

The e1071 Package

April 6, 2005

Version 1.5-8

Date 2005-04-06

Title Misc Functions of the Department of Statistics (e1071), TU Wien

Depends R (>= 1.9.0), class

Imports graphics, stats

Suggests cluster, mlbench, nnet, randomForest, rpart, SparseM, xtable

Author Evgenia Dimitriadou, Kurt Hornik, Friedrich Leisch, David Meyer, and Andreas Weingessel

Maintainer Friedrich Leisch <Friedrich.Leisch@ci.tuwien.ac.at>

Description Functions for latent class analysis, short time Fourier transform, fuzzy clustering, support vector machines,

License GPL version 2. See COPYRIGHT.svm.cpp for the copyright of the svm C++ code.

SaveImage no

LazyLoad yes

R topics documented:

Discrete	2
Glass	3
HouseVotes84	4
Ozone	5
bclust	6
bincombinations	8
bootstrap.lca	8
boxplot.bclust	10
classAgreement	11
cmeans	12
countpattern	14
cshell	15
e1071-deprecated	17
element	17
fclustIndex	18
hamming.distance	20

hamming.window	21
hanning.window	22
ica	23
impute	24
interpolate	24
kurtosis	25
lca	26
matchClasses	27
matchControls	28
moment	30
naiveBayes	31
permutations	32
plot.stft	33
plot.svm	33
plot.tune	35
predict.naiveBayes	36
predict.svm	37
rbridge	39
read.matrix.csr	40
rectangle.window	41
rwiener	41
allShortestPaths	42
sigmoid	43
skewness	44
stft	44
svm	45
tune	50
tune.control	52
tune.wrapper	53
write.svm	54
Index	56

 Discrete

Discrete Distribution

Description

These functions provide information about the discrete distribution where the probability of the elements of `values` is proportional to the values given in `probs`, which are normalized to sum up to 1. `ddiscrete` gives the density, `pdiscrete` gives the distribution function, `qdiscrete` gives the quantile function and `rdiscrete` generates random deviates.

Usage

```
ddiscrete(x, probs, values = 1:length(probs))
pdiscrete(q, probs, values = 1:length(probs))
qdiscrete(p, probs, values = 1:length(probs))
rdiscrete(n, probs, values = 1:length(probs), ...)
```

Arguments

<code>x, q</code>	vector or array of quantiles.
<code>p</code>	vector or array of probabilities.
<code>n</code>	number of observations.
<code>probs</code>	probabilities of the distribution.
<code>values</code>	values of the distribution.
<code>...</code>	ignored (only there for backwards compatibility)

Details

The random number generator is simply a wrapper for `sample` and provided for backwards compatibility only.

Author(s)

Andreas Weingessel and Friedrich Leisch

Examples

```
## a vector of length 30 whose elements are 1 with probability 0.2
## and 2 with probability 0.8.
rdiscrete (30, c(0.2, 0.8))

## a vector of length 100 whose elements are A, B, C, D.
## The probabilities of the four values have the relation 1:2:3:3
rdiscrete (100, c(1,2,3,3), c("A","B","C","D"))
```

Glass

Glass Identification Database

Description

A data frame with 214 observation containing examples of the chemical analysis of 7 different types of glass. The problem is to forecast the type of class on basis of the chemical analysis. The study of classification of types of glass was motivated by criminological investigation. At the scene of the crime, the glass left can be used as evidence (if it is correctly identified!).

Usage

```
data(Glass)
```

Format

A data frame with 214 observations on 10 variables:

[,1]	RI	refractive index
[,2]	Na	Sodium
[,3]	Mg	Magnesium
[,4]	Al	Aluminum
[,5]	Si	Silicon

[,6]	K	Potassium
[,7]	Ca	Calcium
[,8]	Ba	Barium
[,9]	Fe	Iron
[,10]	Type	Type of glass (class attribute)

Source

- Creator: B. German, Central Research Establishment, Home Office Forensic Science Service, Aldermaston, Reading, Berkshire RG7 4PN
- Donor: Vina Spiehler, Ph.D., DABFT, Diagnostic Products Corporation

These data have been taken from the UCI Repository Of Machine Learning Databases at

- <ftp://ftp.ics.uci.edu/pub/machine-learning-databases>
- <http://www.ics.uci.edu/~mlearn/MLRepository.html>

and were converted to R format by Friedrich.Leisch@ci.tuwien.ac.at.

HouseVotes84

United States Congressional Voting Records 1984

Description

This data set includes votes for each of the U.S. House of Representatives Congressmen on the 16 key votes identified by the CQA. The CQA lists nine different types of votes: voted for, paired for, and announced for (these three simplified to yea), voted against, paired against, and announced against (these three simplified to nay), voted present, voted present to avoid conflict of interest, and did not vote or otherwise make a position known (these three simplified to an unknown disposition).

Usage

```
data(HouseVotes84)
```

Format

A data frame with 435 observations on 17 variables:

- 1 Class Name: 2 (democrat, republican)
- 2 handicapped-infants: 2 (y,n)
- 3 water-project-cost-sharing: 2 (y,n)
- 4 adoption-of-the-budget-resolution: 2 (y,n)
- 5 physician-fee-freeze: 2 (y,n)
- 6 el-salvador-aid: 2 (y,n)
- 7 religious-groups-in-schools: 2 (y,n)
- 8 anti-satellite-test-ban: 2 (y,n)
- 9 aid-to-nicaraguan-contras: 2 (y,n)
- 10 mx-missile: 2 (y,n)
- 11 immigration: 2 (y,n)
- 12 synfuels-corporation-cutback: 2 (y,n)
- 13 education-spending: 2 (y,n)
- 14 superfund-right-to-sue: 2 (y,n)
- 15 crime: 2 (y,n)
- 16 duty-free-exports: 2 (y,n)
- 17 export-administration-act-south-africa: 2 (y,n)

Source

- Source: Congressional Quarterly Almanac, 98th Congress, 2nd session 1984, Volume XL: Congressional Quarterly Inc., ington, D.C., 1985
- Donor: Jeff Schlimmer (Jeffrey.Schlimmer@a.gp.cs.cmu.edu)

These data have been taken from the UCI Repository Of Machine Learning Databases at

- <ftp://ftp.ics.uci.edu/pub/machine-learning-databases>
- <http://www.ics.uci.edu/~mlearn/MLRepository.html>

and were converted to R format by Friedrich.Leisch@ci.tuwien.ac.at.

Ozone

Los Angeles ozone pollution data, 1976

Description

A data frame with 366 observations on 13 variables, each observation is one day

Usage

```
data(Ozone)
```

Format

- 1 Month: 1 = January, ..., 12 = December
- 2 Day of month
- 3 Day of week: 1 = Monday, ..., 7 = Sunday
- 4 Daily maximum one-hour-average ozone reading
- 5 500 millibar pressure height (m) measured at Vandenberg AFB
- 6 Wind speed (mph) at Los Angeles International Airport (LAX)
- 7 Humidity (%) at LAX
- 8 Temperature (degrees F) measured at Sandburg, CA
- 9 Temperature (degrees F) measured at El Monte, CA
- 10 Inversion base height (feet) at LAX
- 11 Pressure gradient (mm Hg) from LAX to Daggett, CA
- 12 Inversion base temperature (degrees F) at LAX
- 13 Visibility (miles) measured at LAX

Details

The problem is to predict the daily maximum one-hour-average ozone reading (V4).

Source

Leo Breiman, Department of Statistics, UC Berkeley. Data used in Leo Breiman and Jerome H. Friedman (1985), Estimating optimal transformations for multiple regression and correlation, JASA, 80, pp. 580-598.

bclust

*Bagged Clustering***Description**

Cluster the data in `x` using the bagged clustering algorithm. A partitioning cluster algorithm such as `kmeans` is run repeatedly on bootstrap samples from the original data. The resulting cluster centers are then combined using the hierarchical cluster algorithm `hclust`.

Usage

```
bclust(x, centers=2, iter.base=10, minsize=0,
      dist.method="euclidian",
      hclust.method="average", base.method="kmeans",
      base.centers=20, verbose=TRUE,
      final.kmeans=FALSE, docmdscale=FALSE,
      resample=TRUE, weights=NULL, maxcluster=base.centers, ...)
hclust.bclust(object, x, centers, dist.method=object$dist.method,
             hclust.method=object$hclust.method, final.kmeans=FALSE,
             docmdscale = FALSE, maxcluster=object$maxcluster)
## S3 method for class 'bclust':
plot(x, maxcluster=x$maxcluster, main, ...)
centers.bclust(object, k)
clusters.bclust(object, k, x=NULL)
```

Arguments

<code>x</code>	Matrix of inputs (or object of class "bclust" for plot).
<code>centers, k</code>	Number of clusters.
<code>iter.base</code>	Number of runs of the base cluster algorithm.
<code>minsize</code>	Minimum number of points in a base cluster.
<code>dist.method</code>	Distance method used for the hierarchical clustering, see <code>dist</code> for available distances.
<code>hclust.method</code>	Linkage method used for the hierarchical clustering, see <code>hclust</code> for available methods.
<code>base.method</code>	Partitioning cluster method used as base algorithm.
<code>base.centers</code>	Number of centers used in each repetition of the base method.
<code>verbose</code>	Output status messages.
<code>final.kmeans</code>	If TRUE, a final kmeans step is performed using the output of the bagged clustering as initialization.
<code>docmdscale</code>	Logical, if TRUE a <code>cmdscale</code> result is included in the return value.
<code>resample</code>	Logical, if TRUE the base method is run on bootstrap samples of <code>x</code> , else directly on <code>x</code> .
<code>weights</code>	Vector of length <code>nrow(x)</code> , weights for the resampling. By default all observations have equal weight.
<code>maxcluster</code>	Maximum number of clusters memberships are to be computed for.

object	Object of class "bclust".
main	Main title of the plot.
...	Optional arguments top be passed to the base method in <code>bclust</code> , ignored in <code>plot</code> .

Details

First, `iter.base` bootstrap samples of the original data in `x` are created by drawing with replacement. The base cluster method is run on each of these samples with `base.centers` centers. The `base.method` must be the name of a partitioning cluster function returning a list with the same components as the return value of `kmeans`.

This results in a collection of `iter.base * base.centers` centers, which are subsequently clustered using the hierarchical method `hclust`. Base centers with less than `minsize` points in there respective partitions are removed before the hierarchical clustering.

The resulting dendrogram is then cut to produce `centers` clusters. Hence, the name of the argument `centers` is a little bit misleading as the resulting clusters need not be convex, e.g., when single linkage is used. The name was chosen for compatibility with standard partitioning cluster methods such as `kmeans`.

A new hierarchical clustering (e.g., using another `hclust.method`) re-using previous base runs can be performed by running `hclust.bclust` on the return value of `bclust`.

Value

`bclust` and `hclust.bclust` return objects of class "bclust" including the components

<code>hclust</code>	Return value of the hierarchical clustering of the collection of base centers (Object of class "hclust").
<code>cluster</code>	Vector with indices of the clusters the inputs are assigned to.
<code>centers</code>	Matrix of centers of the final clusters. Only useful, if the hierarchical clustering method produces convex clusters.
<code>allcenters</code>	Matrix of all <code>iter.base * base.centers</code> centers found in the base runs.

