

Package ‘analogue’

April 5, 2012

Type Package

Title Analogue and weighted averaging methods for palaeoecology

Version 0.8-2

Date \$Date: 2012-04-03 21:01:40 +0100 (Tue, 03 Apr 2012) \$

Depends R (>= 2.15.0), stats, graphics, vegan (>= 1.17-12), lattice,grid, MASS, printrcurve

Author Gavin L. Simpson, Jari Oksanen

Maintainer Gavin L. Simpson <gavin.simpson@ucl.ac.uk>

Description Fits Modern Analogue Technique and Weighted Averaging transfer function models for prediction of environmental data from species data, and related methods used in palaeoecology.

License GPL-2

ByteCompile true

URL <http://analogue.r-forge.r-project.org>

Repository CRAN

Date/Publication 2012-04-05 05:37:04

R topics documented:

analogue-package	3
abernethy	4
analog	5
bayesF	7
bootstrap	10
bootstrap.wa	13
bootstrapObject	14
chooseTaxa	15
cma	16

crossval	19
densityplot.residLen	21
deshrink	22
dissimilarities	23
distance	25
fuse	28
getK	30
gradientDist	32
hist.residLen	33
histogram.residLen	34
ImbrieKipp	36
join	38
logitreg	40
mat	43
mcarlo	48
minDC	50
optima	52
panel.Loess	54
panel.Stratplot	55
pcr	57
performance	60
plot.dissimilarities	61
plot.logitreg	62
plot.mat	63
plot.mcarlo	66
plot.minDC	67
plot.pcurve	69
plot.residLen	71
plot.roc	72
plot.wa	75
Pollen	76
pcurve	78
predict.mat	81
predict.wa	84
reconPlot	86
residLen	88
rlgh	90
RMSEP	91
roc	93
screepLOT	95
stdError	97
Stratplot	99
summary.analog	102
summary.bootstrap.mat	104
summary.cma	106
summary.mat	107
summary.predict.mat	109
swapdiat	110

swappH	111
timetrack	111
tran	114
varExpl	117
wa	118
weightedCor	121

Index	124
--------------	------------

analogue-package	<i>Analogue methods for palaeoecology</i>
------------------	---

Description

Fits Modern Analogue Technique transfer function models for prediction of environmental data from species data. Also performs analogue matching, a related technique used in palaeo ecological restoration.

Details

Package: analogue
 Type: Package
 Version: 0.4-3
 Date: 2007-09-15
 Depends: R (>= 2.5.0), stats, graphics, vegan
 License: GPL Version 2
 Built: R 2.5.1; ; 2007-08-03 13:44:47; unix

Index:

RMSEP	Root mean square error of prediction
analog	Analogue matching
analogue-package	Analogue methods for palaeoecology
bayesF	Bayes factors
bootstrap	Bootstrap estimation and errors
bootstrapObject	Bootstrap object description
cma	Close modern analogues
dissimilarities	Extract dissimilarity coefficients from models
distance	Flexibly calculate dissimilarity or distance measures
getK	Extract and set the number of analogues
join	Merge species data sets on common columns (species)
mat	Modern Analogue Technique transfer function models
mcarlo	Monte Carlo simulation of dissimilarities

minDC	Extract minimum dissimilarities
plot.bayesF	Bayes factor plots
plot.dissimilarities	Plots the distribution of extracted dissimilarities
plot.mat	Plot diagnostics for a mat object
plot.mcarlo	Plott Monte Carlo simulated dissimilarity distributions
plot.minDC	Plot of minimum dissimilarity per sample
plot.roc	Plot ROC curves and associated diagnostics
predict.mat	Predict method for Modern Analogue Technique models
reconPlot	Stratigraphic plots of palaeoenvironmental reconstructions
rlgh	Round Loch of Glenhead Diatoms
roc	ROC curve analysis
screeplot.mat	Screeplots of model results
summary.analog	Summarise analogue matching results
summary.bootstrap.mat	Summarise bootstrap resampling for MAT models
summary.cma	Summarise the extraction of close modern analogues
summary.mat	Summarise Modern Analogue Technique models
summary.predict.mat	Summarise MAT model predictions
swapdiat	SWAP sub-fossil diatom training set
swappH	SWAP pH training set

Author(s)

Gavin L. Simpson.

Maintainer: Gavin L. Simpson <gavin.simpson@ucl.ac.uk>

abernethy

Abernethy Forest Pollen Sequence

Description

The classic pollen data set from Abernethy Forest in the Scottish highlands, UK. The data originate from the work of Hilary Birks and Rolf Mathewes (1978) and have been analysed in several texts on quantitative numerical palaeoecology.

