

# proBatch

*Jelena Čuklina, Chloe H. Lee, Patrick Pedrioli and Ruedi Aebersold*

*Institute of Molecular Systems Biology, Department of Biology, ETH Zurich, Switzerland*

*2019-05-03*

## Abstract

This vignette describes how to apply different functions from the **proBatch** package to diagnose and correct for batch effects. Most of the functions are applicable any “omic” data, however, the package has a number of functions, designed specifically for mass spectrometry-based proteomics, and has been tested on SWATH data.

The **proBatch** package provides a complete functionality for batch correction workflow: to prepare the data for analysis, diagnose and correct batch effects and finally, to evaluate the correction with quality control metrics.

The **proBatch** package was programmed and intended for use by researchers without extensive programming skills, but with basic R knowledge.

## Contents

<b>1</b>	<b>Introduction</b>	<b>2</b>
1.1	Batch effects analysis in large-scale data . . . . .	2
<b>2</b>	<b>Preparation for the analysis</b>	<b>3</b>
2.1	Installing dependencies and <b>proBatch</b> . . . . .	3
2.2	Preparing the data for analysis . . . . .	4
2.2.1	Loading the libraries . . . . .	4
2.2.2	Input data formats . . . . .	4
2.2.3	Example dataset . . . . .	5
2.2.4	Preparing sample and peptide annotations . . . . .	6
2.2.5	Other utility functions . . . . .	7
<b>3</b>	<b>Step-by-step workflow</b>	<b>8</b>
3.1	Initial assessment of the raw data matrix . . . . .	8
3.1.1	Plotting the sample mean . . . . .	8
3.1.2	Plotting boxplots . . . . .	9
3.2	Normalization . . . . .	10
3.2.1	Median normalization . . . . .	10
3.2.2	Quantile normalization . . . . .	10
3.3	Diagnostics of batch effects in normalized data . . . . .	11
3.3.1	Hierarchical clustering . . . . .	11
3.3.2	Principal component analysis (PCA) . . . . .	13
3.3.3	Principal variance component analysis (PVCA) . . . . .	14
3.3.4	Peptide-level diagnostics and spike-ins . . . . .	15
3.4	Correction of batch effects . . . . .	16
3.4.1	Continuous drift correction . . . . .	16
3.4.2	Discrete batch correction: combat or peptide-level median centering . . . . .	17
3.4.3	Correct batch effects: universal function . . . . .	19
3.5	Quality control on batch-corrected data matrix . . . . .	20
3.5.1	Heatmap of selected replicate samples . . . . .	20
3.5.2	Correlation distribution of samples . . . . .	21
3.5.3	Correlation of peptide distributions within and between proteins . . . . .	22

4	SessionInfo	23
5	Citation	24
6	References	26

# 1 Introduction

## 1.1 Batch effects analysis in large-scale data

Recent advances in mass-spectrometry enabled fast and near-exhaustive identification and quantification of proteins in complex biological samples [1], allowing for the profiling of large-scale datasets. Obtaining a sufficiently large dataset is, however, associated with considerable logistics efforts. Often multiple handlers at the sample preparation and data acquisition steps are involved e.g. protein extraction, peptide digestion, instrument cleaning. This introduces systematic technical variations known as batch effects.

Batch effects can alter or obscure the biological signal in the data [2, 3]. Thus, the presence and severity of batch effects should be assessed, and, if necessary, corrected.

The fundamental objective of the batch effects adjustment procedure is to make all measurements of samples comparable for a meaningful biological analysis. Normalization brings the measurements into the same scale. Bias in the data, however, can persist even after normalization, as batch effects might affect specific features (peptides, genes) thus requiring additional batch correction procedures. This means, that the correction of technical bias has often two steps: normalization and batch effects correction.

The improvement of the data is best assessed visually at each step of the correction procedure. The initial assessment sets the baseline, before any correction is executed. After normalization, batch effects diagnostics allow to determine the severity of the remaining bias. Finally, the quality control step allows to determine whether the correction improved the quality of the data.

The pipeline, summarizing this workflow, is shown in Fig.1.

##Analysis of large-scale data: steps before and after batch correction

We recommend users to follow this batch correction workflow to ensure all measurements are comparable for downstream analysis. We provide step-by-step illustrations to implement this workflow in the next sections.

Before starting the description, we give a few hints about the steps preceding and following batch effects analysis and correction.

- It is assumed that the initial data processing is completed. In mass spectrometry- based proteomics, this involves primarily peptide-spectrum matching [4, 5] and FDR control [6].
- Data filtering is commonly the next step of data processing. In the context of batch effects correction, both peptide and sample filtering need to be approached with caution. First of all, decoy measurements should be filtered out to ensure correct sample intensity distribution alignment. However, non-proteotypic peptides should be retained. Filtering out low-quality samples, also substantially alters normalization and batch effects correction. The “bad” samples, usually identified by the total intensity of identified peptides or correlation of samples, can be removed either before or after the correction for technical bias. Which option is best for a given dataset, should be decided in each case individually.
- We strongly advocate not to impute missing values before correction and to exclude “requant” values, inferred from SWATH data. Two common strategies to impute values use either “average” measurements, or random noise-level measurements. Both strategies bias the mean/median estimate of the peptide and are detrimental to both normalization and batch effects correction.
- We suggest to perform protein quantification after batch effects correction, as the correction procedure alters the abundances of peptides and peptide transitions, and these abundances are critical for protein

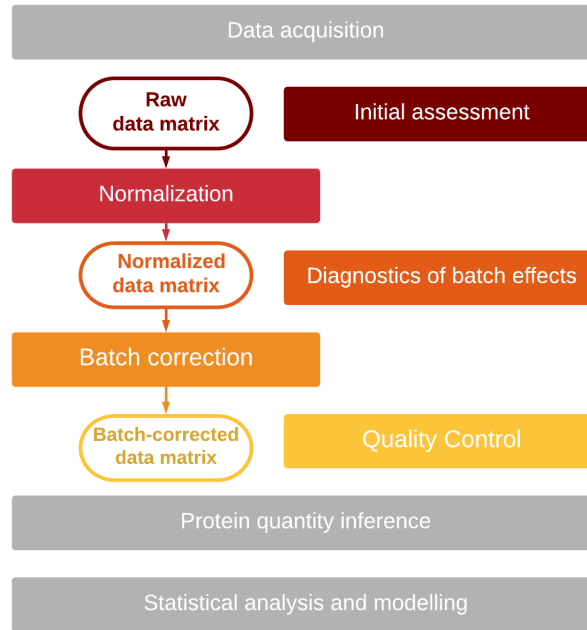


Figure 1: proBatch in batch correction workflow

quantity inference. Instead, we do recommend to correct the technical noise at the level, which is used to infer the proteins (thus, fragment-level for inference tools such as aLFQ or MSstats).

