

Package ‘parallelpam’

November 24, 2022

Type Package

Title Parallel Partitioning-Around-Medoids (PAM) for Big Sets of Data

Version 1.0

Date 2022-10-30

Author Juan Domingo [aut, cre] (<<https://orcid.org/0000-0003-4728-6256>>),
Guillermo Ayala [ctb] (<<https://orcid.org/0000-0002-6231-2865>>),
Spanish Ministry of Science and Innovation, MCIN/AEI
<[doi:10.13039/501100011033](https://doi.org/10.13039/501100011033)> [fnd]

Maintainer Juan Domingo <Juan.Domingo@uv.es>

Description Application of the Partitioning-Around-Medoids (PAM) clustering algorithm described in Schubert, E. and Rousseeuw, P.J.: “Fast and eager k-medoids clustering: O(k) runtime improvement of the PAM, CLARA, and CLARANS algorithms.” Information Systems, vol. 101, p. 101804, (2021). <[doi:10.1016/j.is.2021.101804](https://doi.org/10.1016/j.is.2021.101804)>. It uses a binary format for storing and retrieval of matrices developed for the 'jmatrix' package but the functionality of 'jmatrix' is included here, so you do not need to install it. Also, it is used by package 'scellpam', so if you have installed it, you do not need to install this package. PAM can be applied to sets of data whose dissimilarity matrix can be very big. It has been tested with up to 100.000 points. It does this with the help of the code developed for other package, 'jmatrix', which allows the matrix not to be loaded in 'R' memory (which would force it to be of double type) but it gets from disk, which allows using float (or even smaller data types). Moreover, the dissimilarity matrix is calculated in parallel if the computer has several cores so it can open many threads. The initial part of the PAM algorithm can be done with the BUILD or LAB algorithms; the BUILD algorithm has been implemented in parallel. The optimization phase implements the FastPAM1 algorithm, also in parallel. Finally, calculation of silhouette is available and also implemented in parallel.

License GPL (>= 2)

Imports Rcpp (>= 1.0.8), memuse (>= 4.2.1)

LinkingTo Rcpp

RoxygenNote 7.2.1**Encoding** UTF-8**Suggests** knitr, cluster**VignetteBuilder** knitr**NeedsCompilation** yes**Repository** CRAN**Date/Publication** 2022-11-24 12:10:04 UTC**R topics documented:**

ApplyPAM	2
CalcAndWriteDissimilarityMatrix	5
CalculateSilhouette	6
ClassifAsDataFrame	7
FilterBySilhouetteQuantile	9
FilterBySilhouetteThreshold	11
GetJCol	13
GetJColByName	14
GetJColNames	14
GetJManyCols	15
GetJManyColsByNames	16
GetJManyRows	16
GetJManyRowsByNames	17
GetJNames	18
GetJRow	19
GetJRowByName	19
GetJRowNames	20
GetSubdiag	21
GetTD	21
JMatInfo	22
JWriteBin	23
NumSilToClusterSil	24
ParallelpamSetDebug	25
Index	27

Description

A function to implement the Partitioning-around-medoids algorithm described in Schubert, E. and Rousseeuw, P.J.: "Fast and eager k-medoids clustering: O(k) runtime improvement of the PAM, CLARA, and CLARANS algorithms."

Information Systems, vol. 101, p. 101804, 2021.

doi: <https://doi.org/10.1016/j.is.2021.101804>

Notice that the actual values of the vectors (instances) are not needed. To recover them, look at the data matrix used to generate the distance matrix.

The number of instances, N, is not passed since dissimilarity matrix is NxN and therefore its size indicates the N value.

Usage

```
ApplyPAM(
    dissim_file,
    k,
    init_method = "BUILD",
    initial_med = NULL,
    max_iter = 1000L,
    nthreads = 0L
)
```

Arguments

dissim_file	A string with the name of the binary file that contains the symmetric matrix of dissimilarities. Such matrix should have been generated by CalcAndWriteDissimilarityMatrix and it is a matrix of type 'disttype' (in this type defined as float).
k	A positive integer (the desired number of medoids).
init_method	One of the strings 'PREV', 'BUILD' or 'LAB'. See meaning of initialization algorithms BUILD and LAB in the original paper. 'PREV' should be used exclusively to start the second part of the algorithm (optimization) from a initial set of medoids generated by a former call. Default: BUILD.
initial_med	A vector with initial medoids to start optimization. It is to be used only by the 'PREV' method and it will have been obtained as the first element (L\$med) of the two-element list returned by a previous call to this function used in just-initialize mode (max_iter=0). Default: empty vector.
max_iter	The maximum number of allowed iterations. 0 means stop immediately after finding initial medoids. Default: 1000
nthreads	For the BUILD initialization algorithm (the only part currently implemented in parallel), the number of used threads. -1 means don't use threads (serial implementation). 0 means let the program choose according to the number of cores and of points.

Any other number forces this number of threads. Choosing more than the number of available cores is allowed, but discouraged.

Default: 0

Details

With respect to the returned value, `L$med` has as many components as requested medoids and `L$clasif` has as many components as instances.