The data set consists of 36 pollen taxa from 49 levels, with two additional variables; Age, the age of each sample, and Depth the depth (in cm) below the surface of the peat sequence from which the core was taken.

Usage

data(abernethy)

Format

abernethy is a data frame with 49 samples on 36 species plus sample Age and Depth (in cm).

Source

These data were provided in electronic format by Prof. H. John B. Birks. The original source is Birks and Mathewes (1978).

References

Birks, H.H. and Mathewes, R.W. (1978) Studies in the vegetational history of Scotland. *New Phytologist* **80**, 455-484.

Examples

```
data(abernethy)
head(abernethy)

(plt <- Stratiplot(Age ~ . - Depth,
                  data = chooseTaxa(abernethy, n.occ = 5,
                                    max.abun = 10),
                  type = "poly"))
```

analog

Analogue matching

Description

Analogue matching is a more general implementation of the modern analogue methodology than MAT, where we are only interested in identifying sufficiently similar samples from a modern training as being suitable modern analogues for one or more fossil samples.

Usage

```
analog(x, ...)

## Default S3 method:
analog(x, y,
       method = c("euclidean", "SQeuclidean", "chord", "SQchord",
                  "bray", "chi.square", "SQchi.square",
                  "information", "chi.distance", "manhattan",
                  "kendall", "gower", "alt.gower", "mixed"),
       keep.train = TRUE, ...)
```

Arguments

<code>x, y</code>	data frames with same columns. <code>x</code> is training data and <code>y</code> , the test data.
<code>method</code>	character string naming the dissimilarity methods to be used. See Details below.
<code>keep.train</code>	logical; should the dissimilarity matrix for the training set be stored?
<code>...</code>	arguments passed to or from other methods.

Details

`analog` implements analogue matching *sensu* Flower et al (1997) and Simpson et al (2005), where the aim is to identify suitable close analogues of fossil samples from a modern training set. These results are generally used within ecological restoration, but the identification of close modern analogues for fossil samples is also used as a technique for assessing transfer function reconstructions.

`analog` is a simple and very general function that generates a pairwise dissimilarity matrix for the modern training set, and a second matrix containing the pairwise dissimilarities between each fossil sample and each sample in the training set. These results can then be assessed using other functions and to extract the close modern analogues using function `cma`. See the See Also section below.

Analysis of the pairwise dissimilarity matrix for the modern training set can be used to decide on a suitable dissimilarity threshold for defining close modern analogues. By default this matrix is returned as part of the output from the `analog` function.

Value

A list of class "analog" with the following components:

<code>analog</code>	matrix of pairwise dissimilarities between each fossil sample (<code>y</code>) and each sample in the modern training set (<code>x</code>).
<code>train</code>	if argument <code>keep.train</code> is TRUE then a pairwise dissimilarity matrix for the modern training set.
<code>call</code>	the matched function call.
<code>method</code>	character; the dissimilarity coefficient used.

Author(s)

Gavin L. Simpson

References

Flower, R.J., Juggins, S. and Battarbee, R.W. (1997) Matching diatom assemblages in lake sediment cores and modern surface sediment samples: the implications for lake conservation and restoration with special reference to acidified systems. *Hydrobiologia* **344**; 27–40.

Simpson, G.L., Shilland, E.M., Winterbottom, J. M. and Keay, J. (2005) Defining reference conditions for acidified waters using a modern analogue approach. *Environmental Pollution* **137**; 119–133.

See Also

[distance](#) for the function that calculates the dissimilarity matrices. [cma](#) for extraction of close modern analogues. [dissimilarities](#) and [plot.dissimilarities](#) for analysis of distribution of pairwise dissimilarity matrix for modern training set.

Examples

```
## Imbrie and Kipp example
## load the example data
data(ImbrieKipp)
data(SumSST)
data(V12.122)

## merge training and test set on columns
dat <- join(ImbrieKipp, V12.122, verbose = TRUE)

## extract the merged data sets and convert to proportions
ImbrieKipp <- dat[[1]] / 100
V12.122 <- dat[[2]] / 100

## Imbrie and Kipp foraminifera sea-surface temperature

## analog matching between SWAP and RLGH core
ik.analog <- analog(ImbrieKipp, V12.122, method = "chord")
ik.analog
summary(ik.analog)
```

 bayesF

Bayes factors

Description

Calculates Bayes factors or likelihood ratios of analogue and no-analogue results.

