

Generating heatmaps for Nonnegative Matrix Factorization

Package *NMF* - Version 0.16.4

Renaud Gaujoux

April 16, 2013

Abstract

This vignette describes how to produce different informative heatmaps from NMF objects, such as returned by the function `nmf` in the *NMF* package¹ (Gaujoux et al. 2010). The main drawing engine is implemented by the function `aheatmap`, which is a highly enhanced modification of the function `pheatmap` from the *pheatmap* package², and provides convenient and quick ways of producing high quality and customizable annotated heatmaps. Currently this function is part of the package *NMF*, but may eventually compose a separate package on its own.

Contents

1 Preliminaries	1	4 Consensus matrix: consensusmap	7
1.1 Quick reminder on NMF models . . .	1	4.1 Single fit	8
1.2 Heatmaps for NMF	1	4.2 Single method over a range of ranks	9
1.3 Heatmap engine	2	4.3 Single rank over a range of methods	10
1.4 Data and model	2	5 Generic heatmap engine: aheatmap	11
2 Mixture Coefficient matrix: coefmap	4	6 Session Info	11
3 Basis matrix: basismap	6	6 References	12

1 Preliminaries

1.1 Quick reminder on NMF models

Given a nonnegative target matrix X of dimension $n \times p$, NMF algorithms aim at finding a rank k approximation of the form:

$$X \approx WH,$$

where W and H are nonnegative matrices of dimensions $n \times k$ and $k \times p$ respectively.

The matrix W is the basis matrix, whose columns are the basis components. The matrix H is the mixture coefficient or weight matrix, whose columns contain the contribution of each basis component to the corresponding column of X . We call the rows of H the basis profiles.

1.2 Heatmaps for NMF

Because NMF objects essentially wrap up a pair of matrices, heatmaps are convenient to visualise the results of NMF runs. The package *NMF* provides several specialised heatmap functions, designed to produce heatmaps with sensible default configurations according to the data being drawn. Being all based on a common drawing engine, they share almost identical interfaces and capabilities. The following specialised functions are currently implemented:

¹<http://cran.r-project.org/package=NMF>

²<http://cran.r-project.org/package=pheatmap>

`basismap` draws heatmaps of the basis matrix

`coefmap` draws heatmaps of the mixture coefficient matrix

`consensusmap` draws heatmaps of the consensus matrix, for results of multiple NMF runs.

1.3 Heatmap engine

All the above functions eventually call a common heatmap engine, with different default parameters, chosen to be relevant for the given underlying data. The engine is implemented by the function `aheatmap`. Its development started as modification of the function `pheatmap` from the *pheatmap* package. The initial objective was to improve and increase its capabilities, as well as defining a simplified interface, more consistent with the R core function `heatmap`. We eventually aim at providing a general, flexible, powerful and easy to use engine for drawing annotated heatmaps.

The function `aheatmap` has many advantages compared to other heatmap functions such as `heatmap`, `gplots::heatmap2`, `heatmap.plus::heatmap.plus`, or even `pheatmap`:

- Annotations: unlimited number of annotation tracks can be added to *both* columns and rows, with automated colouring for categorical and numeric variables.
- Compatibility with both base and grid graphics: the function can be directly called in drawing contexts such as `grid`, `mfrow` or `layout`. This is a feature many R users were looking for, and that was strictly impossible with base heatmaps.
- Legends: default automatic legend and colouring;
- Customisation: clustering methods, annotations, colours and legend can all be customised, even separately for rows and columns;
- Convenient interface: many arguments provide multiple ways of specifying their value(s), which speeds up developping/writing and reduce the amount of code required to generate customised plots (e.g. see ??).
- Aesthetics: the heatmaps look globally cleaner, the image and text components are by default well proportioned relatively to each other, and all fit within the graphic device.