Medoids are expressed in `L$med` by its number in the array of points (row in the dissimilarity matrix) starting at 1 (R convention).

`L$clasif` contains the number of the medoid (i.e.: the cluster) to which each instance has been assigned, according to their order in

`L$med` (also from 1).

This means that if `L$clasif[p]` is `m`, the point `p` belongs to the class grouped around medoid `L$med[m]`.

Moreover, if the dissimilarity matrix contains as metadata the point names, the returned are R-named vector with such names.

Value

`L["med","clasif"]` A list of two numeric vectors. See section Details for more information

Examples

```
# Synthetic problem: 10 random seeds with coordinates in [0..20]
# to which random values in [-0.1..0.1] are added
M<-matrix(0,100,500)
rownames(M)<-paste0("rn",c(1:100))
for (i in (1:10))
{
  p<-20*runif(500)
  Rf <- matrix(0.2*(runif(5000)-0.5),nrow=10)
  for (k in (1:10))
  {
    M[10*(i-1)+k,]=p+Rf[k,]
  }
}
JWriteBin(M,"pamtest.bin",dtype="float",dmtpe="full")
CalcAndWriteDissimilarityMatrix("pamtest.bin","pamDL2.bin",distype="L2",restype="float",nthreads=0)
L <- ApplyPAM("pamDL2.bin",10,init_method="BUILD")
# Final value of sum of distances to closest medoid
GetTD(L,"pamDL2.bin")
# Medoids:
L$med
# Medoid in which each individual has been classified
n<-names(L$med)
n[L$clasif]
file.remove("pamtest.bin")
file.remove("pamDL2.bin")
```

CalcAndWriteDissimilarityMatrix
CalcAndWriteDissimilarityMatrix

Description

Writes a binary symmetric matrix with the dissimilarities between ROWS of the data stored in a binary matrix in the scellpam package format.

Notice that, differently from the common practice in single cell, the rows represent cells. This is for efficiency reasons and it is transparent to the user, as long as he/she has generated the binary matrix (with `CsvToBinMat`, `dgCMatToBinMat` or `SceToBinMat`) using the option `transpose=TRUE`.

The input matrix of vectors can be a full or a sparse matrix. Output matrix type can be float or double type (but look at the comments in 'Details').

Usage

```
CalcAndWriteDissimilarityMatrix(  
  ifname,  
  ofname,  
  distype = "L2",  
  restype = "float",  
  comment = "",  
  nthreads = 0L  
)
```

Arguments

<code>ifname</code>	A string with the name of the file containing the counts as a binary matrix, as written by <code>CsvToBinMat</code> , <code>dgCMatToBinMat</code> or <code>SceToBinMat</code>
<code>ofname</code>	A string with the name of the binary output file to contain the symmetric dissimilarity matrix.
<code>distype</code>	The dissimilarity to be calculated. It must be one of these strings: 'L1', 'L2' or 'Pearson'. Default: 'L2'.
<code>restype</code>	The data type of the result. It can be one of the strings 'float' or 'double'. Default: float (and don't change it unless you REALLY need to...).
<code>comment</code>	Comment to be added to the dissimilarity matrix. Default: "" (no comment)
<code>nthreads</code>	Number of threads to be used for the parallel calculations with this meaning: -1: don't use threads. 0: let the function choose according to the number of individuals (cells) and to the number of available cores. Any positive number > 1: use that number of threads. You can use even more than cores, but this is discouraged and raises a warning. Default: 0.

Details

The parameter `restype` forces the output to be a matrix of either floats or doubles. Precision of float is normally good enough; but if you need double precision (may be because you expect your results to be in a large range, two to three orders of magnitude), change it.

Nevertheless, notice that this is at the expense of double memory usage, which is QUADRATIC with the number of individuals (rows) in your input matrix.

Value

No return value, called for side effects (creates a file)

Examples

```
Rf <- matrix(runif(50000),nrow=100)
JWriteBin(Rf,"Rfullfloat.bin",dtype="float",dmtpe="full",
          comment="Full matrix of floats, 100 rows, 500 columns")
JMatInfo("Rfullfloat.bin")
CalcAndWriteDissimilarityMatrix("Rfullfloat.bin","RfullfloatDis.bin",distype="L2",
                                restype="float",comment="L2 distance matrix from full",nthreads=0)
JMatInfo("RfullfloatDis.bin")
JWriteBin(Rf,"Rsparsefloat.bin",dtype="float",dmtpe="sparse",
          comment="Sparse matrix of floats, 100 rows, 500 columns")
JMatInfo("Rsparsefloat.bin")
CalcAndWriteDissimilarityMatrix("Rsparsefloat.bin","RsparsefloatDis.bin",distype="L2",
                                restype="float",comment="L2 distance matrix from sparse",nthreads=0)
JMatInfo("RsparsefloatDis.bin")
Dfu<-GetJManyRows("RfullfloatDis.bin",c(1:nrow(Rf)))
Dsp<-GetJManyRows("RsparsefloatDis.bin",c(1:nrow(Rf)))
max(Dfu-Dsp)
file.remove("Rfullfloat.bin")
file.remove("Rsparsefloat.bin")
file.remove("RfullfloatDis.bin")
file.remove("RsparsefloatDis.bin")
```

CalculateSilhouette *CalculateSilhouette*

Description

Calculates the silhouette of each point of those classified by a clustering algorithm.

