Package: proxyC (via r-universe)

October 30, 2026

Type Package

Title Computes Proximity in Large Matrices

Version 0.6.0

Description Computes proximity between rows or columns of large matrices efficiently in C++. Functions are optimised for large matrices using the Armadillo and Intel TBB libraries. Among various built-in similarity/distance measures, computation of correlation, cosine similarity, Dice coefficient and Euclidean distance is particularly fast.

URL https://github.com/koheiw/proxyC, https://koheiw.github.io/proxyC/

BugReports https://github.com/koheiw/proxyC/issues

License GPL-3

Depends R (>= 3.1.0), methods

Imports Matrix (>= 1.2), Rcpp (>= 0.12.12)

Suggests testthat, entropy, proxy, knitr, rmarkdown

LinkingTo Rcpp, RcppArmadillo (>= 0.7.600.1.0)

NeedsCompilation yes

Encoding UTF-8

VignetteBuilder knitr

Config/Needs/website coop, Rfast, parallelDist, distances, proxy, microbenchmark, ggplot2

Config/roxygen2/version 8.1.0

Author Kohei Watanabe [cre, aut, cph] (ORCID: <https://orcid.org/0000-0001-6519-5265>), Robrecht Cannoodt [aut] (ORCID: <https://orcid.org/0000-0003-3641-729X>)

Maintainer Kohei Watanabe <watanabe.kohei@gmail.com>

Repository https://cran.r-universe.dev

Date/Publication 2026-10-30 07:17:43 UTC

RemoteUrl https://github.com/cran/proxyC

RemoteRef HEAD

RemoteSha 9550c4c4b5d7544ae13f59f642582ce6181e0711

Contents

colMaxs 2
colSds 3
colZeros 3
crossprod 4
mask 5
maxCol 5
simil 6

Index **10**

`colMaxs` *Maximum/minimum values in columns or rows*

Description

Produce the same result as `apply(x, 2, max)`, `apply(x, 1, max)`, `apply(x, 2, min)` and `apply(x, 1, min)`, respectively, but ignore `NA`.

Usage

```
colMaxs(x, use_nan = FALSE)

rowMaxs(x, use_nan = FALSE)

colMins(x, use_nan = FALSE)

rowMins(x, use_nan = FALSE)
```

Arguments

`x`	a base::matrix or Matrix::Matrix object.
`use_nan`	if `TRUE`, return `NaN` for columns or rows with all zero. Only used when `x` is a sparse matrix.

Value

a numeric vector

Examples

```
mt <- Matrix::rsparsematrix(100, 100, 0.9)
colMaxs(mt)
apply(mt, 2, max) # the same
```

colSds *Standard deviation of columns and rows of large matrices*

Description

Produce the same result as `apply(x, 2, sd)` and `apply(x, 1, sd)`, respectively. without coercing matrix to dense matrix. Values are not identical to `sd()` because of the floating point precision issue in C++.

Usage

```
colSds(x, use_nan = FALSE)

rowSds(x, use_nan = FALSE)
```

Arguments

x	a base::matrix or Matrix::Matrix object.
use_nan	if `TRUE`, return `NaN` for columns or rows with all zero. Only used when `x` is a sparse matrix.

Value

a numeric vector

Examples

```
mt <- Matrix::rsparsematrix(100, 100, 0.9)
colSds(mt)
apply(mt, 2, sd) # the same
```

colZeros *Number of zeros in columns and rows of large matrices*

Description

Produce the same result as applying `sum(x == 0)` to each row or column.

Usage

```
colZeros(x)

rowZeros(x)
```

Arguments

x	a base::matrix or Matrix::Matrix object.

Value

an integer vector

Examples

```
mt <- Matrix::rsparsematrix(100, 100, 0.9)
colZeros(mt)
apply(mt, 2, function(x) sum(x == 0)) # the same
```

`crossprod` *Cross-product of large sparse matrices*

Description

Compute the (transposed) cross-product of large sparse matrices using the same infrastructure as `simil()` and `dist()`.