1.4 Data and model

For the purpose of illustrating the use of each heatmap function, we generate a random target matrix, as well as some annotations or covariates:

```
# random data that follow an 3-rank NMF model (with quite some noise:
# sd=2)
X <- syntheticNMF(100, 3, 20, noise = 2)

# row annotations and covariates
n <- nrow(X)
d <- rnorm(n)
e <- unlist(mapply(rep, c("X", "Y", "Z"), 10))
e <- c(e, rep(NA, n - length(e)))
rdata <- data.frame(Var = d, Type = e)

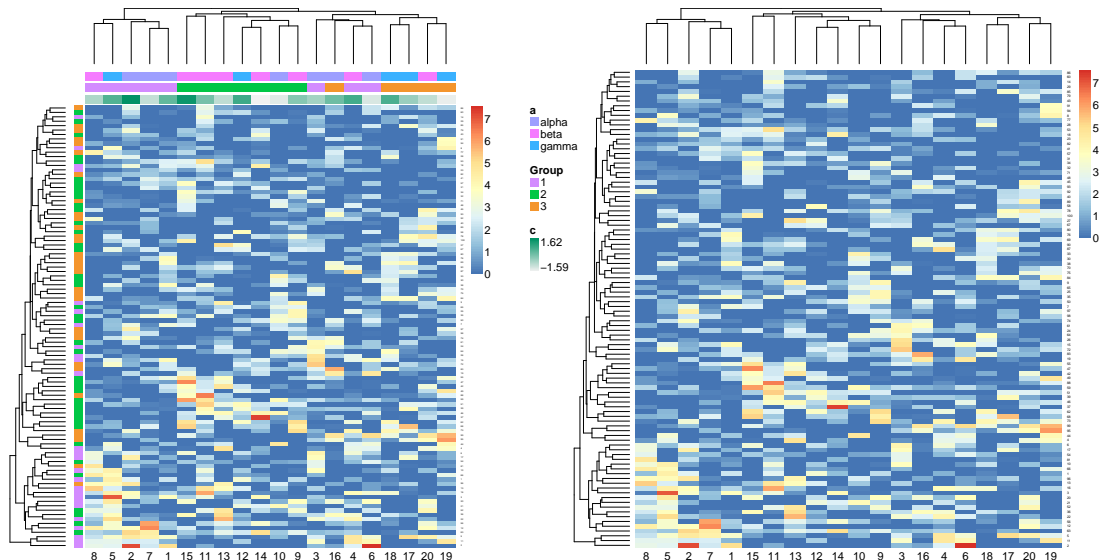
# column annotations and covariates
p <- ncol(X)
a <- sample(c("alpha", "beta", "gamma"), p, replace = TRUE)
```

```
# define covariates: true groups and some numeric variable
c <- rnorm(p)
# gather them in a data.frame
covariates <- data.frame(a, X$pData, c)
```

Note that in the code above, the object `X` returned by `syntheticNMF` *really is* a matrix object, but wrapped through the function `ExposedAttribute` object, which exposes its attributes via a more friendly and access controlled interface `$`. Of particular interests are attributes `'pData'` and `'fData'`, which are lists that contain a factor named `'Group'` that indicates the true underlying clusters. These are respectively defined as each sample's most contributing basis component and the basis component to which each feature contributes the most. They are useful to annotate heatmaps and assess the ability of NMF methods to recover the true clusters.

As an example, one can conveniently visualize the target matrix as a heatmap, with or without the relevant sample and feature annotations, using simple calls to the `aheatmap` function:

```
par(mfrow = c(1, 2))
aheatmap(X, annCol = covariates, annRow = X$fData)
aheatmap(X)
```



Then, we fit an NMF model using multiple runs, that will be used throughout this vignette to illustrate the use of NMF heatmaps:

```
res <- nmf(X, 3, nrun = 10)
res

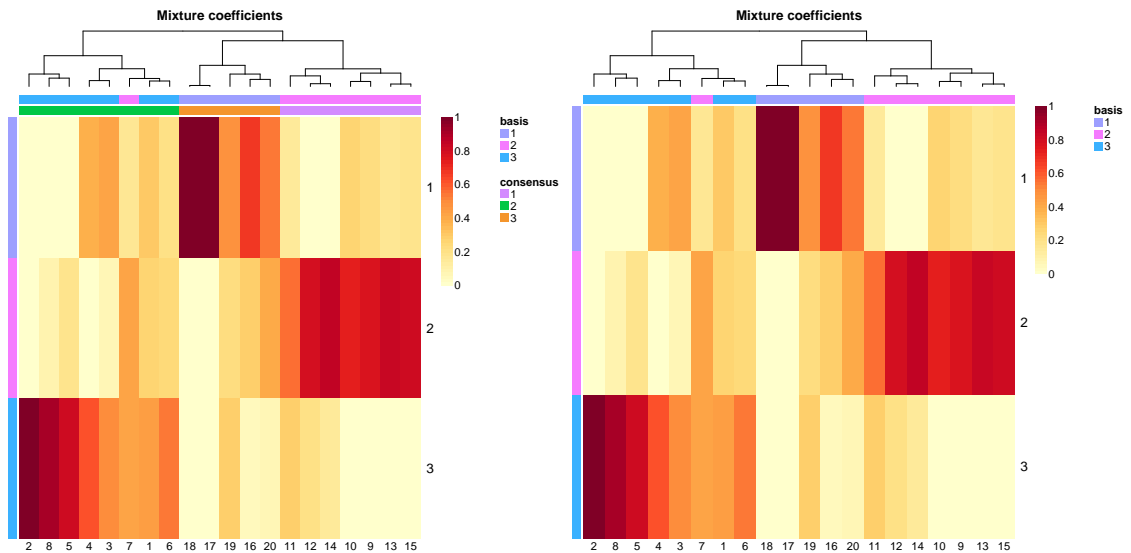
## <Object of class: NMFFitX1 >
##   Method: brunet
##   Runs:   10
##   RNG:
##   407L, 1836648330L, -690549885L, 21815016L, -792180791L, -1546669834L, 1782788863L
##   Total timing:
##   user  system elapsed
##   4.476   0.240   2.275
```