Usage

```
CalculateSilhouette(cl, fdist, nthreads = 0L)
```

Arguments

<code>cl</code>	The array of classification with the number of the class to which each point belongs to. This number must be in $1..number_of_classes$. This function takes something like the <code>L\$clasif</code> array which is the second element of the list returned by <code>ApplyPAM</code>
<code>fdist</code>	The binary file containing the symmetric matrix with the dissimilarities between points (usually, generated by a call to <code>CalcAndWriteDissimilarityMatrix</code>)
<code>nthreads</code>	The number of used threads. -1 means don't use threads (serial implementation). 0 means let the program choose according to the number of cores and of points. Any other number forces this number of threads. Choosing more than the number of available cores is allowed, but discouraged. Default: 0

Value

`sil` Numeric vector with the values of the silhouette for each point, in the same order in which points are in `cl`.
If `cl` is a named vector `sil` will be a named vector, too, with the same names.

Examples

```
# Synthetic problem: 10 random seeds with coordinates in [0..20]
# to which random values in [-0.1..0.1] are added
M<-matrix(0,100,500)
rownames(M)<-paste0("rn",c(1:100))
for (i in (1:10))
{
  p<-20*runif(500)
  Rf <- matrix(0.2*(runif(5000)-0.5),nrow=10)
  for (k in (1:10))
  {
    M[10*(i-1)+k,]=p+Rf[k,]
  }
}
JWriteBin(M,"pamtest.bin",dtype="float",dmtpe="full")
CalcAndWriteDissimilarityMatrix("pamtest.bin","pamDL2.bin",distype="L2",restype="float",nthreads=0)
L <- ApplyPAM("pamDL2.bin",10,init_method="BUILD")
sil <- CalculateSilhouette(L$clasif,"pamDL2.bin")
# Histogram of the silhouette. In this synthetic problem, almost 1 for all points
hist(sil)
file.remove("pamtest.bin")
file.remove("pamDL2.bin")
```

Description

Returns the results of the classification returned by ApplyPAM as a R dataframe

Usage

```
ClassifAsDataFrame(L, fdist)
```

Arguments

L	The list returned by ApplyPAM with fields L\$med and L\$clasif with the numbers of the medoids and the classification of each point
fdist	The binary file containing the symmetric matrix with the dissimilarities between points (usually, generated by a call to CalcAndWriteDissimilarityMatrix or to CalcAndWriteDissimilarityMatrixDouble)

Details

The dataframe has three columns: PointName (name of each point), NNPointName (name of the point which is the center of the cluster to which PointName belongs to) and NNDistance (distance between the points PointName and NNPointName). Medoids are identified by the fact that PointName and NNPointName are equal, or equivalently, NNDistance is 0.

Value

Df Dataframe with columns PointName, NNPointName and NNDistance. See Details for description.

Examples

```
# Synthetic problem: 10 random seeds with coordinates in [0..20]
# to which random values in [-0.1..0.1] are added
M<-matrix(0,100,500)
rownames(M)<-paste0("rn",c(1:100))
for (i in (1:10))
{
  p<-20*runif(500)
  Rf <- matrix(0.2*(runif(5000)-0.5),nrow=10)
  for (k in (1:10))
  {
    M[10*(i-1)+k,]=p+Rf[k,]
  }
}
JWriteBin(M,"pamtest.bin",dtype="float",dmtpe="full")
CalcAndWriteDissimilarityMatrix("pamtest.bin","pamDL2.bin",distype="L2",restype="float",nthreads=0)
L <- ApplyPAM("pamDL2.bin",10,init_method="BUILD")
df <- ClassifAsDataFrame(L,"pamDL2.bin")
df
# Identification of medoids:
which(df[,3]==0)
# Verification they are the same as in L (in different order)
L$med
```



```
file.remove("pamtest.bin")
file.remove("pamDL2.bin")
```

FilterBySilhouetteQuantile

FilterBySilhouetteQuantile

Description

Takes a silhouette, as returned by CalculateSilhouette, the list of medoids and class assignments, as returned by ApplyPam, a quantile and the matrices of counts and dissimilarities and constructs the corresponding matrices clearing off the points whose silhouette is below the lower quantile, except if they are medoids.