Author(s)

Friedrich Leisch

References

Friedrich Leisch. Bagged clustering. Working Paper 51, SFB "Adaptive Information Systems and Modeling in Economics and Management Science", August 1999. <http://www.ci.tuwien.ac.at/~leisch>

See Also

[hclust](#), [kmeans](#), [boxplot.bclust](#)

Examples

```
data(iris)
bc1 <- bclust(iris[,1:4], 3, base.centers=5)
plot(bc1)

table(clusters.bclust(bc1, 3))
centers.bclust(bc1, 3)
```

bincombinations *Binary Combinations*

Description

Returns a matrix containing the 2^p vectors of length p .

Usage

```
bincombinations(p)
```

Arguments

p Length of binary vectors

Author(s)

Friedrich Leisch

Examples

```
bincombinations(2)
bincombinations(3)
```

bootstrap.lca *Bootstrap Samples of LCA Results*

Description

This function draws bootstrap samples from a given LCA model and refits a new LCA model for each sample. The quality of fit of these models is compared to the original model.

Usage

```
bootstrap.lca(l, nsamples=10, lcaiter=30, verbose=FALSE)
```

Arguments

l An LCA model as created by [lca](#)
 nsamples Number of bootstrap samples
 lcaiter Number of LCA iterations
 verbose If TRUE some output is printed during the computations.

Details

From a given LCA model `l`, `nsamples` bootstrap samples are drawn. For each sample a new LCA model is fitted. The goodness of fit for each model is computed via Likelihood Ratio and Pearson's Chisquare. The values for the fitted models are compared with the values of the original model `l`. By this method it can be tested whether the data to which `l` was originally fitted come from an LCA model.

Value

An object of class `bootstrap.lca` is returned, containing

<code>logl</code> , <code>loglsat</code>	The LogLikelihood of the models and of the corresponding saturated models
<code>lratio</code>	Likelihood quotient of the models and the corresponding saturated models
<code>lrationmean</code> , <code>lratiosd</code>	Mean and Standard deviation of <code>lratio</code>
<code>lratioorg</code>	Likelihood quotient of the original model and the corresponding saturated model
<code>zratio</code>	Z-Statistics of <code>lratioorg</code>
<code>pvalzratio</code> , <code>pvalratio</code>	P-Values for <code>zratio</code> , computed via normal distribution and empirical distribution
<code>chisq</code>	Pearson's Chisq of the models
<code>chisqmean</code> , <code>chisqsd</code>	Mean and Standard deviation of <code>chisq</code>
<code>chisqorg</code>	Pearson's Chisq of the original model
<code>zchisq</code>	Z-Statistics of <code>chisqorg</code>
<code>pvalzchisq</code> , <code>pvalchisq</code>	P-Values for <code>zchisq</code> , computed via normal distribution and empirical distribution
<code>nsamples</code>	Number of bootstrap samples
<code>lcaiter</code>	Number of LCA Iterations

Author(s)

Andreas Weingessel

References

Anton K. Formann: "Die Latent-Class-Analysis", Beltz Verlag 1984

See Also

[lca](#)

Examples

```
## Generate a 4-dim. sample with 2 latent classes of 500 data points each.
## The probabilities for the 2 classes are given by type1 and type2.
type1 <- c(0.8,0.8,0.2,0.2)
type2 <- c(0.2,0.2,0.8,0.8)
x <- matrix(runif(4000),nr=1000)
x[1:500,] <- t(t(x[1:500,])<type1)*1
x[501:1000,] <- t(t(x[501:1000,])<type2)*1

l <- lca(x, 2, niter=5)
bl <- bootstrap.lca(l,nsamples=3,lcaiter=5)
bl
```

boxplot.bclust *Boxplot of cluster profiles*

Description

Makes boxplots of the results of a bagged clustering run.

Usage

```
## S3 method for class 'bclust':
boxplot(x, n=nrow(x$centers), bycluster=TRUE,
        main=deparse(substitute(x)), oneplot=TRUE,
        which=1:n, ...)
```

Arguments

x	Clustering result, object of class "bclust".
n	Number of clusters to plot, by default the number of clusters used in the call of <code>bclust</code> .
bycluster	If TRUE (default), a boxplot for each cluster is plotted. If FALSE, a boxplot for each variable is plotted.
main	Main title of the plot, by default the name of the cluster object.
oneplot	If TRUE, all boxplots appear on one screen (using an appropriate rectangular layout).
which	Number of clusters which should be plotted, default is all clusters.
...	Additional arguments for <code>boxplot</code> .

Author(s)

Friedrich Leisch

Examples

```
data(iris)
bc1 <- bclust(iris[,1:4], 3, base.centers=5)
boxplot(bc1)
```

classAgreement	<i>Coefficients comparing classification agreement</i>
----------------	--

Description

`classAgreement()` computes several coefficients of agreement between the columns and rows of a 2-way contingency table.

Usage

```
classAgreement(tab, match.names=FALSE)
```

Arguments

`tab` A 2-dimensional contingency table.
`match.names` Flag whether row and columns should be matched by name.

Details

Suppose we want to compare two classifications summarized by the contingency table $T = [t_{ij}]$ where $i, j = 1, \dots, K$ and t_{ij} denotes the number of data points which are in class i in the first partition and in class j in the second partition. If both classifications use the same labels, then obviously the two classification agree completely if only elements in the main diagonal of the table are non-zero. On the other hand, large off-diagonal elements correspond to smaller agreement between the two classifications. If `match.names` is TRUE, the class labels as given by the row and column names are matched, i.e. only columns and rows with the same dimnames are used for the computation.

If the two classification do not use the same set of labels, or if identical labels can have different meaning (e.g., two outcomes of cluster analysis on the same data set), then the situation is a little bit more complicated. Let A denote the number of all pairs of data points which are either put into the same cluster by both partitions or put into different clusters by both partitions. Conversely, let D denote the number of all pairs of data points that are put into one cluster in one partition, but into different clusters by the other partition. Hence, the partitions disagree for all pairs D and agree for all pairs A . We can measure the agreement by the Rand index $A/(A + D)$ which is invariant with respect to permutations of the columns or rows of T .

Both indices have to be corrected for agreement by chance if the sizes of the classes are not uniform.

Value

A list with components

<code>diag</code>	Percentage of data points in the main diagonal of <code>tab</code> .
<code>kappa</code>	<code>diag</code> corrected for agreement by chance.
<code>rand</code>	Rand index.
<code>crand</code>	Rand index corrected for agreement by chance.

Author(s)

Friedrich Leisch

References

J. Cohen. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20, 37–46, 1960.

Lawrence Hubert and Phipps Arabie. Comparing partitions. *Journal of Classification*, 2, 193–218, 1985.

See Also

[matchClasses](#)

Examples

```
## no class correlations: both kappa and crand almost zero
g1 <- sample(1:5, size=1000, replace=TRUE)
g2 <- sample(1:5, size=1000, replace=TRUE)
tab <- table(g1, g2)
classAgreement(tab)

## let pairs (g1=1,g2=1) and (g1=3,g2=3) agree better
k <- sample(1:1000, size=200)
g1[k] <- 1
g2[k] <- 1

k <- sample(1:1000, size=200)
g1[k] <- 3
g2[k] <- 3

tab <- table(g1, g2)
## both kappa and crand should be significantly larger than before
classAgreement(tab)
```

cmeans

Fuzzy C-Means Clustering

Description

The fuzzy version of the known *k*means clustering algorithm as well as an on-line variant (Unsupervised Fuzzy Competitive learning).

Usage

```
cmeans(x, centers, iter.max = 100, verbose = FALSE,
       dist = "euclidean", method = "cmeans", m = 2,
       rate.par = NULL, weights = 1, control = list())
```

Arguments

<code>x</code>	The data matrix where columns correspond to variables and rows to observations.
<code>centers</code>	Number of clusters or initial values for cluster centers.
<code>iter.max</code>	Maximum number of iterations.

verbose	If TRUE, make some output during learning.
dist	Must be one of the following: If "euclidean", the mean square error, if "manhattan", the mean absolute error is computed. Abbreviations are also accepted.
method	If "cmeans", then we have the <i>c</i> -means fuzzy clustering method, if "ufcl" we have the on-line update. Abbreviations are also accepted.
m	A number greater than 1 giving the degree of fuzzification.
rate.par	A number between 0 and 1 giving the parameter of the learning rate for the on-line variant. The default corresponds to 0.3.
weights	a numeric vector with non-negative case weights. Recycled to the number of observations in <code>x</code> if necessary.
control	a list of control parameters. See Details .

Details

The data given by `x` is clustered by generalized versions of the fuzzy *c*-means algorithm, which use either a fixed-point or an on-line heuristic for minimizing the objective function

$$\sum_i \sum_j w_i u_{ij}^m d_{ij},$$

where w_i is the weight of observation i , u_{ij} is the membership of observation i in cluster j , and d_{ij} is the distance (dissimilarity) between observation i and center j . The dissimilarities used are the sums of squares ("euclidean") or absolute values ("manhattan") of the elementwise differences.

If `centers` is a matrix, its rows are taken as the initial cluster centers. If `centers` is an integer, `centers` rows of `x` are randomly chosen as initial values.

The algorithm stops when the maximum number of iterations (given by `iter.max`) is reached, or when the algorithm is unable to reduce the current value `val` of the objective function by `reltol * (abs(val) * reltol)` at a step. The relative convergence tolerance `reltol` can be specified as the `reltol` component of the list of control parameters, and defaults to `sqrt(.Machine$double.eps)`.

If `verbose` is TRUE, each iteration displays its number and the value of the objective function.

If `method` is "cmeans", then we have the *c*-means fuzzy clustering method, see for example Bezdek (1981). If "ufcl", we have the On-line Update (Unsupervised Fuzzy Competitive Learning) method due to Chung and Lee (1992), see also Pal et al (1996). This method works by performing an update directly after each input signal (i.e., for each single observation).

The parameter `m` defines the degree of fuzzification. It is defined for real values greater than 1 and the bigger it is the more fuzzy the membership values of the clustered data points are.

Value

An object of class "fclust" which is a list with components:

centers	the final cluster centers.
size	the number of data points in each cluster of the closest hard clustering.
cluster	a vector of integers containing the indices of the clusters where the data points are assigned to for the closest hard clustering, as obtained by assigning points to the (first) class with maximal membership.

iter the number of iterations performed.
 membership a matrix with the membership values of the data points to the clusters.
 withinerror the value of the objective function.
 call the call used to create the object.

Author(s)

Evgenia Dimitriadou and Kurt Hornik

References

J. C. Bezdek (1981). *Pattern recognition with fuzzy objective function algorithms*. New York: Plenum.
 Fu Lai Chung and Tong Lee (1992). Fuzzy competitive learning. *Neural Networks*, **7**(3), 539–551.
 Nikhil R. Pal, James C. Bezdek, and Richard J. Hathaway (1996). Sequential competitive learning and the fuzzy c-means clustering algorithms. *Neural Networks*, **9**(5), 787–796.

Examples

```

# a 2-dimensional example
x<-rbind(matrix(rnorm(100,sd=0.3),ncol=2),
          matrix(rnorm(100,mean=1,sd=0.3),ncol=2))
cl<-cmeans(x,2,20,verbose=TRUE,method="cmeans",m=2)
print(cl)

# a 3-dimensional example
x<-rbind(matrix(rnorm(150,sd=0.3),ncol=3),
          matrix(rnorm(150,mean=1,sd=0.3),ncol=3),
          matrix(rnorm(150,mean=2,sd=0.3),ncol=3))
cl<-cmeans(x,6,20,verbose=TRUE,method="cmeans")
print(cl)

```

countpattern

Count Binary Patterns

Description

Every row of the binary matrix `x` is transformed into a binary pattern and these patterns are counted.