Usage

```
bayesF(x, prior = rep(0.5, 2))

## S3 method for class 'bayesF'
plot(x, group = "all", xlab = NULL, ylab = "Pr (A+ | d)",
     col = "red", abline.col = "lightgrey", abline.lty = "dashed", ...)
```

Arguments

x for bayesF an object of class roc. For the plot method, an object of class bayesF, usually the result of a call to bayesF.

prior	numeric; the prior probabilities of analogue and no-analogue, provided as a vector of length 2 whose elements sum to 1. If not provided, the function will use the relative occurrences of analogue and no analogue situations used to evaluate the ROC curve.
group	character vector of length 1 giving the name of the group to plot, or "all" to plot all groups in x.
xlab,ylab	the x- and y-axis labels for the plot.
col	colour of the line used to draw the posterior probability.
abline.col	colour of the vertical line drawn to indicate the optimal dissimilarity determined from the ROC curve.
abline.lty	Line type for indicator of optimal ROC dissimilarity threshold. See par for the allowed line types.
...	other plot arguments passed to plotting functions. Currently ignored.

Details

LR(+), is the likelihood ratio of a positive test result, that the value of d assigns the sample to the group it belongs to. LR(-) is the likelihood ratio of a negative test result, that the value of d assigns the sample to the wrong group.

LR(+) is defined as $LR(+) = TPF/FPF$ (or sensitivity / (1 - specificity)), and LR(-) is defined as $LR(-) = FPF/TNF$ (or (1 - sensitivity) / specificity), in Henderson (1993).

The posterior probability of analogue given a dissimilarity is the LR(+) likelihood ratio values multiplied by the prior odds of analogue, for given values of the dissimilarity, and is then converted to a probability.

The plotting function currently only draws the posterior probability of analogue based on the Bayes factor or likelihood ratio of a positive event (analogue).

Value

For `plot.bayesF` a plot on the currently active device.

For `bayesF`, a list containing the results of computing Bayes factors for each group in x. Each component of this list is itself a list with the following components:

bayesF, posterior.odds, posterior.probs, prior.prob	Bayes factors, posterior odds and probabilities and prior probabilities of true analogue and true non-analogue events. Each components is a list with two components; pos (for true analogue events) and neg (for true non-analogue events). The components of <code>prior.prob</code> are vectors of length 1, whilst components of the other lists are numeric vectors.
roc.points	numeric; the points at which the ROC curve was evaluated.
optimal	numeric; the optimal dissimilarity as assessed by the ROC curve.
max.roc	numeric; the position along the ROC curve at which the slope of the ROC curve is maximal. This is the index of this point on the curve, and can be used to extract the element of <code>bayesF</code> , <code>posterior.odds</code> and <code>posterior.probs</code> for the optimal dissimilarity.

Author(s)

Gavin L. Simpson

References

- Brown, C.D., and Davis, H.T. (2006) Receiver operating characteristics curves and related decision measures: A tutorial. *Chemometrics and Intelligent Laboratory Systems* **80**, 24–38.
- Gavin, D.G., Oswald, W.W., Wahl, E.R. and Williams, J.W. (2003) A statistical approach to evaluating distance metrics and analog assignments for pollen records. *Quaternary Research* **60**, 356–367.
- Henderson, A.R. (1993) Assessing test accuracy and its clinical consequences: a primer for receiver operating characteristic curve analysis. *Annals of Clinical Biochemistry* **30**, 834–846.

See Also

[roc](#) and [plot.bayesF](#).

Examples

```
## load the example data
data(swapdiat, swappH, rlgh)

## merge training and test set on columns
dat <- join(swapdiat, rlgh, verbose = TRUE)

## extract the merged data sets and convert to proportions
swapdiat <- dat[[1]] / 100
rlgh <- dat[[2]] / 100

## fit an analogue matching (AM) model using the squared chord distance
## measure - need to keep the training set dissimilarities
swap.ana <- analog(swapdiat, rlgh, method = "SQchord",
                  keep.train = TRUE)

## fit the ROC curve to the SWAP diatom data using the AM results
## Generate a grouping for the SWAP lakes
clust <- hclust(as.dist(swap.ana$train), method = "ward")
grps <- cutree(clust, 12)

## fit the ROC curve
swap.roc <- roc(swap.ana, groups = grps)
swap.roc

## calculate the Bayes factors of analogue and no-analogue
## (uses observed probabilities of analogue/no-analogue
swap.bayes <- bayesF(swap.roc)
swap.bayes

## plot the probability of analogue
plot(swap.bayes)

## Not run:
```

```
## calculate the Bayes factors of analogue and no-analogue
## with prior probabilities c(0.5, 0.05)
swap.bayes2 <- bayesF(swap.roc, prior = c(0.5, 0.05))
swap.bayes

## plot the probability of analogue
plot(swap.bayes2)

## End(Not run)
```

