

Filter W4M data by sample class or feature attributes

Purpose of the `w4mclassfilter` package

The purpose of the `w4mclassfilter` R package is to provide the computational back-end of a Galaxy tool for inclusion in Workflow4Metabolomics (W4M).

This package (and the Galaxy tool) perform several steps, either to reduce the number samples or features to be analyzed, or to address several data issues that may impede downstream statistical analysis:

- Samples that are missing from either `sampleMetadata` or `dataMatrix` are eliminated.
- Samples may also be eliminated by a “sample class” column in `sampleMetadata`.
- Features that are missing from either `variableMetadata` or `dataMatrix` are eliminated.
- Features may be eliminated by specifying minimum or maximum value (or both) allowable in columns of `variableMetadata`.
- Features may be eliminated by specifying minimum or maximum intensity (or both) allowable in columns of `dataMatrix` for at least one sample for each feature (“range of row-maximum for each feature”).
- Missing values in `dataMatrix` are imputed either to zero (by default) or to the result of a user-supplied imputation function.
- Features and samples that have zero variance are eliminated.
- Samples and features are sorted alphabetically in rows and columns of `variableMetadata`, `sampleMetadata`, and `dataMatrix`
- By default, the names of the first columns of `variableMetadata` and `sampleMetadata` are set respectively to “`variableMetadata`” and “`sampleMetadata`”

How the `w4m_filter_by_sample_class` function is used

Ordinarily, a Galaxy tool wrapper invokes `w4m_filter_by_sample_class`.

For exploratory or debugging purposes, the package may be installed loaded into R and help may then be obtained with the following command:

```
?w4mclassfilter::w4m_filter_by_sample_class
```

W4M uses the XCMS and CAMERA packages to preprocess GC-MS or LC-MS data, producing three files that are documented in detail on the Workflow4Metabolomics (W4M) web site.

In summary:

1. `sampleMetadata.tsv`: a tab-separated file with metadata for the samples, one line per sample
 - One column of this file indicates the *class* of the sample.
 - It is the class that is used by this function to determine whether to include the sample in, or exclude the sample from, further analysis.
2. `variableMetadata.tsv`: a tab-separated file with metadata for the features detected, one line per feature
 - A feature is a location in the two dimensional space defined by the GC-MS or LC-MS data set, which corresponds to a compound or a group of compounds.
 - One dimension is the mass-to-charge ratio, *m/z*.
 - The other dimension is the *retention time*, i.e., how long until the solvent gradient eluted the compound(s) from the column.
3. `dataMatrix.tsv`: a tab separated file with the MS intensities for each sample for each feature:
 - There is one column per sample.

- There is one row per feature.
- If a feature is missing for a sample, the intensity value is NA.
- For numerical reasons, intensities may be negative, but this has no meaning in the real world.

Input- and Output-Format

Ordinary usage of the `w4mclassfilter::w4m_filter_by_sample_class` method is to read from and write to tab-delimited flat files (TSVs) because Galaxy presents datasets to tools as files. However, because general-purpose R packages usually use data structures in memory for their input and output, this function can accept not only with TSVs but also with data structures (`data.frame`, `matrix`, `list`, `env`); see ‘New in release v0.98.3’ below for details.

For all inputs and outputs that are file paths, those paths must be unique.

Feature- and Sample-Elimination

When `w4m_filter_by_sample_class` is invoked:

- an array of class names may be supplied in the `classes` argument. If the `include` argument is true, then only samples whose `class` column in `sampleMetadata` (as named in the `class_column` argument) will be *included* in the output; by contrast, if the `include` argument is false, then only samples whose `class` column in `sampleMetadata` will be *excluded* from the output.
 - Class names may be explicit, or regular expressions may be used.
 - See “New in release v0.98.2” below for details on using regular expressions.
- an array of range specification strings may be supplied in the `variable_range_filter` argument. If supplied, only features having numerical values in the specified column of `variableMetadata` that fall within the specified ranges will be retained in the output. Each range is a string of three colon-separated values (e.g., “mz:200:800”) in the following order:
 - the name of a column of `variableMetadata` which must have numerical data (e.g., “mz”);
 - the minimum allowed value in that column for the feature to be retained (e.g., 200);
 - the maximum allowed value (e.g., 800).
 - if the “maximum” is less than the “minimum”, then the range is exclusive (e.g., “mz:800:200” means retain only features whose mz is NOT in the range 200-800)
- note for the range specification strings: if the name supplied in the first field is ‘FEATMAX’, then the string is defining the minimum (and possibly, though less useful, maximum) intensity for each feature in the `dataMatrix`. For example, “FEATMAX:1e6:” would specify that any feature would be excluded if no sample had an intensity for that feature greater than 1000000.
 - note that if the “maximum” is greater than the “minimum” then the FEATMAX range specification is ignored.

Note that even when no rows or columns of the input `dataMatrix` input have zero variance, there is the possibility that eliminating samples or features may result in some rows or columns having zero variance, adversely impacting downstream statistical analysis. Consequently, `w4m_filter_by_sample_class` eliminates these rows or columns and the corresponding rows from `sampleMetadata`, `variableMetadata`, and `dataMatrix`.

