

HRM: An R Package for Analysing High-dimensional Multi-factor Repeated Measures

by Martin Happ, Solomon W. Harrar, Arne C. Bathke

Abstract High-dimensional longitudinal data pose a serious challenge for statistical inference as many test statistics cannot be computed for high-dimensional data, or they do not maintain the nominal type-I error rate, or have very low power. Therefore, it is necessary to derive new inference methods capable of dealing with high dimensionality, and to make them available to statistics practitioners. One such method is implemented in the package **HRM** described in this article. This new method uses a similar approach as the Welch-Satterthwaite *t*-test approximation and works very well for high-dimensional data as long as the data distribution is not too skewed or heavy-tailed. The package also provides a GUI to offer an easy way to apply the methods.

Introduction

Repeated measures designs appear in many different situations. In clinical studies, physicians might be interested in the effect of different treatments over time. Typically, a univariate or multivariate measurement is observed over a period of time for each subject. The subjects can be separated into different disjoint groups and can often be assumed not to influence one another. That is, observations on different subjects are considered independent. We refer to factors whose levels contain different sets of subjects as whole-plot or between-subjects factors. On the other hand, the time factor is a typical sub-plot or within-subjects factor because it structures the observations within individual subjects. Observations on the same subject may be dependent, and for different level combinations of the whole-plot factors, different types of dependence may be assumed. However, within a group defined by a factor level combination or *cell*, the observation vectors are modeled as coming from the same multivariate distribution. In particular, the dependence structure is the same for all subjects within such a cell.

Let us first consider a simple, additive model for univariate repeated measures data with different groups. This model already allows for different variances in different cells, as described above, but it is not sufficiently general regarding the dependence structures that are encompassed, and therefore a more general model will be considered later. An additive model for repeated measures could be written as

$$Y_{ijk} = \mu_{ij} + A_{ik} + \epsilon_{ijk},$$

where Y_{ijk} is the observation on subject (k) in group (i) at time point (j). It is decomposed additively into the population mean μ_{ik} in group (i) for time point (k), and two random components. Namely, $A_{ik} \sim N(0, \tau_i^2)$ is the specific effect for subject (k) in group (i) and $\epsilon_{ijk} \sim N(0, \sigma_i^2)$ is an additive random error term whose variance may depend on the group. All random variables ϵ_{ijk} and A_{ik} shall be assumed independent. Such a model formulation results in a covariance structure where observations on different subjects, that is, with different index pair (i, k), are independent, and observations Y_{ij_1k} and Y_{ij_2k} on the same subject (k) in group (i) all have the same variance $\sigma_i^2 + \tau_i^2$, as well as the same covariance τ_i^2 , no matter the value of their indices j_1 and j_2 . In other words, the correlation $\rho_{ij_1k, ij_2k} = \tau_i^2 / (\sigma_i^2 + \tau_i^2)$ between two observations does not depend on their time distance $|j_1 - j_2|$.

For other dependence structures the correlation may decrease if two observations are further apart. Such dependence models are certainly possible and sometimes justified. By inspecting the compound symmetry structure above more closely, we observe that also differences of two observations have the same variance independent of the choice of time points. Indeed, this property is generally referred to as sphericity. Therefore, the compound symmetry covariance structure is a special case of sphericity (see e.g., Bathke et al. (2009)).

Box (1954a,b) (or in more detail Huynh and Feldt (1970)) showed that the simple *F*-test is no longer valid when sphericity of the covariance structure cannot be assumed. In lower-dimensional situations, classical corrections such as the Greenhouse-Geisser (Geisser and Greenhouse, 1958; Greenhouse and Geisser, 1959) or Huynh-Feldt (Huynh and Feldt, 1976; Lecoutre, 1991) methods for the repeated measures ANOVA can be used when the covariance matrix sphericity assumption is violated. However, the performance of tests using the Greenhouse-Geisser or Huynh-Feldt corrections heavily depends on the actual underlying covariance structures as both have been derived under the assumption of

homoscedasticity across the levels of the between-subjects factor, which is very often not justifiable in practice. In other words, they assume that the variance-covariance structure between the repeated observations is the same for each cell. This assumption is rather restrictive, in particular in the presence of high dimensionality of the data vectors, and therefore a more general model is needed. In a classical MANOVA-design or repeated measures design, this problem was already tackled by [Harrar \(2009\)](#); [Chen et al. \(2010\)](#); [Konietschke et al. \(2015\)](#); [Pauly et al. \(2015\)](#); [Harrar and Kong \(2016\)](#) and [Bathke et al. \(2018\)](#) among others.

Here and in the following, high-dimensionality of repeated measures means that the number of repeated measurements (d) per subject is larger than the sample sizes (n_i) in group (i) ($d > n_i$). Methods that rely on the inverse of an empirical variance-covariance matrix cannot be calculated in this situation. A well-known example is Hotelling's T^2 test for multivariate data which can be used to test for simple treatment effects, that is, for the hypothesis that expectation vectors are the same for all groups. In a high-dimensional setting, the empirical variance-covariance matrix is singular, rendering Hotelling's T^2 unusable. Alternative test statistics that can in principle still be calculated for high-dimensional data include the ANOVA type statistic ([Brunner et al., 1997](#); [Brunner and Puri, 2001](#)). Also, the Greenhouse-Geisser and the Huynh-Feldt correction for the repeated measures ANOVA can technically be computed, but they do suffer from the limitations mentioned before, and the performance of all three approaches deteriorates with increasing dimensionality (see, e.g., the simulation results in [Happ et al. \(2016\)](#)). Another possible approach for high-dimensional data is presented by the procedure proposed by [Secchi et al. \(2013\)](#), however requiring strong assumptions on the covariance matrices. Specifically, when we denote with Σ_i the $d \times d$ dimensional variance-covariance matrix for the i -th group, then the assumption from [Secchi et al. \(2013\)](#) can be formulated as

$$0 < \lim_{d \rightarrow \infty} \frac{\text{tr}(\Sigma_i^k)}{d} < \infty$$

for $k = 1, 2$, and 4 and for all groups $i = 1, \dots, a$. Similar assumptions have also been used by [Srivastava \(2007\)](#); [Srivastava and Yanagihara \(2010\)](#); [Chen et al. \(2010\)](#) or [Harrar and Kong \(2016\)](#) and these assumptions were discussed in [Pauly et al. \(2015\)](#). For practical applications, it may be difficult to verify this assumption.