bootstrap

Bootstrap estimation and errors

Description

Function to calculate bootstrap statistics for transfer function models such as bootstrap estimates, model RMSEP, sample specific errors for predictions and summary statistics such as bias and R^2 between observed and estimated environment.

[residuals](#) method for objects of class "bootstrap.mat".

Usage

```
bootstrap(object, ...)

## Default S3 method:
bootstrap(object, ...)

## S3 method for class 'mat'
bootstrap(object, newdata, newenv, k,
          weighted = FALSE, n.boot = 1000, ...)

## S3 method for class 'bootstrap.mat'
fitted(object, k, ...)

## S3 method for class 'bootstrap.mat'
residuals(object, which = c("model", "bootstrap"), ...)
```

Arguments

object	an R object of class "mat" for which bootstrap statistics are to be generated, or an object of class "bootstrap.mat" from which fitted values or residuals are extracted.
newdata	a data frame containing samples for which bootstrap predictions and sample specific errors are to be generated. May be missing — See Details. "newdata" must have the same number of columns as the training set data.

newenv	a vector containing environmental data for samples in "newdata". Used to calculate full suite of errors for new data such as a test set with known environmental values. May be missing — See Details. "newenv" must have the same number of rows as "newdata".
k	numeric; how many modern analogues to use to generate the bootstrap statistics (and, if requested, the predictions), fitted values or residuals.
weighted	logical; should the weighted mean of the environment for the "k" modern analogues be used instead of the mean?
n.boot	Number of bootstrap samples to take.
which	character; which set of residuals to return, the model residuals or the residuals of the bootstrap-derived estimates?
...	arguments passed to other methods.

Details

bootstrap is a fairly flexible function, and can be called with or without arguments newdata and newenv.

If called with only object specified, then bootstrap estimates for the training set data are returned. In this case, the returned object will not include component predictions.

If called with both object and newdata, then in addition to the above, bootstrap estimates for the new samples are also calculated and returned. In this case, component predictions will contain the apparent and bootstrap derived predictions and sample-specific errors for the new samples.

If called with object, newdata and newenv, then the full bootstrap object is returned (as described in the Value section below). With environmental data now available for the new samples, residuals, RMSE(P) and R^2 and bias statistics can be calculated.

The individual components of predictions are the same as those described in the components relating to the training set data. For example, returned.object\$predictions\$bootstrap contains the components as returned.object\$bootstrap.

It is not usual for environmental data to be available for the new samples for which predictions are required. In normal palaeolimnological studies, it is more likely that newenv will not be available as we are dealing with sediment core samples from the past for which environmental data are not available. However, if sufficient training set samples are available to justify producing a training and a test set, then newenv will be available, and bootstrap can accommodate this extra information and calculate apparent and bootstrap estimates for the test set, allowing an independent assessment of the RMSEP of the model to be performed.

Typical usage of residuals is

```
resid(object, which = c("model", "bootstrap"), ...)
```

Value

For bootstrap.mat an object of class "bootstrap.mat" is returned. This is a complex object with many components and is described in [bootstrapObject](#).

For residuals, a list containing the requested residuals and metadata, with the following components:

model	Leave one out residuals for the MAT-estimated model.
bootstrap	residuals for the bootstrapped MAT model.
k	numeric; indicating the size of model used in estimates and predictions.
n.boot	numeric; the number of bootstrap samples taken.
auto	logical; whether "k" was chosen automatically or user-selected.
weighted	logical; whether the weighted mean was used instead of the mean of the environment for <i>k</i> -closest analogues.

Author(s)

Gavin L. Simpson

References

Birks, H.J.B., Line, J.M., Juggins, S., Stevenson, A.C. and ter Braak, C.J.F. (1990). Diatoms and pH reconstruction. *Philosophical Transactions of the Royal Society of London; Series B*, **327**; 263–278.