Usage

```
FilterBySilhouetteQuantile(
  s,
  L,
  fallcounts,
  ffilcounts,
  falldissim,
  ffildissim,
  q = 0.2,
  addcom = TRUE
)
```

Arguments

s	A numeric vector with the silhouette coefficient of each point in a classification, as returned by CalculateSilhouette.
L	A list of two numeric vectors, L\$med and L\$clasif, obtained normally as the object returned by ApplyPAM.
fallcounts	A string with the name of the binary file containing the matrix of individuals. It can be either a full or a sparse matrix.
ffilcounts	A string with the name of the binary file that will contain the selected individuals. It will have the same character (full/sparse) and type of the complete file.
falldissim	A string with the name of the binary file containing the dissimilarity matrix of the complete set of individuals. It must be a symmetric matrix of floats.
ffildissim	A string with the name of the binary file that will contain the dissimilarity matrix for the remaining individuals. It will be a symmetric matrix of floats.
q	Quantile to filter. All points whose silhouette is below this quantile will be filtered out. Default: 0.2

`addcom` Boolean to indicate if a comment must be appended to the current comment of counts and dissimilarity matrices to indicate that they are the result of a filtering process. This comment is automatically generated and contains the value of quantile `q`. Successive applications add comments at the end of those already present. Default: TRUE

Details

The renumbering of indices in the returned cluster may seem confusing at first but it was the way of fitting this with the rest of the package. Anyway, notice that if the numeric vectors in the input parameter `L` were named vectors, the point names are appropriately kept in the result so point identity is preserved. Moreover, if the counts and dissimilarity input matrices had row and/or column names, they are preserved in the filtered matrices, too.

Value

`Lr["med", "clasif"]` A list of two numeric vectors.
`Lr$med` is a modification of the corresponding first element of the passed `L` parameter.
`Lr$clasif` has as many components as remaining instances.
 Since points will have been removed, medoid numbering is modified. Therefore, `Lr$med` has the NEW index of each medoid in the filtered set.
`Lr$clasif` contains the number of the medoid (i.e.: the cluster) to which each instance has been assigned, and therefore does not change.
 All indexes start at 1 (R convention). Please, see Details section

Examples

```
# Synthetic problem: 10 random seeds with coordinates in [0..20]
# to which random values in [-0.1..0.1] are added
M<-matrix(0,100,500)
rownames(M)<-paste0("rn",c(1:100))
for (i in (1:10))
{
  p<-20*runif(500)
  Rf <- matrix(0.2*(runif(5000)-0.5),nrow=10)
  for (k in (1:10))
  {
    M[10*(i-1)+k,]=p+Rf[k,]
  }
}
JWriteBin(M,"pamtest.bin",dtype="float",dmtpe="full")
CalcAndWriteDissimilarityMatrix("pamtest.bin","pamDL2.bin",distype="L2",restype="float",nthreads=0)
L <- ApplyPAM("pamDL2.bin",10,init_method="BUILD")
# Which are the medoids
L$med
sil <- CalculateSilhouette(L$clasif,"pamDL2.bin")
Lf<-FilterBySilhouetteQuantile(sil,L,"pamtest.bin","pamtestfilt.bin","pamDL2.bin","pamDL2filt.bin",
                             q=0.4,addcom=TRUE)
# The new medoids are the same points but renumbered, since the L$clasif array has less points
Lf$med
```

```

file.remove("pamtest.bin")
file.remove("pamDL2.bin")
file.remove("pamtestfilt.bin")
file.remove("pamDL2filt.bin")

```

FilterBySilhouetteThreshold

FilterBySilhouetteThreshold

Description

Takes a silhouette, as returned by CalculateSilhouette, the list of medoids and class assignments, as returned by ApplyPam, a threshold and the matrices of counts and dissimilarities and constructs the corresponding matrices clearing off the points whose silhouette is below the threshold, except if they are medoids.

Usage

```

FilterBySilhouetteThreshold(
  s,
  L,
  fallcounts,
  ffilcounts,
  falldissim,
  ffildissim,
  thres = 0,
  addcom = TRUE
)

```

Arguments

s	A numeric vector with the silhouette coefficient of each point in a classification, as returned by CalculateSilhouette.
L	A list of two numeric vectors, L\$med and L\$clasif, obtained normally as the object returned by ApplyPAM.
fallcounts	A string with the name of the binary file containing the matrix of counts per individuals. It can be either a full or a sparse matrix.
ffilcounts	A string with the name of the binary file that will contain the selected individuals. It will have the same character (full/sparse) and type of the complete file.
falldissim	A string with the name of the binary file containing the dissimilarity matrix of the complete set of individuals. It must be a symmetric matrix of floats.
ffildissim	A string with the name of the binary file that will contain the dissimilarity matrix for the remaining individuals. It will be a symmetric matrix of floats.
thres	Threshold to filter. All points whose silhouette is below this threshold will be filtered out. Default: 0.0 (remember that silhouette is in [-1..1])

`addcom` Boolean to indicate if a comment must be appended to the current comment of counts and dissimilarity matrices to indicate that they are the result of a filtering process. This comment is automatically generated and contains the value of threshold `t`. Successive applications add comments at the end of those already present. Default: TRUE

Details

The renumbering of indices in the returned cluster may seem confusing at first but it was the way of fitting this with the rest of the package. Anyway, notice that if the numeric vectors in the input parameter `L` were named vectors, the point names are appropriately kept in the result so point identity is preserved. Moreover, if the individuals and dissimilarity input matrices had row and/or column names, they are preserved in the filtered matrices, too.

Value

`Lr["med", "clasif"]` A list of two numeric vectors.

`Lr$med` is a modification of the corresponding first element of the passed `L` parameter.

`Lr$clasif` has as many components as remaining instances.

Since points will have been removed, medoid numbering is modified. Therefore, `Lr$med` has the NEW index of each medoid in the filtered set.

