{T}|\}$
2. if $E_{\alpha}(T)=E_{\alpha}\{T(\alpha)\}$ then $T(\alpha) \leq T$.

This result is a proof of existence, but also a proof of uniqueness: consider another subtree $T^{\prime}$ so that $T$ and $T^{\prime}$ both minimise $E_{\alpha}$ and are not nested, then $T(\alpha)$ does not exist in accordance with second condition.

The idea of introducing cost-complexity function at this stage is to check only a subset of different subtrees of $T_{\text {MAX }}$ : optimal subtrees for different values of $\alpha$. The starting point is to define the first optimal subtree in the sequence so that $E\left(T_{1}\right)=$ $E\left(T_{\mathrm{MAX}}\right)$ and the size of $T_{1}$ is minimum among other subtrees with the same cost level. To get $T_{1}$ out of $T_{\mathrm{MAX}}$ for each terminal node of $T_{\mathrm{MAX}}$ it is necessary to verify the condition $E(t)=E\left(t_{L}\right)+E\left(t_{R}\right)$ and if it is fulfilled—node $t$ is pruned. The process is looped until no extra pruning is available-the resulting tree $T(0)$ becomes $T_{1}$.

Define a node $t$ as an ancestor of $t^{\prime}$ and $t^{\prime}$ as descendant of $t$ if there is a connected path down the tree leading from $t$ to $t^{\prime}$. Consider Fig. 20.26 where nodes $t_{4}, t_{5}, t_{8}, t_{9}, t_{10}$ and $t_{11}$ are descendants of $t_{2}$ while nodes $t_{6}$ and $t_{7}$ are not descendants of $t_{2}$ although they are positioned lower since it is not possible to connect them with a path from $t_{2}$ to these nodes without engaging $t_{1}$. Nodes $t_{4}, t_{2}$ and $t_{1}$ are ancestors of $t_{9}$ and $t_{3}$ is not ancestor of $t_{9}$.

Define the branch $T_{t}$ of the tree $T$ as a subtree based on node $t$ and all its descendants. An example is given in Fig. 20.27. Pruning a branch $T_{t}$ from a tree $T$ means deleting all descendant nodes of $t$. Denote the transformed tree as $T-T_{t}$. Pruning the branch $T_{t_{2}}$ results in the tree described in Fig. 20.28.

For any branch $T_{t}$ define the internal misclassification estimate as:

$$
\begin{equation*}
E\left(T_{t}\right)=\sum_{t^{\prime} \in \tilde{T}_{t}} E\left(t^{\prime}\right) \tag{20.43}
\end{equation*}
$$

where $\tilde{T}_{t}$ is the set of terminal nodes of $T_{t}$. Hence for an arbitrary node $t$ of $T_{t}$ :


Fig. 20.26 Decision tree hierarchy

Fig. 20.27 The branch $T_{t_{2}}$ of the original tree $T$


Fig. 20.28 $T-T_{t_{2}}$ the pruned tree $T$


$$
\begin{equation*}
E(t)>E\left(T_{t}\right) \tag{20.44}
\end{equation*}
$$

Consider now the cost-complexity misclassification estimate for branches or single nodes. Define for a single node $\{t\}$ :

$$
\begin{equation*}
E(\{t\})=E(t)+\alpha \tag{20.45}
\end{equation*}
$$

and for a branch:

$$
\begin{equation*}
E_{\alpha}\left(T_{t}\right)=E\left(T_{t}\right)+\alpha\left|\tilde{T}_{t}\right| \tag{20.46}
\end{equation*}
$$

When $E_{\alpha}\left(T_{t}\right)<E_{\alpha}(\{t\})$ the branch $T_{t}$ is preferred to a single node $\{t\}$ according to cost-complexity. For some $\alpha$ both (20.45) and (20.46) will become equal. This critical value of $\alpha$ can be determined from:

$$
\begin{equation*}
E_{\alpha}\left(T_{t}\right)<E_{\alpha}(\{t\}) \tag{20.47}
\end{equation*}
$$

which is equivalent to

$$
\begin{equation*}
\alpha<\frac{E(t)-E\left(T_{t}\right)}{\left|\tilde{T}_{t}\right|-1} \tag{20.48}
\end{equation*}
$$

where $\alpha>0$ since $E(t)>E\left(T_{t}\right)$

To obtain the next member of the subtrees sequence, i.e. $T_{2}$ out of $T_{1}$ a special node called weak link is determined. For this purpose a function $g_{1}(t), t \in T_{1}$ is defined as

$$
g_{1}(t)=\left\{\begin{array}{l}
\frac{E(t)-E\left(T_{t}\right)}{\left|\tilde{T}_{t}\right|-1}, t \notin \tilde{T}_{1}  \tag{20.49}\\
+\infty, t \in \tilde{T}_{1}
\end{array}\right.
$$

Node $\bar{t}_{1}$ is a weak link in $T_{1}$ if

$$
\begin{equation*}
g_{1}\left(\bar{t}_{1}\right)=\min _{t \in T_{1}} g_{1}(t) \tag{20.50}
\end{equation*}
$$

and a new value for $\alpha_{2}$ is defined as

$$
\begin{equation*}
\alpha_{2}=g_{1}\left(\bar{t}_{1}\right) \tag{20.51}
\end{equation*}
$$

A new tree $T_{2} \prec T_{1}$ in the sequence is obviously defined by pruning the branch $T_{\bar{t}_{1}}$, i.e.

$$
\begin{equation*}
T_{2}=T_{1}-T_{\bar{t}_{1}} \tag{20.52}
\end{equation*}
$$

The process is looped until root node $\left\{t_{0}\right\}$-the final member of sequence-is reached. When there are multiple weak links detected, for instance $g_{k}\left(\bar{t}_{k}\right)=g_{k}\left(\bar{t}_{k}^{\prime}\right)$, then both branches are pruned, i.e. $T_{k+1}=T_{k}-T_{\bar{t}_{k}}-T_{\bar{t}_{k}^{\prime}}$.

In this way it is possible to get the sequence of optimal subtrees $T_{\mathrm{MAX}} \succ T_{1} \succ$ $T_{2} \succ T_{3} \succ \cdots \succ\left\{t_{0}\right\}$ for which it is possible to prove that the sequence $\left\{\alpha_{k}\right\}$ is increasing, i.e. $\alpha_{k}<\alpha_{k+1}, k \geq 1$ and $\alpha_{1}=0$. For $k \geq 1: \alpha_{k} \leq \alpha<\alpha_{k+1}$ and $T(\alpha)=T\left(\alpha_{k}\right)=T_{k}$.

Practically this tells us how to implement the search algorithm. First, the maximum tree $T_{\text {MAX }}$ is taken, then $T_{1}$ is found and a weak link $\bar{t}_{1}$ is detected and branch $T_{\bar{t}_{1}}$ is pruned off, $\alpha_{2}$ is calculated and the process is continued.

When the algorithm is applied to $T_{1}$, the number of pruned nodes is usually quite significant. For instance, consider the following typical empirical evidence (see Table 20.6). When the trees become smaller, the difference in the number of terminal nodes also gets smaller.

Finally, it is worth mentioning that the sequence of optimally pruned subtrees is a subset of trees which might be constructed using direct method of internal misclassification estimator minimisation given a fixed number of terminal nodes.