See Also

mat, plot.mat, summary.bootstrap.mat, [residuals](#)

Examples

```
## Imbrie and Kipp example
## load the example data
data(ImbrieKipp)
data(SumSST)
data(V12.122)

## merge training and test set on columns
dat <- join(ImbrieKipp, V12.122, verbose = TRUE)

## extract the merged data sets and convert to proportions
ImbrieKipp <- dat[[1]] / 100
V12.122 <- dat[[2]] / 100

## Imbrie and Kipp foraminifera sea-surface temperature
## fit the MAT model using the squared chord distance measure
ik.mat <- mat(ImbrieKipp, SumSST, method = "SQchord")

## bootstrap training set
ik.boot <- bootstrap(ik.mat, n.boot = 100)
ik.boot
summary(ik.boot)

## Bootstrap fitted values for training set
fitted(ik.boot)

## residuals
resid(ik.boot) # uses abbreviated form
```

`bootstrap.wa`*Bootstrap estimation and errors for WA models*

Description

Function to calculate bootstrap statistics for transfer function models such as bootstrap estimates, model RMSEP, sample specific errors for predictions and summary statistics such as bias and R^2 between observed and estimated environment.

Usage

```
## S3 method for class 'wa'  
bootstrap(object, n.boot = 1000, verbose = TRUE, ...)
```

Arguments

<code>object</code>	an R object of class "wa" for which bootstrap statistics are to be generated.
<code>n.boot</code>	numeric; the number of bootstrap samples to draw.
<code>verbose</code>	logical; should bootstrap progress be printed to the console?
<code>...</code>	arguments passed to other methods.

Details

See [bootstrap.mat](#) for further details. This method is not as feature packed as [bootstrap.mat](#) but can be used to evaluate the model performance of WA transfer function models.

Value

An object with the same components as [predict.wa](#).

Author(s)

Gavin L. Simpson

References

Birks, H.J.B., Line, J.M., Juggins, S., Stevenson, A.C. and ter Braak, C.J.F. (1990). Diatoms and pH reconstruction. *Philosophical Transactions of the Royal Society of London; Series B*, **327**; 263–278.

See Also

`wa`, `plot.wa`.

Examples

```
## Imbrie and Kipp
data(ImbrieKipp)
data(SumSST)
ik.wa <- wa(SumSST ~ ., data = ImbrieKipp, tol.dw = TRUE,
           min.tol = 2, small.tol = "min")
ik.wa

## compare actual tolerances to working values
with(ik.wa, rbind(tolerances, model.tol))

## bootstrap the WA model
ik.boot <- bootstrap(ik.wa, n.boot = 100)

## performance statistics
performance(ik.boot)
```

bootstrapObject	<i>Bootstrap object description</i>
-----------------	-------------------------------------

Description

Objects of class `bootstrap.mat` are a complex containing many sub-components. This object is described here in more detail.

Details

A large object is returned with some or all of the following depending on whether newdata and newenv are supplied or not.

observed: vector of observed environmental values.

model: a list containing the apparent or non-bootstrapped estimates for the training set. With the following components:

estimated: estimated values for the response

residuals: model residuals

r.squared: Apparent R^2 between observed and estimated values of response

avg.bias: Average bias of the model residuals

max.bias: Maximum bias of the model residuals

rmse: Apparent error (RMSE) for the model.

k: numeric; indicating the size of model used in estimates and predictions

bootstrap: a list containing the bootstrap estimates for the training set. With the following components:

estimated: Bootstrap estimates for the response

residuals: Bootstrap residuals for the response

r.squared: Bootstrap derived R^2 between observed and estimated values of the response

`avg.bias`: Average bias of the bootstrap derived model residuals
`max.bias`: Maximum bias of the bootstrap derived model residuals
`rmsep`: Bootstrap derived RMSEP for the model
`s1`: Bootstrap derived S1 error component for the model
`s2`: Bootstrap derived S2 error component for the model
`k`: numeric; indicating the size of model used in estimates and predictions
`sample.errors`: a list containing the bootstrap-derived sample specific errors for the training set.
 With the following components:
`rmsep`: Bootstrap derived RMSEP for the training set samples
`s1`: Bootstrap derived S1 error component for training set samples
`s2`: Bootstrap derived S2 error component for training set samples
`weighted`: logical; whether the weighted mean was used instead of the mean of the environment
 for k -closest analogues
`auto`: logical; whether "k" was chosen automatically or user-selected
`n.boot`: numeric; the number of bootstrap samples taken
`call`: the matched call
`type`: model type
`predictions`: a list containing the apparent and bootstrap-derived estimates for the new data, with
 the following components:
`observed`: the observed values for the new samples — only if `newenv` is provided
`model`: a list containing the apparent or non-bootstrapped estimates for the new samples. A
 list with the same components as `model`, above
`bootstrap`: a list containing the bootstrap estimates for the new samples, with some or all
 of the same components as `bootstrap`, above
`sample.errors`: a list containing the bootstrap-derived sample specific errors for the new
 samples, with some or all of the same components as `sample.errors`, above