## 2 Preparation for the analysis

### 2.1 Installing dependencies and proBatch

proBatch is primarily a wrapper of functions from other packages, therefore it has numerous dependencies. If some of these dependencies are not installed, you will need to do that before running proBatch.

```

bioc_deps <- c("GO.db", "impute", "preprocessCore", "pvca", "sva" )
cran_deps <- c("corrplot", "data.table", "ggplot2", "ggfortify", "lazyeval", "pheatmap", "reshape2", "rlang",
              "tibble", "dplyr", "tidyr", "wesanderson", "WGCNA")

if (!requireNamespace("BiocManager", quietly = TRUE))
  install.packages("BiocManager")
BiocManager::install(bioc_deps)
install.packages(cran_deps)
  
```

To install the proBatch package, the following commands can be executed in R:

```

# Once the proBatch package is in Bioconductor, can easily install by:
install.packages("proBatch")

# Alternatively, install the development version from GitHub:
install.packages("devtools")
devtools::install_github("sympioticMe/proBatch", build_vignettes = TRUE)
  
```

## 2.2 Preparing the data for analysis

### 2.2.1 Loading the libraries

In this vignette, we use functions from `dplyr`, `tibble`, `ggplot2` and other `tidyverse` package family to transform some data frames

```
require(dplyr)
require(tibble)
require(ggplot2)
```

### 2.2.2 Input data formats

To analyze an experiment for batch effects, three tables need to be loaded into environment:

The package typically requires three datasets: 1) measurement (data matrix), 2) sample annotation, and 3) feature annotation (optional) tables. If you are familiar with the `Biobase` package, these correspond to 1) `assayData`, 2) joined `phenoData` and `protocolData`, and 3) `featureData`.

#### 1) Measurement table

Either a wide data matrix or long format data frame. In the wide (matrix) format, referred in this vignette as `data_matrix`, rows represent features (for proteomics, peptides/fragments) and columns represent samples. In the long format, referred in this vignette as `df_long`, each row is a measurement of a specific feature (peptide, fragment) in the specific sample. At least three columns are required: feature ID column, measurement (intensity) column and sample ID column. For batch correction, we also assume, that the imputed values, e.g. requants in SWATH, are flagged in quality column, such as `m_score`, so that they can be filtered out (see below).

In this vignette, the essential columns have the following names:

```
feature_id_col = 'peptide_group_label'
measure_col = 'Intensity'
sample_id_col = 'FullRunName'
essential_columns = c(feature_id_col, measure_col, sample_id_col)
```

The names of the columns can be technology-specific. These column names are specific to the OpenSWATH tsv output format.

In the package, we provide the functionality to convert from long to matrix format (see section 2.2.4.1 “Utility functions”).

Note that the sample IDs (column names in `data_matrix`) should match the values of the sample ID column in `sample_annotation` and the feature ID column values should match the feature annotation table (here - `peptide_annotation`). For OpenSWATH tsv file, which is a long format data frame, `peptide_annotation` can be generated in the beginning of the analysis.

#### 2) Sample annotation

A data frame, where one row corresponds to one sample (run/file), and the columns contain information on biological and technical factors. Minimally, sample annotation has to contain a sample ID column, at least one technical and one biological factor column, and a biological ID column (unique ID for the biological replicate, which is repeated for each technical replicate).

In our example data, these columns are:

1. `sample_id_col = 'FullRunName'`
2. technical covariates:
  - `SacrificeDate` - date when tissues were extracted

- ProteinPrepDate - date when samples were prepared
  - RunDate (and RunTime, if available) - will be used to determine run order
  - MS\_batch - number of MS batches (in this case, sets of runs between machine cleaning)
3. biological covariates:
- Strain
  - Diet
  - Sex
4. biospecimen\_id\_col = "EarTag"

Thus, technical and biological factors are:

```
technical_covariates = c('MS_batch', 'digestion_batch', 'RunDate', 'RunTime')
biological_covariates = c('Strain', 'Diet', 'Sex', 'Age_Days')
biospecimen_id_col = "EarTag"
```

In the example dataset, you will also find `order` column, which is, however, best inferred from the run date/time with the corresponding function (see Utility functions below).

### 3) Feature (peptide) annotation

A dataframe, where one row corresponds to one feature (in MS proteomics - peptide or fragment), and the columns are names of proteins and corresponding genes. Thus, the minimum columns are feature ID (`peptide_group_id`) and name of corresponding protein (in this vignette, we use `Gene` name).

#### 2.2.3 Example dataset

The `proBatch` package can be applied to any dataset, for which an intensity matrix and a sample annotation tables are available. However, the package was primarily designed with proteomic data in mind, and thoroughly tested on SWATH data. Thus, as an example dataset we include a reduced SWATH-MS measurement file, generated from a BXD mouse aging study. In this study, the liver proteome of mouse from BXD reference population have been profiled to identify proteome changes associated with aging. The animals of each strain were subjected to Chow and High-Fat Diet and sacrificed at different ages (the age factor is excluded from the example data as age-related differences are the focus of an unpublished manuscript).

This dataset has a few features, that make it a good illustrative example: 1. This is a large dataset of 371 samples, that was affected by multiple technical factors, described above in the `sample_annotation` subsection. Specifically, 7 MS batches drive the similarity of the samples. 2. The technical factors bias the data in at least two ways: discrete shifts (affecting different peptides in a batch-specific way), and continuous shifts from MS drift associated with sample running order. We will illustrate, how such biases can be corrected. 3. Replicate structure: samples from two animals were injected in the MS instrument every 10-15 samples. Additionally, several samples were repeated back-to-back in the end and in the beginning of two consecutive batches. This replication scheme allows to evaluate the coefficient of variation and is highly beneficial for assessment of sample correlation.

The example SWATH data and annotation files can be loaded from the package with the function `data()`.

```
library(proBatch)
data("example_proteome", "example_sample_annotation", "example_peptide_annotation",
     package = "proBatch")
```

## 2.2.4 Preparing sample and peptide annotations

proBatch provides utility functions to facilitate the preparation of sample and peptide annotation. Feel free to skip this section if you don't require them.

### 2.2.4.1 Defining the order of samples from running date and time

In proteomics, sequential measurement of samples may introduce order-related effects. To facilitate the examination of such effects, it is necessary to define an order column in the sample annotation. Using the `date_to_sample_order()` function one can infer sample order from the date and time of the measurements.