Table 20.6 Typical pruning speed

| Tree | $T_{1}$ | $T_{2}$ | $T_{3}$ | $T_{4}$ | $T_{5}$ | $T_{6}$ | $T_{7}$ | $T_{8}$ | $T_{9}$ | $T_{10}$ | $T_{11}$ | $T_{12}$ | $T_{13}$ |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| $\left\|\tilde{T}_{k}\right\|$ | 71 | 63 | 58 | 40 | 34 | 19 | 10 | 9 | 7 | 6 | 5 | 2 | 1 |

Consider an example of tree $T(\alpha)$ with 7 terminal nodes, then there is no other subtree $T$ with 7 terminal nodes having lower $E(T)$. Otherwise

$$
E_{\alpha}(T)=E(T)+7 \alpha<E_{\alpha}\{T(\alpha)\}=\min _{T \leq T_{\mathrm{MAX}}} E_{\alpha}(T)
$$

which is impossible by definition.
Applying the method of $V$-fold cross-validation to the sequence $T_{\mathrm{MAX}} \succ T_{1} \succ$ $T_{2} \succ T_{3} \succ \cdots \succ\left\{t_{0}\right\}$, an optimal tree is determined. On the other hand it is frequently pointed out that choice of tree with minimum value of $E^{\mathrm{CV}}(T)$ is not always adequate since $E^{\mathrm{CV}}(T)$ is not too robust, i.e. there is a whole range of values $E^{\mathrm{CV}}(T)$ satisfying $E^{\mathrm{CV}}(T)<E_{\mathrm{MIN}}^{\mathrm{CV}}(T)+\varepsilon$ for small $\varepsilon>0$. Moreover, when $V<N$ a simple change of random generator seed will definitely result in changed values of $\left|\tilde{T}_{k}\right|$ minimising $\hat{E}\left(T_{K}\right)$. Hence a so-called one standard error empirical rule is applied which states that if $T_{k_{0}}$ is the tree minimising $E^{\mathrm{CV}}\left(T_{k_{0}}\right)$ from the sequence $T_{\mathrm{MAX}} \succ T_{1} \succ T_{2} \succ T_{3} \succ \cdots \succ\left\{t_{0}\right\}$, then a value $k_{1}$ and a correspondent tree $T_{k_{1}}$ are selected so that

$$
\begin{equation*}
\underset{k_{1}}{\operatorname{argmax}} \hat{E}\left(T_{k_{1}}\right) \leq \hat{E}\left(T_{k_{0}}\right)+\sigma\left\{\hat{E}\left(T_{k_{0}}\right)\right\} \tag{20.53}
\end{equation*}
$$

where $\sigma(\cdot)$ denotes sample estimate of standard error and $\hat{E}(\cdot)$-the relevant sample estimators.

The dotted line in Fig. 20.29 shows the area where the values of $\hat{E}\left(T_{k}\right)$ only slightly differ from $\min _{\mid \tilde{\sim}} \mid \leq\left(T_{k}\right)$. The left edge which is roughly equivalent to 16 $\left|\tilde{T}_{k}\right|$ terminal nodes shows the application of one standard error rule. The use of one


Fig. 20.29 The example of relationship between $\hat{E}\left(T_{k}\right)$ and number of terminal nodes
standard error rule allows not only to achieve more robust results but also to get trees of lower complexity given the error comparable with $\min \hat{E}\left(T_{k}\right)$.

## Regression Trees

Up to now we concentrate on classification trees. Although regression trees share a similar logical framework, there are some differences which need to be addressed. The important difference between classification and regression trees is the type of dependent variable $Y$. When $Y$ is discrete, a decision tree is called a classification tree, a regression tree is a decision tree with a continuous dependent variable.

Gini index and twoing rule discussed in previous sections assume that the number of classes is finite and hence introduce some measures based mainly on $p(j \mid t)$ for arbitrary class $j$ and node $t$. But since in case of continuous dependent variable there are no more classes, this approach cannot be used anymore unless groups of continuous values are effectively substituted with artificial classes. Since there are no classes anymore-how can be the maximum regression tree determined? Analogously with discrete case, absolute homogeneity can be then described only after some adequate impurity measure for regression trees is introduced.

Recall the idea of Gini index, then it becomes quite natural to use the variance as impurity indicator. Since for each node data variance can be easily computed, then splitting criterion for an arbitrary node $t$ can be written as

$$
\begin{equation*}
s^{*}=\underset{s}{\operatorname{argmax}}\left[p_{L} \operatorname{var}\left\{t_{L}(s)\right\}+p_{R} \operatorname{var}\left\{t_{R}(s)\right\}\right] \tag{20.54}
\end{equation*}
$$

where $t_{L}$ and $t_{R}$ are emerging child nodes which are, of course, directly dependent on the choice of $s^{*}$.

Hence the maximum regression tree can be easily defined as a structure where each node has only the same predicted values. It is important to point out that since continuous data have much higher chances to take different values comparing with discrete ones, the size of maximum regression tree is usually very big.

When the maximum regression tree is properly defined, it is then of no problem to get an optimally size tree. Like with classification trees, maximum regression tree is usually supposed to be upwardly pruned with the help of cost-complexity function and cross-validation. That is why the majority of results presented above is applied to regression trees as well.

Fig. 20.30 Decision tree for bankruptcy dataset: Gini index, $\bar{N}=30$
a MVACARTBan1


## Bankruptcy Analysis

This section provides a practical study on bankruptcy data involving decision trees. A dataset with 84 observations representing different companies is constituted by three variables:

- net income to total assets ratio
- total liabilities to total assets ratio
- company status ( -1 if bankrupt and 1 if not)

The data is from SEC (2004).
The goal is to predict and describe the company status given the two primary financial ratios. Since no additional information like the functional form of possible relationship is available, the use of a classification tree is an active alternative.

The tree given in Fig. 20.30 was constructed using the Gini index and a $\bar{N}=$ 30 constraint, i.e. the number of points in each of the terminal nodes can not be more than 30 . Numbers in parentheses displayed on terminal nodes are observation quantities belonging to Class 1 and Class -1 .

If we loose the constraint to $\bar{N}=10$, the decision rule changes, see Fig. 20.31. How exactly did the situation change? Consider the Class 1 terminal nodes of the tree on Fig. 20.30. The first one contains 21 observations and thus was split for $\bar{N}=10$. When it was split two new nodes of different classes emerged and for both of them the impurity measure has decreased.

We may conclude that $\bar{N} \approx 10$ is a good choice and analysing the tree produced we can state that for this particular example the net income to total assets $\left(X_{1}\right)$ ratio appears to be an important class indicator. The successful classification ratio


Fig. 20.31 Decision tree for bankruptcy dataset: Gini index, $\bar{N}=10 Q_{\text {MVACARTBan2 }}$


Fig. 20.32 Successful classification ratio dynamic over the number of terminal nodes: crossvalidation Q MVAbancrupcydis
dynamic over the number of terminal nodes is shown in Fig. 20.32. It is chosen by cross-validation method.

For this example with relatively small sample size we construct two maximum trees-using the Gini and twoing rules, see Figs. 20.33 and 20.34. Looking at both decision trees we see that the choice of impurity measure is not so important as the right choice of tree size.



| $\hookrightarrow$ | CART is a tree based method splitting the data sequentially into a <br> binary tree |
| :--- | :--- |
| $\hookrightarrow$ | CART determined the nodes by minimising an impurity measure at <br> each mode |
|  | CART is non-parametric: <br>  <br> When no data structure hypotheses are available, non-parametric <br> analysis becomes the single effective data mining tool. CART is a <br> flexible nonparametric data mining tool |
| $\hookrightarrow$ | CART does not require variables to be selected in advance: <br>  <br>  <br>  <br>  <br> From a learning sample CART will automatically select the most <br> significant ones |
| $\hookrightarrow$ | CART is very efficient in computational terms: <br> Although all possible data splits are analysed, the CART architec- <br> ture is flexible enough to do all of them quickly |
|  | CART is robust to the effect of outliers: |
|  | Due to data-splitting nature of decision rules creation it is possible <br> to distinguish between datasets with different characteristics and <br> hence to neutralise outliers in separate nodes |
|  | CART can use any combination of continuous and categorical data: <br> Researchers are no longer limited to a particular class of data and <br> will be able to capture more real-life examples |

### 20.6 Boston Housing

Coming back to the Boston Housing data set, we compare the results of EPP on the original data $\mathcal{X}$ and the transformed data $\hat{\mathcal{X}}$ motivated in Sect. 1.9. So we exclude $X_{4}$ (indicator of Charles River) from the present analysis.

The aim of this analysis is to see from a different angle whether our proposed transformations yield more normal distributions and whether it will yield data with less outliers. Both effects will be visible in our projection pursuit analysis.