`Lr$clasif` contains the number of the medoid (i.e.: the cluster) to which each instance has been assigned, and therefore does not change.

All indexes start at 1 (R convention). Please, see Details section

Examples

```
# Synthetic problem: 10 random seeds with coordinates in [0..20]
# to which random values in [-0.1..0.1] are added
M<-matrix(0,100,500)
rownames(M)<-paste0("rn",c(1:100))
for (i in (1:10))
{
  p<-20*runif(500)
  Rf <- matrix(0.2*(runif(5000)-0.5),nrow=10)
  for (k in (1:10))
  {
    M[10*(i-1)+k,]=p+Rf[k,]
  }
}
JWriteBin(M,"pamtest.bin",dtype="float",dmtpe="full")
CalcAndWriteDissimilarityMatrix("pamtest.bin","pamDL2.bin",distype="L2",restype="float",nthreads=0)
L <- ApplyPAM("pamDL2.bin",10,init_method="BUILD")
# Which are the medoids
L$med
sil <- CalculateSilhouette(L$clasif,"pamDL2.bin")
Lf<-FilterBySilhouetteQuantile(sil,L,"pamtest.bin","pamtestfilt.bin","pamDL2.bin","pamDL2filt.bin",
                              q=0.4,addcom=TRUE)
# The new medoids are the same points but renumbered, since the L$clasif array has less points
Lf$med
```

```
file.remove("pamtest.bin")
file.remove("pamDL2.bin")
file.remove("pamtestfilt.bin")
file.remove("pamDL2filt.bin")
```

GetJCol

GetJCol

Description

Returns (as a R numeric vector) the requested column number from the matrix contained in a jmatrix binary file

Usage

```
GetJCol(fname, ncol)
```

Arguments

fname	String with the file name that contains the binary data.
ncol	The number of the column to be returned, in R-numbering (from 1)

Value

A numeric vector with the values of elements in the requested column

Examples

```
Rf <- matrix(runif(48),nrow=6)
rownames(Rf) <- c("A","B","C","D","E","F")
colnames(Rf) <- c("a","b","c","d","e","f","g","h")
JWriteBin(Rf,"Rfullfloat.bin",dtype="float",dmtpe="full",comment="Full matrix of floats")
Rf[,3]
vf<-GetJCol("Rfullfloat.bin",3)
vf
file.remove("Rfullfloat.bin")
```

 GetJColByName

GetJColByName

Description

Returns (as a R numeric vector) the requested named column from the matrix contained in a jmatrix binary file

Usage

```
GetJColByName(fname, colname)
```

Arguments

fname	String with the file name that contains the binary data.
colname	The name of the column to be returned. If the matrix has no column names, or the name is not found, an empty vector is returned

Value

A numeric vector with the values of elements in the requested column

Examples

```
Rf <- matrix(runif(48),nrow=6)
rownames(Rf) <- c("A","B","C","D","E","F")
colnames(Rf) <- c("a","b","c","d","e","f","g","h")
JWriteBin(Rf,"Rfullfloat.bin",dtype="float",dmtpe="full",comment="Full matrix of floats")
Rf[,"c"]
vf<-GetJColByName("Rfullfloat.bin","c")
vf
file.remove("Rfullfloat.bin")
```

 GetJColNames

GetJColNames

Description

Returns a R StringVector with the column names of a matrix stored in the binary format of package jmatrix, if it has them stored.

Usage

```
GetJColNames(fname)
```

Arguments

fname String with the file name that contains the binary data.

Value

A R StringVector with the column names, or the empty vector if the binaryfile has no row column names as metadata.

Examples

```
Rf <- matrix(runif(48),nrow=6)
rownames(Rf) <- c("A","B","C","D","E","F")
colnames(Rf) <- c("a","b","c","d","e","f","g","h")
JWriteBin(Rf,"Rfullfloat.bin",dtype="float",dmtpe="full",comment="Full matrix of floats")
cn<-GetJColNames("Rfullfloat.bin")
cn
file.remove("Rfullfloat.bin")
```

GetJManyCols

GetJManyCols

Description

Returns (as a R numeric matrix) the columns with the requested column numbers from the matrix contained in a jmatrix binary file

Usage

```
GetJManyCols(fname, extcols)
```

Arguments

fname String with the file name that contains the binary data.
 extcols A numeric vector with the indexes of the columns to be extracted, in R-numbering (from 1)

Value

A numeric matrix with the values of elements in the requested columns

Examples

```
Rf <- matrix(runif(48),nrow=6)
rownames(Rf) <- c("A","B","C","D","E","F")
colnames(Rf) <- c("a","b","c","d","e","f","g","h")
JWriteBin(Rf,"Rfullfloat.bin",dtype="float",dmtpe="full",comment="Full matrix of floats")
Rf[,c(1,4)]
vc<-GetJManyCols("Rfullfloat.bin",c(1,4))
vc
file.remove("Rfullfloat.bin")
```

GetJManyColsByNames *GetJManyColsByNames*

Description

Returns (as a R numeric matrix) the columns with the requested column names from the matrix contained in a jmatrix binary file

Usage

```
GetJManyColsByNames(fname, extcolnames)
```

Arguments

fname	String with the file name that contains the binary data.
extcolnames	A numeric vector with the names of the columns to be extracted. If the binary file has no column names, or <code>_any_</code> of the column names is not present, an empty matrix is returned.