You can specify the columns illustrating date and time with `time_column` and their formats with the `dateTimeFormat` parameters (see POSIX date format for reference).

```
generated_sample_annotation <- date_to_sample_order(example_sample_annotation,
  time_column = c('RunDate', 'RunTime'),
  new_time_column = 'generated_DateTime',
  dateTimeFormat = c("%b_%d", "%H:%M:%S"),
  new_order_col = 'generated_order',
  instrument_col = NULL)

library(knitr)
kable(generated_sample_annotation[1:5,] %>%
  select(c("RunDate", "RunTime", "order", "generated_DateTime", "generated_order")))
```

RunDate	RunTime	order	generated_DateTime	generated_order
Oct_05	18:35:00	1	2019-10-05 18:35:00	1
Oct_05	20:12:00	2	2019-10-05 20:12:00	2
Oct_05	21:50:00	3	2019-10-05 21:50:00	3
Oct_05	23:28:00	4	2019-10-05 23:28:00	4
Oct_06	01:51:00	5	2019-10-06 01:51:00	5

The new time and order columns have been generated. Note that the `generated_order` has the same order as the manually annotated order column.

### 2.2.4.2 Generating peptide annotation from OpenSWATH data

From the OpenSWATH output, you can generate peptide annotation using `create_peptide_annotation()` by denoting the peptide ID with the `feature_id_col` and the annotation columns with the `annotation_col` parameters.

```
generated_peptide_annotation <- create_peptide_annotation(example_proteome,
  feature_id_col = 'peptide_group_label',
  annotation_col = c('ProteinName', 'Gene'))
```

In practice, the generation of peptide annotation from proteomic data allows one to remove peptide annotation columns from the intensity dataframe, thereby reducing the memory load, and can be achieved as follows:

```
example_proteome = example_proteome %>% select(one_of(essential_columns))
gc()
#>      used (Mb) gc trigger (Mb) max used (Mb)
#> Ncells 3367187 179.9   6778948 362.1 4185310 223.6
#> Vcells 6044716 46.2    12255594 93.6 10051970 76.7
```

Additionally, smaller peptide annotation matrices allow for faster mapping of UniProt identifiers to gene names and other IDs.

## 2.2.5 Other utility functions

### 2.2.5.1 Transforming the data to long or wide format

Plotting functions accept data in either data matrix or long data frame formats. Our package provides the helper functions `long_to_matrix()` and `matrix_to_long()` to conveniently convert datasets back and forth.

```
example_matrix <- long_to_matrix(example_proteome)
```

### 2.2.5.2 Transforming the data to log scale

Additionally, if the data are expected to be log-transformed, one can:

```
log_transformed_matrix <- log_transform(example_matrix)
```

### 2.2.5.3 Defining the color scheme

For color annotation, you first assign colors to biological and technical covariates with `sample_annotation_to_colors()`. Using this function columns are assigned to factoric, non-factoric or numeric types to specify colors into qualitative or sequential color scales.

```
color_scheme <- sample_annotation_to_colors (example_sample_annotation,  
      factor_columns = c('MS_batch', 'EarTag', "Strain", "Diet", "digestion_batch", "Sex"),  
      not_factor_columns = 'DateTime',  
      numeric_columns = c('order'))  
color_list = color_scheme$list_of_colors
```

## 3 Step-by-step workflow

### 3.1 Initial assessment of the raw data matrix


Before any correction, it is informative to set the baseline of the data quality by examining global quantitative patterns in the raw data matrix. Commonly, batch effects manifest as batch-specific intensity distribution changes.

In proteomics, batch-specific intensity drifts of sample mean can occur. Thus, it is important to carefully keep track of the order of sample measurement. Order inference can be performed as shown in the previous section “Defining the order of samples from running date and time”. If the order column is not available (`order_col = NULL`), the samples order in the sample annotation is used for plotting.

#### 3.1.1 Plotting the sample mean

The `plot_sample_mean()` function illustrates global average vs. sample running order. This can be helpful to visualize the global quantitative pattern and to identify discrepancies within or between batches.

```
batch_col = 'MS_batch'
plot_sample_mean(log_transformed_matrix, example_sample_annotation, order_col = 'order',
                 batch_col = batch_col, color_by_batch = TRUE, ylims = c(12, 16),
                 color_scheme = color_list[[batch_col]])
```




We can clearly see down-sloping trends in the BXD aging dataset. In fact, during the data acquisition, the mass-spectrometer had to be interrupted several times for tuning and/or column exchange as the signal was decreasing.



### 3.1.2 Plotting boxplots

Alternatively, `plot_boxplots()` captures the global distribution vs. the sample running order.

```
log_transformed_long <- matrix_to_long(log_transformed_matrix)
batch_col = 'MS_batch'
plot_boxplot(log_transformed_long, example_sample_annotation,
             batch_col = batch_col, color_scheme = color_list[[batch_col]])
```



In many cases, global quantitative properties such as sample medians or standard deviations won't match. The initial assessment via mean plots or boxplots can capture such information and hint at which normalization method is better suitable. If the distributions are comparable, methods as simple as global median centering can fix the signal shift, while quantile normalization can help in case of divergent distributions.

## 3.2 Normalization

In large-scale experiments, the total intensity of the samples is likely to be different due to a number of reasons, such as different amount of sample loaded or fluctuations in measurement instrument sensitivity. To make samples comparable, they need to be scaled. This process is called normalization. In **proBatch**, two normalization approaches are used: median centering and quantile normalization. The normalization function `normalize_data()` by default takes log-transformed data, and if needed, log-transformation can be done on-the-fly by specifying `log_base = 2` for log2-transformation.

### 3.2.1 Median normalization

Median normalization is a conservative approach that shifts the intensity of the sample to the global median of the experiment. If the distributions of samples are dramatically different and this cannot be explained by non-technical factors, such as heterogeneity of samples, other approaches, such as quantile normalization need to be used.

```
median_normalized_matrix = normalize_data(log_transformed_matrix,
                                         normalizeFunc = "medianCentering")
```

Same result will be achieved with:

```
median_normalized_matrix = normalize_data(example_matrix,
                                         normalizeFunc = "medianCentering", log_base = 2)
```


### 3.2.2 Quantile normalization

Quantile normalization sets different distributions of individual samples to the same quantiles, which forces the distribution of the raw signal intensities to be the same in all samples. This method is computationally effective and has simple assumption that the majority of features (genes, proteins) is constant among the samples, thus also the distribution in principle are identical.

```
quantile_normalized_matrix = normalize_data(log_transformed_matrix,
                                           normalizeFunc = "quantile")
```

After quantile or median normalization, you can easily check if the global pattern improved by generating mean or boxplots and comparing them side by side. Here are the mean plots before and after normalization of the log transformed dataset.

```
plot_sample_mean(quantile_normalized_matrix, example_sample_annotation,
                 color_by_batch = TRUE, ylimits = c(12, 16),
                 color_scheme = color_list[[batch_col]])
```



### 3.3 Diagnostics of batch effects in normalized data

Now is the time to diagnose for batch effects and evaluate to what extent technical variance still exists in the normalized data matrix. The positive effect of normalization is sometimes not sufficient to control for peptide and protein-specific biases associated with a certain batch source. These biases can be identified via diagnostic plots. Here, we describe our essential toolbox of batch effect diagnostic approaches. Note that sample annotation and/or peptide annotation are necessary for the implementation of these plots.