We first apply the Jones and Sibson index to the non-transformed data with 50 randomly chosen 13-dimensional directions. Figure 20.35 displays the results in the following form. In the lower part, we see the values of the Jones and Sibson index. It should be constant for 13-dimensional normal data. We observe that this is clearly not the case. In the upper part of Fig. 20.35 we show the standard normal density as a green curve and two densities corresponding to two extreme index values. The red, slim curve corresponds to the maximal value of the index among the 50 projections. The blue curve, which is close to the normal, corresponds to the minimal value of

Fig. 20.35 Projection Pursuit with the Sibson-Jones index with 13 original variables Q MVAppsib


Fig. 20.36 Projection Pursuit with the Sibson-Jones index with 13 transformed variables Q MVAppsib

the Jones and Sibson index. The corresponding values of the indices have the same colour in the lower part of Fig. 20.35. Below the densities, a jitter plot shows the distribution of the projected points $\alpha^{\top} x_{i}(i=1, \ldots, 506)$. We conclude from the outlying projection in the red distribution that several points are in conflict with the normality assumption.

Figure 20.36 presents an analysis with the same design for the transformed data. We observe in the lower part of the figure values that are much lower for the Jones and Sibson index (by a factor of 10) with lower variability which suggests that the transformed data is closer to the normal. ("Closeness" is interpreted here in the sense of the Jones and Sibson index.) This is confirmed by looking to the upper part of Fig. 20.36 which has a significantly less outlying structure than in Fig. 20.35.

### 20.7 Exercises

Exercise 20.1 Calculate the Simplicial Depth for the Swiss bank notes data set and compare the results to the univariate medians. Calculate the Simplicial Depth again for the genuine and counterfeit bank notes separately.

Exercise 20.2 Construct a configuration of points in $\mathbb{R}^{2}$ such that $x_{\text {med, } j}$ from (20.2) is not in the "centre" of the scatterplot.

Exercise 20.3 Apply the SIR technique to the US companies data with $Y=$ market value and $X=$ all other variables. Which directions do you find?
Exercise 20.4 Simulate a data set with $X \sim N_{4}\left(0, I_{4}\right), Y=\left(X_{1}+3 X_{2}\right)^{2}+\left(X_{3}-\right.$ $\left.X_{4}\right)^{4}+\varepsilon$ and $\varepsilon \sim N\left(0,(0.1)^{2}\right)$. Use SIR and SIR II to find the EDR directions.

Exercise 20.5 Apply the Projection Pursuit technique on the Swiss bank notes data set and compare the results to the PC analysis and the Fisher discriminant rule.

Exercise 20.6 Apply the SIR and SIR II technique on the car data set in Table 22.3 with $Y=$ price.

Exercise 20.7 Generate four regions on the two-dimensional unit square by sequentially cutting parallel to the coordinate axes. Generate 100 two-dimensional Uniform random variables and label them according to their presence in the above regions. Apply the CART algorithm to find the regions bound and to classify the observations.

Exercise 20.8 Modify Exercise 20.7 by defining the regions as lying above and below the main diagonal of the unit square. Make a CART analysis and comment on the complexity of the tree.

Exercise 20.9 Apply the SVM with different radial basis parameter $r$ and different capacity parameter $c$ in order to separate two circular datasets. This example is often called the Orange Peel exercise and involves two Normal distributions $N\left(\mu, \Sigma_{i}\right), i=1,2$, with covariance matrices $\Sigma_{1}=2 \mathcal{I}_{2}$ and $\Sigma_{2}=0.5 \mathcal{I}_{2}$.
Exercise 20.10 The noisy spiral data set consists of two intertwining spirals that need to be separated by a non-linear classification method. Apply the SVM with different radial basis parameter $r$ and capacity parameter $c$ in order to separate the two spiral datasets.

Exercise 20.11 Apply the SVM to separate the bankrupt from the surviving (profitable) companies using the profitability and leverage ratios given in the Bankruptcy data set in Table 22.21.

## Part IV Appendix

## Chapter 21 <br> Symbols and Notations

## Basics

| X, Y | Random variables or vectors |  |
| :---: | :---: | :---: |
| $X_{1}, X_{2}, \ldots, X_{p}$ | Random variables |  |
| $X=\left(X_{1}, \ldots, X_{p}\right)^{\top}$ | Random vector |  |
| $X \sim$ 。 | $X$ has distribution $\cdot$ |  |
| $\mathcal{A}, \mathcal{B}$ | Matrices | 53 |
| $\Gamma, \Delta$ | Matrices | 60 |
| $\mathcal{X}, \mathcal{Y}$ | Data matrices | 81 |
| $\Sigma$ | Covariance matrix | 80 |
| $1_{n}$ | Vector of ones $(\underbrace{1, \ldots, 1})^{\top}$ | 54 |
| $0_{n}$ | Vector of zeros $(\underbrace{0, \ldots, 0})^{\top}$ | 54 |
| I(.) | $n \text {-times }$ <br> Indicator function, i.e. for a set $M$ is $\boldsymbol{I}=1$ on $\boldsymbol{M}, \boldsymbol{I}=0$ otherwise |  |
| i | $\sqrt{-1}$ |  |
| $\Rightarrow$ | Implication |  |
| $\Leftrightarrow$ | Equivalence |  |
| $\approx$ | Approximately equal |  |
| $\otimes$ | Kronecker product |  |
| iff | if and only if, equivalence |  |

## Mathematical Abbreviations

| $\operatorname{tr}(\mathcal{A})$ | Trace of matrix $\mathcal{A}$ |
| :--- | :--- |
| $\operatorname{hull}\left(x_{1}, \ldots, x_{k}\right)$ | Convex hull of points $\left\{x_{1}, \ldots, x_{k}\right\}$ |
| $\operatorname{diag}(\mathcal{A})$ | Diagonal of matrix $\mathcal{A}$ |
| $\operatorname{rank}(\mathcal{A})$ | Rank of matrix $\mathcal{A}$ |
| $\operatorname{det}(\mathcal{A})$ | Determinant of matrix $\mathcal{A}$ |
| $C(\mathcal{A})$ | Column space of matrix $\mathcal{A}$ |

## Samples

| $x, y$ | Observations of $X$ and $Y$ |  |
| :--- | :--- | :--- |
| $x_{1}, \ldots, x_{n}=\left\{x_{i}\right\}_{i=1}^{n}$ | Sample of $n$ observations of $X$ |  |
| $\mathcal{X}=\left\{x_{i j}\right\}_{i=1, \ldots, n ; j=1, \ldots, p}$ | $(n \times p)$ data matrix of observations of $X_{1}, \ldots, X_{p}$ | 81 |
|  | or of $X=\left(X_{1}, \ldots, X_{p}\right)^{T}$ |  |
| $x_{(1)}, \ldots, x_{(n)}$ | The order statistic of $x_{1}, \ldots, x_{n}$ | 5 |
| $\mathcal{H}$ | Centering matrix $\mathcal{H}=\mathcal{I}_{n}-n^{-1} 1_{n} 1_{n}^{\top}$ | 90 |

## Densities and Distribution Functions

| $f(x)$ | Density of $X$ |
| :--- | :--- |
| $f(x, y)$ | Joint density of $X$ and $Y$ |
| $f_{X}(x), f_{Y}(y)$ | Marginal densities of $X$ and $Y$ |
| $f_{X_{1}}\left(x_{1}\right), \ldots, f_{X_{p}}\left(x_{2}\right)$ | Marginal densities of $X_{1}, \ldots, X_{p}$ |
| $\hat{f}_{h}(x)$ | Histogram or kernel estimator of $f(x)$ |
| $F(x)$ | Distribution function of $X$ |
| $F(x, y)$ | Joint distribution function of $X$ and $Y$ |
| $F_{X}(x), F_{Y}(y)$ | Marginal distribution functions of $X$ and $Y$ |
| $F_{X_{1}}\left(x_{1}\right), \ldots, f_{X_{p}}\left(x_{p}\right)$ | Marginal distribution functions of $X_{1}, \ldots, X_{p}$ |
| $\varphi(x)$ | Density of the standard normal distribution |
| $\Phi(x)$ | Standard normal distribution function |
| $\varphi_{X}(t)$ | Characteristic function of $X$ |
| $m_{k}$ | $k$-th moment of $X$ |
| $\kappa_{j}$ | Cumulants or semi-invariants of $X$ |