Value

A numeric matrix with the values of elements in the requested columns

Examples

```
Rf <- matrix(runif(48),nrow=6)
rownames(Rf) <- c("A","B","C","D","E","F")
colnames(Rf) <- c("a","b","c","d","e","f","g","h")
JWriteBin(Rf,"Rfullfloat.bin",dtype="float",dmtpe="full",comment="Full matrix of floats")
Rf[,c(1,4)]
vf<-GetJManyColsByNames("Rfullfloat.bin",c("a","d"))
vf
file.remove("Rfullfloat.bin")
```

GetJManyRows *GetJManyRows*

Description

Returns (as a R numeric matrix) the rows with the requested row numbers from the matrix contained in a jmatrix binary file

Usage

```
GetJManyRows(fname, extrows)
```


Arguments

fname	String with the file name that contains the binary data.
extrows	A numeric vector with the indexes of the rows to be extracted, in R-numbering (from 1)

Value

A numeric matrix with the values of elements in the requested rows

Examples

```
Rf <- matrix(runif(48),nrow=6)
rownames(Rf) <- c("A","B","C","D","E","F")
colnames(Rf) <- c("a","b","c","d","e","f","g","h")
JWriteBin(Rf,"Rfullfloat.bin",dtype="float",dmtpe="full",comment="Full matrix of floats")
Rf[c(1,4),]
vc<-GetJManyRows("Rfullfloat.bin",c(1,4))
vc
file.remove("Rfullfloat.bin")
```

GetJManyRowsByNames *GetJManyRowsByNames*

Description

Returns (as a R numeric matrix) the rows with the requested row names from the matrix contained in a jmatrix binary file

Usage

```
GetJManyRowsByNames(fname, extrownames)
```

Arguments

fname	String with the file name that contains the binary data.
extrownames	A numeric vector with the names of the rows to be extracted. If the binary file has no row names, or <code>_any_</code> of the row names is not present, an empty matrix is returned.

Value

A numeric matrix with the values of elements in the requested rows

Examples

```
Rf <- matrix(runif(48),nrow=6)
rownames(Rf) <- c("A","B","C","D","E","F")
colnames(Rf) <- c("a","b","c","d","e","f","g","h")
JWriteBin(Rf,"Rfullfloat.bin",dtype="float",dmtpe="full",comment="Full matrix of floats")
Rf[c("A","C"),]
vf<-GetJManyRowsByNames("Rfullfloat.bin",c("A","C"))
vf
file.remove("Rfullfloat.bin")
```

 GetJNames

GetJNames

Description

Returns a R list of two elements, rownames and colnames, each of them being a R StringVector with the corresponding names

Usage

```
GetJNames(fname)
```

Arguments

fname String with the file name that contains the binary data.

Value

N["rownames","colnames"]: A list with two elements named rownames and colnames which are R StringVectors. If the binary file has no row or column names as metadata BOTH will be returned as empty vectors, even if one of them exists. If you want to extract only one, use either GetBinRowNames or GetBinColNames, as appropriate.

Examples

```
Rf <- matrix(runif(48),nrow=6)
rownames(Rf) <- c("A","B","C","D","E","F")
colnames(Rf) <- c("a","b","c","d","e","f","g","h")
JWriteBin(Rf,"Rfullfloat.bin",dtype="float",dmtpe="full",comment="Full matrix of floats")
N<-GetJNames("Rfullfloat.bin")
N["rownames"]
N["colnames"]
file.remove("Rfullfloat.bin")
```

GetJRow	<i>GetJRow</i>
---------	----------------

Description

Returns (as a R numeric vector) the requested row number from the matrix contained in a jmatrix binary file

Usage

```
GetJRow(fname, nrow)
```

Arguments

fname	String with the file name that contains the binary data.
nrow	The number of the row to be returned, in R-numbering (from 1)

Value

A numeric vector with the values of elements in the requested row

Examples

```
Rf <- matrix(runif(48),nrow=6)
rownames(Rf) <- c("A","B","C","D","E","F")
colnames(Rf) <- c("a","b","c","d","e","f","g","h")
JWriteBin(Rf,"Rfullfloat.bin",dtype="float",dmtpe="full",comment="Full matrix of floats")
Rf[3,]
vf<-GetJRow("Rfullfloat.bin",3)
vf
file.remove("Rfullfloat.bin")
```

GetJRowByName	<i>GetJRowByName</i>
---------------	----------------------

Description

Returns (as a R numeric vector) the requested named row from the matrix contained in a jmatrix binary file

Usage

```
GetJRowByName(fname, rowname)
```

Arguments

fname	String with the file name that contains the binary data.
rowname	The name of the row to be returned. If the matrix has no row names, or the name is not found, an empty vector is returned

Value

A numeric vector with the values of elements in the requested row

Examples

```
Rf <- matrix(runif(48),nrow=6)
rownames(Rf) <- c("A","B","C","D","E","F")
colnames(Rf) <- c("a","b","c","d","e","f","g","h")
JWriteBin(Rf,"Rfullfloat.bin",dtype="float",dmtype="full",comment="Full matrix of floats")
Rf["C",]
vf<-GetJRowByName("Rfullfloat.bin","C")
vf
file.remove("Rfullfloat.bin")
```

GetJRowNames

GetJRowNames

Description

Returns a R StringVector with the row names of a matrix stored in the binary format of package jmatrix, if it has them stored.