#### 3.3.1 Hierarchical clustering


Hierarchical clustering is an algorithm that groups similar samples into a tree-like structure called dendrogram. Similar samples cluster together and the driving force of this similarity can be visualized by coloring the leaves of the dendrogram by technical and biological variables.

Our package provides `plot_sample_clustering()` and `plot_heatmap()` to plot the dendrogram by itself or with a heatmap. You can easily color annotations on the leaves of the dendrograms or heatmaps to identify what is the driving force of the clustering.

Once your color annotation is ready, for the specific covariates of interest, you can subset the color dataset and feed it into the clustering functions.


```
color_annotation <- color_scheme$color_df
selected_annotations <- c("MS_batch", "digestion_batch", "Diet")
#select only a subset of samples for plotting
color_annotation <- color_annotation %>% select(one_of(selected_annotations))

#Plot clustering between samples
plot_hierarchical_clustering(quantile_normalized_matrix, color_annotation,
                             distance = "euclidean", agglomeration = 'complete',
                             label_samples = FALSE)
```



Similarly, you can plot a heatmap by supplementing the color list. You decide whether to show annotations in the column, row or both by specifying the required covariates with `sample_annotation_col`, `sample_annotation_row`, or both.

```
plot_heatmap(quantile_normalized_matrix, example_sample_annotation,
             sample_annotation_col = selected_annotations,
             cluster_cols = TRUE,
             annotation_color_list = color_scheme$list_of_colors,
             show_rownames = FALSE, show_colnames = FALSE)
```



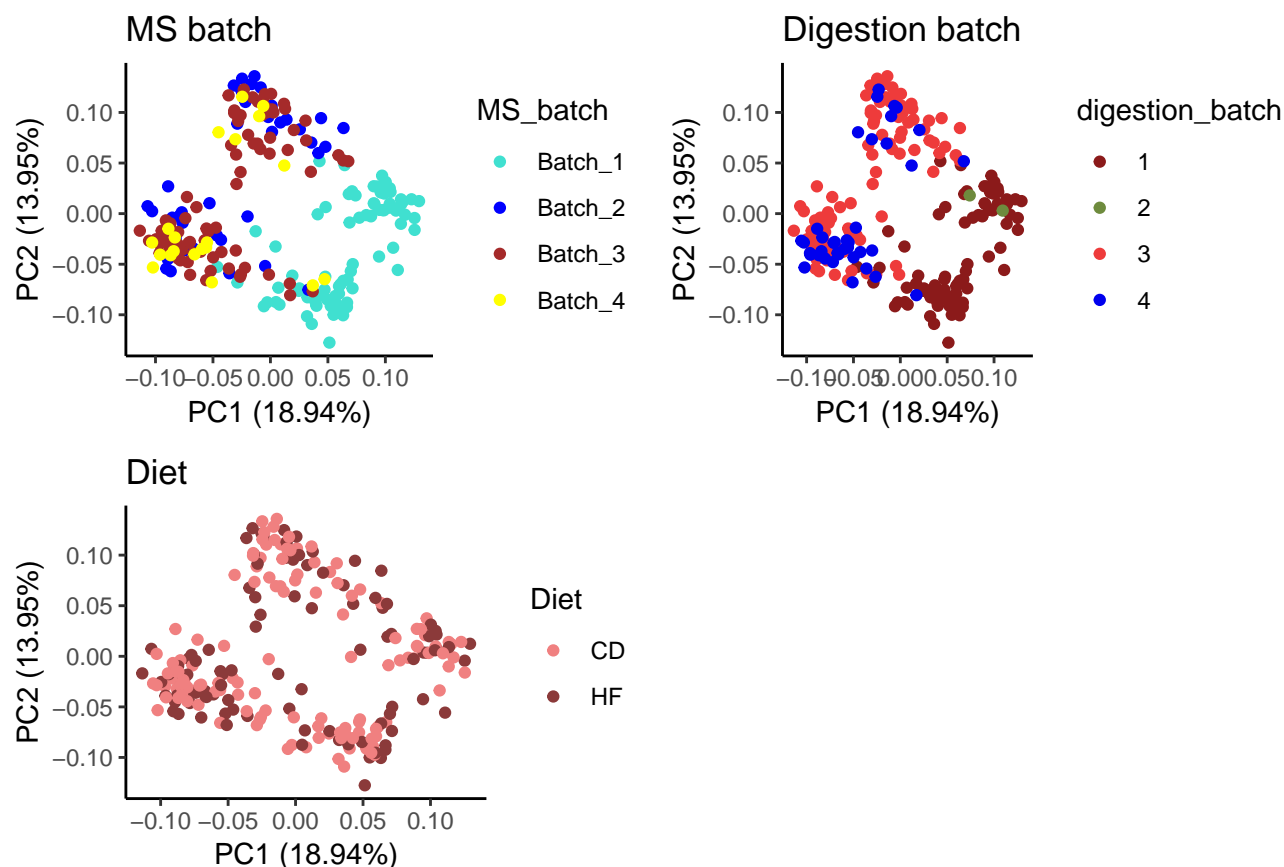
From the clustering analysis, we can clearly see that the driving force behind the sample clustering is the MS batch.

### 3.3.2 Principal component analysis (PCA)

PCA is a technique that identifies the leading directions of variation, known as principal components. The projection of data on two principal components allows to visualize sample proximity. This technique is particularly convenient to assess replicate similarity.

You can identify the covariate leading the direction of variations by coloring potential candidates.

```
plot_PCA(quantile_normalized_matrix, example_sample_annotation, color_by = 'MS_batch',
         plot_title = "MS batch", colors_for_factor = color_list[[batch_col]])
plot_PCA(quantile_normalized_matrix, example_sample_annotation, color_by = "digestion_batch",
         plot_title = "Digestion batch", colors_for_factor = color_list[["digestion_batch"]])
plot_PCA(quantile_normalized_matrix, example_sample_annotation, color_by = "Diet",
         plot_title = "Diet", colors_for_factor = color_list[["Diet"]])
```




By plotting the first two principal components and applying different color overlaps, we see once again that the clusters overlap nicely with the MS batches.

### 3.3.3 Principal variance component analysis (PVCA)

The main advantage of this approach is the quantification of the variance, associated with both technical and biological covariates. Briefly, principal variance component analysis uses a linear model to match each principal component to the sources of variation and weighs the variance of each covariate by the eigenvalue of the PC [7]. Thus, the resulting value reflects the variance explained by that covariate.