## Moments

| $\mathrm{E} X, \mathrm{E} Y$ | Mean values of random variables or vectors $X$ and $Y$ | 80 |
| :--- | :--- | ---: |
| $\sigma_{X Y}=\operatorname{Cov}(X, Y)$ | Covariance between random variables $X$ and $Y$ | 80 |
| $\sigma_{X X}=\operatorname{Var}(X)$ | Variance of random variable $X$ | 80 |
| $\rho_{X Y}=\frac{\operatorname{Cov}(X, Y)}{\sqrt{\operatorname{Var}(X) \operatorname{Var}(Y)}}$ | Correlation between random variables $X$ and $Y$ | 84 |
| $\Sigma_{X Y}=\operatorname{Cov}(X, Y)$ | Covariance between random vectors $X$ <br> i.e., $\operatorname{Cov}(X, Y)=E(X-E X)(Y-E Y)^{\top}$ |  |
| $\Sigma_{X X}=\operatorname{Var}(X)$ | Covariance matrix of the random vector $X$ |  |

## Empirical Moments

$\bar{x}=\frac{1}{n} \sum_{i=1}^{n} x_{i}$
$s_{X Y}=\frac{1}{n} \sum_{i=1}^{n}\left(x_{i}-\bar{x}\right)\left(y_{i}-\bar{y}\right)$
$s_{X X}=\frac{1}{n} \sum_{i=1}^{n}\left(x_{i}-\bar{x}\right)^{2}$
$\mathcal{R}=\left\{r_{X_{i} X_{j}}\right\}=\mathcal{D}^{-1 / 2} \mathcal{S D} \mathcal{D}^{-1 / 2}$

Empirical covariance of random variables $X$ and $Y$80 sampled by $\left\{x_{i}\right\}_{i=1, \ldots, n}$ and $\left\{y_{i}\right\}_{i=1, \ldots, n}$ Empirical variance of random variable $X$ sampled by 80 $\left\{x_{i}\right\}_{i=1, \ldots, n}$
Empirical correlation of $X$ and $Y$ 84

Empirical covariance matrix of $X_{1}, \ldots, X_{p}$ or of the 80,90 random vector $X=\left(X_{1}, \ldots, X_{p}\right)^{\top}$
Average of $X$ sampled by $\left\{x_{i}\right\}_{i=1, \ldots, n}$

Empirical correlation matrix of $X_{1}, \ldots, X_{p}$ or of the 84,91 random vector $X=\left(X_{1}, \ldots, X_{p}\right)^{\top}$

## Distributions

| $\varphi(x)$ | Density of the standard normal distribution |
| :--- | :--- |
| $\Phi(x)$ | Distribution function of the standard normal distribution |
| $N(0,1)$ | Standard normal or Gaussian distribution |
| $N\left(\mu, \sigma^{2}\right)$ | Normal distribution with mean $\mu$ and variance $\sigma^{2}$ |
| $N_{p}(\mu, \Sigma)$ | $p$-Dimensional normal distribution with mean $\mu$ and <br> covariance matrix $\Sigma$ |


| $\xrightarrow{\mathcal{L}}$ | Convergence in distribution <br> CLT | Central Limit Theorem |
| :--- | :--- | ---: |
| $\chi_{p}^{2}$ | $\chi^{2}$ distribution with $p$ degrees of freedom | 143 |
| $\chi_{1-\alpha ; p}^{2}$ | $1-\alpha$ quantile of the $\chi^{2}$ distribution with $p$ degrees of |  |
| $t_{n}$ | $t$-Distribution with $n$ degrees of freedom |  |
| $t_{1-\alpha / 2 ; n}$ | $1-\alpha / 2$ quantile of the $t$-distribution with $n$ d.f. |  |
| $F_{n, m}$ | $F$-Distribution with $n$ and $m$ degrees of freedom |  |
| $F_{1-\alpha ; n, m}$ | $1-\alpha$ quantile of the $F$-distribution with $n$ and $m$ degrees <br> of freedom |  |
| $T_{p, n}^{2}$ | Hotelling $T^{2}$-distribution with $p$ and $n$ degrees of <br> freedom |  |

## Chapter 22 <br> Data

All data sets are available on the Springer webpage or at the authors' home pages.

### 22.1 Boston Housing Data

The Boston housing data set was collected by Harrison and Rubinfeld (1978). It comprise 506 observations for each census district of the Boston metropolitan area. The data set was analysed in Belsley, Kuh, and Welsch (1980).
$X_{1}$ : Per capita crime rate,
$X_{2}$ : Proportion of residential land zoned for large lots,
$X_{3}$ : Proportion of nonretail business acres,
$X_{4}$ : Charles River ( 1 if tract bounds river, 0 otherwise),
$X_{5}$ : Nitric oxides concentration,
$X_{6}$ : Average number of rooms per dwelling,
$X_{7}$ : Proportion of owner-occupied units built prior to 1940,
$X_{8}$ : Weighted distances to five Boston employment centers,
$X_{9}$ : Index of accessibility to radial highways,
$X_{10}$ : Full-value property tax rate per $\$ 10,000$,
$X_{11}$ : Pupil/teacher ratio,
$X_{12}: \quad 1000(B-0.63)^{2} \boldsymbol{I}(B<0.63)$ where $B$ is the proportion of African American,
$X_{13}$ : \% lower status of the population,
$X_{14}$ : Median value of owner-occupied homes in $\$ 1,000$.

### 22.2 Swiss Bank Notes

Six variables measured on 100 genuine and 100 counterfeit old Swiss 1000 -franc bank notes. The data stem from Flury and Riedwyl (1988). The columns correspond to the following six variables.
$X_{1}$ : Length of the bank note,
$X_{2}$ : Height of the bank note, measured on the left,
$X_{3}$ : Height of the bank note, measured on the right,
$X_{4}$ : Distance of inner frame to the lower border,
$X_{5}$ : Distance of inner frame to the upper border,
$X_{6}$ : Length of the diagonal.

Observations 1-100 are the genuine bank notes and the other 100 observations are the counterfeit bank notes.

### 22.3 Car Data

The car data set (Chambers, Cleveland, Kleiner \& Tukey, 1983) consists of 13 variables measured for 74 car types. The abbreviations in this section are as follows:

| $X_{1}:$ | P | Price, |
| :--- | :--- | :--- |
| $X_{2}:$ | M | Mileage (in miles per gallon), |
| $X_{3}:$ | R 78 | Repair record 1978 (rated on a 5-point scale; 5 best, 1 worst), |
| $X_{4}:$ | R 77 | Repair record 1977 (scale as before), |
| $X_{5}:$ | H | Headroom (in inches), |
| $X_{6}:$ | R | Rear seat clearance (distance from front seat back to rear seat, in inches), |
| $X_{7}:$ | Tr | Trunk space (in cubic feet), |
| $X_{8}:$ | W | Weight (in pound), |
| $X_{9}:$ | L | Length (in inches), |
| $X_{10}:$ | T | Turning diameter (clearance required to make a U-turn, in feet), |
| $X_{11}:$ | D | Displacement (in cubic inches), |
| $X_{12}:$ | G | Gear ratio for high gear, |
| $X_{13}:$ | C | Company headquarter (1 for USA, 2 for Japan, 3 for Europe). |

### 22.4 Classic Blue Pullovers Data

This is a data set consisting of ten measurements of four variables. The story: A textile shop manager is studying the sales of "classic blue" pullovers over ten periods. He uses three different marketing methods and hopes to understand his sales as a fit of these variables using statistics. The variables measured are
$X_{1}$ : Numbers of sold pullovers,
$X_{2}$ : Price (in EUR),
$X_{3}$ : Advertisement costs in local newspapers (in EUR),
$X_{4}$ : Presence of a sales assistant (in hours per period).

### 22.5 US Companies Data

The data set consists of measurements for 79 US companies. The abbreviations in this section are as follows:

```
X1: A Assets (USD),
X2: S Sales (USD),
X 3: MV Market value (USD),
X4: P Profits (USD),
X5: CF Cash flow (USD),
X6: E Employees.
```


### 22.6 French Food Data

The data set consists of the average expenditures on food for several different types of families in France (manual workers $=$ MA, employees $=E M$, managers $=C A$ ) with different numbers of children (2, 3, 4 or 5 children). The data is taken from Lebart, Morineau, and Fénelon (1982).