NB: To keep the vignette simple, we always use the default NMF method (i.e. 'brunet'), but all steps could be performed using a different method, or multiple methods in order to compare their performances.

2 Mixture Coefficient matrix: coefmap

The coefficient matrix of the result can be plotted using the function `coefmap`. The default behaviour for multiple NMF runs is to add two annotation tracks that show the clusters obtained by the best fit and the hierarchical clustering of the consensus matrix³. In the legend, these tracks are named *basis* and *consensus* respectively. For single NMF run or NMF model objects, no consensus data are available, and only the clusters from the fit are displayed.

```
opar <- par(mfrow = c(1, 2))
# coefmap from multiple run fit: includes a consensus track
coefmap(res)
# coefmap of a single run fit: no consensus track
coefmap(minfit(res))
```



```
par(opar)
```

NB: Note how both heatmaps were drawn on the same plot, simply using the standard call to `par(mfrow=c(1,2))`. This is impossible to achieve with the R core function `heatmap`. See Section 5 for more details about compatibility with base and grid graphics.

By default:

- the rows are not ordered;

³The hierarchical clustering is computed using the consensus matrix itself as a similarity measure, and average linkage. See `?consensushc`.

- the columns use the default ordering of `aheatmap`, but may easily be ordered according to the clusters defined by the dominant basis component for each column with `Colv="basis"`, or according to those implied by the consensus matrix, i.e. as in `consensusmap`, with `Colv="consensus"`;
- each column is scaled to sum up to one;
- the color palette used is 'YlOrRd' from the *RColorBrewer* package⁴ (Neuwirth 2011), with 50 breaks.

In term of arguments passed to the heatmap engine `aheatmap`, these default settings translate as:


```
Rowv = NA
Colv = TRUE
scale = "c1"
color = "YlOrRd:50"
annCol = predict(object) + predict(object, "consensus")
```

If the ordering does not come from a hierarchical clustering (e.g., if `Colv='basis'`), then no dendrogram is displayed. The default behaviour of `aheatmap` can be obtained by setting arguments `Rowv=TRUE`, `Colv=TRUE`, `scale='none'`.

The automatic annotation tracks can be hidden all together by setting argument `tracks=NA`, displayed separately by passing only one of the given names (e.g. `tracks=':basis'` or `tracks='basis:'` for the row or column respectively), and their legend names may be changed by specifying e.g. `tracks=c(Metagene=':basis', 'consensus')`. Beside this, they are handled by the heatmap engine function `aheatmap` and can be customised as any other annotation tracks – that can be added via the same argument `annCol` (see Section 5 or `?aheatmap` for more details).

```
opar <- par(mfrow=c(1,2))
# removing all automatic annotation tracks
coefmap(res, tracks=NA)
# customized plot
coefmap(res, Colv = 'euclidean'
, main = "Metagene contributions in each sample", labCol = NULL
, annRow = list(Metagene=':basis'), annCol = list(':basis', Class=a, Index=c)
, annColors = list(Metagene='Set2')
, info = TRUE)
```

⁴<http://cran.r-project.org/package=RColorBrewer>



```
par(opar)
```


NB: The feature that allows to display some information about the fit at the bottom of the plot via argument `info=TRUE` is still experimental. It is helpful mostly when developing algorithms or doing an analysis, but would seldom be used in publications.

3 Basis matrix: basismap

The basis matrix can be plotted using the function `basismap`. The default behaviour is to add an annotation track that shows for each row the dominant basis component. That is, for each row, the index of the basis component with the highest loading.