The method implemented in the package **HRM** does not assume any particular covariance structures ([Happ et al., 2016, 2017](#)), neither homoscedasticity across the between-subjects factor levels, nor sphericity regarding the levels of the within-subjects factor. In particular, no compound symmetry is assumed. Also, it does not require a stringent covariance assumption as in [Secchi et al. \(2013\)](#) or [Srivastava and Yanagihara \(2010\)](#). It is therefore rather general. A limitation exists, however, also for this method. The theory has been derived under the assumption of normal errors. While it performs well empirically for a wide range of simulated data distributions, including discrete data, we recommend caution when the data is very skewed or heavy-tailed. For details, see Section 2.6.

In the following subsection, we will describe the model that is underlying the **HRM** package. Here, it will become clear that this model allows for rather general covariance structures, rendering the associated R package applicable for a wide range of high-dimensional data sets.

Statistical model

Consider a setting with two whole- and two sub-plot factors. Models for other designs can be formulated analogously. We refer to the whole-plot factors as A and S (e.g., group and subgroup) and to the sub-plot factors as C and D , respectively. For example, in the EEG data in Section , the sub-plot factors are *variable* and *region*. The observations are represented by independent random vectors

$$\mathbf{X}_{ijk} = (\mathbf{X}'_{ijk11}, \dots, \mathbf{X}'_{ijkbc})' \sim N(\boldsymbol{\mu}_{ij}, \boldsymbol{\Sigma}_{ij}).$$

Here, \mathbf{X}_{ijk} is the $d = b \times c$ dimensional data vector of the k -th subject in the i -th group and j -th subgroup for $i = 1, \dots, a$ and $j = 1, \dots, s$. Note that we do not impose any special dependence structure on the covariance-matrices $\boldsymbol{\Sigma}_{ij}$.

Accordingly, the factors A , S , B , and C have a , s , b , and c factor levels, respectively. The sample size in the i -th group and j -th subgroup is denoted by n_{ij} . Overall, there are $N = \sum_{i,j} n_{ij}$ experimental units. Our hypotheses of interest concern interaction effects between two or more factors, and main effects of a single factor. They can all be formulated as quadratic forms using appropriate projection

matrices K_ϕ which depend on the respective hypothesis of interest ϕ , namely

$$H_0(\phi) : \boldsymbol{\mu}' K_\phi \boldsymbol{\mu} = 0.$$

In a setting with only one whole- and one sub-plot factor with a and d factor levels (*group* and *time*), this projection matrix is simply given by a Kronecker product of two matrices, that is $K_\phi = W_\phi \otimes S_\phi$. If we want to test for the main group effect, we choose $W_\phi = P_a$ and $S_\phi = 1/d J_d$ where P_a is the $a \times a$ dimensional centering matrix and J_d is the $d \times d$ dimensional matrix whose entries are all equal to 1. With the latter matrix, we average over the sub-plot factor, and the former matrix centers a vector by subtracting its mean from all components of the vector. If we want to test for an interaction effect, we use for both W_ϕ and S_ϕ a centering matrix and if we want to test only for the main time effect we simply use $W_\phi = 1/a J_a$ and $S_\phi = P_d$. The corresponding test statistic is then constructed as

$$T(\phi) := \frac{\bar{X}' K_\phi \bar{X}}{\text{tr}(\hat{\Sigma}_\phi)},$$

where

$$\hat{\Sigma}_\phi = K_\phi \left(\bigoplus_{i=1}^a \bigoplus_{j=1}^s \frac{1}{n_{ij}} \hat{\Sigma}_\phi \right),$$

\bar{X} is the $as \times d$ -dimensional vector containing the means of all groups at each time point, and $\hat{\Sigma}_\phi$ is the empirical variance-covariance matrix for the i -th group. In general, we do not know the exact distribution of $T(\phi)$, but we can approximate the sampling distribution of $T(\phi)$ under null hypothesis by an F -distribution with estimated degrees of freedom \hat{f} and \hat{f}_0 . We refer to Brunner et al. (2012); Happ et al. (2016, 2017) for details regarding this so-called Box approximation and in particular the estimation of the degrees of freedom.

Note that this method only provides an approximation to the sampling distribution. Although it generally performs well, it is not asymptotically exact. There are asymptotic tests such as those proposed by Pauly et al. (2015) or Sattler and Pauly (2017), converging under some assumptions such as $\min\{n_i, i = 1, \dots, a\} \rightarrow \infty$ and they also provide a small sample size approximation. Similar to the aforementioned papers, an advantage of the approximative method in HRM is that it is working very well even for rather small sample sizes.

The R package HRM

In the R package HRM, the test statistic is implemented for designs described in Section 2.2. It is possible to use tests with up to four factors (at most two whole-plot or three sub-plot factors). Two S3 methods are provided to facilitate different data input formats. Both S3 methods can be called by the generic function `hrm_test`. The data can either be provided in the wide table format where each row represents the observations from one subject. This means that for each group, all observations have to be stored in a matrix. Then the matrices from all groups need to be elements of a `list`. This method only supports one whole- or one sub-plot factor, that is, a maximum of two factors can be used.

The other way is to provide the data in the long table format as a `data.frame`. Here, all observations are in stored in one column. The other columns of the `data.frame` specify the factor levels to which an observation belongs. For this type of data, at least one whole- and one sub-plot factor have to be used and they support in general up to two whole- and two sub-plot factors. There is also the special case of one whole- and three sub-plot factors implemented.

Depending on the type of data that is given as an argument to `hrm_test`, either the S3 method `hrm_test.list` or `hrm_test.data.frame` is internally called. We will give now a short description of these two S3 methods.

Methods

The S3 method `hrm_test.list` has two parameters, `data` and `alpha`. The second parameter has its default value set to 0.05 and specifies the type-I error level of the procedure which is used for calculating the critical value for the test procedure. The first parameter is a list containing the data in the wide table format for all groups. That is, the list has the structure

```
data <- list(group_1, group_2, ..., group_a),
```

where `group_i` is the data for the i -th group in the wide table format (the repeated measurements are the columns). Clearly, this method only works with one whole- and one sub-plot factor.

The method `hrm_test.data.frame` can be used for designs with up to two whole- and three sub-plot factors, but it is limited to a maximum number of four factors. The parameter data needs to be a `data.frame` containing the data with the aforementioned columns. Similar to the method `hrm_test.list` we can also specify the nominal type-I error rate by using the parameter `alpha`. The method `hrm_test.data.frame` needs a formula object. These are special R objects that allow to write a statistical model in a compact way and is used by many R packages. Let us assume the model

$$Y_{ijk} = \mu_{ij} + \epsilon_{ijk} = \alpha_i + \beta_j + \gamma_{ij} + \epsilon_{ijk},$$

where Y_{ijk} is the observation at time j for subject k in the i -th group and ϵ_{ijk} is the measurement error. The influence of the whole-plot factor is described by α and of the sub-plot factor by β . We denote by γ the interaction between these two factors. To write this model in a compact way in R we can use the formula

```
response ~ whole-plot factor * sub-plot factor.
```

Here the expression on the left hand site of \sim is the response variable. It is explained by the variables on the right hand site. The symbol $*$ means that the interaction term γ_{ij} is included in the model, that is $\mu_{ij} = \alpha_i + \beta_j + \gamma_{ij}$. Therefore main effects and the interaction effect are included in the model. This formula is equivalent to

```
response ~ whole-plot factor + sub-plot factor
+ whole-plot factor : sub-plot factor,
```

where `:` specifically denotes for the interaction effect. Hence the $*$ notation is an abbreviation for adding main and interaction effects. If we are only interested in the main effects then we could use the formula

```
response ~ whole-plot factor + sub-plot factor,
```

or alternatively if we are only interested in the interaction effect, our formula has the form

```
response ~ whole-plot factor : sub-plot factor.
```

For the method `hrm_test.data.frame` we also need to specify the column name for the subjects. This column name needs to be a character. Otherwise an error message is returned by the method.