Usage

```
countpattern(x, matching=FALSE)
```

Arguments

`x` A matrix of binary observations
`matching` If TRUE an additional vector is returned which stores which row belongs to which pattern

Value

A vector of length $2^{\text{ncol}(x)}$ giving the number of times each pattern occurs in the rows of x . The names of this vector are the binary patterns. They are sorted according to their numeric value. If `matching` is `TRUE`, a list of the following two vectors is returned.

`pat` Numbers of patterns as described above.
`matching` Vector giving the position of the pattern of each row of x in `pat`.

Author(s)

Andreas Weingessel

Examples

```
xx <- rbind(c(1,0,0),c(1,0,0),c(1,0,1),c(0,1,1),c(0,1,1))
countpattern(xx)
countpattern(xx, matching=TRUE)
```

cshell

Fuzzy C-Shell Clustering

Description

The *c*-shell clustering algorithm, the shell prototype-based version (ring prototypes) of the fuzzy *k*means clustering method.

Usage

```
cshell(x, centers, iter.max=100, verbose=FALSE, dist="euclidean",
       method="cshell", m=2, radius = NULL)
```

Arguments

`x` The data matrix, where columns correspond to the variables and rows to observations.

`centers` Number of clusters or initial values for cluster centers

`iter.max` Maximum number of iterations

`verbose` If `TRUE`, make some output during learning

`dist` Must be one of the following: If `"euclidean"`, the mean square error, if `"manhattan"`, the mean absolute error is computed. Abbreviations are also accepted.

`method` Currently, only the `"cshell"` method; the *c*-shell fuzzy clustering method

`m` The degree of fuzzification. It is defined for values greater than 1

`radius` The radius of resulting clusters

Details

The data given by `x` is clustered by the fuzzy *c*-shell algorithm.

If `centers` is a matrix, its rows are taken as the initial cluster centers. If `centers` is an integer, `centers` rows of `x` are randomly chosen as initial values.

The algorithm stops when the maximum number of iterations (given by `iter.max`) is reached.

If `verbose` is `TRUE`, it displays for each iteration the number the value of the objective function.

If `dist` is `"euclidean"`, the distance between the cluster center and the data points is the Euclidean distance (ordinary `kmeans` algorithm). If `"manhattan"`, the distance between the cluster center and the data points is the sum of the absolute values of the distances of the coordinates.

If `method` is `"cshell"`, then we have the *c*-shell fuzzy clustering method.

The parameter `m` defines the degree of fuzzification. It is defined for real values greater than 1 and the bigger it is the more fuzzy the membership values of the clustered data points are.

The parameter `radius` is by default set to `0.2` for every cluster.

Value

`cshell` returns an object of class `"cshell"`.

<code>centers</code>	The final cluster centers.
<code>size</code>	The number of data points in each cluster.
<code>cluster</code>	Vector containing the indices of the clusters where the data points are assigned to. The maximum membership value of a point is considered for partitioning it to a cluster.
<code>iter</code>	The number of iterations performed.
<code>membership</code>	a matrix with the membership values of the data points to the clusters.
<code>withinerror</code>	Returns the sum of square distances within the clusters.
<code>call</code>	Returns a call in which all of the arguments are specified by their names.

Author(s)

Evgenia Dimitriadou

References

Rajesh N. Dave. *Fuzzy Shell-Clustering and Applications to Circle Detection in Digital Images*. Int. J. of General Systems, Vol. **16**, pp. 343-355, 1996.

Examples

```
## a 2-dimensional example
x<-rbind(matrix(rnorm(50,sd=0.3),ncol=2),
          matrix(rnorm(50,mean=1,sd=0.3),ncol=2))
cl<-cshell(x,2,20,verbose=TRUE,method="cshell",m=2)
print(cl)

# assign classes to some new data
y<-rbind(matrix(rnorm(13,sd=0.3),ncol=2),
          matrix(rnorm(13,mean=1,sd=0.3),ncol=2))
#      ycl<-predict(cl, y, type="both")
```

e1071-deprecated *Deprecated Functions in Package e1071*

Description

These functions are provided for compatibility with older versions of package **e1071** only, and may be defunct as soon as of the next release.

Usage

```
read.octave(file, quiet = FALSE)
```

Arguments

`file` a character string giving the name of a file in Octave text data format (as created by `save -ascii`) to read.

`quiet` do not print information on type and size of the object read in.

Details

`read.octave` reads only the first object found in the file and ignores the rest.

Value

`read.octave` returns an object of type vector or matrix.

See Also

[Deprecated](#); [read.octave](#)

element *Extract Elements of an Array*

Description

Returns the element of `x` specified by `i`.

Usage

```
element(x, i)
```

Arguments

`x` Array of arbitrary dimensionality.

`i` Vector of the same length as `x` has dimension.

Author(s)

Friedrich Leisch

See Also[Extract](#)**Examples**

```
x <- array(1:20, dim=c(2,5,2))
element(x, c(1,4,2))
```

fclustIndex

*Fuzzy Cluster Indexes (Validity/Performance Measures)***Description**

Calculates the values of several fuzzy validity measures. The values of the indexes can be independently used in order to evaluate and compare clustering partitions or even to determine the number of clusters existing in a data set.

Usage

```
fclustIndex(y, x, index = "all")
```

Arguments

<code>y</code>	An object of a fuzzy clustering result of class "fclust"
<code>x</code>	Data matrix
<code>index</code>	The validity measures used: "gath.geva", "xie.beni", "fukuyama.sugeno", "partition.coefficient", "partition.entropy", "proportion.exponent", "separation.index" and "all" for all the indexes.

Details

The validity measures and a short description of them follows, where N is the number of data points, u_{ij} the values of the membership matrix, v_j the centers of the clusters and k te number of clusters.

gath.geva: Gath and Geva introduced 2 main criteria for comparing and finding optimal partitions based on the heuristics that a better clustering assumes clear separation between the clusters, minimal volume of the clusters and maximal number of data points concentrated in the vicinity of the cluster centroids. These indexes are only for the cmeans clustering algorithm valid. For the first, the “fuzzy hypervolume” we have: $F_{HV} = \sum_{j=1}^c [\det(F_j)]^{1/2}$, where

$$F_j = \frac{\sum_{i=1}^N u_{ij}(x_i - v_j)(x_i - v_j)^T}{\sum_{i=1}^N u_{ij}}$$

for the case when the defuzzification parameter is 2. For the second, the “average partition density”: $D_{PA} = \frac{1}{k} \sum_{j=1}^k \frac{S_j}{[\det(F_j)]^{1/2}}$, where $S_j = \sum_{i=1}^N u_{ij}$.

Moreover, the “partition density” which expresses the general partition density according to the physical definition of density is calculated by: $P_D = \frac{S}{F_{HV}}$, where $S = \sum_{j=1}^k \sum_{i=1}^N u_{ij}$.

xie.beni: This index is a function of the data set and the centroids of the clusters. Xie and Beni explained this index by writing it as a ratio of the total variation of the partition and the centroids (U, V) and the separation of the centroids vectors. The minimum values of this index

under comparison support the best partitions. $u_{XB}(U, V; X) = \frac{\sum_{j=1}^k \sum_{i=1}^N u_{ij}^2 \|x_i - v_j\|^2}{N(\min_{j \neq l} \{\|v_j - v_l\|^2\})}$

fukuyama.sugeno: This index consists of the difference of two terms, the first combining the fuzziness in the membership matrix with the geometrical compactness of the representation of the data set via the prototypes, and the second the fuzziness in its row of the partition matrix with the distance from the i th prototype to the grand mean of the data. The minimum values of this index also propose a good partition. $u_{FS}(U, V; X) = \sum_{i=1}^N \sum_{j=1}^k (u_{ij}^2)^q (\|x_i - v_j\|^2 - \|v_j - \bar{v}\|^2)$

partition.coefficient: An index which measures the fuzziness of the partition but without considering the data set itself. It is a heuristic measure since it has no connection to any property of the data. The maximum values of it imply a good partition in the meaning of a least fuzzy clustering. $F(U; k) = \frac{tr(UU^T)}{N} = \frac{\langle U, U \rangle}{N} = \frac{\|U\|^2}{N}$

- $F(U; k)$ shows the fuzziness or the overlap of the partition and depends on kN elements.
- $1/k \leq F(U; k) \leq 1$, where if $F(U; k) = 1$ then U is a hard partition and if $F(U; k) = 1/k$ then $U = [1/k]$ is the centroid of the fuzzy partition space P_{fk} . The converse is also valid.

partition.entropy: It is a measure that provides information about the membership matrix without also considering the data itself. The minimum values imply a good partition in the meaning of a more crisp partition. $H(U; k) = \sum_{i=1}^N h(u_i)/N$, where $h(u) = -\sum_{j=1}^k u_j \log_a(u_j)$ the Shannon's entropy.

- $H(U; k)$ shows the uncertainty of a fuzzy partition and depends also on kN elements. Specifically, $h(u_i)$ is interpreted as the amount of fuzzy information about the membership of x_i in k classes that is retained by column u_j . Thus, at $U = [1/k]$ the most information is withheld since the membership is the fuzziest possible.
- $0 \leq H(U; k) \leq \log_a(k)$, where for $H(U; k) = 0$ U is a hard partition and for $H(U; k) = \log_a(k)$ $U = [1/k]$.

proportion.exponent: It is a measure $P(U; k)$ of fuzziness adept to detect structural variations in the partition matrix as it becomes more fuzzier. A crisp cluster in the partition matrix can drive it to infinity when the partition coefficient and the partition entropy are more sensitive to small changes when approaching a hard partition. Its evaluation does not also involve the data or the algorithm used to partition them and its maximum implies the optimal partition but without knowing what maximum is a statistically significant maximum.

- $0 \leq P(U; k) < \infty$, since the $[0, 1]$ values explode to $[0, \infty)$ due to the natural logarithm. Specifically, $P = 0$ when and only when $U = [1/k]$, while $P \rightarrow \infty$ when any column of U is crisp.
- $P(U; k)$ can easily explode and it is good for partitions with large column maximums and at detecting structural variations.

separation.index (known as CS Index): This index identifies unique cluster structure with well-defined properties that depend on the data and a measure of distance. It answers the question if the clusters are compact and separated, but it rather seems computationally infeasible for big data sets since a distance matrix between all the data membership values has to be calculated. It also presupposes that a hard partition is derived from the fuzzy one.

$D_1(U; k; X, d) = \min_{i+1 \leq l \leq k-1} \left\{ \min_{1 \leq j \leq k} \left\{ \frac{dis(u_j, u_l)}{\max_{1 \leq m \leq k} \{dia(u_m)\}} \right\} \right\}$, where dia is the diameter of the subset, dis the distance of two subsets, and d a metric. U is a CS partition of $X \Leftrightarrow D_1 > 1$. When this holds then U is unique.

Value

Returns a vector with the validity measures values.