Author(s)

Gavin L. Simpson

See Also

`mat`, `plot.mat`, `summary.bootstrap.mat`, [residuals](#)

chooseTaxa

*Select taxa (variables) on basis of maximum abundance attained and
number of occurrences*

Description

Select taxa (variables) from an object on the basis of one or both of maximum abundance and number of occurrences greater than user-specified values. This is a simple utility function to encapsulate this common task in filtering palaeoecological data sets.

Usage

```
chooseTaxa(object, ...)

## Default S3 method:
chooseTaxa(object, n.occ = 1, max.abun = 0,
           type = c("AND", "OR"), ...)
```

Arguments

<code>object</code>	an R object for which a suitable method exists. The default method assumes a matrix-like object such as a data frame or a numeric matrix.
<code>n.occ</code>	numeric; number of occurrences representing the lower limit for selection. A taxon is included in the returned subset if it is present a total of <code>n.occ</code> times or more. See argument <code>type</code> for a modifier which might exclude the taxon even if it would be included on the basis of <code>n.occ</code> .
<code>max.abun</code>	numeric; maximum abundance representing the lower limit for selection. A taxon is included in the returned subset if it attains abundance equal to or greater than <code>max.abun</code> in one or more sample. See argument <code>type</code> for a modifier which might exclude the taxon even if it would be included on the basis of <code>max.abun</code> .
<code>type</code>	character; one of "AND" or "OR", controlling how the criteria <code>n.occ</code> and <code>max.abun</code> are combined to generate a subset of the variables in <code>object</code> .
<code>...</code>	arguments passed on to subsequent methods.

Value

Returns the supplied data frame or matrix with a subset of columns (taxa) that meet the criteria chosen.

Author(s)

Gavin L. Simpson

Examples

```
data(ImbrieKipp)
IK2 <- chooseTaxa(ImbrieKipp, n.occ = 5)
dim(ImbrieKipp)
dim(IK2)
```

cma

Close modern analogues

Description

Extracts and formats close modern analogue samples from a modern reference set that are closer than a defined cut off threshold.

Usage

```

cma(object, ...)

## Default S3 method:
cma(object, ...)

## S3 method for class 'analog'
cma(object, cutoff, prob = c(0.01, 0.025, 0.05), ...)

## S3 method for class 'mat'
cma(object, k, cutoff, prob = c(0.01, 0.025, 0.05), ...)

## S3 method for class 'predict.mat'
cma(object, k, cutoff, prob = c(0.01, 0.025,
0.05), ...)

## S3 method for class 'cma'
plot(x, method = c("overplot", "jitter", "stack"),
     jitter = 0.1, vertical = FALSE,
     draw.quant = TRUE, xlab = NULL, ylab = "",
     main = "", cex.axis = NULL, ...,
     col.quant = "red", lty.quant = "dashed")

```

Arguments

object	an object for which close modern analogues are to be returned. Currently only for objects of class analog .
k	numeric; the number of analogues to return.
cutoff	numeric; critical value determining level below which samples from the modern reference set are defined as close modern analogues. May be missing, in which case the 2.5% quantile of the training set dissimilarities is used unless <code>object\$train</code> is NULL, in which case "cutoff" must be supplied.
prob	numeric vector of probabilities with values in [0,1], for which quantiles of the distribution of training set dissimilarities will be calculated. See quantile .
...	arguments to be passed to other cma methods or additional arguments passed to stripchart .
x	an object of class "cma".
method	the method to be used to separate coincident points. The default method "overplot" causes such points to be overplotted, but it is also possible to specify "jitter" to jitter the points, or "stack" have coincident points stacked. The last method only makes sense for very granular data.
jitter	when <code>method="jitter"</code> is used, jitter gives the amount of jittering applied.
vertical	when vertical is TRUE the plots are drawn vertically rather than the default horizontal.
draw.quant	logical; should the quantiles be drawn on the stripchart?

`xlab, ylab, main` Graphical parameters

`cex.axis` The magnification to be used for axis annotation relative to the current setting of `cex`. See [par](#).