All these methods previously described return an HRM object which contains a list consisting of the results in a `data.frame`, the formula, the type-I error, the subject column name and the column names of the whole- and sub-plot factors which are used and the data used for testing.

In the special case when only one whole- and one sub-plot factor is used, it is possible to plot the so-called profile curves. In each group $i = 1, \dots, a$, the mean \bar{x}_{it} is calculated at each time point $t = 1, \dots, d$. Then, the points $(1, \bar{x}_{i1}), \dots, (d, \bar{x}_{id})$ are plotted for all groups $i = 1, \dots, a$. An example can be seen in Figure 3. For plotting these curves, the S3 method `plot.HRM` is available. It needs an object of class "HRM" which is returned by the function `hrm_test`. Additionally we can change the labels for the x -axis and the y -axis with the parameters `xlab` and `ylab` respectively. Furthermore, it is possible to disable the legend by setting the argument `legend = FALSE`, the title for the legend can be changed with the argument `legend.title`. By setting `legend.title = NULL`, no legend title is displayed. Internally, this method is only a `ggplot2` wrapper (Wickham, 2009). If there is only one sub-plot but no whole-plot factor used, it is still possible to plot profile curves (see for example Figure 5).

For objects of class "HRM", there are also the methods `print.HRM` and `summary.HRM` available. The first method just reproduces the standard output from `hrm_test` whereas `summary.HRM` lists additionally the whole- and sub-plot factors which were used in the model.

To improve the performance in case of a large dimension, this package uses mainly the dual empirical variance-covariance matrices (Brunner et al., 2012; Happ et al., 2017). Because the trace function is invariant under cyclic permutations, the empirical $d \times d$ variance-covariance matrix can be reduced to an $n_i \times n_i$ matrix. This is especially useful if $d \gg n_i$, where n_i is the sample size in the i -th group. Another improvement is to use

$$\text{tr} \left(\widehat{\Sigma}_i \widehat{\Sigma}_j \right) = \mathbf{1}'_d \left(\widehat{\Sigma}_i * \widehat{\Sigma}_j \right) \mathbf{1}_d$$

for calculating the trace of a matrix product, where $*$ is the Hadamard-Schur product. The effects of these two improvements are only noticeable for calculating a single test statistic if d is very large ($d > 1000$).

An additional enhancement for high-dimensional data is achieved by using the method `hrm_test.list` instead of `hrm_test.data.frame`. Although it can be sometimes convenient to use the long table format in a `data.frame`, for computing the degrees of freedom and the test statistic, the data has to be

dimension d	data.frame	data.table
10^5	224.17	0.64
2×10^5	1411.93	1.07

Table 1: Running times in seconds for converting a wide table into a long table for 100 subjects.

separated internally by the whole-plot factors. This procedure can be quite time consuming for very large datasets (dimension $d > 10^5$ and number of subjects > 100) if the data structure `data.frame` is used. The problem in this case is that the data structure `data.frame` is not suitable for large data sets. But most functions that import data into R return a `data.frame` object. In such a case, it is beneficial to provide the data already separated by the whole-plot factors in a list. This can be done with the method `hrm_test.list`. One disadvantage by using this method is, that it is currently not possible to specify more than one whole- or sub-plot factor. Another way to solve this problem is to convert the `data.frame` into a `data.table` object (Dowle and Srinivasan, 2017) which is specifically developed for large data sets. But the syntax for `data.table` is completely different from `data.frame` and their syntax is also not compatible with each other. In order to demonstrate the effect, we ran a small simulation with $d = 10^5$ and $n = 100$. Here, it takes about 4 minutes to convert a `data.frame` from the wide to the long table format and when we used a `data.table` it takes under 1 second to do the same (see Table 1). By increasing the dimension, the time it takes for the conversion increases exponentially if you use a `data.frame`. For this simulation we used the the function `Sys.time` to measure the calculation time and repeated each calculation five times. Note that the time needed for rearranging the data heavily depends on the computer on which the R script runs. In our case, we employed a computer with an AMD Ryzen 5 1600 CPU (6×3.2 GHz) and 16GB DDR-4 RAM (2800 MHz). Therefore the results from this simulation cannot be generalized but it can be an indication that a `data.frame` should not be rearranged if the dimension is large. Therefore, using the method `hrm_test.list` may help to avoid this additional step of converting your `data.frame` into a `data.table`.

Graphical user interface

For more convenient access to the package’s capabilities, a graphical user interface (GUI) can be used by typing the command `hrm_GUI()` into the R console. A window opens where the data file (long table format) can be loaded and viewed. This window is shown in Figure 1. After typing the formula as described before and selecting the column identifying the subjects, the results can be displayed in a separate window. Optionally, the results can be saved as a \LaTeX -table or just as plain text. For each hypothesis test, the degrees of freedom of the F approximation, the value of the test statistic, the critical value for the test, and the p -value are displayed. An example of the results window for the EEG data is shown in Figure 2. This data set contains of 160 subjects who have been divided into one of four diagnostic groups (whole-plot factor). For each subject 40 measurements from an EEG are available (sub-plot factor). This example is explained in more detail in Section 2.5.1. In the results window are the test statistics along with the degrees of freedom and p -values for main and interaction effects displayed.

If only one whole- and one sub-plot factor is being used, the profile curves of the groups are displayed as a graphic in a separate window. The plot can be saved in the pdf format.

The GUI relies on the function `hrm_test` for doing the analysis and `plot.HRM` for creating the plot as shown in Figure 3. That is, there is no difference between using the GUI or using both generic functions directly. For creating the GUI, the **RGtk2** package is used (Lawrence and Temple Lang, 2010). This package allows us to utilise the cross-platform widget-toolkit GTK+ with R. The current version 2.20.34 of this package cannot be installed on macOS, therefore the GUI for the package **HRM** does not work on macOS.