Usage

```
GetJRowNames(fname)
```

Arguments

fname	String with the file name that contains the binary data.
-------	--

Value

A R StringVector with the row names, or the empty vector if the binary file has no row names as metadata.

Examples

```
Rf <- matrix(runif(48),nrow=6)
rownames(Rf) <- c("A","B","C","D","E","F")
colnames(Rf) <- c("a","b","c","d","e","f","g","h")
JWriteBin(Rf,"Rfullfloat.bin",dtype="float",dmtype="full",comment="Full matrix of floats")
rn<-GetJRowNames("Rfullfloat.bin")
rn
file.remove("Rfullfloat.bin")
```

 GetSubdiag

GetSubdiag

Description

Takes a symmetric matrix and returns a vector with all its elements under the main diagonal (without those at the diagonal itself) Done as an instrumental function to check the PAM in package cluster. To be removed in final version of the package.

Usage

```
GetSubdiag(fname)
```

Arguments

fname The name of the file with the dissimilarity matrix in jmatrix binary format.

Value

The vector with the values under the main diagonal, sorted by columns (i.e.: $m(2,1) \dots m(n,1)$, $m(3,2) \dots m(n,2)$, ..., $m(n-1,n)$)

Examples

```
Rns <- matrix(runif(49),nrow=7)
Rsym <- 0.5*(Rns+t(Rns))
rownames(Rsym) <- c("A","B","C","D","E","F","G")
colnames(Rsym) <- c("a","b","c","d","e","f","g")
JWriteBin(Rsym,"Rsymfloat.bin",dtype="float",dmtpe="symmetric")
d<-GetSubdiag("Rsymfloat.bin")
Rsym
d
file.remove("Rsymfloat.bin")
```

 GetTD

GetTD

Description

Function that takes a PAM classification (as returned by ApplyPAM) and the dissimilarity matrix and returns the value of the TD function (sum of dissimilarities between each point and its closest medoid, divided by the number of points). This function is mainly for debugging/internal use.

Usage

```
GetTD(L, dissim_file)
```

Arguments

- L** A list of two numeric vectors, L["med","clasif"], as returned by ApplyPAM (please, consult the help of ApplyPAM for details)
- dissim_file** A string with the name of the binary file that contains the symmetric matrix of dissimilarities. Such matrix should have been generated by CalcAndWriteDissimilarityMatrix and it is a matrix of type 'disttype' (currently defined as float).

Value

TD The value of the TD function.

Examples

```
# Synthetic problem: 10 random seeds with coordinates in [0..20]
# to which random values in [-0.1..0.1] are added
M<-matrix(0,100,500)
rownames(M)<-paste0("rn",c(1:100))
for (i in (1:10))
{
  p<-20*runif(500)
  Rf <- matrix(0.2*(runif(5000)-0.5),nrow=10)
  for (k in (1:10))
  {
    M[10*(i-1)+k,]=p+Rf[k,]
  }
}
JWriteBin(M,"pamtest.bin",dtype="float",dmtpe="full")
CalcAndWriteDissimilarityMatrix("pamtest.bin","pamDL2.bin",distype="L2",restype="float",nthreads=0)
L <- ApplyPAM("pamDL2.bin",10,init_method="BUILD")
# Final value of sum of distances to closest medoid
GetTD(L,"pamDL2.bin")
```

JMatInfo

JMatInfo

Description

Shows in the screen or writes to a file information about a matrix stored in the binary format of package jmatrix

Usage

```
JMatInfo(fname, fres = "")
```

Arguments

- fname** String with the file name that contains the binary data.
- fres** String with the name of the file to write the information. Default: "" (information is written to the console)

Value

No return value, called for its side effects (writes on screen or creates a file)

Examples

```
Rf <- matrix(runif(48),nrow=6)
rownames(Rf) <- c("A","B","C","D","E","F")
colnames(Rf) <- c("a","b","c","d","e","f","g","h")
JWriteBin(Rf,"Rfullfloat.bin",dtype="float",dmtpe="full",comment="Full matrix of floats")
JMatInfo("Rfullfloat.bin")
file.remove("Rfullfloat.bin")
```

JWriteBin

JWriteBin

Description

Writes a R matrix to a disk file as a binary matrix in the jmatrix format

Usage

```
JWriteBin(M, fname, dtype = "float", dmtpe = "full", comment = "")
```

Arguments

M	The R matrix to be written
fname	The name of the file to write
dtype	The data type of the matrix to be written: one of the strings 'short', 'int', 'long', 'float' or 'double'. Default: 'float'
dmtpe	The matrix type: one of the strings 'full', 'sparse' or 'symmetric'. Default: 'full'
comment	A optional string with the comment to be added as metadata. Default: "" (empty string, no added comment)

Details

Use this function cautiously. Differently to the functions to get one or more rows or columns from the binary file, which book only the memory strictly needed for the vector/matrix and do not load all the binary file in memory, this function books the full matrix in the requested data type and writes it later so with very big matrices you might run out of memory.