Author(s)

Evgenia Dimitriadou

References

- James C. Bezdek, *Pattern Recognition with Fuzzy Objective Function Algorithms*, Plenum Press, 1981, NY.
- L. X. Xie and G. Beni, *Validity measure for fuzzy clustering*, IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. **3**, n. 8, p. 841-847, 1991.
- I. Gath and A. B. Geva, *Unsupervised Optimal Fuzzy Clustering*, IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. **11**, n. 7, p. 773-781, 1989.
- Y. Fukuyama and M. Sugeno, *A new method of choosing the number of clusters for the fuzzy c-means method*, Proc. 5th Fuzzy Syst. Symp., p. 247-250, 1989 (in japanese).

See Also

[cmeans](#)

Examples

```
# a 2-dimensional example
x<-rbind(matrix(rnorm(100,sd=0.3),ncol=2),
           matrix(rnorm(100,mean=1,sd=0.3),ncol=2))
cl<-cmeans(x,2,20,verbose=TRUE,method="cmeans")
resultindexes <- fclustIndex(cl,x, index="all")
resultindexes
```

hamming.distance *Hamming Distances of Vectors*

Description

If both `x` and `y` are vectors, `hamming.distance` returns the Hamming distance (number of different bytes) between this two vectors. If `x` is a matrix, the Hamming distances between the rows of `x` are computed and `y` is ignored.

Usage

```
hamming.distance(x, y)
```

Arguments

<code>x</code>	a vector or matrix.
<code>y</code>	an optional vector.

Examples

```
x <- c(1, 0, 0)
y <- c(1, 0, 1)
hamming.distance(x, y)
z <- rbind(x,y)
rownames(z) <- c("Fred", "Tom")
hamming.distance(z)
```

hamming.window *Computes the Coefficients of a Hamming Window.*

Description

The filter coefficients w_i of a Hamming window of length n are computed according to the formula

$$w_i = 0.54 - 0.46 \cos \frac{2\pi i}{n - 1}$$

Usage

```
hamming.window(n)
```

Arguments

`n` The length of the window.

Value

A vector containing the filter coefficients.

Author(s)

Andreas Weingessel

References

For a definition of the Hamming window, see for example
Alan V. Oppenheim and Roland W. Schaffer: "Discrete-Time Signal Processing", Prentice-Hall,
1989.

See Also

stft, hanning.window

Examples

```
hamming.window(10)

x<-rnorm(500)
y<-stft(x, wtype="hamming.window")
plot(y)
```

hanning.window *Computes the Coefficients of a Hanning Window.*

Description

The filter coefficients w_i of a Hanning window of length n are computed according to the formula

$$w_i = 0.5 - 0.5 \cos \frac{2\pi i}{n-1}$$

Usage

```
hanning.window(n)
```

Arguments

`n` The length of the window.

Value

A vector containing the filter coefficients.

Author(s)

Andreas Weingessel

References

For a definition of the Hanning window, see for example
Alan V. Oppenheim and Roland W. Schaffer: "Discrete-Time Signal Processing", Prentice-Hall,
1989.

See Also

stft, hamming.window

Examples

```
hanning.window(10)

x<-rnorm(500)
y<-stft(x, wtype="hanning.window")
plot(y)
```

impute *Replace Missing Values*

Description

Replaces missing values of a matrix or dataframe with the medians (`what="median"`) or means (`what="mean"`) of the respective columns.

Usage

```
impute(x, what = c("median", "mean"))
```

Arguments

`x` A matrix or dataframe.
`what` What to impute.

Value

A matrix or dataframe.

Author(s)

Friedrich Leisch

Examples

```
x<- matrix(1:10, ncol=2)
x[c(1,3,7)] <- NA
print(x)
print(impute(x))
```

interpolate *Interpolate Values of Array*

Description

For each row in matrix `x`, the hypercube of `a` containing this point is searched. The corners of the hypercube are linearly interpolated. By default, `dimnames(a)` is taken to contain the coordinate values for each point in `a`. This can be overridden using `adims`. If `method="constant"`, the value of the "lower left" corner of the hypercube is returned.

Usage

```
interpolate(x, a, adims=lapply(dimnames(a), as.numeric),
            method="linear")
```

Arguments

<code>x</code>	Matrix of values at which interpolation shall take place.
<code>a</code>	Array of arbitrary dimension.
<code>adims</code>	List of the same structure as <code>dimnames(a)</code> .
<code>method</code>	Interpolation method, one of "linear" or "constant".

Author(s)

Friedrich Leisch

See Also

[approx](#), [spline](#)

Examples

```
x <- seq(0,3,0.2)
z <- outer(x,x, function(x,y) sin(x*y))
dimnames(z) <- list(x,x)
sin(1.1*2.1)
interpolate(c(1.1, 2.1),z)
```

kurtosis

Kurtosis

Description

Computes the kurtosis.

Usage

```
kurtosis(x, na.rm=FALSE)
```

Arguments

<code>x</code>	a numeric vector containing the values whose kurtosis is to be computed.
<code>na.rm</code>	a logical value indicating whether NA values should be stripped before the computation proceeds.

Details

If $N = \text{length}(x)$, then the kurtosis of x is defined as

$$N^{-1} \text{sd}(x)^{-4} \sum_i (x_i - \text{mean}(x))^4 - 3.$$

Value

The kurtosis of `x`.

Examples

```
x <- rnorm(100)
kurtosis(x)
```

lca *Latent Class Analysis (LCA)*

Description

A latent class analysis with k classes is performed on the data given by x .

Usage

```
lca(x, k, niter=100, matchdata=FALSE, verbose=FALSE)
```

Arguments

<code>x</code>	Either a data matrix of binary observations or a list of patterns as created by countpattern
<code>k</code>	Number of classes used for LCA
<code>niter</code>	Number of Iterations
<code>matchdata</code>	If TRUE and x is a data matrix, the class membership of every data point is returned, otherwise the class membership of every pattern is returned.
<code>verbose</code>	If TRUE some output is printed during the computations.

Value

An object of class "lca" is returned, containing

<code>w</code>	Probabilities to belong to each class
<code>p</code>	Probabilities of a '1' for each variable in each class
<code>matching</code>	Depending on <code>matchdata</code> either the class membership of each pattern or of each data point
<code>logl, loglsat</code>	The LogLikelihood of the model and of the saturated model
<code>bic, bicsat</code>	The BIC of the model and of the saturated model
<code>chisq</code>	Pearson's Chisq
<code>lhquot</code>	Likelihood quotient of the model and the saturated model
<code>n</code>	Number of data points.
<code>np</code>	Number of free parameters.

Author(s)

Andreas Weingessel

References

Anton K. Formann: "Die Latent-Class-Analysis", Beltz Verlag 1984

See Also

[countpattern](#), [bootstrap.lca](#)

Examples

```
## Generate a 4-dim. sample with 2 latent classes of 500 data points each.
## The probabilities for the 2 classes are given by type1 and type2.
type1 <- c(0.8,0.8,0.2,0.2)
type2 <- c(0.2,0.2,0.8,0.8)
x <- matrix(runif(4000),nr=1000)
x[1:500,] <- t(t(x[1:500,])<type1)*1
x[501:1000,] <- t(t(x[501:1000,])<type2)*1

l <- lca(x, 2, niter=5)
print(l)
summary(l)
p <- predict(l, x)
table(p, c(rep(1,500),rep(2,500)))
```

matchClasses

Find similar classes in two-way contingency tables

Description

Try to find a mapping between the two groupings, such that as many cases as possible are in one of the matched pairs.

Usage

```
matchClasses(tab, method="rowmax", iter=1, maxexact=9, verbose=TRUE)
compareMatchedClasses(x, y, method="rowmax", iter=1,
                      maxexact=9, verbose=FALSE)
```

Arguments

tab	Two-way contingency table of class memberships
method	One of "rowmax", "greedy" or "exact".
iter	Number of iterations used in greedy search.
verbose	If TRUE, display some status messages during computation.
maxexact	Maximum number of variables for which all possible permutations are computed.
x, y	Vectors or matrices with class memberships.

Details

If method="rowmax", then each class defining a row in the contingency table is mapped to the column of the corresponding row maximum. Hence, some columns may be mapped to more than one row (while each row is mapped to a single column).

If method="greedy" or method="exact", then the contingency table must be a square matrix and a unique mapping is computed. This corresponds to a permutation of columns and rows, such that sum of the main diagonal, i.e., the trace of the matrix, gets as large as possible. For both methods, first all pairs where row and columns maxima correspond and are bigger than the sum of all other elements in the corresponding columns and rows together are located and fixed (this is a necessary condition for maximal trace).

If `method="exact"`, then for the remaining rows and columns, all possible permutations are computed and the optimum is returned. This can get computationally infeasible very fast. If more than `maxexact` rows and columns remain after applying the necessary condition, then `method` is reset to `"greedy"`. If `method="greedy"`, then a greedy heuristic is tried `iter` times. Repeatedly a row is picked at random and matched to the free column with the maximum value.

`compareMatchedClasses()` computes the contingency table for each combination of columns from `x` and `y` and applies `matchClasses` to that table. The columns of the table are permuted accordingly and then the table is passed to `classAgreement`. The resulting agreement coefficients (diag, kappa, ...) are returned. The return value of `compareMatchedClasses()` is a list containing a matrix for each coefficient; with element (k,l) corresponding to the k-th column of `x` and l-th column of `y`. If `y` is missing, then the columns of `x` are compared with each other.

Author(s)

Friedrich Leisch

See Also

[classAgreement](#)

Examples

```
## a stupid example with no class correlations:
g1 <- sample(1:5, size=1000, replace=TRUE)
g2 <- sample(1:5, size=1000, replace=TRUE)
tab <- table(g1, g2)
matchClasses(tab, "exact")

## let pairs (g1=1,g2=4) and (g1=3,g2=1) agree better
k <- sample(1:1000, size=200)
g1[k] <- 1
g2[k] <- 4

k <- sample(1:1000, size=200)
g1[k] <- 3
g2[k] <- 1

tab <- table(g1, g2)
matchClasses(tab, "exact")

## get agreement coefficients:
compareMatchedClasses(g1, g2, method="exact")
```

matchControls *Find matched control group*

Description

Finds controls matching the cases as good as possible.

Usage

```
matchControls(formula, data = list(), subset, contlabel = "con",
              caselabel = NULL, dogrep = TRUE, replace = FALSE)
```

Arguments

formula	A formula indicating cases, controls and the variables to be matched. Details are described below.
data	an optional data frame containing the variables in the model. By default the variables are taken from the environment which <code>matchControls</code> is called from.
subset	an optional vector specifying a subset of observations to be used in the matching process.
contlabel	A string giving the label of the control group.
caselabel	A string giving the labels of the cases.
dogrep	If <code>TRUE</code> , then <code>contlabel</code> and <code>caselabel</code> are matched using <code>grep</code> , else string comparison (exact equality) is used.
replace	If <code>FALSE</code> , then every control is used only once.

Details

The left hand side of the `formula` must be a factor determining whether an observation belongs to the case or the control group. By default, all observations where a `grep` of `contlabel` matches, are used as possible controls, the rest is taken as cases. If `caselabel` is given, then only those observations are taken as cases. If `dogrep = TRUE`, then both `contlabel` and `caselabel` can be regular expressions.