### 22.7 Car Marks

The data are averaged marks for 23 car types from a sample of 40 persons. The marks range from 1 (very good) to 6 (very bad) like German school marks. The variables are:
$X_{1}$ : A Economy,
$X_{2}$ : B Service,
$X_{3}$ : C Non-depreciation of value,
$X_{4}$ : D Price, Mark 1 for very cheap cars,
$X_{5}$ : E Design,
$X_{6}: ~ F ~ S p o r t y ~ c a r, ~$
$X_{7}: \quad \mathrm{G}$ Safety,
$X_{8}: ~ H \quad$ Easy handling.

### 22.8 French Baccalauréat Frequencies

The data consist of observations of 202,100 baccalauréats from France in 1976 and give the frequencies for different sets of modalities classified into regions. For a reference see Bouroche and Saporta (1980). The variables (modalities) are:
$X_{1}$ : A Philosophy-Letters,
$X_{2}$ : B Economics and Social Sciences,
$X_{3}$ : C Mathematics and Physics,
$X_{4}$ : D Mathematics and Natural Sciences,
$X_{5}$ : E Mathematics and Techniques,
$X_{6}: \mathrm{F}$ Industrial Techniques,
$X_{7}: \quad \mathrm{G}$ Economic Techniques,
$X_{8}: \quad \mathrm{H}$ Computer Techniques.

### 22.9 Journaux Data

This is a data set that was created from a survey completed in the 1980s in Belgium questioning people's reading habits. They were asked where they live (10 regions comprised of 7 provinces and 3 regions around Brussels) and what kind of newspaper they read on a regular basis. The 15 possible answers belong to 3
classes: Flemish newspapers (first letter v), French newspapers (first letter f) and both languages (first letter b).

| $X_{1}:$ | WaBr | Walloon Brabant |
| :--- | :--- | :--- |
| $X_{2}:$ | Brar | Brussels area |
| $X_{3}:$ | Antw | Antwerp |
| $X_{4}:$ | FlBr | Flemish Brabant |
| $X_{5}:$ | OcFl | Occidental Flanders |
| $X_{6}:$ | OrFl | Oriental Flanders |
| $X_{7}:$ | Hain | Hainaut |
| $X_{8}:$ | Lièg | Liège |
| $X_{9}:$ | Limb | Limburg |
| $X_{10}:$ | Luxe | Luxembourg |

### 22.10 US Crime Data

This is a data set consisting of 50 measurements of 7 variables. It states for 1 year (1985) the reported number of crimes in the 50 states of the US classified according to 7 categories ( $X_{3}-X_{9}$ ).
$X_{1}$ : Land area (land)
$X_{2}$ : Population 1985 (popu 1985)
$X_{3}$ : Murder (murd)
$X_{4}$ : Rape
$X_{5}$ : Robbery (robb)
$X_{6}$ : Assault (assa)
$X_{7}$ : Burglary (burg)
$X_{8}$ : Larcery (larc)
$X_{9}$ : Autothieft (auto)
$X_{10}$ : US states region number (reg)
$X_{11}$ : US states division number (div)

| Division numbers |  | Region numbers |  |
| :--- | :---: | :--- | :--- |
| New England | 1 | Northeast | 1 |
| Mid Atlantic | 2 | Midwest | 2 |
| E N Central | 3 | South | 3 |
| W N Central | 4 | West | 4 |
| S Atlantic | 5 |  |  |
| E S Central | 6 |  |  |
| W S Central | 7 |  |  |
| Mountain | 8 |  |  |
| Pacific | 9 |  |  |

### 22.11 Plasma Data

In Olkin and Veath (1980), the evolution of citrate concentration in the plasma is observed at three different times of day, $X_{1}(8 \mathrm{am}), X_{2}(11 \mathrm{am})$ and $X_{3}(3 \mathrm{pm})$, for two groups of patients. Each group follows a different diet.
$X_{1}: 8 \mathrm{am}$
$X_{2}: 11 \mathrm{am}$
$X_{3}: 3 \mathrm{pm}$

### 22.12 WAIS Data

Morrison (1990) compares the results of four subtests of the Wechsler Adult Intelligence Scale (WAIS) for two categories of people: in group one are $n_{1}=37$ people who do not present a senile factor, group two are those ( $n_{2}=12$ ) presenting a senile factor.

WAIS subtests:

| $X_{1}:$ | Information |
| :--- | :--- |
| $X_{2}:$ | Similarities |
| $X_{3}:$ | Arithmetic |
| $X_{4}:$ | Picture completion |

### 22.13 ANOVA Data

The yields of wheat have been measured in 30 parcels which have been randomly attributed to 3 lots prepared by one of 3 different fertilisers A, B and C.
$X_{1}$ : Fertiliser A
$X_{2}$ : Fertiliser B
$X_{3}$ : Fertiliser C

### 22.14 Timebudget Data

In Volle (1985), we can find data on 28 individuals identified according to sex, country where they live, professional activity and matrimonial status, which indicates the amount of time each person spent on ten categories of activities over 100 days ( $100 \cdot 24 \mathrm{~h}=2,400 \mathrm{~h}$ total in each row) in the year 1976.
$\begin{array}{lll}X_{1}: & \text { prof : } & \text { Professional activity } \\ X_{2}: & \text { tran : } & \text { Transportation linked to professional activity } \\ X_{3}: & \text { hous : } & \text { Household occupation } \\ X_{4}: & \text { kids : } & \text { Occupation linked to children } \\ X_{5}: & \text { shop : } & \text { Shopping } \\ X_{6}: & \text { pers : } & \text { Time spent for personal care } \\ X_{7}: & \text { eat : } & \text { Eating } \\ X_{8}: & \text { slee : } & \text { Sleeping } \\ X_{9}: & \text { tele : } & \text { Watching television } \\ X_{10}: & \text { leis : } & \text { Other leisures }\end{array}$
maus: Active men in the USA
waus: Active women in the USA
wnus: Nonactive women in the USA
mmus: Married men in USA
wmus: Married women in USA
msus: Single men in USA
wsus: Single women in USA
mawe: Active men from Western countries
wawe: Active women from Western countries
wnwe: Nonactive women from Western countries
mmwe: Married men from Western countries
wmwe: Married women from Western countries
mswe: Single men from Western countries
wswe: Single women from Western countries
mayo: Active men from Yugoslavia
wayo: Active women from Yugoslavia
wnyo: Nonactive women from Yugoslavia
mmyo: Married men from Yugoslavia
wmyo: Married women from Yugoslavia
msyo: $\quad$ Single men from Yugoslavia
wsyo: Single women from Yugoslavia
maes: Active men from Eastern countries
waes: Active women from Eastern countries
wnes: Nonactive women from Eastern countries
mmes: Married men from Eastern countries
wmes: Married women from Eastern countries
mses: Single men from Eastern countries
wses: Single women from Eastern countries

### 22.15 Geopol Data

This data set contains a comparison of 41 countries according to 10 different political and economic parameters.

| $X_{1}:$ | popu | Population |
| :--- | :--- | :--- |
| $X_{2}:$ | giph | Gross Internal Product per habitant |
| $X_{3}:$ | ripo | Rate of increase of the population |
| $X_{4}:$ | rupo | Rate of urban population |
| $X_{5}:$ | rlpo | Rate of illiteracy in the population |
| $X_{6}:$ | rspo | Rate of students in the population |
| $X_{7}:$ | eltp | Expected lifetime of people |
| $X_{8}:$ | rnnr | Rate of nutritional needs realised |
| $X_{9}:$ | nunh | Number of newspapers and magazines per 1,000 habitants |
| $X_{10}:$ | nuth | Number of television per 1,000 habitants |


| AFS | South Africa | DAN | Denmark | MAR | Marocco |
| :--- | :--- | :--- | :--- | :--- | :--- |
| ALG | Algeria | EGY | Egypt | MEX | Mexico |
| BRD | Germany | ESP | Spain | NOR | Norway |
| GBR | Great Britain | FRA | France | PER | Peru |
| ARS | Saudi Arabia | GAB | Gabun | POL | Poland |
| ARG | Argentine | GRE | Greece | POR | Portugal |
| AUS | Australia | HOK | Hong Kong | SUE | Sweden |
| AUT | Austria | HON | Hungary | SUI | Switzerland |
| BEL | Belgium | IND | India | THA | Tailand |
| CAM | Cameroon | IDO | Indonesia | URS | USSR |
| CAN | Canada | ISR | Israel | USA | USA |
| CHL | Chile | ITA | Italia | VEN | Venezuela |
| CHN | China | JAP | Japan | YOU | Yugoslavia |
| CUB | Cuba | KEN | Kenia |  |  |