Type 'int' is really long int (8-bytes in most modern machines) so using 'int' or 'long' is equivalent. Type is coerced from double (the internal type of R matrices) to the requested type, which may provoke a loose of precision.

If M is a named-R matrix, row and column names are written as metadata, too.

Also, if you write as symmetric a matrix which is not such, only the lower-diagonal part will be written. The rest of the data will be lost. In this case, if the matrix has row and column names, only row names are written.

Value

No return value, called for side effects (creates a file)

Examples

```
Rf <- matrix(runif(48),nrow=6)
rownames(Rf) <- c("A","B","C","D","E","F")
colnames(Rf) <- c("a","b","c","d","e","f","g","h")
JWriteBin(Rf,"Rfullfloat.bin",dtype="float",dmtpe="full",comment="Full matrix of floats")
file.remove("Rfullfloat.bin")
```

NumSilToClusterSil *NumSilToClusterSil*

Description

Takes a silhouette in the form of a NumericVector, as returned by CalculateSilhouette, and returns it as a numeric matrix appropriate to be plotted by the package 'cluster'

Usage

```
NumSilToClusterSil(cl, s)
```

Arguments

- | | |
|----|--|
| cl | The array of classification with the number of the class to which each point belongs to. This number must be in 1..number_of_classes.
This function takes something like the L\$clasif array which is the second element of the list returned by ApplyPAM |
| s | The numeric value of the silhouette for each point, with points in the same order as they appear in cl.
This is the vector returned by after a call to CalculateSilhouette with the same value of parameter cl. |

Value

sp A silhouette in the format of the cluster package which is a NumericMatrix of as many rows as points and three columns: cluster, neighbor and sil_width.
Its structure and dimension names are as in package 'cluster', which allows to use it with the silhouette plotting functions of such package
This means you can do library(cluster) followed by plot(NumSilToClusterSil(cl,s)) to get a beautiful plot.

Examples

```

# Synthetic problem: 10 random seeds with coordinates in [0..20]
# to which random values in [-0.1..0.1] are added
M<-matrix(0,100,500)
rownames(M)<-paste0("rn",c(1:100))
for (i in (1:10))
{
  p<-20*runif(500)
  Rf <- matrix(0.2*(runif(5000)-0.5),nrow=10)
  for (k in (1:10))
  {
    M[10*(i-1)+k,]=p+Rf[k,]
  }
}
JWriteBin(M,"pamtest.bin",dtype="float",dmtpe="full")
CalcAndWriteDissimilarityMatrix("pamtest.bin","pamDL2.bin",distype="L2",restype="float",nthreads=0)
L <- ApplyPAM("pamDL2.bin",10,init_method="BUILD")
sil <- CalculateSilhouette(L$classif,"pamDL2.bin")
sp <- NumSilToClusterSil(L$classif,sil)
# Here you can install and load package cluster and make plot(sp)
file.remove("pamtest.bin")
file.remove("pamDL2.bin")

```

ParallelpamSetDebug *ParallelpamSetDebug*

Description

Sets debugging in parallelpam package to ON (with TRUE) or OFF (with FALSE) for both parts of it.

On package load the default status is OFF.

Setting debugging of any part to ON shows a message. Setting to OFF does not show anything (since debugging is OFF...)

Usage

```
ParallelpamSetDebug(deb = TRUE, debjmat = FALSE)
```

Arguments

deb	boolean, TRUE to generate debug messages for the PAM algorithm and FALSE to turn them off. Default: true.
debjmat	boolean, TRUE to generate debug messages for the jmatrix part inside this package and FALSE to turn them off. Default: false

Value

No return value, called for side effects (internal boolean flag changed)

Examples

```
ParallelpamSetDebug(TRUE, debjmat=TRUE)  
ParallelpamSetDebug(TRUE, debjmat=FALSE)
```

Index

ApplyPAM, [2](#)

CalcAndWriteDissimilarityMatrix, [5](#)

CalculateSilhouette, [6](#)

ClassifAsDataFrame, [7](#)

FilterBySilhouetteQuantile, [9](#)

FilterBySilhouetteThreshold, [11](#)

GetJCol, [13](#)

GetJColByName, [14](#)

GetJColNames, [14](#)

GetJManyCols, [15](#)

GetJManyColsByNames, [16](#)

GetJManyRows, [16](#)

GetJManyRowsByNames, [17](#)

GetJNames, [18](#)

GetJRow, [19](#)

GetJRowByName, [19](#)

GetJRowNames, [20](#)

GetSubdiag, [21](#)

GetTD, [21](#)

JMatInfo, [22](#)

JWriteBin, [23](#)

NumSilToClusterSil, [24](#)

ParallelpamSetDebug, [25](#)