For opening data files, we rely on the functions provided in the package **tcltk** to ensure that this procedure works on multiple computer environments. For viewing the data we utilise the package **RGtk2Extras** (Taverner et al., 2012) which provides a very easy way do display and manipulate data in a GUI. The packages **cairoDevice** and **ggplot2** are used for plotting (Lawrence, 2017; Wickham, 2009) profile curves. Similar as before, we have to avoid functions that only work for a specific computer environment. Therefore we work with the package **cairoDevice** which is capable of displaying a graphic in a window and does not depend on a specific computer environment. For saving the results of the function `hrm_test` as a \LaTeX -table, we are using the package **xtable** (Dahl, 2016) which converts a `data.frame` into the corresponding \LaTeX code for a table.

Because of the limitation of **RGtk2**, the GUI does not work on macOS (OS X 10.11.6). But the package **HRM** itself can be installed on macOS and the function `hrm_test` can be run regardless.

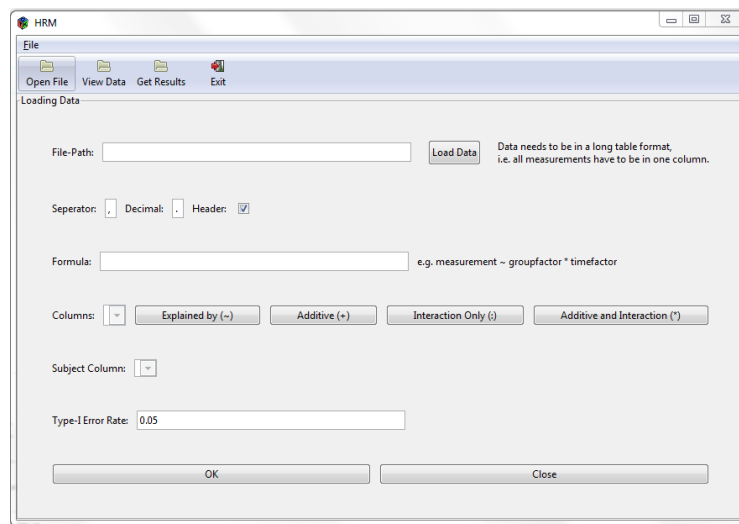


Figure 1: GUI for the **HRM** package. The data can be loaded either by clicking 'Open File' in the toolbar or by clicking on the button 'Load Data' next to the textfield for the file path. The formula has to be specified in the appropriate textfield and the subject column needs to be chosen in the dropdown list.

Apart from that, the GUI works on all other newer operating systems (for example Debian 7.3.0-12 or Microsoft Windows 7, 8, 10).

Examples

EEG Data Example

In order to demonstrate how to work with the package, we use the EEG data from [Staffen et al. \(2014\)](#) which is included in the **HRM** package. There are 160 subjects who have been diagnosed with Alzheimer's disease (AD), mild cognitive impairments (MCI), or subjective cognitive complaints (SCC+ or SCC-). For each subject, four variables (activity, complexity, mobility and brain rate) are measured at 10 different brain regions (frontal left/right, parietal left/right, central left/right, temporal left/right, and occipital left/right). We are using the diagnostic groups and sex as whole-plot factors, and the region and variable as sub-plot factors. Sample sizes are given in Table 2. Here in all combinations of diagnostic group and sex the dimension is larger than the sample size, therefore this is considered a high-dimensional setting.

Sex	AD	MCI	SCC+	SCC-
male	12	27	14	6
female	24	30	31	16
sum	36	57	45	22

Table 2: Sample sizes for the EEG data.

To perform the analysis we are using the graphical user interface and the R console. The commands

```
# Save EEG data from package HRM
write.table(EEG,
  file = "EEG.csv",
  sep = ",",
  dec = ".",
  col.names = TRUE)
# Open the GUI
hrm_GUI()
```

will load the package and save the EEG data in a csv file in the current working directory. The last line launches the GUI (see Figure 1). By clicking on the 'Open File' button we can select the EEG.csv file we previously saved. Depending on the data, we may have to change the separator mark, decimal

wholeplot-factors: group,	subplot-factors: dimension,					
hypothesis	df1	df2	crit	test	p.value	sign.code
group (weighted)	3.0	130	2.7	1.9	0.136	
group	2.8	95	2.8	1.7	0.185	
dimension	2.5	269	2.8	3056.6	<0.001	***
group : dimension	7.2	269	2.0	2.4	0.022	*

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Figure 2: Results for the EEG data. The results can be saved as a \LaTeX table by clicking on the button on the bottom or in the toolbar.

mark, or deselect that the file contains header. This can be done in the fields below the textfield for the file path. We can look at the data by clicking on the button 'View Data' in the toolbar. A new window is created displaying the data we have just loaded (seen in Figure 4). The data consists of columns specifying the whole-plot (group, sex) and sub-plot factors (variable, region). Additionally, one column contains the response variable from each person (value), and the subject column identifies the subjects. A unique identification of the subjects is necessary to determine automatically within the function `hrm_test` which factors are whole- and which are sub-plot factors. By combining the factors region and variable into just one factor, we obtain a new factor which we will simply call dimension. It is provided in the column with the same name.

Another way to look at the data is by using the command

```
View(EEG)
```

which will also display the complete data set. After we have loaded and looked at the data we need to specify the formula in the GUI. In this example, the formula is given by

```
value ~ group * sex * region * variable
```

which means that for the four factors tests for main and interaction effects will be performed. If we were only interested in testing the main effects we could use the formula

```
value ~ group + sex + region + variable
```

where we are using `+` instead of `*`. If we simply want to test interaction effects we could use `:` instead. In addition to specifying the formula, we need to select which column identifies the subjects. This is done by selecting the 'subject' column in the dropdown list. To perform the tests, we simply click on the 'Ok' button on the bottom of the window or the 'Get Results' button in the toolbar. A new window (see Figure 2) will open which contains the results of the performed tests. For each tested hypothesis, the degrees of freedom, the value of the test statistic, and the p-value are displayed. Above the results, a summary about the factors used in the analysis is provided. To perform these tests manually, we may use the commands

```
r <- hrm_test(formula = value ~ group*sex*region*variable,
              data = EEG,
              subject = "subject")
library(xtable)
print(xtable(r$result), include.rownames = FALSE)

% latex table generated in R 3.4.1 by xtable 1.8-2 package
\begin{table}[ht]
  \centering
  \begin{tabular}{lrrrrrl}
    \hline
```