**NB:** PVCA calculation is a computationally demanding procedure. For a data matrix of several hundred samples and several thousands of peptides it can easily take several hours. So it is generally a good idea to run this analysis as a stand-alone script on a powerful machine.

```
pvca <- plot_PVCA(quantile_normalized_matrix, example_sample_annotation,
                  technical_covariates = c('MS_batch', 'digestion_batch'),
                  biological_covariates = c(biological_covariates, biospecimen_id_col))
```



The biggest proportion of variance in the peptide measurement was derived from mass spectrometry batches. In a typical experiment, the overall magnitude of variances coming from biological factors should be high while technical variance should be kept at minimum<sup>1</sup>.

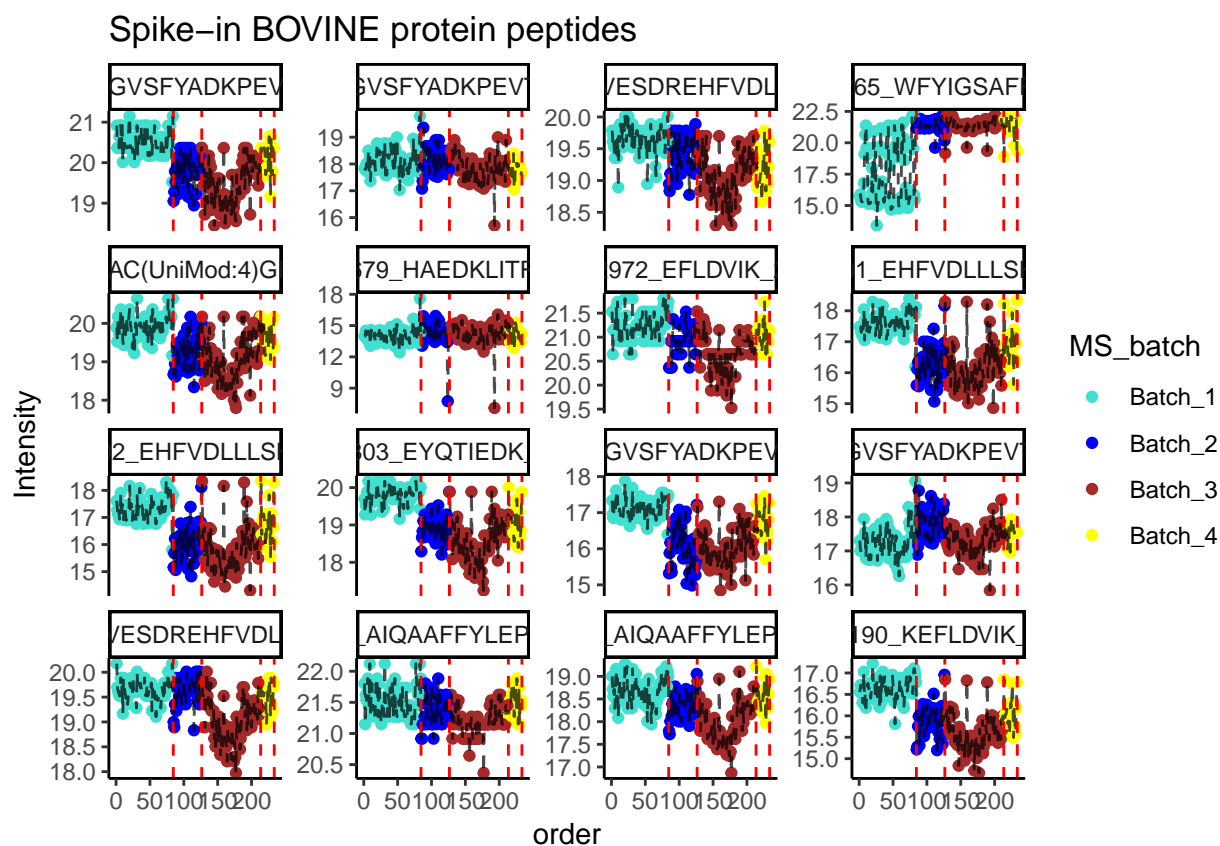
<sup>1</sup>Application of hierarchical clustering, PCA and PVCA in their classical implementation is not possible if missing values are present in the matrix. It has been noticed previously that missing values can be associated with technical bias [8], and most commonly, it is suggested that missing values need to be imputed [8-9]. However, we would like to suggest to use missing value imputation with extreme caution. First of all, missing value imputation alters the sample proximity. Additionally, imputed missing values, which can be obtained for SWATH data, can alter the correction of the batches.

### 3.3.4 Peptide-level diagnostics and spike-ins

Feature-level diagnostics are very informative for batch effect correction. To assess the bias in the data, one can choose a feature (peptide, protein, gene), the quantitative behaviour of which is known. In our package, `plot_peptides_of_one_protein()` allows plotting peptides of interest e.g. from biologically well understood protein. If spike-in proteins or peptides have been added to the mixture, one can use the `plot_spike_ins()` function instead. In most DIA datasets iRT peptides [10] are added in controlled quantities and can be visualized with the `plot_iRT()` function.

In mass-spectrometry, also the trends associated with this order can be assessed for a few representative peptides, thus the `order` column is also important for this diagnostics.

```
quantile_normalized_long <- matrix_to_long(quantile_normalized_matrix, example_sample_annotation)
plot_spike_in(quantile_normalized_long, example_sample_annotation,
  peptide_annotation = generated_peptide_annotation,
  protein_col = 'Gene', spike_ins = "BOVINE_Alag",
  plot_title = 'Spike-in BOVINE protein peptides',
  color_by_batch = TRUE, color_scheme = color_list[[batch_col]])
```



It is clear that while the pre-determined quantities of spike-ins or peptides of known biology have their expected intensities, the trend is dominated by mass spectrometry signal drift. After confirming either continuous or discrete batch effects exist in a dataset, by one or more of these methods, proceed by selecting a batch correction method.

### 3.4 Correction of batch effects

Depending on the type of batch effects, different batch correction methods should be implemented. In most cases, batch-specific signal shift needs to be corrected. For this case, feature median-centering can be applied, or, to use across-feature information in a Bayesian framework, the ComBat approach can be used. If there is continuous drift in the data, one has to start from continuous drift correction.