This track can be disabled by setting `tracks=NA`, and extra row annotations can be added using the same argument `annRow`.

```
opar <- par(mfrow = c(1, 2))
# default plot
basismap(res)
# customized plot: only use row special annotation track.
basismap(res, main = "Metagenes", annRow = list(d, e), tracks = c(Metagene = ":basis"))
```



```
par(opar)
```

By default:

- the columns are not ordered;
- the rows are ordered by hierarchical clustering using default distance and linkage methods ('euclidean' and 'complete');
- each row is scaled to sum up to one;
- the color palette used is 'YlOrRd' from the *RColorBrewer* package⁵ (Neuwirth 2011), with 50 breaks.

In term of arguments passed to the heatmap engine `aheatmap`, these default settings translate as:

```
Colv = NA
scale = "r1"
color = "YlOrRd:50"
annRow = predict(object, "features")
```

4 Consensus matrix: consensusmap

When doing clustering with NMF, a common way of assessing the stability of the clusters obtained for a given rank is to consider the consensus matrix computed over multiple independent NMF runs, which is the average of the connectivity matrices of each separate run⁶. This procedure is usually repeated over a certain range of factorization ranks, and the results are compared to identify which rank gives the best clusters, possibly in the light of some extra knowledge one could have about the samples (e.g. covariates). The functions `nmf` and `consensusmap` make it easy to implement this whole process.

⁵<http://cran.r-project.org/package=RColorBrewer>

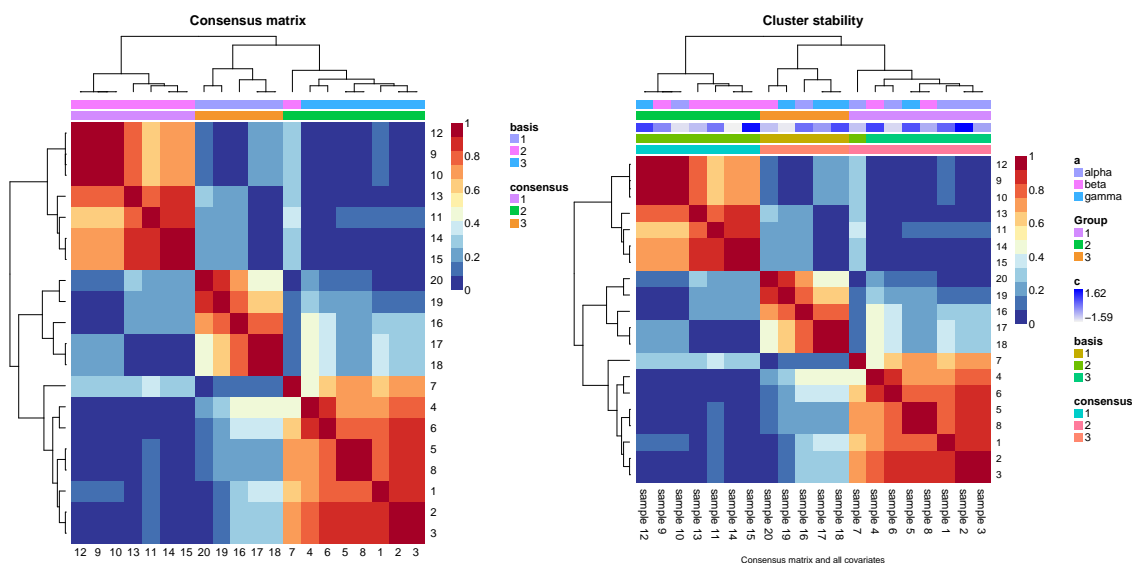
⁶Hence, stability here means robustness with regards to the initial starting point, and shall not be interpreted as in e.g. cross-validation/bootstrap analysis. However, one can argue that having very consistent clusters across runs somehow supports for a certain regularity or the presence of an underlying pattern in the data.

NB: The consensus plots can also be generated for fits obtained from single NMF runs, in which case the consensus matrix simply reduces to a single connectivity matrix. This is a binary matrix (i.e. entries are either 0 or 1), that will always produce a bi-colour heatmap, and by default clear blocks for each cluster.