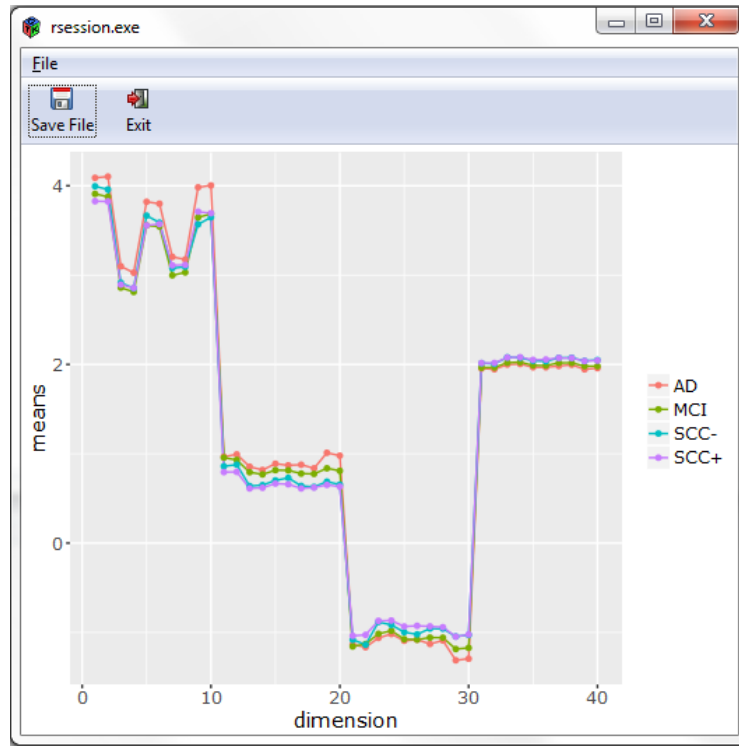


Figure 3: Profile Plot for the EEG data. The image can be saved by clicking on 'Save File'.

```

hypothesis & df1 & df2 & crit & test & p.value & sign.code \\
\hline
group & 3.04 & 116.29 & 2.67 & 1.66 & 0.18 & \\
sex & 1.00 & 116.29 & 3.92 & 6.05 & 0.02 & * \\
region & 5.59 & 246.82 & 2.18 & 196.60 & 0.00 & *** \\
variable & 1.51 & 178.72 & 3.36 & 4930.69 & 0.00 & *** \\
group:sex & 3.04 & 116.29 & 2.67 & 0.18 & 0.91 & \\
group:region & 14.09 & 246.82 & 1.73 & 0.80 & 0.67 & \\
group:variable & 4.52 & 178.72 & 2.33 & 3.34 & 0.01 & ** \\
sex:region & 5.59 & 246.82 & 2.18 & 0.99 & 0.43 & \\
sex:variable & 1.51 & 178.72 & 3.36 & 4.89 & 0.02 & * \\
region:variable & 7.76 & 299.89 & 1.99 & 143.58 & 0.00 & *** \\
group:sex:region & 14.09 & 246.82 & 1.73 & 0.58 & 0.88 & \\
group:sex:variable & 4.52 & 178.72 & 2.33 & 0.26 & 0.92 & \\
group:region:variable & 18.75 & 299.89 & 1.63 & 1.04 & 0.41 & \\
    
```

	group	value	sex	subject	variable	region	dimension
1	SCC+	2.80418	W	1	1	1	1
2	SCC+	2.285004	W	1	1	2	2
3	SCC+	1.473094	W	1	1	3	3
4	SCC+	1.548939	W	1	1	4	4
5	SCC+	2.160346	W	1	1	5	5
6	SCC+	2.143525	W	1	1	6	6
7	SCC+	1.366017	W	1	1	7	7
8	SCC+	1.678226	W	1	1	8	8
9	SCC+	1.928796	W	1	1	9	9
10	SCC+	2.067021	W	1	1	10	10
11	SCC+	0.8473393	W	1	2	1	11
12	SCC+	0.9197796	W	1	2	2	12

Figure 4: EEG data used in the example.


```
sex:region:variable & 7.76 & 299.89 & 1.99 & 1.26 & 0.27 & \\
group:sex:region:variable & 18.75 & 299.89 & 1.63 & 0.74 & 0.77 & \\
\hline
\end{tabular}
\end{table}
```

to obtain the results stated here. The command `xtable` is used to convert a data frame into a \LaTeX table. This is also possible with the GUI. In the toolbar of the results window there is a button labeled with 'Save as \LaTeX table'. By clicking on it we can select a path to save a tex file containing the results.

As an alternative formula, we may use the following.

```
value ~ group * dimension
```

Here, only one whole- and one sub-plot factor are specified. After clicking on the 'Ok' button, two windows will open. In the first window, the results of the test statistics are displayed. The other window shows a plot of the profiles (see Figure 3). We can also create these plots by using the function

```
plot(hrm_test(value ~ group*dimension, data = EEG, subject = "subject"))
```

Here the function `hrm_test` returns an object from class "HRM" and the function `plot` is applied to an object from such a class. If we calculate the test statistics in this two-factor example and save the results again as a table, we obtain the following:

```
hrm_test(formula = value ~ group*dimension,
          data = EEG,
          subject = "subject")
```

Call:

```
value ~ group * dimension
```

	hypothesis	df1	df2	crit	test	p.value	sign.code
group (weighted)		2.993584	130.10764	2.676095	1.881486	0.13600798	
group		2.779981	95.00304	2.768161	1.657329	0.18482432	
dimension		2.504068	268.94339	2.803050	3056.569162	0.00000000	***
group : dimension		7.155131	268.94339	2.031723	2.376905	0.02164946	*

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

T-cell Activation Example

It is also possible to use the function `hrm_test` with only one sub-plot factor and no whole-plot factor. A data example with such a design is included in the package `longitudinal` (Opgen-Rhein and Strimmer, 2015) and is originally taken from Rangel et al. (2004). The data are from a microarray time series for T-lymphocytes activation. Here, the expression levels of $n = 58$ genes are measured at $d = 100$ time points ($d > n$), thus presenting a high-dimensional situation. In order to analyse the data, we need to load the necessary packages first and then convert the data into the long table format. At the end, we can use the function `hrm_test` which detects a significant time effect. Alternatively, we could have used the GUI. But for this approach we would have needed to save the data as a csv file first because the GUI only supports loading data sets from external files.