Imputation of Missing Values

The `w4mclassfilter::w4m_filter_imputation` function is the default imputation method used by `w4m_filter_by_sample_class`; if other methods are to be used in the Galaxy tool, they might best eventually be incorporated into the `w4mclassfilter` R package, although they could be implemented in any R code to be used by the Galaxy tool wrapper.

By default, the `w4m_filter_by_sample_class` function imputes negative and NA intensity values as zero using the following `w4m_filter_imputation` function.

```
w4m_filter_imputation <-
function(m) {
  # replace NA values with zero
  m[is.na(m)] <- 0
  # replace negative values with zero, if applicable
  m[m<0] <- 0
  # return matrix as the result
  return (m)
}
```

New in release v0.98.2 - support regular expressions

Beginning with v0.98.2, `w4mclassfilter` supports use of R regular expression patterns to select class-names.

The R `base::grep1` function (at the core of this functionality) uses POSIX 1003.2 standard regular expressions, which allow precise pattern-matching and are exhaustively defined at:

http://pubs.opengroup.org/onlinepubs/9699919799/basedefs/V1_chap09.html

However, only a few basic building blocks of regular expressions need to be mastered for most cases:

Within square brackets:

- “^” as the first character after the left bracket specifies that none listed characters should be matched
- “-” separates a range of characters, e.g., “4-7” or “b-f”.

Outside of square brackets:

- “^” matches the beginning of a class-name
- “\$” matches the end of a class-name
- “.” matches a single character
- “*” matches the character specified immediately before zero or more times

Caveat: The tool wrapper uses the comma (“,”) to split a list of sample-class names, so *commas may not be used within regular expressions for this tool*

First Example: Consider a field of class-names consisting of

marq3,marq6,marq9,marq12,front3,front6,front9,front12

this regular expression	matches this set of sample-class names
<code>^front[0-9][0-9]*\$</code>	“front3,front6,front9,front12”
<code>^[a-z][a-z]3\$</code>	“front3,marq3”
<code>^[a-z][a-z]12\$</code>	“front12,marq12”
<code>^[a-z][a-z][0-9]\$</code>	“front3,front6,front9,marq3,marq6,marq9”

Second Example: Consider these regular expression patterns as possible matches to a sample-class name

AB0123

this regular expression	matches this set of sample-class names
<code>^[A-Z][A-Z][0-9][0-9]*\$</code>	AB0123
<code>^[A-Z][A-Z]*[0-9][0-9]*\$</code>	AB0123
<code>^[A-Z][0-9]*</code>	AB0123, see Note 1.
<code>^[A-Z][A-Z][0-9]</code>	AB0123, see Note 2.

this regular expression	matches this set of sample-class names
<code>^[A-Z][A-Z]*[0-9][0-9]\$</code>	NO MATCH, see Note 3.
<code>^[A-Z][0-9]*\$</code>	NO MATCH, see Note 4.

- Note 1. The first character is a letter, “*” can specify zero characters, and end of line did not need to be matched.
- Note 2. The first two characters are letters and the third is a digit.
- Note 3. The name does not end with the pattern “`[A-Z][0-9][0-9]$`”, i.e., it ends with four digits, not two.
- Note 4. The pattern specifies that second character and all those that follow, if present, must be digits.

New in release v0.98.3 - more flexible input and output

R package v0.98.3 - add support for env, list, data.frame, and matrix I/O * To support XCMS outside the context of Galaxy, `w4m_filter_by_sample_class` now supports input from and output to data structures as follows:

Inputs:

1. `dataMatrix_in`

- if a string, treat as a file path (i.e., same as in previous releases)
- else if an env or list, read `dataMatrix_in$dataMatrix`
- else if a data.frame, use `as.matrix(dataMatrix_in)`
- else if a matrix, use directly
- else error

2. `sampleMetadata_in`

- if a string, treat as a file path (i.e., same as in previous releases)
- else if an env or list, read `sampleMetadata_in$sampleMetadata`
- else if a data.frame, use directly
- else error

3. `variableMetadata_in`

- if a string, treat as a file path (i.e., same as in previous releases)
- else if an env or list, read `sampleMetadata_in$sampleMetadata`
- else if a data.frame, use directly
- else error

Outputs:

1. `dataMatrix_out`

- if a string, treat as a file path (i.e., same as in previous releases)
- else if an env or list, write `dataMatrix_out$dataMatrix`
- else error

2. `sampleMetadata_out`

- if a string, treat as a file path (i.e., same as in previous releases)
- else if an env or list, write `sampleMetadata_out$sampleMetadata`
- else error

3. `variableMetadata_out`

- if a string, treat as a file path (i.e., same as in previous releases)
- else if an env or list, write `sampleMetadata_out$sampleMetadata`
- else error

New in release v0.98.6 - feature-attribute range filtering

R package v0.98.6

- Support eliminating features whose attributes fall outside specified ranges. For more detail, see “Feature- and Sample-Elimination” above.
- Sort sample names and feature names because some statistical tools expect the same order in `dataMatrix` row and column names as in the corresponding metadata files.

New in release v0.98.7 - rename first column of sample metadata

R package v0.98.7

- First column of sample metadata is by default renamed to “sampleMetadata” unless argument ‘name_smplmetadata_col1’ is supplied and set to FALSE.