`col.quant, lty.quant` colour and line type in which to draw the quantile lines.

Details

The `plot` method is simply a wrapper to [stripchart](#).

The methods for `mat` and `predict.mat` objects allow the user to select the `k`-closest analogues (argument `k`) or those samples as close or closer than a stated threshold of dissimilarity (argument `cutoff`). Only one of `k` and `cutoff` may be specified. If neither is specified, `getK` is used to extract the value for `k` stored within object. As such, the default is to return the automatically selected set of `k` closest samples, behaviour that is consistent with other functions in the package.

Value

For the `plot` method, a plot on the current device. Invisibly the plotted data are returned; see Note for further details.

A list of class "cma" with the following components:

`close` a named list of named vectors of close modern analogues and their dissimilarities. The names of the list components are the names of the fossil samples. The named vector in each component of `close` is the distances for the close modern analogues from the training set that are as close as `cutoff`, or closer, to the fossil sample.

`call` the matched call.

`cutoff` the cutoff threshold used to define close modern analogues.

`quant` numeric vector of the requested quantiles. Note returned by the [predict.mat](#) method.

`probs` the probabilities of the requested quantiles.

`method` character; the dissimilarity coefficient used

`n.analogs` numeric vector of the number of analogues per fossil sample.

Note

Only objects of classes [analog](#), [mat](#), and [predict.mat](#) are supported.

The `plot` method invisibly returns a list with the following components:

- `distances` a vector of stacked distances extracted from object.
- `groups` a factor listing the fossil sample for which the distances are the distances to the close modern analogues for the training set.

Author(s)

Gavin L. Simpson

References

Flower, R.J., Juggins, S. and Battarbee, R.W. (1997) Matching diatom assemblages in lake sediment cores and modern surface sediment samples: the implications for lake conservation and restoration with special reference to acidified systems. *Hydrobiologia* **344**; 27–40.

Simpson, G.L., Shilland, E.M., Winterbottom, J. M. and Keay, J. (2005) Defining reference conditions for acidified waters using a modern analogue approach. *Environmental Pollution* **137**; 119–133.

See Also

[analog](#), [stripchart](#), or [boxplot](#) for an alternative representation.

Examples

```
## Imbrie and Kipp example
## load the example data
data(ImbrieKipp)
data(SumSST)
data(V12.122)

## merge training and test set on columns
dat <- join(ImbrieKipp, V12.122, verbose = TRUE)

## extract the merged data sets and convert to proportions
ImbrieKipp <- dat[[1]] / 100
V12.122 <- dat[[2]] / 100

## analog matching between SWAP and RLGH reference samples
(ik.ana <- analog(ImbrieKipp, V12.122, method = "chord"))

## close modern analogues --- no analogues are "close"
(ik.cma <- cma(ik.ana, cutoff = 0.91))
summary(ik.cma)

## plot the results
plot(ik.cma)
```

Description

Performs leave-one-out, k -fold, n k -fold and bootstrap cross-validation of palaeoecological transfer function models.

Usage

```
crossval(obj, ...)

## S3 method for class 'wa'
crossval(obj, method = c("L00", "kfold", "bootstrap"),
         nboot = 100, nfold = 10, folds = 5,
         verbose = getOption("verbose"), ...)
```

Arguments

obj	A fitted transfer function model. Currently, only objects of class "wa" are supported.
method	character; type of cross-validation.
nboot	numeric; number of bootstrap samples.
nfold	numeric; number of chunks into which the training data are split. The k in k -fold.
folds	numeric; the number of times k -fold CV is performed.
verbose	logical; should progress of the CV be displayed?
...	Arguments passed to other methods.

Value

Returns an object of class "crossval", a list with the following components:

fitted.values	numeric vector; the cross-validated estimates of the response.
residuals	numeric vector; residuals computed from the cross-validated estimates of the response.
performance	data frame; cross-validation performance statistics for the model.
CVparams	list; parameters holding details of the cross-validation process.
call	the matched call.

Author(s)

Gavin L. Simpson

See Also

[wa](#)

Examples

```
## Load the Imbrie & Kipp data and
## summer sea-surface temperatures
data(ImbrieKipp)
data(SumSST)

## fit the WA model
```

```

mod <- wa(SumSST ~., data = ImbrieKipp)
mod

## Leave one out CV
cv.loo <- crossval(mod)
cv.loo

## k-fold CV (k == 10)
cv.kfold <- crossval(mod, kfold = 10, folds = 1, method = "kfold")
cv.kfold

## n k-fold CV (k == 10, n = 10)
cv.nkfold <- crossval(mod, kfold = 10, folds = 10, method = "kfold")
cv.nkfold

## bootstrap with 250 bootstrap samples
cv.boot <- crossval(mod, method = "bootstrap", nboot = 250)
cv.boot

## extract fitted values and residuals
fitted(cv.boot)
resid(cv.boot)

```

densityplot.residLen *Lattice density plot for residual lengths*

Description

Lattice [densityplot](#) method for [residLen](#) objects.