This data example is just used as a demonstration for the package. From a statistical point of view, the gene expressions do not necessarily have to be independent, as different genes may have an influence on each other. For example in eukaryotes, gene expression is highly regulated. Some genes encode for transcription factors (for example zinc-fingers) which can up- or downregulate the expression of other genes. With regard to this data example the gene FYB seems to influence the expression of several interleukin receptor genes and GATA-3 (a zinc-finger transcription factor) (Rangel et al., 2004). Therefore, it is very plausible that the assumption of independent genes is violated for this data example. However, for the illustration of the method we are making the simplifying assumption that observations on different 'subjects' (genes, in this case) are independent from each other.

```
library(longitudinal)
library(tidyr)
library(HRM)

# transforming the data into the long table format
data(tcell)
```




Figure 5: Profile curve.

```
data <- t(print(tcell.10))
data <- as.data.frame(data)
data$subject = 1:58
data_long <- gather(data, time, measurement, 1:100, factor_key = TRUE)
```

```
hrm_test(data_long, measurement ~ time, subject = "subject")
```

```
Call:
measurement ~ time
```

hypothesis	df1	df2	crit	test	p.value	sign.code
time	5.514332	314.3169	2.179376	6.268063	6.538363e-06	***

```
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

To plot the profile curve for this data set, we simply apply the `plot` function to an "HRM" object returned by the function `hrm_test`. By setting the optional argument `legend` to `FALSE`, no legend is added to the plot (see Figure 5).

```
object_hrm <- hrm_test(measurement ~ time, data = data_long, subject = "subject")
plot(object_hrm, legend = FALSE, xlab = "time", ylab = "mean expression level")
```

For the `plot.HRM` method it is also possible to use additional arguments which will be added to the "ggplot2" object. As an example we want to use the `theme_bw` style with no legend. Here we need to take into consideration that `theme_bw` overwrites everything related to the legend already specified in `plot.HRM`. Therefore, we need to specify separately that a legend should not be displayed. The code

```
plot(object_hrm, ... = theme_bw() +
      theme(legend.title = element_blank(), legend.position = "none") +
      theme(axis.text = element_text(size=18), axis.title = element_text(size = 25)) )
```

produces Figure 6 with a white background and no legend. Because `plot.HRM` simply returns an "ggplot2" object, we also could have used the following code.