Usage

```
crossprod(x, y = NULL, min_prod = NULL, sparse = TRUE, digits = 14)

tcrossprod(x, y = NULL, min_prod = NULL, sparse = TRUE, digits = 14)
```

Arguments

`x`	a base::matrix or Matrix::Matrix object. Only sparse matrices are covered to the Matrix::CsparseMatrix internally.
`y`	if a base::matrix or Matrix::Matrix object is provided, proximity between documents or features in `x` and `y` is computed.
`min_prod`	the minimum product to be recorded.
`sparse`	if `TRUE`, returns Matrix::sparseMatrix object. When neither `min_simil` nor `rank` is used, dense matrices require less space in RAM.
`digits`	determines rounding of small values towards zero. Use primarily to correct floating point errors. Rounding is performed in C++ in a similar way as base::zapsmall.

Value

a Matrix::Matrix object

mask *Create a pattern matrix for masking*

Description

Create a pattern matrix for `simil()` or `dist()` to enable masked similarity computation. If the matrix is passed to the function, it computes similarity scores only for cells with `TRUE`.

Usage

```
mask(x, y = NULL)
```

Arguments

`x`	a numeric or character vector matched against each other.
`y`	a numeric or character vector matched against `x` if provided.

Value

a sparse logical matrix with `TRUE` for matched pairs.

Examples

```
mt1 <- Matrix::rsparsematrix(100, 6, 1.0)
colnames(mt1) <- c("a", "a", "d", "d", "e", "e")
mt2 <- Matrix::rsparsematrix(100, 5, 1.0)
colnames(mt2) <- c("a", "b", "c", "d", "e")

(msk <- mask(colnames(mt1), colnames(mt2)))
simil(mt1, mt2, margin = 2, mask = msk, drop0 = TRUE)
```

maxCol *Column or row indices of maximum values*

Description

Find column or row indices of maximum values in matrices. Works in a similar way as `base::max.col()` but without conversion of sparse matrices to dense matrices.

Usage

```
maxCol(x, ties_method = c("random", "first", "last"), use_nan = FALSE)

maxRow(x, ties_method = c("random", "first", "last"), use_nan = FALSE)
```

Arguments

x	a base::matrix or Matrix::Matrix object.
ties_method	specify how to handle ties. If "random", choose an index randomly within the ties; if "first" or "last" returns the index of the first or the last tie elements, respectively.
use_nan	if TRUE, return NA for columns or rows with all zero. Applies only when x is a sparse matrix.

Value

an integer vector

Examples

```
mt <- Matrix::rsparsematrix(100, 100, 0.9)
maxCol(mt, ties_method = "first")
max.col(mt, ties.method = "first") # the same
```

---

simil *Compute similarity/distance between rows or columns of large matrices*

Description

Fast similarity/distance computation function for large matrices using multiple processors. To save computation time and storage, it can also apply arbitrary threshold (min_simil) or rank (rank) to the similarity/distance scores and return them in a sparse matrix.

Usage

```
simil(
  x,
  y = NULL,
  margin = 1,
 method = c("cosine", "correlation", "dice", "edice", "jaccard", "ejaccard", "fjaccard",
    "hamann", "faith", "simple matching"),
  mask = NULL,
  min_simil = NULL,
  rank = NULL,
  drop0 = FALSE,
  diag = FALSE,
  use_nan = NULL,
  sparse = TRUE,
  digits = 14
)

dist(
```

```
  x,
  y = NULL,
  margin = 1,
 method = c("euclidean", "chisquared", "kullback", "jeffreys", "jensen", "manhattan",
    "maximum", "canberra", "minkowski", "hamming"),
  mask = NULL,
  p = 2,
  smooth = 0,
  drop0 = FALSE,
  diag = FALSE,
  use_nan = NULL,
  sparse = TRUE,
  digits = 14
)
```