4.1 Single fit

In section Section 1.4, the NMF fit `res` was computed with argument `nrun=10`, and therefore contains the best fit over 10 runs, as well as the consensus matrix computed over all the runs⁷. This can be plotted using the function `consensusmap`, which allows for the same kind of customization as the other NMF heatmap functions:

```
opar <- par(mfrow = c(1, 2))
# default plot
consensusmap(res)
# customized plot
consensusmap(res, annCol = covariates, annColors = list(c = "blue"), labCol = "sample ",
  main = "Cluster stability", sub = "Consensus matrix and all covariates")
```



```
par(opar)
```

By default:

- the rows and columns of the consensus heatmap are symmetrically ordered by hierarchical clustering using the consensus matrix as a similarity measure and average linkage, and the associated dendrogram is displayed;
- the color palette used is the reverse of 'RdYlBu' from the *RColorBrewer* package⁸ (Neuwirth 2011).

⁷If one were interested in keeping the fits from all the runs, the function `nmf` should have been called with argument `.options='k'`. See section *Options* in `?nmf`. The downstream handling of the result would remain identical.

⁸<http://cran.r-project.org/package=RColorBrewer>

In term of arguments passed to the heatmap engine `aheatmap`, these default settings translate as:

```
distfun = function(x) as.dist(1 - x) # x being the consensus matrix
hclustfun = "average"
Rowv = TRUE
Colv = "Rowv"
color = "-RdYlBu"
```

4.2 Single method over a range of ranks

The function `nmf` accepts a range of value for the rank (argument `rank`), making it fit NMF models for each value in the given range⁹:

```
res2_7 <- nmf(X, 2:7, nrun = 10, .options = "v")


## Compute NMF rank= 2 ... + measures ... OK
## Compute NMF rank= 3 ... + measures ... OK
## Compute NMF rank= 4 ... + measures ... OK
## Compute NMF rank= 5 ... + measures ... OK
## Compute NMF rank= 6 ... + measures ... OK
## Compute NMF rank= 7 ... + measures ... OK

class(res2_7)

## [1] "NMF.rank"
```

The result `res2_7` is an S3 object of class `'NMF.rank'`, that contains – amongst other data – a list of the best fits obtained for each value of the rank in range `[2, 7]`. The method of `consensusmap` defined for class `'NMF.rank'`, which plots all the consensus matrices on the same plot:

```
consensusmap(res2_7)
```



⁹Before version 0.6, this feature was provided by the function `nmfEstimateRank`. From version 0.6, the function `nmf` accepts ranges of ranks, and internally calls the function `nmfEstimateRank` – that remains exported and can still be called directly. See documentation `?nmfEstimateRank` for more details on the returned value.

NB: The main title of each consensus heatmap can be customized by passing to argument `main` a character vector or a list whose elements specify each title. All other arguments are used in each internal call to `consensusmap`, and will therefore affect all the plots simultaneously. The layout can be specified via argument `layout` as a numeric vector giving the number of rows and columns in a `mfrow`-like way, or as a matrix that will be passed to R core function `layout`. See `?consensusmap` for more details and example code.