### 22.16 US Health Data

This is a data set consisting of 50 measurements of 13 variables. It states for 1 year (1985) the reported number of deaths in the 50 states of the US classified according to 7 categories.
$X_{1}$ : Land area (land)
$X_{2}$ : Population 1985 (popu)
$X_{3}$ : Accident (acc)
$X_{4}$ : Cardiovascular (card)
$X_{5}$ : Cancer (canc)
$X_{6}$ : Pulmonar (pul)
$X_{7}$ : Pneumonia flu (pnue)
$X_{8}: \quad$ Diabetis (diab)
$X_{9}$ : Liver (liv)
$X_{10}$ : Doctors (doc)
$X_{11}$ : Hospitals (hosp)
$X_{12}$ : US states region number (r)
$X_{13}$ : US states division number (d)

| Division numbers |  | Region numbers |  |
| :--- | :--- | :--- | :--- |
| New England | 1 | Northeast | 1 |
| Mid Atlantic | 2 | Midwest | 2 |
| E N Central | 3 | South | 3 |
| W N Central | 4 | West | 4 |
| S Atlantic | 5 |  |  |
| E S Central | 6 |  |  |
| W S Central | 7 |  |  |
| Mountain | 8 |  |  |
| Pacific | 9 |  |  |

### 22.17 Vocabulary Data

This example of the evolution of the vocabulary of children can be found in Bock (1975). Data are drawn from test results on file in the Records Office of the Laboratory School of the University of Chicago. They consist of scores, obtained from a cohort of pupils from the eighth through eleventh grade levels, on alternative forms of the vocabulary section of the Cooperative Reading Test. It provides the following scaled scores shown for the sample of 64 subjects (the origin and units are fixed arbitrarily).

### 22.18 Athletic Records Data

This data set provides data on Men's athletic records for 55 countries in 1984 Olympic Games.

### 22.19 Unemployment Data

This data set provides unemployment rates in all federal states of Germany in November 2005.

### 22.20 Annual Population Data

The data shows yearly average population rates for Former territory of the Federal Republic of Germany incl. Berlin-West (given in 1,000 inhabitants).

### 22.21 Bankruptcy Data I

The data are the profitability, leverage, and bankruptcy indicators for 84 companies.
The data set contains information on 42 of the largest companies that filed for protection against creditors under Chap. 11 of the US Bankruptcy Code in 20012002 after the stock market crash of 2000 . The bankrupt companies were matched with 42 surviving companies with the closest capitalisations and the same US industry classification codes available through the Division of Corporate Finance of the Securities and Exchange Commission (SEC, 2004).

The information for each company was collected from the annual reports for 1998-1999 (SEC, 2004), i.e. 3 years prior to the defaults of the bankrupt companies. The following data set contains profitability and leverage ratios calculated, respectively, as the ratio of net income (NI) and total assets (TA) and the ratio of total liabilities (TL) and total assets (TA).

### 22.22 Bankruptcy Data II

Altman (1968), quoted by Morrison (1990), reports financial data on 66 banks.
$\mathrm{X} 1=($ Working capital)/(total assets)
$\mathrm{X} 2=($ Retained earnings) $/$ (total assets)
$\mathrm{X} 3=($ Earnings before interest and taxes $) /$ (total assets)
$\mathrm{X} 4=($ Market value equity $) /($ book value of total liabilities)
X5 $=($ Sales $) /($ total assets $)$
The first 33 observations correspond to bankrupt banks and the last 33 for solvent banks as indicated by the last columns: values of $y$

## References

ALLBUS. (2006). Germany general social survey 1980-2004.
Altman, E. (1968). Financial ratios, discriminant analysis and the prediction of corporate bankruptcy. The Journal of Finance, 23, 589-609.
Andrews, D. (1972). Plots of high-dimensional data. Biometrics, 28, 125-136.
Backhaus, K., Erichson, B., Plinke, W., \& Weiber, R. (1996). Multivariate analysemethoden. Berlin: Springer.
Bartlett, M. S. (1939). A note on tests of significance in multivariate analysis. Proceedings of the Cambridge Philosophical Society, 35, 180-185.
Bartlett, M. S. (1954). A note on multiplying factors for various chi-squared approximations. Journal of the Royal Statistical Society: Series B, 16, 296-298.
Belsley, D. A., Kuh, E., \& Welsch, R. E. (1980). Regression diagnostics. New York: Wiley.
Blæsild, P., \& Jensen, J. (1981). Multivariate distributions of hyperbolic type. In Statistical distributions in scientific work-Proceedings of the NATO Advanced Study Institute held at the Università degli studi di Trieste (Vol. 4, pp. 45-66).
Bock, R. D. (1975). Multivariate statistical methods in behavioral research. New York: Mc GrawHill.
Bouroche, J.-M., \& Saporta, G. (1980). L’analyse des données. Paris: Presses Universitaires de France.
Breiman, L. (1973). Statistics: With a view towards application. Boston: Houghton Mifflin Company.
Breiman, L., Friedman, J. H., Olshen, R., \& Stone, C. J. (1984). Classification and regression trees. Belmont: Wadsworth.
Chambers, J. M., Cleveland, W. S., Kleiner, B., \& Tukey, P. A. (1983). Graphical methods for data analysis. Boston: Duxbury Press.
Chen, Y., Härdle, W., \& Jeong, S.-O. (2008). Nonparametric risk management with generalized hyperbolic distributions. Journal of the American Statistical Association, 103, 910-923.
Chernoff, H. (1973). Using faces to represent points in k-dimensional space graphically. Journal of the American Statistical Association, 68, 361-368.
Cook, R. D., \& Weisberg, S. (1991). Comment on "sliced inverse regression for dimension reduction". Journal of the American Statistical Association, 86(414), 328-332.
Dillon, W. R., \& Goldstein, M. (1984). Multivariate analysis. New York: Wiley.
Duan, N., \& Li, K.-C. (1991). Slicing regression: A link-free regression method. Annals of Statistics, 19(2), 505-530.
Efron, B., Hastie, T., Johnstone, I., \& Tibshirani, R. (2004). Least angle regression (with discussion). Annals of Statistics, 32(2), 407-499.