```
plot(object_hrm) + theme_bw() +
  theme(legend.title = element_blank(), legend.position = "none") +
  theme(axis.text = element_text(size = 18), axis.title = element_text(size = 25))
```

That is we can pass additional arguments through `"..."` in the method `hrm.plot` or, even easier, just add the additional arguments directly to the object to manipulate the graphic.

Limitations

In Section 2.2, we have already stated the mathematical assumptions for the methods implemented in the HRM package. We are assuming normally distributed observations from independent subjects. As




Figure 6: Profile curve generated with additional arguments for `ggplot2`.

long as the data distribution is not ‘too far’ from a normal distribution, the test works fine. However, for very skewed or heavy-tailed distributions, the test may in general not maintain the type-I error rate. This could especially be a problem in combination with heteroscedastic data and unbalanced groups. Note that this limitation still holds for large sample sizes because the described method is only an approximation and not an asymptotically exact test. This limitation also applies in particular to the case of multivariate normally distributed observations. Therefore it is important to check the assumptions or do some simulations to see whether the test is robust against some violations of the assumptions. As we see in the following simulations, the test statistic is mainly affected by skewed or heavy-tailed distributions but works well for data generated by a binomial distribution. For our two examples used in Section 2.5 they do not satisfy the assumptions. For the genetic data, the gene expressions (subjects) are in general correlated. The EEG data seems to violate the normality assumption as some QQ-plots from marginal distributions show larger deviations from a straight line. But it is very difficult to check these distributional assumptions in general for high-dimensional data sets. In our case we only use those data sets for demonstration purposes but for real analyses one should be careful about the assumptions as the type-I error can be quite inflated as we can see in the subsequent simulations.

In order to illustrate the behaviour of the proposed method when its formal assumptions are violated, we have performed several simulations for testing for the main group effect. For all simulations we used the dimension $d = 20$ and performed 10^4 simulation runs. We have already observed that the test statistic does not perform well for skewed data. Therefore we concentrated in this simulation study on the exponential and log-normal distribution. But we have also considered a discrete distribution as in such a case the test seems to perform very well.

For the first simulation, we defined the random variables $X_{1jk} = \epsilon_{1j} + \delta_1$ for $j = 1, \dots, 20$ and $k = 1, \dots, n_1 = 20$. In the second group we chose $X_{2jk} = \epsilon_{2j} + \delta_2$ for $j = 1, \dots, 20$ and $k = 1, \dots, n_1 = 30$. Here, $\delta_1 \sim \mathcal{E}(1)$, $\delta_2 \sim \mathcal{E}(2)$, $\epsilon_{1j} \sim \mathcal{E}(1)$ and $\epsilon_{2j} \sim \mathcal{E}(2)$, where \mathcal{E} refers to the exponential distribution. Then we shifted the second group equally for all 20 time points such that both groups have the same mean which means that the null hypothesis of no main group effect holds. For this setting, the simulated type-I error rate was 0.0688. In the second simulation we used for the second group $\mathcal{E}(4)$ distributions. In this case, the simulated type-I error rate was 0.0748. Additionally, we used different sample sizes to see the impact of negative or positive pairing. The results are displayed in Table 3. For all these settings, the simulated type-I error was slightly inflated. We performed similar simulations using log-normal distributions. Here, we considered $\delta_1 \sim LN(0, 1)$, $\delta_2 \sim LN(0, \theta)$, $\epsilon_{1j} \sim LN(0, 1)$ and $\epsilon_{2j} \sim LN(0, \theta)$ with $\theta = \frac{1}{2}, 2$. Both groups have been shifted such that they have the same mean. The results in Table 4 show heavily inflated type-I error rates (up to over 30%).

In a third simulation, we considered binomial distributions as an example for discrete data. Here, we defined $\delta_1, \epsilon_{2j} \sim Bin(m, 0.5)$ and $\delta_1, \epsilon_{2j} \sim Bin(m, 0.9)$ for $m = 5, 10, 15$ and for sample sizes $n_1 = 20, 15$ and $n_2 = 15, 20$. The other simulation parameters were the same as for the other two simulations. That is, the dimension was 20 and all random variables have been shifted to mean zero such that the null hypothesis holds. The results are shown in Table 5. For binomial data, the method seemed to work very well even for heteroscedastic and unbalanced groups. Furthermore, the value of m did not seem to influence the result very much and even in the case of Bernoulli distributions

($m = 1$) the simulated type-I error rate was close to the nominal level of 5%.

Therefore the package HRM should not be used with distributions similar to log-normal or exponential distributions, that is very skewed or heavy-tailed distributions. Instead, in those cases, we recommend using a nonparametric test such as the rank-based ANOVA type statistic or a similar method.

θ	sample sizes	type-I error rate
2	20, 30	0.0688
4	20, 30	0.0748
2	30, 20	0.0543
4	30, 20	0.0685
2	25, 25	0.0616
4	25, 25	0.0702

Table 3: Type-I error rate simulation with different settings for testing for the main group effect. We used centred exponential distributions with parameter $\theta = 1$ in the first group and $\theta = 2, 4$ for the second group, and the dimension was $d = 20$.

θ	sample sizes	type-I error rate
0.5	20, 30	0.1157
2	20, 30	0.2927
0.5	30, 20	0.0917
2	30, 20	0.3126
0.5	25, 25	0.1032
2	25, 25	0.2974

Table 4: Type-I error rate simulation using log-normal distributions for testing for the main group effect. For the first group we used centred $LN(0, 1)$ and for the second group centred $LN(0, \theta)$ distributions, and the dimension was $d = 20$.

m	sample sizes	type-I error rate
1	20, 15	0.0488
5	20, 15	0.0476
10	20, 15	0.0503
15	20, 15	0.0467
1	15, 20	0.0515
5	15, 20	0.0485
10	15, 20	0.0479
15	15, 20	0.0489

Table 5: Type-I error rate simulation using binomial distributions for testing for the main group effect. For the first group we used centred $Bin(m, 0.5)$ and for the second group we used centred $Bin(m, 0.9)$ random variables, and the dimension was $d = 20$.

Another drawback of our method is that it is not invariant under scale transformations of individual variables. This means that changing the scale in one dimension (using m/s instead of km/h) may lead to a different result because the observations are not standardized to avoid the problem of singular empirical covariance matrices. Therefore the same scale should be used for all dimensions.

The package `MANOVA.RM` by Friedrich et al. (2017b) provides several asymptotically exact tests that do not have these limitations (see for example Friedrich and Pauly (2018)). However, not all have been extended to high-dimensional settings (Friedrich et al., 2017a) and some may also require a resampling step which may take a while to compute for large data sets. But this might only become a real problem if you want to simulate the power of the test for multiple situations. The ANOVA type statistic mentioned before has the problem that it gets quite conservative when the dimension is large (Brunner et al., 2002; Bathke et al., 2009; Friedrich et al., 2017a), and additionally, the test is also not invariant under scale transformations. Friedrich et al. (2017a) proposed a potential solution for this problem in the low-dimensional case.

Conclusions

The aim of the **HRM** package is to provide an easy to use way to do inference on high-dimensional repeated measures. To that end, a graphical user interface has also been implemented in this package. With the GUI, it is possible to load and view the data, plot the profile curves and save the results of hypothesis tests as a table which can be inserted very easily into a \LaTeX document. The GUI is only optional, that is, all of its functions can be used directly. There are different functions to allow for data in the long or wide table format. Although both formats can be transformed easily into each other, not having to do it is more convenient for users. This applies especially to statistics practitioners who are not expert R users.

The package currently supports up to four factors. If data in the wide table format are used, the maximum number of factors, that can be used, is two. For the long table format, it is possible to use up to two whole-plot factors and up to three sub-plot factors, but the maximum number of factors is four. Table 6 summarizes the minimum and maximum number of factors that can be used.

Table Format	Whole-plot	Sub-plot	Total
long table	0–2	0–3	1–4
wide table	1	1	2

Table 6: Minimum and maximum number of factors which can be used with the package **HRM**.