4.3 Single rank over a range of methods

If one is interested in comparing methods, for a given factorization rank, then one can fit an NMF model for each method by providing the function `nmf` with a `list` in argument `method`:

```
res_methods <- nmf(X, 3, list("lee", "brunet", "nsNMF"), nrun = 10)

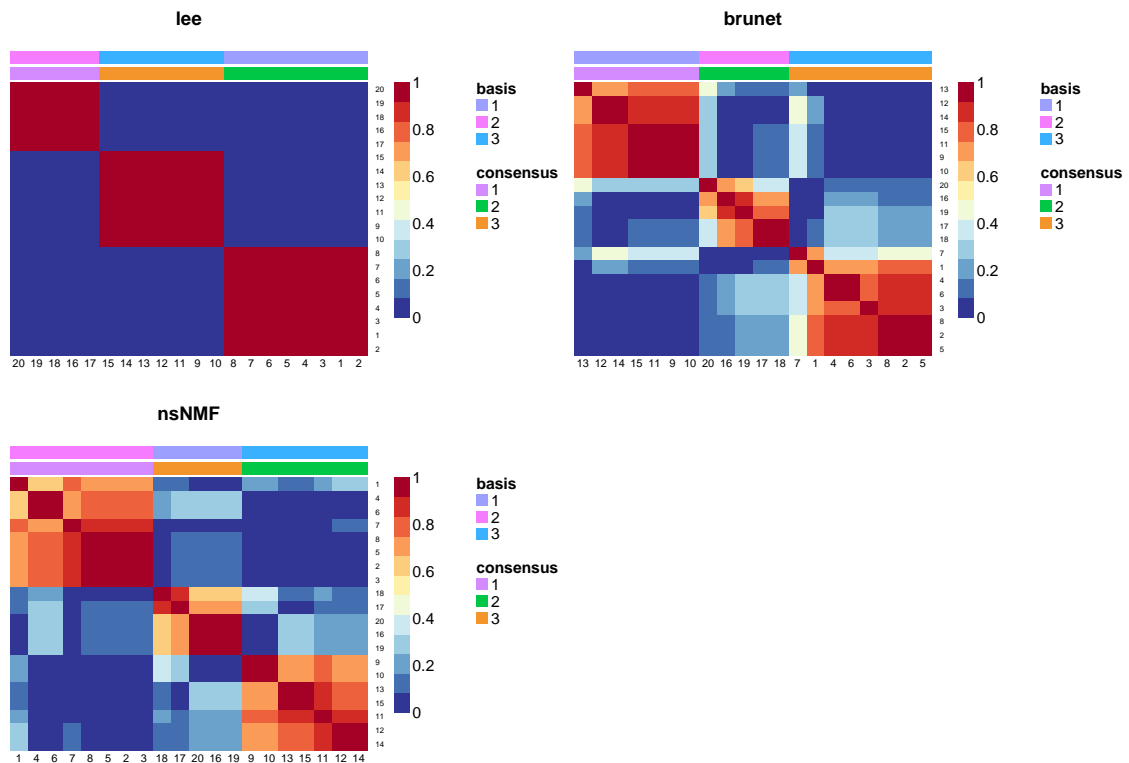
## Compute NMF method 'lee' [1/3] ... OK
## Compute NMF method 'brunet' [2/3] ... OK
## Compute NMF method 'nsNMF' [3/3] ... OK

class(res_methods)

## [1] "NMFList"
## attr(,"package")
## [1] "NMF"
```

The result `res_methods` is an S4 object of class `NMFList`, which is essentially a named list, that contains each fits and the CPU time required by the whole computation. As previously, the sequence of consensus matrices is plotted with `consensusmap`:

```
consensusmap(res_methods)
```



5 Generic heatmap engine: aheatmap

This section still needs to be written, but many examples of annotated heatmaps can be found in the demos 'aheatmap' and 'heatmaps':

```
demo("aheatmap")
# or
demo("heatmaps")
```

These demos and the plots they generate can also be browsed online at http://nmf.r-forge.r-project.org/_DEMOS.html.

6 Session Info

- R version 2.15.3 (2013-03-01), i686-pc-linux-gnu
- Locale: LC_CTYPE=en_US.UTF-8, LC_NUMERIC=C, LC_TIME=en_US.UTF-8, LC_COLLATE=en_US.UTF-8, LC_MONETARY=en_US.UTF-8, LC_MESSAGES=en_US.UTF-8, LC_PAPER=C, LC_NAME=C, LC_ADDRESS=C, LC_TELEPHONE=C, LC_MEASUREMENT=en_US.UTF-8, LC_IDENTIFICATION=C
- Base packages: base, datasets, graphics, grDevices, grid, methods, stats, utils
- Other packages: BH 1.51.0-0, bigmemory 4.4.3, bigmemory.sri 0.1.2, Biobase 2.18.0, BiocGenerics 0.4.0, digest 0.6.3, knitr 1.1, NMF 0.16.4, pkgmaker 0.15.4, RColorBrewer 1.0-5, registry 0.2, rngtools 1.2, synchronicity 1.1.0

- Loaded via a namespace (and not attached): codetools 0.2-8, colorspace 1.2-1, doParallel 1.0.1, evaluate 0.4.3, foreach 1.4.0, formatR 0.7, gridBase 0.4-6, iterators 1.0.6, parallel 2.15.3, stringr 0.6.2, tools 2.15.3, xtable 1.7-1

References

- [1] Renaud Gaujoux and Cathal Seoighe. “A flexible R package for nonnegative matrix factorization”. In: *BMC Bioinformatics* 11.1 (2010), p. 367. ISSN: 1471-2105. DOI: [10.1186/1471-2105-11-367](https://doi.org/10.1186/1471-2105-11-367). URL: <http://www.biomedcentral.com/1471-2105/11/367> (see p. 1).
- [2] Erich Neuwirth. *RColorBrewer: ColorBrewer palettes*. R package version 1.0-5. 2011. URL: <http://CRAN.R-project.org/package=RColorBrewer> (see pp. 5, 7, 8).