Arguments

x	a base::matrix or Matrix::Matrix object. Only sparse matrices are covered to the Matrix::CsparseMatrix internally.
y	if a base::matrix or Matrix::Matrix object is provided, proximity between documents or features in x and y is computed.
margin	integer indicating margin of similarity/distance computation. 1 indicates rows or 2 indicates columns.
method	method to compute similarity or distance
mask	a pattern matrix created using mask() for masked similarity/distance computation. The shape of the matrix must be the same as the resulting matrix. Cannot be used when diag = TRUE.
min_simil	the minimum similarity value to be recorded.
rank	an integer value specifying top-n most similarity values to be recorded.
drop0	if TRUE, removes zero values to make the similarity/distance matrix sparse. It has no effect when dense = TRUE.
diag	if TRUE, only compute diagonal elements of the similarity/distance matrix; useful when comparing corresponding rows or columns of x and y.
use_nan	if TRUE, returns NaN if the standard deviation of a vector is zero when method is "correlation"; if all the values are zero in a vector when method is "cosine", "chisquared", "kullback", "jeffreys" or "jensen". Note that use of NaN makes the similarity/distance matrix denser and therefore larger in RAM. If FALSE, return zero in same use situations as above. If NULL, will also return zero but also generate a warning (default).
sparse	if TRUE, returns Matrix::sparseMatrix object. When neither min_simil nor rank is used, dense matrices require less space in RAM.
digits	determines rounding of small values towards zero. Use primarily to correct floating point errors. Rounding is performed in C++ in a similar way as base::zapsmall.
p	weight for Minkowski distance.
smooth	adds a fixed value to all the cells to avoid division by zero. Only used when method is "chisquared", "kullback", "jeffreys" or "jensen".

Details

Available Methods:

Similarity:

- `cosine`: cosine similarity
- `correlation`: Pearson's correlation
- `jaccard`: Jaccard coefficient
- `ejaccard`: the real value version of `jaccard`
- `fjaccard`: Fuzzy Jaccard coefficient
- `dice`: Dice coefficient
- `edice`: the real value version of `dice`
- `hamann`: Hamann similarity
- `faith`: Faith similarity
- `simple matching`: the percentage of common elements

Distance:

- `euclidean`: Euclidean distance
- `chisquared`: chi-squared distance
- `kullback`: Kullback–Leibler divergence
- `jeffreys`: Jeffreys divergence
- `jensen`: Jensen–Shannon divergence
- `manhattan`: Manhattan distance
- `maximum`: the largest difference between values
- `canberra`: Canberra distance
- `minkowski`: Minkowski distance
- `hamming`: Hamming distance

See the vignette for how the similarity and distance are computed: `vignette("measures", package = "proxyC")`

Parallel Computing:

It performs parallel computing using Intel oneAPI Threads Building Blocks. The number of threads for parallel computing should be specified via `options(proxyC.threads)` before calling the functions. If the value is -1, all the available threads will be used. Unless the option is used, the number of threads will be limited by the environmental variables (`OMP_THREAD_LIMIT` or `RCPP_PARALLEL_NUM_THREADS`) to comply with CRAN policy and offer backward compatibility.

Value

a Matrix::Matrix object

See Also

zapsmall

Examples

```
mt <- Matrix::rsparsematrix(100, 100, 0.01)
simil(mt, method = "cosine")[1:5, 1:5]
mt <- Matrix::rsparsematrix(100, 100, 0.01)
dist(mt, method = "euclidean")[1:5, 1:5]
```

Index

base::matrix, *2–4*, *6*, *7*
base::zapsmall, *4*, *7*

colMaxs, 2
colMins *(*colMaxs*)*, 2
colSds, 3
colZeros, 3
crossprod, 4

dist *(*simil*)*, 6
dist(), *4*, *5*

mask, 5
mask(), *7*
Matrix::CsparseMatrix, *4*, *7*
Matrix::Matrix, *2–4*, *6–8*
Matrix::sparseMatrix, *4*, *7*
maxCol, 5
maxRow *(*maxCol*)*, 5

rowMaxs *(*colMaxs*)*, 2
rowMins *(*colMaxs*)*, 2
rowSds *(*colSds*)*, 3
rowZeros *(*colZeros*)*, 3

simil, 6
simil(), *4*, *5*

tcrossprod *(*crossprod*)*, 4