#### 3.4.1 Continuous drift correction


Continuous drifts are specific to mass-spectrometry, thus, the user-friendly methods for its correction have not been implemented before. In this package, we suggest a novel procedure to correct for MS signal drift. We developed a procedure based on nonlinear LOESS fitting. For each peptide and each batch, a non-linear trend is fitted to the normalized data and this trend is subtracted to correct for within-batch variation. Note, that the resulting data are not batch-free as within-batch means and variances are batch-dependent. However, now the batches are discrete and thus can be corrected using discrete methods such as median-centering or ComBat.

```
loess_fit <- adjust_batch_trend(quantile_normalized_matrix, example_sample_annotation)
loess_fit_matrix <- loess_fit$data_matrix
```

One important parameter in LOESS fitting is span, which determines the degree of smoothing. The LOESS span ranges from 0 to 1: the greater the value of span, the smoother is the fitted curve. Since we want the curve to reflect signal drift, we want to avoid overfitting but be sensitive to fit the trend accurately. Currently, we suggest to evaluate several peptides to determine the best smoothing degree for a given dataset.


```
loess_fit_30 <- adjust_batch_trend(quantile_normalized_matrix, example_sample_annotation, span = 0.3)

quantile_normalized_long <- matrix_to_long(quantile_normalized_matrix)
plot_with_fitting_curve(pep_name = "10231_QDVDVWLWQQEGSSK_2",
  df_long = quantile_normalized_long, example_sample_annotation,
  color_by_batch = TRUE, color_scheme = color_list[[batch_col]],
  fit_df = loess_fit_30$fit_df, plot_title = "Span = 30%")
```



```
loess_fit_70 <- adjust_batch_trend(quantile_normalized_matrix, example_sample_annotation, span = 0.7)
plot_with_fitting_curve(pep_name = "10231_QDVDVWLWQQEGSSK_2",
  df_long = quantile_normalized_long, example_sample_annotation,
  color_by_batch = TRUE, color_scheme = color_list[[batch_col]],
  fit_df = loess_fit_70$fit_df, plot_title = "Span = 70%")
```





Curve fitting is largely dependent on the number of consecutive measurements. In proteomics, missing values are not uncommon. If too many points are missing, a curve cannot be fit accurately. This is especially common for small batches. In this case, we suggest to not fit the curve to the specific peptide within the specific batch, and proceed directly to discrete correction methods. To identify such peptides, absolute and relative thresholds (`abs.threshold` and `pct.threshold`) on the number of missing values for each peptide can be passed as parameters to `normalize_custom_fit()`.

### 3.4.2 Discrete batch correction: combat or peptide-level median centering

Once the data are normalized and corrected for continuous drift, only discrete batch effects is left to be corrected.


Currently, both batch correction procedures accept long format of the data:

```
loess_fit_df <- matrix_to_long(loess_fit_matrix)
```

#### 3.4.2.1 Feature-level median centering

Feature-level median centering is the simplest approach for batch effects correction. However, if the variance is different between batches, other approaches need to be used.

```
peptide_median_df <- center_peptide_batch_medians(loess_fit_df, example_sample_annotation)
plot_single_feature(pep_name = "10231_QDQDVWLWQQEGSSK_2", df_long = peptide_median_df,
  example_sample_annotation, color_by_col = NULL, measure_col = 'Intensity_normalized',
  plot_title = "Feature-level Median Centered")
```




#### 3.4.2.2 ComBat

ComBat is well-suited for batches with distinct distributions, but restricted to peptides that don't have missing batch measurements. ComBat, uses parametric and non-parametric empirical Bayes framework for adjusting data for batch effects [11]. The function `correct_with_ComBat()` can incorporate several covariates and make data comparable across batches.

```
comBat_matrix <- correct_with_ComBat(loess_fit_matrix, example_sample_annotation)
#> Standardizing Data across genes
```

To illustrate the correction we use the “10231\_QDVDVWLWQQEGSSK\_2” spike-in peptide.

```
combat_df <- matrix_to_long(comBat_matrix)
plot_single_feature (pep_name = "10231_QDVDVWLWQQEGSSK_2", loess_fit_df,
                     example_sample_annotation, plot_title = "Loess Fitted", color_by_col = NULL)
plot_single_feature (pep_name = "10231_QDVDVWLWQQEGSSK_2", combat_df,
                     example_sample_annotation, plot_title = "ComBat corrected", color_by_col = NULL)
```



ComBat fixed the discrete batch effects and also made the distributions between batches similar to one another.

### 3.4.3 Correct batch effects: universal function

We provide a convenient all-in-one function for batch correction. The function `correct_batch_effects()` corrects MS signal drift and discrete shift in a single function call. Simply specify which discrete correction method is preferred at `discreteFunc` either “ComBat” or “MedianCentering” and supplement other arguments such as `span`, `abs.threshold` or `pct.threshold` as in `normalize_custom_fit()`.

```
batch_corrected_matrix <- correct_batch_effects(data_matrix = quantile_normalized_matrix,
                                                example_sample_annotation, discreteFunc = 'ComBat',
                                                abs.threshold = 5, pct.threshold = 0.20)
#> Standardizing Data across genes
```

### 3.5 Quality control on batch-corrected data matrix

In most cases, the batch effects correction method is evaluated by its ability to remove technical confounding, visible on hierarchical clustering or PCA. However, it is rarely shown whether the biological signal is not destroyed, or, better even, improved. Often, and increase in the number of differentially expressed genes is presented as an improvement. However, every reasonably designed experiment has replicates that can serve as an excellent control. In addition, peptides within a given protein should behave similarly and correlation of these peptides should improve after batch correction.

#### 3.5.1 Heatmap of selected replicate samples

In this study, 10 samples were run in the same order before and after the tuning of the mass-spectrometer, which marks the boundary between batches 2 and 3. The correlation between these replicates can be illustrated by `corrplot` (`flavor = 'corrplot'`) or by “pretty heatmap”, of “pheatmap” (`flavor = 'pheatmap'`)

First, we specify, which samples we want to correlate

```
earTags <- c("ET1524", "ET2078", "ET1322", "ET1566", "ET1354", "ET1420", "ET2154",
            "ET1515", "ET1506", "ET2577", "ET1681", "ET1585", "ET1518", "ET1906")

# Prepare color annotation
factors_to_show = c("MS_batch", "EarTag")
replicate_annotation <- example_sample_annotation %>%
  filter(MS_batch == 'Batch_2' | MS_batch == "Batch_3") %>%
  filter(EarTag %in% earTags) %>%
  remove_rownames() %>%
  column_to_rownames(var="FullRunName") %>%
  select(factors_to_show) # Annotate MS_batch and EarTag on pheatmap


# sample ID of biological replicates
replicate_filenames = replicate_annotation %>%
  rownames()

breaksList <- seq(0.7, 1, by = 0.01) # color scale of pheatmap
heatmap_colors = colorRampPalette(
  rev(RColorBrewer::brewer.pal(n = 7, name = "RdYlBu")))(length(breaksList))

# Plot the heatmap
plot_sample_corr_heatmap(quantile_normalized_matrix, samples_to_plot = replicate_filenames,
  flavor = 'pheatmap', plot_title = 'Quantile Normalized',
  annotation_colors = color_list[factors_to_show],
  annotation_col = replicate_annotation,
  color = heatmap_colors, breaks = breaksList,
  cluster_rows= FALSE, cluster_cols=FALSE, fontsize = 4,
  annotation_names_col = TRUE, annotation_legend = FALSE,
  show_colnames = FALSE)

plot_sample_corr_heatmap(batch_corrected_matrix, samples_to_plot = replicate_filenames,
  flavor = 'pheatmap', plot_title = 'Batch Corrected',
  annotation_colors = color_list[factors_to_show],
  annotation_col = replicate_annotation,
  color = heatmap_colors, breaks = breaksList,
  cluster_rows=FALSE, cluster_cols=FALSE, fontsize = 4,
  annotation_names_col = TRUE, annotation_legend = FALSE,
```

```
show_colnames = FALSE)
```



Before the correction, samples from one batch correlate better and of ten higher than the replicates. However, after the correction, the correlation between replicates becomes higher than the correlation between non-related samples regardless of the batch.