Embrechts, P., McNeil, A. J., \& Straumann, D. (1999, May). Correlation and dependence in risk management: Properties and pitfalls. RISK Magazine, 69-71.
EUNITE. (2001). Electricity load forecast competition of the European network on intelligent technologies for smart adaptive systems. http://neuron.tuke.sk/competition/
Everitt, B., \& Dunn, G. (1998). Applied multivariate data analysis. London: Edward Arnold.
Fahrmeir, L., \& Hamerle, A. (1984). Multivariate statistische verfahren. Berlin: De Gruyter.
Fang, K. T., Kotz, S., \& Ng, K. W. (1990). Symmetric multivariate and related distributions. London: Chapman and Hall.
Fengler, M. R., Härdle, W., \& Villa, C. (2003). The dynamics of implied volatilities: A common principal components approach. Review of Derivative Research, 6, 179-202.
Flury, B. (1988). Common principle components analysis and related multivariate models. New York: Wiley.
Flury, B., \& Gautschi, W. (1986). An algorithm for simultaneous orthogonal transformation of several positive definite symmetric matrices to nearly diagonal form. SIAM Journal on Scientific and Statistical Computing, 7, 169-184.
Flury, B., \& Riedwyl, H. (1988). Multivariate statistics: A practical approach. Cambridge: Cambridge University Press.
Franke, J., Härdle, W., \& Hafner, C. (2011). Introduction to statistics of financial markets (3rd ed.). Heidelberg: Springer.
Friedman, J. H., Hastie, T., \& Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. Journal of Statistical Software, 33(1), 1-22.
Friedman, J. H., \& Stuetzle, W. (1981). Projection pursuit classification. Unpublished manuscript.
Friedman, J. H., \& Tukey, J. W. (1974). A projection pursuit algorithm for exploratory data analysis. IEEE Transactions on Computers, C 23, 881-890.
Gale, D., Kuhn, H. W., \& Tucker, A. W. (1951). Linear programming and the theory of games. In T. C. Koopmans (Ed.), Activity analysis of production and allocation (pp. 317-329). New York: John Wiley and Sons.
Gibbins, R. (1985). Canonical analysis. A review with application in ecology. Berlin: Springer.
Giri, N. C. (1996). Multivariate statistical analysis. New York: Marcel Dekker.
Gosset, W. S. (1908). The probable error of a mean. Biometrika, 6, 1-25.
Graham, M., \& Kennedy, J. (2003). Using curves to enhance parallel coordinate visualisations. In Proceedings of seventh international conference on information visualization, 2003 (IV 2003) (pp. 10-16).
Hall, P. (1992). The bootstrap and edgeworth expansion. Statistical Series. New York: Springer.
Hall, P., \& Li, K.-C. (1993). On almost linearity of low dimensional projections from high dimensional data. Annals of Statistics, 21(2), 867-889.
Härdle, W. (1991). Smoothing techniques, with implementations in S. New York: Springer.
Härdle, W., Hautsch, N., \& Overbeck, L. (2009). Applied quantitative finance (2nd ed.). Heidelberg: Springer.
Härdle, W., Müller, M., Sperlich, S., \& Werwatz, A. (2004). Non- and semiparametric models. Heidelberg: Springer.
Härdle, W., \& Scott, D. (1992). Smoothing by weighted averaging of rounded points. Computational Statistics, 7, 97-128.
Harrison, D., \& Rubinfeld, D. L. (1978). Hedonic prices and the demand for clean air. Journal of Environment Economics \& Management, 5, 81-102.
Hoaglin, W., Mosteller, F., \& Tukey, J. (1983). Understanding robust and exploratory data analysis. New York: Wiley.
Hodges, J. L., \& Lehman, E. L. (1956). The efficiency of some non-parametric competitors of the $t$-test. Annals of Mathematical Statistics, 27, 324-335.
Hotelling, H. (1935). The most predictable criterion. Journal of Educational Psychology, 26, 139-142.
Hotelling, H. (1953). New light on the correlation coefficient and its transform. Journal of the Royal Statistical Society, Series B, 15, 193-232.
Huber, P. (1985). Projection pursuit. Annals of Statistics, 13(2), 435-475.

Inselberg, A. (1985). A goodness of fit test for binary regression models based on smoothing methods. The Visual Computer, 1, 69-91.
Johnson, R. A., \& Wichern, D. W. (1998). Applied multivariate analysis (4th ed.). Englewood Cliffs, NJ: Prentice Hall.
Jones, M. C., \& Sibson, R. (1987). What is projection pursuit? (with discussion). Journal of the Royal Statistical Society, Series A, 150(1), 1-36.
Kaiser, H. F. (1985). The varimax criterion for analytic rotation in factor analysis. Psychometrika, 23, 187-200.
Kendall, K., \& Stuart, S. (1977). Distribution theory. In The advanced theory of statistics (Vol. 1). London: Griffin.
Klinke, S., \& Polzehl, J. (1995). Implementation of kernel based indices in XGobi. Discussion paper 47, SFB 373. Humboldt-University of Berlin.
Kruskal, J. B. (1965). Analysis of factorial experiments by estimating a monotone transformation of data. Journal of the Royal Statistical Society, Series B, 27, 251-263.
Kruskal, J. B. (1969). Toward a practical method which helps uncover the structure of a set of observations by finding the line tranformation which optimizes a new "index of condensation". In R. C. Milton \& J. A. Nelder (Eds.), Statistical computation (pp. 427-440). New York: Academic Press.
Kruskal, J. B. (1972). Linear transformation of multivariate data to reveal clustering. In R. N. Shepard, A. K. Romney, \& S. B. Nerlove (Eds.), Multidimensional scaling: Theory and applications in the behavioural sciences (Vol. 1, pp. 179-191). London: Seminar Press.
Lachenbruch, P. A., \& Mickey, M. R. (1968). Estimation of error rates in discriminant analysis. Technometrics, 10, 1-11.
Laplace, P.-S. (1774). Mémoire sur la probabilité des causes par les événements. Savants étranges, 6, 621-656.
Lawson, C., \& Hansen, R. (1974). Solving least square problems. Englewood Cliffs: Prentice Hall.
Lebart, L., Morineau, A. and Fénelon, J. P. (1982). Traitement des données statistiques, Dunod, Paris.
Lewin-Koh, N. (2006). Hexagon binnning. Technical Report.
Li, K.-C. (1991). Sliced inverse regression for dimension reduction (with discussion). Journal of the American Statistical Association, 86(414), 316-342.
Li, K.-C. (1992). On principal Hessian directions for data visualization and dimension reduction: Another application of Stein's lemma. Journal of the American Statistical Association, 87, 1025-1039.
Mardia, K. V., Kent, J. T., \& Bibby, J. M. (1979). Multivariate analysis. Duluth/London: Academic Press.
Meier, L., van de Geer, S., \& Bühlmann, P. (2008). The group lasso for logistic regression. Journal of the Royal Statistical Society, Series B, 70, 53-71.
Mercer, J. (1909). Functions of positive and negative type and their connection with the theory of integral equations. Philosophical Transactions of the Royal Society of London, 209, 415-446.
Morrison, D. (1990). Multivariate statistical methods (3rd ed., 495 pp.). New York: McGraw-Hill.
Muirhead, R. J. (1982). Aspects of multivariate statistics. New York: Wiley.
Nelsen, R. B. (1999). An introduction to copulas. New York: Springer.
Olkin, I., \& Veath, M. (1980). Maximum likelihood estimation in a two-way analysis with correlated errors in one classification. Biometrika, 68, 653-660.
Osborne, M., Presnell, B., \& Turlach, B. (2000). On the lasso and its dual. Journal of Computational and Graphical Statistics, 9(2), 319-337.
Parzen, E. (1962). On estimating of a probability density and mode. Annals of Mathematical Statistics, 35, 1065-1076.
Rosenblatt, M. (1956). Remarks on some nonparametric estimates of a density function. Annals of Mathematical Statistics, 27, 832-837.
Schott, J. R. (1994). Determining the dimensionality in sliced inverse regression. Journal of the American Statistical Association, 89(425), 141-148.

Scott, D. (1985). Averaged shifted histograms: Effective nonparametric density estimation in several dimensions. Annals of Statistics, 13, 1024-1040.
SEC (2004). Securities and exchange commission: Archive of historical documents. http://www. sec.gov/cgi-bin/srch-edgar
Shevade, S. K., \& Keerthi, S. S. (2003). A simple and efficient algorithm for gene selection using sparse logistic regression. Bioinformatics, 19, 2246-2253.
Silverman, B. W. (1986). Density estimation for statistics and data analysis. Monographs on statistics and applied probability (Vol. 26). London: Chapman and Hall.
Simon, N., Friedman, J. H., Hastie, T., \& Tibshirani, R. (2013). A sparse-group lasso. Journal of Computational and Graphical Statistics, 22(2), 231-245.
Sklar, A. (1959). Fonctions de répartition à $n$ dimensions et leurs marges. Publications de l'Institut de Statistique de L'Université de Paris, 8, 229-231.
Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society, Series B, 58, 267-288.
Tufte, E. (1983). The visual display of quantitative information. Cheshire: Graphics Press.
Vapnik, V. (1995). The nature of statistical learning theory. New York: Springer.
Volle, V. (1985). Analyse des Données. Paris: Economica.
Whittle, P. (1958). On the smoothing of probability density functions. Journal of the Royal Statistical Society, Series B, 55, 549-557.
Yuan, M., \& Lin, Y. (2006). Model selection and estimation in regression with grouped variables. Journal of the Royal Statistical Society, Series B, 68, 49-67.
Zou, H., \& Hastie, T. (2005). Regularization and variable selection via the elastic net. Journal of the Royal Statistical Society, Series B, 67, 301-320.