The right hand side of the `formula` gives the variables that should be matched. The matching is done using the `daisy` distance from the `cluster` package, i.e., a model frame is built from the formula and used as input for `daisy`. For each case, the nearest control is selected. If `replace = FALSE`, each control is used only once.

Value

Returns a list with components

cases	Row names of cases.
controls	Row names of matched controls.
factor	A factor with 2 levels indicating cases and controls (the rest is set to <code>NA</code>).

Author(s)

Friedrich Leisch

Examples

```
Age.case <- 40 + 5 * rnorm(50)
Age.cont <- 45 + 10 * rnorm(150)
Age <- c(Age.case, Age.cont)

Sex.case <- sample(c("M", "F"), 50, prob = c(.4, .6), replace = TRUE)
Sex.cont <- sample(c("M", "F"), 150, prob = c(.6, .4), replace = TRUE)
Sex <- as.factor(c(Sex.case, Sex.cont))

casecont <- as.factor(c(rep("case", 50), rep("cont", 150)))

## now look at the group properties:
```

```

boxplot(Age ~ casecont)
barplot(table(Sex, casecont), beside = TRUE)

m <- matchControls(casecont ~ Sex + Age)

## properties of the new groups:
boxplot(Age ~ m$factor)
barplot(table(Sex, m$factor))

```

moment

Statistical Moment

Description

Computes the (optionally centered and/or absolute) sample moment of a certain order.

Usage

```
moment(x, order=1, center=FALSE, absolute=FALSE, na.rm=FALSE)
```

Arguments

x	a numeric vector containing the values whose moment is to be computed.
order	order of the moment to be computed, the default is to compute the first moment, i.e., the mean.
center	a logical value indicating whether centered moments are to be computed.
absolute	a logical value indicating whether absolute moments are to be computed.
na.rm	a logical value indicating whether NA values should be stripped before the computation proceeds.

Details

When `center` and `absolute` are both `FALSE`, the moment is simply $\text{sum}(x^{\text{order}}) / \text{length}(x)$.

Author(s)

Kurt Hornik and Friedrich Leisch

See Also

[mean](#), [var](#)

Examples

```

x <- rnorm(100)

## Compute the mean
moment(x)
## Compute the 2nd centered moment (!= var)
moment(x, order=2, center=TRUE)

## Compute the 3rd absolute centered moment
moment(x, order=3, center=TRUE, absolute=TRUE)

```

naiveBayes	<i>Naive Bayes Classifier</i>
------------	-------------------------------

Description

Computes the conditional a-posterior probabilities of a categorical class variable given independent predictor variables using the Bayes rule.

Usage

```
## S3 method for class 'formula':
naiveBayes(formula, data, ..., subset, na.action = na.pass)
## Default S3 method:
naiveBayes(x, y, ...)
```

Arguments

<code>x</code>	A numeric matrix, or a data frame of categorical and/or numeric variables.
<code>y</code>	Class vector.
<code>formula</code>	A formula of the form <code>class ~ x1 + x2 + ...</code> . Interactions are not allowed.
<code>data</code>	Either a data frame of predictors (caegorical and/or numeric) or a contingency table.
<code>...</code>	Currently not used.
<code>subset</code>	For data given in a data frame, an index vector specifying the cases to be used in the training sample. (NOTE: If given, this argument must be named.)
<code>na.action</code>	A function to specify the action to be taken if NAs are found. The default action is not to count them for the computation of the probability factors. An alternative is <code>na.omit</code> , which leads to rejection of cases with missing values on any required variable. (NOTE: If given, this argument must be named.)

Details

The standard naive Bayes classifier (at least this implementation) assumes independence of the predictor variables, and gaussian distribution (given the target class) of metric predictors.

Value

An object of class "naiveBayes" including components:

<code>apriori</code>	Class distribution for the dependent variable.
<code>tables</code>	A list of tables, one for each predictor variable. For each categorical variable a table giving, for each attribute level, the conditional probabilities given the target class. For each numeric variable, a table giving, for each target class, mean and standard deviation of the (sub-)variable.

Author(s)

David Meyer (David.Meyer@R-project.org)

See Also[predict.naiveBayes](#)**Examples**

```
## Categorical data only:
data(HouseVotes84)
model <- naiveBayes(Class ~ ., data = HouseVotes84)
predict(model, HouseVotes84[1:10,-1])
predict(model, HouseVotes84[1:10,-1], type = "raw")

pred <- predict(model, HouseVotes84[,-1])
table(pred, HouseVotes84$Class)

## Example of using a contingency table:
data(Titanic)
m <- naiveBayes(Survived ~ ., data = Titanic)
m
predict(m, as.data.frame(Titanic)[,1:3])

## Example with metric predictors:
data(iris)
m <- naiveBayes(Species ~ ., data = iris)
## alternatively:
m <- naiveBayes(iris[,-5], iris[,5])
m
table(predict(m, iris[,-5]), iris[,5])
```

permutations

All permutations of integers 1:n

Description

Returns a matrix containing all permutations of the integers 1:n (one permutation per row).

Usage

```
permutations(n)
```

Arguments

n Number of element to permute.

Author(s)

Friedrich Leisch

Examples

```
permutations(3)
```

`plot.stft`*Plot Short Time Fourier Transforms*

Description

An object of class "stft" is plotted as a gray scale image. The x-axis corresponds to time, the y-axis to frequency. If the default colormap is used, dark regions in the plot correspond to high values at the particular time/frequency location.

Usage

```
## S3 method for class 'stft':  
plot(x, col = gray(63:0/63), ...)
```

Arguments

<code>x</code>	An object of class "stft" as obtained by the function <code>stft</code> .
<code>col</code>	An optional colormap. By default 64 gray values are used, where white corresponds to the minimum value and black to the maximum.
<code>...</code>	further arguments to be passed to or from methods.

Value

No return value. This function is only for plotting.

Author(s)

Andreas Weingessel

See Also

`stft`

Examples

```
x<-rnorm(500)  
y<-stft(x)  
plot(y)
```

`plot.svm`*Plot svm objects*

Description

Generates a scatter plot of the input data of a svm fit for classification models by highlighting the classes and support vectors. Optionally, draws a filled contour plot of the class regions.

Usage

```
plot.svm(x, data, formula, fill = TRUE, grid = 50, slice = list(),  
symbolPalette = palette(), svSymbol = "x", dataSymbol = "o", ...)
```

Arguments

<code>x</code>	An object of class <code>svm</code>
<code>data</code>	data to visualize. Should be the same used for fitting.
<code>formula</code>	formula selecting the visualized two dimensions. Only needed if more than two input variables are used.
<code>fill</code>	switch indicating whether a contour plot for the class regions should be added.
<code>grid</code>	granularity for the contour plot.
<code>slice</code>	a list of named numeric values for the dimensions held constant (only needed if more than two variables are used). Dimensions not specified are fixed at 0.
<code>symbolPalette</code>	Color palette used for the class the data points and support vectors belong to.
<code>svSymbol</code>	Symbol used for support vectors.
<code>dataSymbol</code>	Symbol used for data points (other than support vectors).
<code>...</code>	additional graphics parameters passed to <code>filled.contour</code> and <code>plot</code> .

Author(s)

David Meyer
(David.Meyer@R-project.org)

See Also

[svm](#)

Examples

```
## a simple example
library(MASS)
data(cats)
m <- svm(Sex~., data = cats)
plot(m, cats)

## more than two variables: fix 2 dimensions
data(iris)
m2 <- svm(Species~., data = iris)
plot(m2, iris, Petal.Width ~ Petal.Length,
      slice = list(Sepal.Width = 3, Sepal.Length = 4))

## plot with custom symbols and colors
plot(m, cats, svSymbol = 1, dataSymbol = 2, symbolPalette = rainbow(4),
      color.palette = terrain.colors)
```

plot.tune	<i>Plot tuning object</i>
-----------	---------------------------

Description

Visualizes the results of parameter tuning.

Usage

```
plot.tune(x, type = c("contour", "perspective"), theta = 60,
          col = "lightblue", main = NULL, xlab = NULL, ylab = NULL,
          swapxy = FALSE, transform.x = NULL, transform.y = NULL,
          transform.z = NULL, color.palette = topo.colors,
          nlevels = 20, ...)
```

Arguments

<code>x</code>	an object of class <code>tune</code>
<code>type</code>	choose whether a contour plot or a perspective plot is used if two parameters are to be visualized. Ignored if only one parameter has been tuned.
<code>theta</code>	angle of azimuthal direction.
<code>col</code>	the color(s) of the surface facets. Transparent colours are ignored.
<code>main</code>	main title
<code>xlab, ylab</code>	titles for the axes. N.B. These must be character strings; expressions are not accepted. Numbers will be coerced to character strings.
<code>swapxy</code>	if TRUE, the parameter axes are swapped (only used in case of two parameters).
<code>transform.x, transform.y, transform.z</code>	functions to transform the parameters (<code>x</code> and <code>y</code>) and the error measures (<code>z</code>). Ignored if NULL.
<code>color.palette</code>	color palette used in contour plot.
<code>nlevels</code>	number of levels used in contour plot.
<code>...</code>	Further graphics parameters.

Author(s)

David Meyer (based on C/C++-code by Chih-Chung Chang and Chih-Jen Lin)
(David.Meyer@R-project.org)

See Also

[tune](#)

Examples

```
data(iris)
obj <- tune.svm(Species~., data = iris, sampling = "fix",
               gamma = 2^c(-8,-4,0,4), cost = 2^c(-8,-4,-2,0))
plot(obj, transform.x = log2, transform.y = log2)
plot(obj, type = "perspective", theta = 120, phi = 45)
```

predict.naiveBayes *Naive Bayes Classifier*

Description

Computes the conditional a-posterior probabilities of a categorical class variable given independent predictor variables using the Bayes rule.

Usage

```
predict.naiveBayes(object, newdata, type = c("class", "raw"), threshold = 0.001,
```

Arguments

object	An object of class "naiveBayes".
newdata	A dataframe with new predictors.
type	see value.
threshold	Value replacing cells with 0 probabilities.
...	Currently not used.

Details

The standard naive Bayes classifier (at least this implementation) assumes independence of the predictor variables, and gaussian distribution (given the target class) of metric predictors. For attributes with missing values, the corresponding table entries are omitted for prediction.

Value

If `type = "raw"`, the conditional a-posterior probabilities for each class are returned, and the class with maximal probability else.

Author(s)

David Meyer <David.Meyer@R-project.org>

See Also

[naiveBayes](#)

Examples

```
## Categorical data only:
data(HouseVotes84)
model <- naiveBayes(Class ~ ., data = HouseVotes84)
predict(model, HouseVotes84[1:10,-1])
predict(model, HouseVotes84[1:10,-1], type = "raw")

pred <- predict(model, HouseVotes84[,-1])
table(pred, HouseVotes84$Class)

## Example of using a contingency table:
data(Titanic)
```

```

m <- naiveBayes(Survived ~ ., data = Titanic)
m
predict(m, as.data.frame(Titanic)[,1:3])

## Example with metric predictors:
data(iris)
m <- naiveBayes(Species ~ ., data = iris)
## alternatively:
m <- naiveBayes(iris[,-5], iris[,5])
m
table(predict(m, iris[,-5]), iris[,5])

```

predict.svm

Predict method for Support Vector Machines

Description

This function predicts values based upon a model trained by `svm`.