There are still a few limitations to the method implemented in the package. Improvements to the test statistic to get rid of these limitations are part of future work and will also be implemented in the package. For example, a nonparametric version of the test statistic will be provided in the future by setting an additional argument `nonparametric = TRUE` for the function `hrm_test`. This new test will use (pseudo)-ranks and can be applied to metric or non-metric data. Additionally, the package will be further improved to increase the performance for large data.

Funding

The research was supported by Austrian Science Fund (FWF) I 2697-N31.

Bibliography

- A. C. Bathke, O. Schabenberger, R. D. Tobias, and L. V. Madden. Greenhouse-geisser adjustment and the anova-type statistic: Cousins or twins? *The American Statistician*, 63(3):239–246, 2009. URL <https://doi.org/10.1198/tast.2009.08187>. [p1, 12]
- A. C. Bathke, S. Friedrich, M. Pauly, F. Konietschke, W. Staffen, N. Strobl, and Y. Höller. Testing mean differences among groups: Multivariate and repeated measures analysis with minimal assumptions. *Multivariate Behavioral Research*, 0(0):0–0, 2018. URL <https://doi.org/10.1080/00273171.2018.1446320>. PMID: 29565679. [p2]
- G. Box. Some theorems on quadratic forms applied in the study of analysis of variance problems, ii. effects of inequality of variance and of correlation between errors in the two-way classification. *The Annals of Mathematical Statistics*, 25(3):484–498, 1954a. URL <https://doi.org/10.1214/aoms/1177728717>. [p1]
- G. E. Box. Some theorems on quadratic forms applied in the study of analysis of variance problems, i. effect of inequality of variance in the one-way classification. *The Annals of Mathematical Statistics*, 25(2):290–302, 1954b. URL <https://doi.org/10.1214/aoms/1177728786>. [p1]
- E. Brunner and M. L. Puri. Nonparametric methods in factorial designs. *Statistical Papers*, 42(1):1–52, 2001. URL <https://doi.org/10.1007/s003620000039>. [p2]
- E. Brunner, H. Dette, and A. Munk. Box-type approximations in nonparametric factorial designs. *Journal of the American Statistical Association*, 92(440):1494–1502, 1997. URL <https://doi.org/10.1080/01621459.1997.10473671>. [p2]
- E. Brunner, S. Domhof, and F. Langer. *Nonparametric Analysis of Longitudinal Data in Factorial Designs*. John Wiley & Sons, 2002. [p12]

- E. Brunner, A. C. Bathke, and M. Placzek. Estimation of box's ϵ for low-and high-dimensional repeated measures designs with unequal covariance matrices. *Biometrical Journal*, 54(3):301–316, 2012. URL <https://doi.org/10.1002/bimj.201100160>. [p3, 4]
- S. X. Chen, Y.-L. Qin, and others. A two-sample test for high-dimensional data with applications to gene-set testing. *The Annals of Statistics*, 38(2):808–835, 2010. URL <https://dx.doi.org/10.1214/09-AOS716>. [p2]
- D. B. Dahl. *Xtable: Export Tables to LaTeX or HTML*, 2016. URL <https://CRAN.R-project.org/package=xtable>. R package version 1.8-2. [p5]
- M. Dowle and A. Srinivasan. *Data.table: Extension of 'data.frame'*, 2017. URL <https://CRAN.R-project.org/package=data.table>. R package version 1.10.4-3. [p5]
- S. Friedrich and M. Pauly. Mats: Inference for potentially singular and heteroscedastic manova. *Journal of Multivariate Analysis*, 165:166 – 179, 2018. ISSN 0047-259X. URL <https://doi.org/10.1016/j.jmva.2017.12.008>. [p12]
- S. Friedrich, E. Brunner, and M. Pauly. Permuting longitudinal data in spite of the dependencies. *Journal of Multivariate Analysis*, 153:255 – 265, 2017a. ISSN 0047-259X. URL <https://doi.org/10.1016/j.jmva.2016.10.004>. [p12]
- S. Friedrich, F. Konietschke, and M. Pauly. *MANOVA.RM: Analysis of Multivariate Data and Repeated Measures Designs*, 2017b. URL <https://CRAN.R-project.org/package=MANOVA.RM>. R package version 0.1.2. [p12]
- S. Geisser and S. W. Greenhouse. An extension of box's results on the use of the f distribution in multivariate analysis. *The Annals of Mathematical Statistics*, 29(3):885–891, 1958. URL <https://doi.org/10.1214/aoms/1177706545>. [p1]
- S. W. Greenhouse and S. Geisser. On methods in the analysis of profile data. *Psychometrika*, 24(2): 95–112, 1959. URL <https://doi.org/10.1007/BF02289823>. [p1]
- M. Happ, S. W. Harrar, and A. C. Bathke. Inference for low-and high-dimensional multigroup repeated measures designs with unequal covariance matrices. *Biometrical Journal*, 58(4):810–830, 2016. URL <https://doi.org/10.1002/bimj.201500064>. [p2, 3]
- M. Happ, S. W. Harrar, and A. C. Bathke. High-dimensional repeated measures. *Journal of Statistical Theory and Practice*, 11(3):468–477, 2017. URL <https://doi.org/10.1080/15598608.2017.1307792>. [p2, 3, 4]
- S. W. Harrar. Asymptotics for tests on mean profiles, additional information and dimensionality under non-normality. *Journal of Statistical Planning and Inference*, 139(8):2685 – 2705, 2009. ISSN 0378-3758. URL <https://doi.org/10.1016/j.jspi.2008.12.008>. [p2]
- S. W. Harrar and X. Kong. High-dimensional multivariate repeated measures analysis with unequal covariance matrices. *Journal of Multivariate Analysis*, 145:1 – 21, 2016. ISSN 0047-259X. URL <https://doi.org/10.1016/j.jmva.2015.11.012>. [p2]
- H. Huynh and L. S. Feldt. Conditions under which mean square ratios in repeated measurements designs have exact f -distributions. *Journal of the American Statistical Association*, 65(332):1582–1589, 1970. URL <https://doi.org/10.1080/01621459.1970.10481187>. [p1]
- H. Huynh and L. S. Feldt. Estimation of the box correction for degrees of freedom from sample data in randomized block and split-plot designs. *Journal of Educational Statistics*, 1(1):69–82, 1976. URL <https://doi.org/10.3102/10769986001001069>. [p1]
- F. Konietschke, A. C. Bathke, S. W. Harrar, and M. Pauly. Parametric and nonparametric bootstrap methods for general manova. *Journal of Multivariate Analysis*, 140:291 – 301, 2015. ISSN 0047-259X. URL <https://doi.org/10.1016/j.jmva.2015.05.001>. [p2]
- M. Lawrence. *cairoDevice: Embeddable Cairo Graphics Device Driver*, 2017. URL <https://CRAN.R-project.org/package=cairoDevice>. R package version 2.24. [p5]
- M. Lawrence and D. Temple Lang. RGtk2: A graphical user interface toolkit for R. *Journal of Statistical Software*, 37(8):1–52, 2010. URL <http://www.jstatsoft.org/v37/i08/>. [p5]
- B. Lecoutre. A correction for the $\bar{\epsilon}$ approximate test in repeated measures designs with two or more independent groups. *Journal of Educational Statistics*, 16(4):371–372, 1991. URL <https://doi.org/10.3102/10769986016004371>. [p1]

- R. Opgen-Rhein and K. Strimmer. *Longitudinal: Analysis of Multiple Time Course Data*, 2015. URL <https://CRAN.R-project.org/package=longitudinal>. R package version 1.1.12. [p9]
- M. Pauly, D. Ellenberger, and E. Brunner. Analysis of high-dimensional one group repeated measures designs. *Statistics*, 49(6):1243–1261, 2015. URL <https://doi.org/10.1080/02331888.2015.1050022>. [p2, 3]
- C. Rangel, J. Angus, Z. Ghahramani, M. Lioumi, E. Sotheran, A. Gaiba, D. L. Wild, and F. Falciani. Modeling t-cell activation using gene expression profiling and state-space models. *Bioinformatics*, 20(9):1361–1372, 2004. URL <https://doi.org/10.1093/bioinformatics/bth093>. [p9]
- P. Sattler and M. Pauly. Inference for high-dimensional split-plot-designs: A unified approach for small to large numbers of factor levels. *arXiv preprint arXiv:1706.02592*, 2017. [p3]
- P. Secchi, A. Stamm, and S. Vantini. Inference for the mean of large p small n data: A finite-sample high-dimensional generalization of hotelling’s theorem. *Electronic Journal of Statistics*, 7:2005–2031, 2013. URL <https://doi.org/10.1214/13-EJS833>. [p2]
- M. S. Srivastava. Multivariate theory for analyzing high dimensional data. *JOURNAL OF THE JAPAN STATISTICAL SOCIETY*, 37(1):53–86, 2007. URL <https://doi.org/10.14490/jjss.37.53>. [p2]
- M. S. Srivastava and H. Yanagihara. Testing the equality of several covariance matrices with fewer observations than the dimension. *Journal of Multivariate Analysis*, 101(6):1319 – 1329, 2010. ISSN 0047-259X. URL <https://doi.org/10.1016/j.jmva.2009.12.010>. [p2]
- W. Staffen, N. Strobl, H. Zauner, Y. Höller, J. Dobesberger, and E. Trinka. Combining spect and eeg analysis for assessment of disorders with amnesic symptoms to enhance accuracy in early diagnostics. *Poster A19 Presented at the 11th Annual Meeting of the Austrian Society of Neurology. 26th-29th March 2014, Salzburg, Austria.*, 2014. [p6]
- T. Taverner, with contributions from John Verzani, I. Conde, and L. Andronic. *RGtk2Extras: Data Frame Editor and Dialog Making Wrapper for RGtk2*, 2012. URL <https://CRAN.R-project.org/package=RGtk2Extras>. R package version 0.6.1. [p5]
- H. Wickham. *Ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag, 2009. ISBN 978-0-387-98140-6. URL <http://ggplot2.org>. [p4, 5]

Martin Happ
University of Salzburg
Department of Mathematics
5020 Salzburg
Austria
martin.happ@sbg.ac.at

Solomon W. Harrar
University of Kentucky
Department of Statistics
Lexington, KY 40536
USA
solomon.harrar@uky.edu

Arnce C. Bathke
University of Salzburg
Department of Mathematics
5020 Salzburg
Austria
arne.bathke@sbg.ac.at