## Index

Actual error rate (AER), 417
Admissible, 413
Agglomerative techniques, 393
Allocation rules, 407
Analysis of variance (ANOVA), 72
Andrews' Curves, 29
Angle between two vectors, 72
Apparent error rate (APER), 417

Bayes discriminant rule, 413
Bernoulli distributions, 144
Best line, 308
Binary structure, 387
Binomial sampling, 273
Biplots, 439
Bootstrap, 176
sample, 179
Bootstrap sample, 179
Boston Housing, 40, 110, 242, 262, 346, 378, 400, 552
Boxplot, 4, 5
construction, 7

Canonical correlation, 443
analysis, 443
coefficient, 445
variable, 445
vector, 445
Capital asset pricing model (CAPM), 497
Cauchy distribution, 154
Centering matrix, 90
Central limit theorem (CLT), 143-145
Centroid, 396
Characteristic functions, 123, 130

Classic blue pullovers, 82
Cluster algorithms, 392
Cluster analysis, 385
Cobb-Douglas production function, 254
Cochran theorem, 192
Coefficient of determination, 96, 107
adjusted, 107
Column space, 74, 306
Common factors, 361
Common principal components analysis (CPCA), 342
Communality, 362
Complete linkage, 396
Computationally intensive techniques, 501
Concentration ellipsoid, 139
Conditional approximations, 188
Conditional covariance, 513
Conditional density, 119
Conditional distribution, 186
Conditional expectation, 126, 512, 513
Conditional pdf, 119
Confidence interval, 146
Confusion matrix, 417
Contingency tables, 264, 273, 425
Contrast, 227
Convex hull, 504
Copula, 120, 168
Copulae, 166
Correlation, 84
multiple, 188
Correspondence analysis, 425
Covariance, 80
Covariance matrix
decomposition, 320
properties, 124

Cramer-Rao, 208
Cramer-Rao lower bound, 206
Cramer-Wold, 131
Cumulant, 132
Cumulative distribution function (cdf), 118
Curse of dimensionality, 511

Data depth, 504
Degrees of freedom, 102
Dendrogram, 394
Density estimates, 11
Density functions, 118
Determinant, 56
Deviance, 269
Diagonal matrix, 54
Dice, 388
Discriminant analysis, 407
Discriminant rule, 408
in practice, 415
Dissimilarity of cars, 458
Distance
d, 68
Euclidean, 68
iso-distance curves, 69
matrix, 460
measures, 389
Distribution, 118
Draftman's plot, 21
Duality relations, 313
Duality theorem, 464

Effective dimension reduction directions, 511, 513
Effective dimension reduction space, 511
Efficient portfolio, 488
Eigenvalues, 57
Eigenvectors, 57
Elastic net, 297
Elliptical distribution, 196
Elliptically symmetric distribution, 511
Existence of a riskless asset, 492
Expected cost of misclassification, 409
Explained variation, 96
Exploratory projection pursuit, 506
Extremes, 7

F-spread, 7
F-test, 104
Faces, 23
Factorial
axis, 308, 309
method, 336
representation, 314, 317
variable, 308, 316
Factors, 306
analysis model, 359, 360
model, 367
scores, 376
Farthest Neighbor, 396
Fisher information, 209
Fisher information matrix, 207, 208
Fisher's linear discrimination function, 418
Five-number summary, 5
Flury faces, 24
Fourths, 6
French food expenditure, 340
Full model, 103

G-inverse, 57
non-uniqueness, 61
General multinormal distribution, 193
Gradient, 65
Group-building algorithm, 386

Heavy-tailed distributions, 149
Hessian, 65
Hexagon, 37
binning algorithm, 37, 50
plot, 38
Hierarchical algorithm, 393
Histograms, 11
Hotelling $T^{2}$-distribution, 193
Hyperbolic, 151
Idempotent matrix, 54
Identity matrix, 54
Independence copula, 122
Independent, 85, 119
Inertia, 315, 317
Information matrix, 208
Interpretation of the factors, 363
Interpretation of the principal components, 327
Invariance of scale, 363
Inverse, 56
Inverse regression, 511, 513

Jaccard, 388
Jacobian, 135
Jordan decomposition, 60, 61

Kernel
densities, 15
estimator, 15
Kulczynski, 388

Laplace distribution, 154
Lasso, 282
Likelihood function, 202
Likelihood ratio test, 214
Limit theorems, 142
Linear discriminant analysis (LDA), 412
Linear regression, 93
Linear transformation, 91
Link function, 511
Loadings, 361, 362
non-uniqueness, 364
Log-likelihood function, 202
Log-linear, 264
Logit models, 272

Mahalanobis distance, 412
Mahalanobis transformation, 93, 137, 138
Marginal densities, 119
Marketing strategies, 101
Maximum likelihood discriminant rule, 408
Maximum likelihood estimator, 202
MDS direction, 459
Mean-variance, 487, 488
Median, 5, 502
Metric methods, 459
Mixture model, 156
Model with interactions, 258
Moments, 123
Multidimentional scaling (MDS), 455
Multinormal, 139, 183
distribution, 137
Multivariate deneralised hyperbolic distribution, 160
Multivariate laplace distribution, 163
Multivariate median, 504
Multivariate t-distribution, 163, 196

Nearest neighbor, 395
Non-metric solution, 479
Nonexistence of a riskless asset, 491
Nonhomogeneous, 92
Nonmetric methods of MDS, 459
Norm of a vector, 72
Normal distribution, 203
Normal-inverse Gaussian, 151
Normalized principal components (NPCs), 335
Null space, 74

Odds, 273
Order statistics, 5
Orthogonal complement, 75
Orthogonal matrix, 54

Orthonormed, 309
Outliers, 3
Outside bars, 7

Parallel profiles, 238
Partitioned covariance matrix, 184
Partitioned matrixes, 66
PAV algorithm, 466, 484
Pearson chi-square, 269
Pearson chi-square test for independence, 269
Pool-adjacent violators algorithm, 466, 484
Portfolio
analysis, 487
choice, 487
Positive
definite, 62
definiteness, 65
or negative dependence, 22
semidefinite, 62, 90
Principal
axes, 70
components, 324
method, 372
in practice, 324
technique, 324
transformation, 321, 323
factors, 370
Principal components analysis (PCA), 512, 515
Profile
analysis, 238
method, 476
Projection
matrix, 75
pursuit, 505
pursuit regression, 508
vector, 511
Proximity between objects, 387
Proximity measure, 386
p-value, 269

Quadratic discriminant analysis, 414
Quadratic forms, 62
Quadratic response model, 253
Quality of the representations, 339

Randomized discriminant rule, 413
Rank, 55
Reduced model, 103
Rotations, 73, 374
Row space, 306
Russel and Rao (RR), 388

Sampling distributions, 142
Scatterplot matrix, 21
Separation line, 21
Similarity of objects, 387
Simple analysis of variance (ANOVA), 100
Simple Matching, 388, 389
Single linkage, 395
Singular normal distribution, 140
Singular value decomposition (SVD), 61, 313
Sliced inverse regression, 511, 516 algorithm, 512
Sliced inverse regression II, 513, 514, 516, 517
algorithm, 514
Solution
nonmetric, 483
Specific factors, 361
Specific variance, 362
Spectral decompositions, 60
Spherical distribution, 195
Standardized linear combinations (SLC), 320
Statistics, 142
Stimulus, 475
Student's $t$ with $n, 152$
Student's t-distribution, 94
Sum of squares, 102
Summary statistics, 89
Support vector machines (SVMs), 519
Swiss bank data, 4
Symmetric matrix, 54

T-test, 94
Tanimoto, 388
Three-way tables, 267
Total variation, 96
Trace, 56
Trade-off analysis, 476
Transformations, 135
Transpose, 56
Two factor method, 476

Unbiased estimator, 208
Uncorrelated factors, 361
Unexplained variation, 96
Unit vector, 72
Upper triangular matrix, 54

Variance explained by PCs, 333
Varimax
criterion, 374
method, 374
rotation method, 374

Ward clustering, 396
Wishart
density, 192
distribution, 191, 192