Usage

```

## S3 method for class 'svm':
predict(object, newdata, decision.values = FALSE,
        probability = FALSE, ..., na.action = na.omit)

```

Arguments

<code>object</code>	Object of class "svm", created by <code>svm</code> .
<code>newdata</code>	A matrix containing the new input data. A vector will be transformed to a $n \times 1$ matrix.
<code>decision.values</code>	Logical controlling whether the decision values of all binary classifiers computed in multiclass classification shall be computed and returned.
<code>probability</code>	Logical indicating whether class probabilities should be computed and returned. Only possible if the model was fitted with the <code>probability</code> option enabled.
<code>na.action</code>	A function to specify the action to be taken if 'NA's are found. The default action is <code>na.omit</code> , which leads to rejection of cases with missing values on any required variable. An alternative is <code>na.fail</code> , which causes an error if NA cases are found. (NOTE: If given, this argument must be named.)
<code>...</code>	Currently not used.

Value

A vector of predicted values (for classification: a vector of labels, for density estimation: a logical vector). If `decision.value` is `TRUE`, the vector gets a "decision.values" attribute containing a $n \times c$ matrix (n number of predicted values, c number of classifiers) of all c binary classifiers' decision values. There are $k * (k - 1) / 2$ classifiers (k number of classes). The colnames of the matrix indicate the labels of the two classes. If `probability` is `TRUE`, the vector gets a "probabilities" attribute containing a $n \times k$ matrix (n number of predicted values, k number of classes) of the class probabilities.

Note

If the training set was scaled by `svm` (done by default), the new data is scaled accordingly using scale and center of the training data.

Author(s)

David Meyer (based on C++-code by Chih-Chung Chang and Chih-Jen Lin)
(David.Meyer@R-project.org)

See Also

[svm](#)

Examples

```
data(iris)
attach(iris)

## classification mode
# default with factor response:
model <- svm(Species ~ ., data = iris)

# alternatively the traditional interface:
x <- subset(iris, select = -Species)
y <- Species
model <- svm(x, y, probability = TRUE)

print(model)
summary(model)

# test with train data
pred <- predict(model, x)
# (same as:)
pred <- fitted(model)

# compute decision values and probabilities
pred <- predict(model, x, decision.values = TRUE, probability = TRUE)
attr(pred, "decision.values")[1:4,]
attr(pred, "probabilities")[1:4,]

## try regression mode on two dimensions

# create data
x <- seq(0.1, 5, by = 0.05)
y <- log(x) + rnorm(x, sd = 0.2)

# estimate model and predict input values
m <- svm(x, y)
new <- predict(m, x)

# visualize
plot(x, y)
points(x, log(x), col = 2)
points(x, new, col = 4)

## density-estimation
```

```
# create 2-dim. normal with rho=0:
X <- data.frame(a = rnorm(1000), b = rnorm(1000))
attach(X)

# traditional way:
m <- svm(X, gamma = 0.1)

# formula interface:
m <- svm(~., data = X, gamma = 0.1)
# or:
m <- svm(~ a + b, gamma = 0.1)

# test:
newdata <- data.frame(a = c(0, 4), b = c(0, 4))
predict (m, newdata)

# visualize:
plot(X, col = 1:1000 %in% m$index + 1, xlim = c(-5,5), ylim=c(-5,5))
points(newdata, pch = "+", col = 2, cex = 5)
```

rbridge

Simulation of Brownian Bridge

Description

rwiener returns a time series containing a simulated realization of the Brownian bridge on the interval $[0, \text{end}]$. If $W(t)$ is a Wiener process, then the Brownian bridge is defined as $W(t) - t W(1)$.

Usage

```
rbridge(end = 1, frequency = 1000)
```

Arguments

end	the time of the last observation.
frequency	the number of observations per unit of time.

See Also

rwiener

Examples

```
# simulate a Brownian bridge on [0,1] and plot it

x <- rbridge()
plot(x, type="l")
```

read.matrix.csr *read/write sparse data*

Description

reads and writes a file in sparse data format.

Usage

```
read.matrix.csr(file, fac = TRUE, ncol = NULL)
write.matrix.csr(x, file = "out.dat", y = NULL)
```

Arguments

x	An object of class <code>matrix.csr</code>
y	A vector (either numeric or a factor)
file	The filename.
fac	If TRUE and y-values are stored in the file, the values are interpreted as factor levels.
ncol	Number of columns, detected automatically. Can be used to add empty columns (these are not stored in the sparse format).

Value

If the data file includes no y variable, `read.matrix.csr` returns an object of class `matrix.csr`, else a list with components:

x	object of class <code>matrix.csr</code>
y	vector of numeric values or factor levels, depending on <code>fac</code> .

Author(s)

David Meyer (based on C/C++-code by Chih-Chung Chang and Chih-Jen Lin)
(David.Meyer@R-project.org)

See Also

[matrix.csr](#)

Examples

```
## Not run:
library(methods)
if (require(SparseM)) {
  data(iris)
  x <- as.matrix(iris[,1:4])
  y <- iris[,5]
  xs <- as.matrix.csr(x)
  write.matrix.csr(xs, y = y, file="iris.dat")
  xs2 <- read.matrix.csr("iris.dat")$x
  if (!all(as.matrix(xs) == as.matrix(xs2)))
```

```
    stop("Error: objects are not equal!")
  }
## End(Not run)
```

rectangle.window *Computes the Coefficients of a Rectangle Window.*

Description

Returns the filter coefficients of a rectangle window. That is a vector of n 1.

The purpose of this function is just to have a name for the R command `rep(1, n)`.

Usage

```
rectangle.window(n)
```

Arguments

`n` The length of the window.

Value

A vector of length n filled with 1.

Author(s)

Andreas Weingessel

See Also

`stft`

Examples

```
x<-rnorm(500)
y<-stft(x, wtype="rectangle.window")
plot(y)
```

rwiener *Simulation of Wiener Process*

Description

`rwiener` returns a time series containing a simulated realization of the Wiener process on the interval $[0, \text{end}]$

Usage

```
rwiener(end = 1, frequency = 1000)
```

Arguments

end the time of the last observation.
 frequency the number of observations per unit of time.

Examples

```
# simulate a Wiener process on [0,1] and plot it

x <- rwiener()
plot(x,type="l")
```

allShortestPaths *Find Shortest Paths Between All Nodes in a Directed Graph*

Description

allShortestPaths finds all shortest paths in a directed (or undirected) graph using Floyd's algorithm. extractPath can be used to actually extract the path between a given pair of nodes.

Usage

```
allShortestPaths(x)
extractPath(obj, start, end)
```

Arguments

x matrix or distance object
 obj return value of allShortestPaths
 start integer, starting point of path
 end integer, end point of path

Details

If x is a matrix, then $x[i, j]$ has to be the length of the direct path from point i to point j . If no direct connection from point i to point j exist, then $x[i, j]$ should be either NA or Inf. Note that the graph can be directed, hence $x[i, j]$ need not be the same as $x[j, i]$. The main diagonal of x is ignored. Alternatively, x can be a distance object as returned by `dist` (corresponding to an undirected graph).

Value

allShortestPaths returns a list with components

length A matrix with the total lengths of the shortest path between each pair of points.
 middlePoints A matrix giving a point in the middle of each shortest path (or 0 if the direct connection is the shortest path), this is mainly used as input for extractPath.
 extractPath returns a vector of node numbers giving with the shortest path between two points.

Author(s)

Friedrich Leisch

References

Kumar, V., Grama, A., Gupta, A. and Karypis, G. Introduction to Parallel Programming - Design and Analysis of Algorithms, Benjamin Cummings Publishing, 1994, ISBN 0-8053-3170-0

Examples

```
## build a graph with 5 nodes
x <- matrix(NA, 5, 5)
diag(x) <- 0
x[1,2] <- 30; x[1,3] <- 10
x[2,4] <- 70; x[2,5] <- 40
x[3,4] <- 50; x[3,5] <- 20
x[4,5] <- 60
x[5,4] <- 10
print(x)

## compute all path lengths
z <- allShortestPaths(x)
print(z)

## the following should give 1 -> 3 -> 5 -> 4
extractPath(z, 1, 4)
```

sigmoid

The logistic function and derivatives

Description

Sigmoid $1/(1 + \exp(-x))$, first and second derivative.

Usage

```
sigmoid(x)
dsigmoid(x)
d2sigmoid(x)
```

Arguments

x a numeric vector

Author(s)

Friedrich Leisch

Examples

```
plot(sigmoid, -5, 5, ylim = c(-.2, 1))
plot(dsigmoid, -5, 5, add = TRUE, col = 2)
plot(d2sigmoid, -5, 5, add = TRUE, col = 3)
```

 skewness

Skewness

Description

Computes the skewness.

Usage

```
skewness(x, na.rm=FALSE)
```

Arguments

`x` a numeric vector containing the values whose skewness is to be computed.
`na.rm` a logical value indicating whether NA values should be stripped before the computation proceeds.

Details

If $N = \text{length}(x)$, then the skewness of x is defined as

$$N^{-1} \text{sd}(x)^{-3} \sum_i (x_i - \text{mean}(x))^3.$$

Value

The skewness of `x`.

Examples

```
x <- rnorm(100)
skewness(x)
```

 stft

Computes the Short Time Fourier Transform of a Vector

Description

This function computes the Short Time Fourier Transform of a given vector `X`.

First, time-slices of length `win` are extracted from the vector. The shift of one time-slice to the next one is given by `inc`. The values of these time-slices are smoothed by multiplying them with a window function specified in `wtype`. For the thus obtained windows, the Fast Fourier Transform is computed.

Usage

```
stft(X, win=min(80,floor(length(X)/10)), inc=min(24,
floor(length(X)/30)), coef=64, wtype="hanning.window")
```

Arguments

<code>x</code>	The vector from which the stft is computed.
<code>win</code>	Length of the window. For long vectors the default window size is 80, for short vectors the window size is chosen so that 10 windows fit in the vector.
<code>inc</code>	Increment by which the window is shifted. For long vectors the default increment is 24, for short vectors the increment is chosen so that 30 increments fit in the vector.
<code>coef</code>	Number of Fourier coefficients
<code>wtype</code>	Type of window used

Value

Object of type `stft`. Contains the values of the stft and information about the parameters.

<code>values</code>	A matrix containing the results of the stft. Each row of the matrix contains the <code>coef</code> Fourier coefficients of one window.
<code>windowsize</code>	The value of the parameter <code>win</code>
<code>increment</code>	The value of the parameter <code>inc</code>
<code>windowtype</code>	The value of the parameter <code>wtype</code>

Author(s)

Andreas Weingessel

See Also

`plot.stft`

Examples

```
x<-rnorm(500)
y<-stft(x)
plot(y)
```

Description

`svm` is used to train a support vector machine. It can be used to carry out general regression and classification (of `nu` and `epsilon`-type), as well as density-estimation. A formula interface is provided.