### 3.5.2 Correlation distribution of samples

For example, in the mice aging experiment, biological replicates, ET1506 and ET1524, were injected every 30-40 MS runs. The correlation between these biological replicates should improve after normalization and batch correction.

The `plot_sample_corr_distribution()` function plots correlation distribution between biological replicates and non-replicates in the same or different batches by `plot_param = 'batch_replicate'`. Alternatively, you can compute the correlation between different batches by `plot_param = 'batches'`.

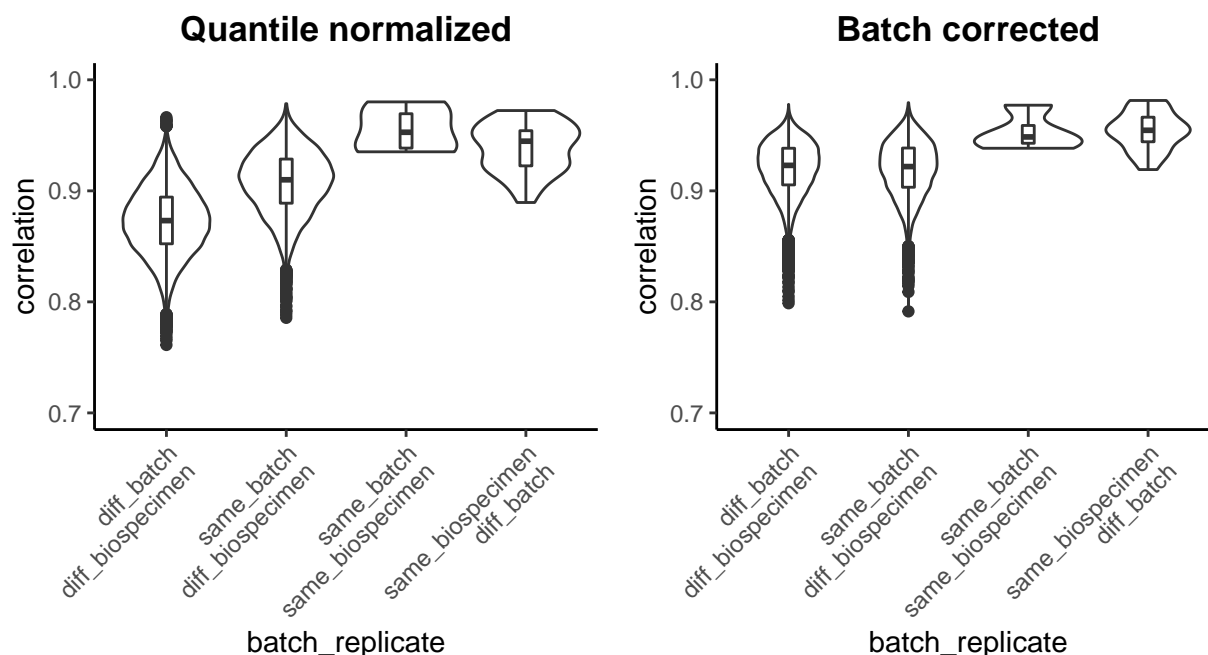
It should be noted, however, that the comparison of sample correlation should not be approached by evaluating the individual examples of within-replicate vs within-batch corrections, but rather by comparing the distribution. Unless these examples are shown in the context of the whole distribution structure, they can lead to erroneous conclusion. The sample correlation is often used to prove the quality of the measurement, as it is typically very high (examples of the replicate correlation above .95 are common for mass spectrometry).

```
sample_cor_norm <- plot_sample_corr_distribution(quantile_normalized_matrix,
  example_sample_annotation,
  batch_col = 'MS_batch',
  biospecimen_id_col = "EarTag",
  plot_title = 'Quantile normalized',
  plot_param = 'batch_replicate')

sample_cor_norm + theme(axis.text.x = element_text(angle = 45, hjust = 1)) + ylim(0.7,1)

sample_cor_batchCor <- plot_sample_corr_distribution(batch_corrected_matrix,
  example_sample_annotation,
  batch_col = 'MS_batch',
  plot_title = 'Batch corrected',
  plot_param = 'batch_replicate')

sample_cor_batchCor + theme(axis.text.x = element_text(angle = 45, hjust = 1)) + ylim(0.7, 1)
```



### 3.5.3 Correlation of peptide distributions within and between proteins

Peptides of the same protein are likely to correlate. Therefore, we can compare within- vs between-protein peptide correlation before and after batch correction to check if the correlation of peptides between the same proteins increases while that of different proteins stays the same.


**NB:** For a data matrix containing several thousands of peptides, calculation of peptide correlation is a computationally demanding procedure. It can easily take several hours. Therefore, it is generally recommended to run this analysis as a stand-alone script on a powerful machine.

The `plot_peptide_corr_distribution()` function plots correlation distribution between peptides of the same protein. However, an improvement of peptide correlation may not be clearly exhibited for a reduced dataset.

```
peptide_cor_norm <- plot_peptide_corr_distribution(quantile_normalized_matrix,
  generated_peptide_annotation, protein_col = 'Gene', plot_title = 'Quantile normalized')
peptide_cor_norm + geom_hline(yintercept=0, linetype="dashed", color = "grey")

peptide_cor_batchCor <- plot_peptide_corr_distribution(batch_corrected_matrix,
  generated_peptide_annotation, protein_col = 'Gene', plot_title = 'Batch corrected')
peptide_cor_batchCor + geom_hline(yintercept=0, linetype="dashed", color = "grey")
```

Here shows peptide correlation distribution generated from a complete dataset of the BXD mouse aging study:



## 4 SessionInfo

```
sessionInfo()
#> R version 3.6.0 (2019-04-26)
#> Platform: x86_64-w64-mingw32/x64 (64-bit)
#> Running under: Windows Server 2012 R2 x64 (build 9600)
#>
#> Matrix products: default
#>
#> locale:
#> [1] LC_COLLATE=C
#> [2] LC_CTYPE=English_United States.1252
#> [3] LC_MONETARY=English_United States.1252
```

```

#> [4] LC_NUMERIC=C
#> [5] LC_TIME=English_United_States.1252
#>
#> attached base packages:
#> [1] stats      graphics  grDevices  utils      datasets  methods    base
#>
#> other attached packages:
#> [1] knitr_1.22      proBatch_1.0.0 ggplot2_3.1.1  tibble_2.1.1
#> [5] dplyr_0.8.0.1
#>
#> loaded via a namespace (and not attached):
#> [1] nlme_3.1-139      bitops_1.0-6      matrixStats_0.54.0
#> [4] lubridate_1.7.4   robust_0.4-18     fit.models_0.5-14
#> [7] bit64_0.9-7       doParallel_1.0.14 RColorBrewer_1.1-2
#> [10] dynamicTreeCut_1.63-1 tools_3.6.0       backports_1.1.4
#> [13] R6_2.4.0          rpart_4.1-15      mgcv_1.8-28
#> [16] Hmisc_4.2-0       DBI_1.0.0          lazyeval_0.2.2
#> [19] BiocGenerics_0.30.0 colorspace_1.4-1  nnet_7.3-12
#> [22] withr_2.1.2       tidymodels_0.2.5  gridExtra_2.3
#> [25] bit_1.1-14        compiler_3.6.0    preprocessCore_1.46.0
#> [28] WGCNA_1.67        Biobase_2.44.0    htmlTable_1.13.1
#> [31] labeling_0.3       scales_1.0.0      checkmate_1.9.1
#> [34] DEoptimR_1.0-8    mutnorm_1.0-10    robustbase_0.93-4
#> [37] genefilter_1.66.0 readr_1.3.1        stringr_1.4.0
#> [40] digest_0.6.18     foreign_0.8-71    rmarkdown_1.12
#> [43] rrcov_1.4-7       base64enc_0.1-3   pkgconfig_2.0.2
#> [46] htmltools_0.3.6   limma_3.40.0      highr_0.8
#> [49] htmlwidgets_1.3   rlang_0.3.4       rstudioapi_0.10
#> [52] RSQLite_2.1.1     impute_1.58.0     BiocParallel_1.18.0
#> [55] acepack_1.4.1     RCurl_1.95-4.12   magrittr_1.5
#> [58] GO.db_3.8.2       Formula_1.2-3     Matrix_1.2-17
#> [61] Rcpp_1.0.1        munsell_0.5.0     S4Vectors_0.22.0
#> [64] viridis_0.5.1     stringi_1.4.3     yaml_2.2.0
#> [67] MASS_7.3-51.4     plyr_1.8.4        grid_3.6.0
#> [70] blob_1.1.1        parallel_3.6.0    crayon_1.3.4
#> [73] lattice_0.20-38   splines_3.6.0     annotate_1.62.0
#> [76] hms_0.4.2         pillar_1.3.1      fastcluster_1.1.25
#> [79] reshape2_1.4.3    codetools_0.2-16  stats4_3.6.0
#> [82] XML_3.98-1.19     glue_1.3.1        evaluate_0.13
#> [85] latticeExtra_0.6-28 data.table_1.12.2  png_0.1-7
#> [88] foreach_1.4.4     gtable_0.3.0      purrr_0.3.2
#> [91] tidyr_0.8.3       assertthat_0.2.1  xfun_0.6
#> [94] xtable_1.8-4       survival_2.44-1.1 pcaPP_1.9-73
#> [97] viridisLite_0.3.0 pheatmap_1.0.12   iterators_1.0.10
#> [100] AnnotationDbi_1.46.0 memoise_1.1.0     IRanges_2.18.0
#> [103] cluster_2.0.9     corrplot_0.84     sva_3.32.0
#> [106] ggfortify_0.4.6

```

## 5 Citation

To cite this package, please use:



```

citation('proBatch')
#>
#> To cite proBatch in publications use:
#>
#> Cuklina J., Lee C.H., Williams E.G., Collins B., Sajic T.,
#> Pedrioli P., Rodriguez-Martinez M., Aebersold R. Chapter 3:
#> Systematic overview of batch effects in proteomics Doctoral
#> thesis, ETH Zurich, 2018
#> https://doi.org/10.3929/ethz-b-000307772
#>
#> A BibTeX entry for LaTeX users is
#>
#> @PhdThesis{
#>   title = {Computational challenges in biomarker discovery from high-throughput proteomic data},
#>   author = {Jelena Cuklina and Chloe H. Lee and Evan G. Williams and Ben Collins and Tatjana Sajic},
#>   url = {https://doi.org/10.3929/ethz-b-000307772},
#>   school = {ETH Zurich},
#>   year = {2018},
#> }

```

## 6 References

- [1] O. T. Schubert, H. L. Röst, B. C. Collins, G. Rosenberger, and R. Aebersold. «Quantitative proteomics: challenges and opportunities in basic and applied research». *Nature Protocols* 12:7 (2017), pp. 1289–1294.
- [2] E. G. Williams, Y. Wu, P. Jha et al. «Systems proteomics of liver mitochondria function». *Science* 352:6291 (2016), aad0189.
- [3] Y. Liu, A. Buil, B. C. Collins et al. «Quantitative variability of 342 plasma proteins in a human twin population». *Molecular Systems Biology* 11:2 (2015), pp. 786–786.
- [4] H. L. Rost, G. Rosenberger, P. Navarro et al. «OpenSWATH enables automated, targeted analysis of data-independent acquisition MS data.» *Nature biotechnology* 32:3 (2014), pp. 219–23.
- [5] P. G. Pedrioli, «Trans-Proteomic Pipeline: A Pipeline for Proteomic Analysis.» *Methods in Molecular Biology Proteome Bioinformatics*, May 2009, pp. 213–238.
- [6] G. Rosenberger et al. «Statistical control of peptide and protein error rates in large-scale targeted data-independent acquisition analysis.» *Nature Methods* 14:9 (2017), pp. 921–927.
- [7] P. R. Bushel. *pvca: Principal Variance Component Analysis (PVCA)*. Package version 1.18.0. 2013.
- [8] Y. V. Karpievitch, A. R. Dabney, and R. D. Smith. «Normalization and missing value imputation for label-free LC-MS analysis». *BMC Bioinformatics* 13:Suppl 16 (2012), S5.
- [9] S. Tyanova, T. Temu, P. Sinitcyn et al. «The Perseus computational platform for comprehensive analysis of (prote)omics data». *Nat Methods* 13:9 (2016), pp. 731–740.
- [10] C. Escher, L. Reiter, B. Maclean et al. «Using iRT, a normalized retention time for more targeted measurement of peptides». *Proteomics* 12:8 (2012), pp. 1111–1121.
- [11] A. W. B. Johnston, Y. Li, and L. Ogilvie. «Metagenomic marine nitrogen fixation—feast or famine?» *Trends in microbiology* 13:9 (2005), pp. 416–20.