Usage

```
## S3 method for class 'formula':
svm(formula, data = NULL, ..., subset, na.action =
na.omit, scale = TRUE)
## Default S3 method:
svm(x, y = NULL, scale = TRUE, type = NULL, kernel =
"radial", degree = 3, gamma = 1 / ncol(as.matrix(x)), coef0 = 0, cost = 1, nu =
class.weights = NULL, cachesize = 40, tolerance = 0.001, epsilon = 0.1,
shrinking = TRUE, cross = 0, probability = FALSE, fitted = TRUE,
..., subset, na.action = na.omit)
```

Arguments

formula	a symbolic description of the model to be fit.
data	an optional data frame containing the variables in the model. By default the variables are taken from the environment which ‘svm’ is called from.
x	a data matrix, a vector, or a sparse matrix (object of class <code>matrix.csr</code> as provided by the package SparseM).
y	a response vector with one label for each row/component of x. Can be either a factor (for classification tasks) or a numeric vector (for regression).
scale	A logical vector indicating the variables to be scaled. If <code>scale</code> is of length 1, the value is recycled as many times as needed. Per default, data are scaled internally (both x and y variables) to zero mean and unit variance. The center and scale values are returned and used for later predictions.
type	svm can be used as a classification machine, as a regresson machine, or for novelty detection. Depending of whether y is a factor or not, the default setting for type is C-classification or eps-regression, respectively, but may be overwritten by setting an explicit value. Valid options are: <ul style="list-style-type: none"> • C-classification • nu-classification • one-classification (for novelty detection) • eps-regression • nu-regression
kernel	the kernel used in training and predicting. You might consider changing some of the following parameters, depending on the kernel type. <p>linear: $u'v$</p> <p>polynomial: $(\gamma u'v + coef0)^{degree}$</p> <p>radial basis: $e^{-(\gamma u-v ^2)}$</p> <p>sigmoid: $\tanh(\gamma u'v + coef0)$</p>
degree	parameter needed for kernel of type polynomial (default: 3)
gamma	parameter needed for all kernels except linear (default: 1/(data dimension))
coef0	parameter needed for kernels of type polynomial and sigmoid (default: 0)
cost	cost of constraints violation (default: 1)—it is the ‘C’-constant of the regularization term in the Lagrange formulation.
nu	parameter needed for nu-classification, nu-regression, and one-classification

<code>class.weights</code>	a named vector of weights for the different classes, used for asymmetric class sizes. Not all factor levels have to be supplied (default weight: 1). All components have to be named.
<code>cache.size</code>	cache memory in MB (default 40)
<code>tolerance</code>	tolerance of termination criterion (default: 0.001)
<code>epsilon</code>	epsilon in the insensitive-loss function (default: 0.1)
<code>shrinking</code>	option whether to use the shrinking-heuristics (default: TRUE)
<code>cross</code>	if a integer value $k > 0$ is specified, a k -fold cross validation on the training data is performed to assess the quality of the model: the accuracy rate for classification and the Mean Squared Error for regression
<code>fitted</code>	logical indicating whether the fitted values should be computed and included in the model or not (default: TRUE)
<code>probability</code>	logical indicating whether the model should allow for probability predictions.
<code>...</code>	additional parameters for the low level fitting function <code>svm.default</code>
<code>subset</code>	An index vector specifying the cases to be used in the training sample. (NOTE: If given, this argument must be named.)
<code>na.action</code>	A function to specify the action to be taken if NAs are found. The default action is <code>na.omit</code> , which leads to rejection of cases with missing values on any required variable. An alternative is <code>na.fail</code> , which causes an error if NA cases are found. (NOTE: If given, this argument must be named.)

Details

For multiclass-classification with k levels, $k > 2$, `libsvm` uses the ‘one-against-one’-approach, in which $k(k-1)/2$ binary classifiers are trained; the appropriate class is found by a voting scheme. `libsvm` internally uses a sparse data representation, which is also high-level supported by the package **SparseM**.

If the predictor variables include factors, the formula interface must be used to get a correct model matrix.

`plot.svm` allows a simple graphical visualization of classification models.

The probability model for classification fits a logistic distribution using maximum likelihood to the decision values of all binary classifiers, and computes the a-posteriori class probabilities for the multi-class problem using quadratic optimization. The probabilistic regression model assumes (zero-mean) laplace-distributed errors for the predictions, and estimates the scale parameter using maximum likelihood.

Value

An object of class "svm" containing the fitted model, including:

<code>SV</code>	The resulting support vectors (possibly scaled).
<code>index</code>	The index of the resulting support vectors in the data matrix. Note that this index refers to the preprocessed data (after the possible effect of <code>na.omit</code> and <code>subset</code>)
<code>coefs</code>	The corresponding coefficients times the training labels.
<code>rho</code>	The negative intercept.
<code>sigma</code>	In case of a probabilistic regression model, the scale parameter of the hypothesized (zero-mean) laplace distribution estimated by maximum likelihood.

probA, probB numeric vectors of length $k(k-1)/2$, k number of classes, containing the parameters of the logistic distributions fitted to the decision values of the binary classifiers ($1 / (1 + \exp(a x + b))$).

Note

Data are scaled internally, usually yielding better results.

Author(s)

David Meyer (based on C/C++-code by Chih-Chung Chang and Chih-Jen Lin)
(David.Meyer@R-project.org)

References

- Chang, Chih-Chung and Lin, Chih-Jen:
LIBSVM: a library for Support Vector Machines
<http://www.csie.ntu.edu.tw/~cjlin/libsvm>
- Exact formulations of models, algorithms, etc. can be found in the document:
Chang, Chih-Chung and Lin, Chih-Jen:
LIBSVM: a library for Support Vector Machines
<http://www.csie.ntu.edu.tw/~cjlin/papers/libsvm.ps.gz>
- More implementation details and speed benchmarks can be found on: Rong-En Fan and Pai-Hsune Chen and Chih-Jen Lin:
Working Set Selection Using the Second Order Information for Training SVM
<http://www.csie.ntu.edu.tw/~cjlin/papers/quadworkset.pdf>

See Also

`predict.svm`, `plot.svm`, `matrix.csr` (in package **SparseM**)

Examples

```
data(iris)
attach(iris)

## classification mode
# default with factor response:
model <- svm(Species ~ ., data = iris)

# alternatively the traditional interface:
x <- subset(iris, select = -Species)
y <- Species
model <- svm(x, y)

print(model)
summary(model)

# test with train data
pred <- predict(model, x)
# (same as:)
pred <- fitted(model)

# Check accuracy:
table(pred, y)
```

```

# compute decision values and probabilities:
pred <- predict(model, x, decision.values = TRUE)
attr(pred, "decision.values")[1:4,]

# visualize (classes by color, SV by crosses):
plot(cmdscale(dist(iris[,-5])),
     col = as.integer(iris[,5]),
     pch = c("o","+") [1:150 %in% model$index + 1])

## try regression mode on two dimensions

# create data
x <- seq(0.1, 5, by = 0.05)
y <- log(x) + rnorm(x, sd = 0.2)

# estimate model and predict input values
m <- svm(x, y)
new <- predict(m, x)

# visualize
plot(x, y)
points(x, log(x), col = 2)
points(x, new, col = 4)

## density-estimation

# create 2-dim. normal with rho=0:
X <- data.frame(a = rnorm(1000), b = rnorm(1000))
attach(X)

# traditional way:
m <- svm(X, gamma = 0.1)

# formula interface:
m <- svm(~., data = X, gamma = 0.1)
# or:
m <- svm(~ a + b, gamma = 0.1)

# test:
newdata <- data.frame(a = c(0, 4), b = c(0, 4))
predict (m, newdata)

# visualize:
plot(X, col = 1:1000 %in% m$index + 1, xlim = c(-5,5), ylim=c(-5,5))
points(newdata, pch = "+", col = 2, cex = 5)

# weights: (example not particularly sensible)
i2 <- iris
levels(i2$Species)[3] <- "versicolor"
summary(i2$Species)
wts <- 100 / table(i2$Species)
wts
m <- svm(Species ~ ., data = i2, class.weights = wts)

```

tune

Parameter tuning of functions using grid search

Description

This generic function tunes hyperparameters of statistical methods using a grid search over supplied parameter ranges.

Usage

```
tune(method, train.x, train.y = NULL, data = list(), validation.x =
      NULL, validation.y = NULL, ranges = NULL, predict.func = predict,
      tunecontrol = tune.control(), ...)
best.tune(...)
```

Arguments

<code>method</code>	function to be tuned.
<code>train.x</code>	either a formula or a matrix of predictors.
<code>train.y</code>	the response variable if <code>train.x</code> is a predictor matrix. Ignored if <code>train.x</code> is a formula.
<code>data</code>	data, if a formula interface is used. Ignored, if predictor matrix and response are supplied directly.
<code>validation.x</code>	an optional validation set. Depending on whether a formula interface is used or not, the response can be included in <code>validation.x</code> or separately specified using <code>validation.y</code> .
<code>validation.y</code>	if no formula interface is used, the response of the (optional) validation set.
<code>ranges</code>	a named list of parameter vectors spanning the sampling space. The vectors will usually be created by <code>seq</code> .
<code>predict.func</code>	optional predict function, if the standard <code>predict</code> behaviour is inadequate.
<code>tunecontrol</code>	object of class "tune.control", as created by the function <code>tune.control()</code> . If omitted, <code>tune.control()</code> gives the defaults.
<code>...</code>	Further parameters passed to the training functions.

Details

As performance measure, the classification error is used for classification, and the mean squared error for regression. It is possible to specify only one parameter combination (i.e., vectors of length 1) to obtain an error estimation of the specified type (bootstrap, cross-classification, etc.) on the given data set. For convenience, there are several `tune.foo()` wrappers defined, e.g., for `nnet()`, `randomForest()`, `rpart()`, `svm()`, and `knn()`.

Value

For `tune`, an object of class `tune`, including the components:

```
best.parameters
  a 1 x k data frame, k number of parameters.
```

`best.performance`
 best achieved performance.

`performances` if requested, a data frame of all parameter combinations along with the corresponding performance results.

if requested, the model trained on the complete training data using the best parameter combination.

`best.tune` returns the best model detected by `tune`.

Author(s)

David Meyer
 (David.Meyer@R-project.org)

See Also

[tune.control](#), [plot.tune](#), [tune.svm](#), [tune.wrapper](#)

Examples

```
data(iris)
## tune `svm` for classification with RBF-kernel (default in svm),
## using one split for training/validation set

obj <- tune(svm, Species~., data = iris,
           ranges = list(gamma = 2^(-1:1), cost = 2^(2:4)),
           tunecontrol = tune.control(sampling = "fix")
           )

## alternatively:
## obj <- tune.svm(Species~., data = iris, gamma = 2^(-1:1), cost = 2^(2:4))

summary(obj)
plot(obj)

## tune `knn` using a convenience function; this time with the
## conventional interface and bootstrap sampling:
x <- iris[,-5]
y <- iris[,5]
obj2 <- tune.knn(x, y, k = 1:5, tunecontrol = tune.control(sampling = "boot"))
summary(obj2)
plot(obj2)

## tune `rpart` for regression, using 10-fold cross validation (default)
data(mtcars)
obj3 <- tune.rpart(mpg~., data = mtcars, minsplit = c(5,10,15))
summary(obj3)
plot(obj3)

## simple error estimation for lm using 10-fold cross validation
tune(lm, mpg~., data = mtcars)
```

tune.control *control parameters for the tune function*

Description

Creates an object of class `tune.control` to be used with the `tune` function, containing various control parameters.

Usage

```
tune.control(random = FALSE, nrepeat = 1, repeat.aggregate = min,
             sampling = c("cross", "fix", "bootstrap"), sampling.aggregate = mean,
             cross = 10, fix = 2/3, nboot = 10, boot.size = 9/10, best.model = TRUE,
             performances = TRUE)
```

Arguments

<code>random</code>	if an integer value is specified, random parameter vectors are drawn from the parameter space.
<code>nrepeat</code>	specifies how often training shall be repeated.
<code>repeat.aggregate</code>	function for aggregating the repeated training results.
<code>sampling</code>	sampling scheme. If <code>sampling = "cross"</code> , a cross-times cross validation is performed. If <code>sampling = "boot"</code> , <code>nboot</code> training sets of size <code>boot.size</code> (part) are sampled from the supplied data. If <code>sampling = "fix"</code> , a single split into training/validation set is used, the training set containing a <code>fix</code> part of the supplied data. Note that a separate validation set can be supplied via <code>validation.x</code> and <code>validation.y</code> . It is only used for <code>sampling = "boot"</code> and <code>sampling = "fix"</code> ; in the latter case, <code>fix</code> is set to 1.
<code>sampling.aggregate</code>	function for aggregating the training results on the generated training samples.
<code>cross</code>	number of partitions for cross-validation.
<code>fix</code>	part of the data used for training in fixed sampling.
<code>nboot</code>	number of bootstrap replications.
<code>boot.size</code>	size of the bootstrap samples.
<code>best.model</code>	if <code>TRUE</code> , the best model is trained and returned (the best parameter set is used for training on the complete training set).
<code>performances</code>	if <code>TRUE</code> , the performance results for all parameter combinations are returned.

Value

An object of class `"tune.control"` containing all the above parameters (either the defaults or the user specified values).

Author(s)

David Meyer
(David.Meyer@R-project.org)

See Also[tune](#)

tune.wrapper	<i>Convenience tuning wrapper functions</i>
--------------	---

Description

Convenience tuning wrapper functions, using `tune`.

Usage

```
tune.svm(x, y = NULL, data = NULL, degree = NULL, gamma = NULL, coef0 = NULL, cost = NULL,
         nu = NULL, ...)
best.svm(x, tunecontrol = tune.control(), ...)

tune.nnet(x, y = NULL, data = NULL, size = NULL, decay = NULL, trace = FALSE,
          tunecontrol = tune.control(nrepeat = 5), ...)
best.nnet(x, tunecontrol = tune.control(nrepeat = 5), ...)

tune.rpart(formula, data, na.action = na.omit, minsplit = NULL,
            minbucket = NULL, cp = NULL, maxcompete = NULL, maxsurrogate = NULL,
            usesurrogate = NULL, xval = NULL, surrogatestyle = NULL, maxdepth = NULL,
            predict.func = NULL, ...)
best.rpart(formula, tunecontrol = tune.control(), ...)
rpart.wrapper(formula, minsplit=20, minbucket=round(minsplit/3), cp=0.01,
              maxcompete=4, maxsurrogate=5, usesurrogate=2, xval=10,
              surrogatestyle=0, maxdepth=30, ...)

tune.randomForest(x, y = NULL, data = NULL, nodesize = NULL, mtry = NULL, ntree = NULL,
                  tunecontrol = tune.control(), ...)
best.randomForest(x, tunecontrol = tune.control(), ...)

tune.knn(x, y, k = NULL, l = NULL, ...)
knn.wrapper(x, y, k = 1, l = 0, ...)
```

Arguments

`formula, x, y, data` formula and data arguments of function to be tuned.

`predict.func` predicting function.

`na.action` function handling missingness.

`minsplit, minbucket, cp, maxcompete, maxsurrogate, usesurrogate, xval, surrogatestyle` rpart parameters.

`degree, gamma, coef0, cost, nu` svm parameters.

`k, l` knn parameters.

`mtry, nodesize, ntree` randomForest parameters.

size, decay, trace
 parameters passed to nnet.
 tunecontrol object of class "tune.control" containing tuning parameters.
 ... Further parameters passed to tune.

Details

For examples, see the help page of `tune()`.

Value

`tune.foo()` returns a tuning object including the best parameter set obtained by optimizing over the specified parameter vectors. `best.foo()` directly returns the best model, i.e. the fit of a new model using the optimal parameters found by `tune.foo`.

Author(s)

David Meyer
 (David.Meyer@R-project.org)

See Also

[tune](#)

<code>write.svm</code>	<i>Write SVM object to file</i>
------------------------	---------------------------------

Description

This function exports an SVM object (trained by `svm`) to two specified files. One is in the format that the function `'svm_load_model'` of `libsvm` can read. The other is for scaling data, containing a data with centers and scales for all variables.

Usage

```
write.svm(object, svm.file, scale.file)
```

Arguments

`object` Object of class "svm", created by `svm`.
`svm.file` filename to export the svm object to.
`scale.file` filename to export the scaling data of the svm object to.

Details

This function is useful when SVM models trained in R shall be used in other environments. The SVM model is saved in the standard format of `libsvm`. The scaling data is written to a separate file because scaling data is not included in the standard format of `libsvm`. The format of the scaling data file is a n times 2 matrix: the n -th row corresponds to the n -th dimension of the data, the columns being formed of the corresponding mean and scale.

Author(s)

Tomomi TAKASHINA (based on 'predict.svm' by David Meyer) <t.takashina@computer.org>

See Also

[svm](#)

Examples

```
data(iris)
attach(iris)

## classification mode
# default with factor response:
model <- svm (Species~., data=iris)

# export SVM object to file
write.svm(model, svm.file = "iris-classifier.svm", scale.file = "iris-classifier.scale")

# read scale file
# the n-th row is corresponding to n-th dimension. The 1st column contains the
# center value, the 2nd column is the scale value.

read.table("iris-classifier.scale")
```

Index

- *Topic **IO**
 - `read.matrix.csr`, 39
 - *Topic **arith**
 - `interpolate`, 23
 - *Topic **array**
 - `element`, 16
 - *Topic **category**
 - `classAgreement`, 10
 - `matchClasses`, 26
 - `naiveBayes`, 30
 - `predict.naiveBayes`, 35
 - *Topic **classif**
 - `naiveBayes`, 30
 - `plot.svm`, 32
 - `predict.naiveBayes`, 35
 - `predict.svm`, 36
 - `svm`, 44
 - `write.svm`, 53
 - *Topic **cluster**
 - `bclust`, 5
 - `cmeans`, 11
 - `cshell`, 14
 - `fclustIndex`, 17
 - `lca`, 25
 - *Topic **datagen**
 - `permutations`, 31
 - *Topic **datasets**
 - `Glass`, 2
 - `HouseVotes84`, 3
 - `Ozone`, 4
 - *Topic **distribution**
 - `Discrete`, 1
 - `rbridge`, 38
 - `rwiener`, 40
 - *Topic **hplot**
 - `boxplot.bclust`, 9
 - *Topic **manip**
 - `impute`, 23
 - `matchControls`, 27
 - *Topic **math**
 - `sigmoid`, 42
 - *Topic **misc**
 - `e1071-deprecated`, 16
 - *Topic **models**
 - `plot.tune`, 34
 - `tune`, 49
 - `tune.control`, 51
 - `tune.wrapper`, 52
 - *Topic **multivariate**
 - `bclust`, 5
 - `bootstrap.lca`, 7
 - `countpattern`, 13
 - `hamming.distance`, 19
 - `ica`, 22
 - `interpolate`, 23
 - `lca`, 25
 - *Topic **neural**
 - `plot.svm`, 32
 - `predict.svm`, 36
 - `svm`, 44
 - `write.svm`, 53
 - *Topic **nonlinear**
 - `plot.svm`, 32
 - `predict.svm`, 36
 - `svm`, 44
 - `write.svm`, 53
 - *Topic **optimize**
 - `allShortestPaths`, 41
 - *Topic **ts**
 - `hamming.window`, 20
 - `hanning.window`, 21
 - `plot.stft`, 32
 - `rectangle.window`, 40
 - `stft`, 43
 - *Topic **univar**
 - `kurtosis`, 24
 - `moment`, 29
 - `skewness`, 43
 - *Topic **utilities**
 - `bincombinations`, 7
- `allShortestPaths`, 41
- `approx`, 24
- `bclust`, 5, 9
- `best.nnet (tune.wrapper)`, 52

- best.randomForest (tune.wrapper), 52
- best.rpart (tune.wrapper), 52
- best.svm (tune.wrapper), 52
- best.tune (tune), 49
- bincombinations, 7
- bootstrap.lca, 7, 25
- boxplot, 9
- boxplot.bclust, 6, 9

- centers.bclust (bclust), 5
- classAgreement, 10, 27
- clusters.bclust (bclust), 5
- cmdscale, 5
- cmeans, 11, 19
- compareMatchedClasses (matchClasses), 26
- countpattern, 13, 25
- cshell, 14

- d2sigmoid (sigmoid), 42
- daisy, 28
- ddiscrete (Discrete), 1
- Deprecated, 16
- Discrete, 1
- dist, 5, 41
- dsigmoid (sigmoid), 42

- e1071-deprecated, 16
- element, 16
- Extract, 17
- extractPath (allShortestPaths), 41

- fclustIndex, 17

- Glass, 2
- grep, 28

- hamming.distance, 19
- hamming.window, 20
- hanning.window, 21
- hclust, 5, 6
- hclust.bclust (bclust), 5
- HouseVotes84, 3

- ica, 22
- impute, 23
- interpolate, 23

- kmeans, 5, 6
- knn.wrapper (tune.wrapper), 52
- kurtosis, 24

- lca, 7, 8, 25
- matchClasses, 11, 26
- matchControls, 27
- matrix.csr, 39, 45, 47
- mean, 29
- moment, 29

- naiveBayes, 30, 35

- Ozone, 4

- pdiscrete (Discrete), 1
- permutations, 31
- plot.bclust (bclust), 5
- plot.ica (ica), 22
- plot.stft, 32
- plot.svm, 32, 47
- plot.tune, 34, 50
- predict.lca (lca), 25
- predict.naiveBayes, 31, 35
- predict.svm, 36, 47
- print.bootstrap.lca (bootstrap.lca), 7
- print.fclust (cmeans), 11
- print.ica (ica), 22
- print.lca (lca), 25
- print.naiveBayes (naiveBayes), 30
- print.summary.lca (lca), 25
- print.summary.svm (svm), 44
- print.summary.tune (tune), 49
- print.svm (svm), 44
- print.tune (tune), 49

- qdiscrete (Discrete), 1

- rbridge, 38
- rdiscrete (Discrete), 1
- read.matrix.csr, 39
- read.octave, 16
- read.octave (e1071-deprecated), 16
- rectangle.window, 40
- rpart.wrapper (tune.wrapper), 52
- rwiener, 40

- sample, 2
- sigmoid, 42
- skewness, 43
- spline, 24
- stft, 43
- summary.lca (lca), 25
- summary.svm (svm), 44
- summary.tune (tune), 49
- svm, 33, 37, 44, 54

- tune, 34, 49, 52, 53

`tune.control`, 50, 51
`tune.knn` (`tune.wrapper`), 52
`tune.nnet` (`tune.wrapper`), 52
`tune.randomForest` (`tune.wrapper`),
52
`tune.rpart` (`tune.wrapper`), 52
`tune.svm`, 50
`tune.svm` (`tune.wrapper`), 52
`tune.wrapper`, 50, 52

`var`, 29

`write.matrix.csr`
(`read.matrix.csr`), 39
`write.svm`, 53