Usage

```

## S3 method for class 'residLen'
densityplot(x, ..., xlab = NULL, ylab = NULL)

```

Arguments

x	Object of class "residLen", the result of a call to residLen .
xlab, ylab	Axis labels. If not supplied, suitable defaults are generated, depending on whether RDA or CCA was used as the underlying ordination model.
...	Additional arguments passed to densityplot .

Value

Returns an object of class "trellis". See [densityplot](#) for details.

Author(s)

Gavin L. Simpson

See Also[residLen](#), [plot.residLen](#), [hist.residLen](#), [histogram.residLen](#).**Examples**

```
## load the Imbrie and Kipp example data
data(ImbrieKipp, SumSST, V12.122)

## squared residual lengths for Core V12.122
rlens <- residLen(ImbrieKipp, SumSST, V12.122)
rlens

## plot the density functions of the residual distances
densityplot(rlens)
```

deshrink

*Deshrinking techniques for WA transfer functions***Description**

In Weighted Averaging models averages are taken twice and thus WA estimates shrink towards the training set mean and need to be deshrunk. `deshrink` performs this deshrinking using several techniques, whilst `deshrinkPred` will deshrink WA estimates for new samples given a set of deshrinking coefficients.

Usage

```
deshrink(env, wa.env, type = c("inverse", "classical",
                              "expanded", "none"))

deshrinkPred(x, coef, type = c("inverse", "classical",
                              "expanded", "none"))
```

Arguments

<code>env</code>	numeric; original environmental values.
<code>wa.env</code>	numeric; initial weighted average estimates.
<code>type</code>	character; the type of deshrinking. One of "inverse", "classical", "expand", "none".
<code>x</code>	numeric; estimates to be deshrunk.
<code>coef</code>	numeric; deshrinking coefficients to use. Currently needs to be a vector of length 2. These should be supplied in the order β_0, β_1 .

Value

For `deshrinkPred` a numeric vector of `deshrunk` estimates.

For an object of class `"deshrink"`, inheriting from class `"list"`, with two components. The type of `deshrinking` performed is stored within attribute `"type"`. The components of the returned object are:

`coefficients` The `deshrinking` coefficients used.
`env` The `deshrunk` WA estimates.

Warning

`deshrinkPred`, does not currently check that the correct coefficients have been supplied in the correct order.

Author(s)

Gavin L. Simpson & Jari Oksanen

References

Birks, H.J.B. (1995) Quantitative environmental reconstructions. In *Statistical modelling of Quaternary science data* (eds. D. Maddy & J.S. Brew). Quaternary Research Association technical guide 5. Quaternary Research Association, Cambridge.

See Also

[wa](#)

dissimilarities *Extract dissimilarity coefficients from models*

Description

Extracts a vector of dissimilarity coefficients from an object for further analysis.

Usage

```
dissimilarities(object, ...)
dissim(object, ...)

## S3 method for class 'analog'
dissimilarities(object, which = c("train", "analogs"),
                ...)

## S3 method for class 'mat'
dissimilarities(object, ...)
```

Arguments

object	an R object from which the dissimilarity values are to be extracted. Currently only for objects of class "analog".
which	character; which set of dissimilarities should be extracted. One of "train" or "analogs".
...	arguments passed to other methods.

Details

The function can be called using the much shorter name "dissim".

Value

A vector of dissimilarities.

Author(s)

Gavin L. Simpson

See Also

[analog](#), [plot.dissimilarities](#)

Examples

```
## Imbrie and Kipp example
## load the example data
data(ImbrieKipp)
data(SumSST)
data(V12.122)

## merge training and test set on columns
dat <- join(ImbrieKipp, V12.122, verbose = TRUE)

## extract the merged data sets and convert to proportions
ImbrieKipp <- dat[[1]] / 100
V12.122 <- dat[[2]] / 100

## analog matching between SWAPImbrie & Kipp and V12.122 core
ik.analog <- analog(ImbrieKipp, V12.122, method = "chord")
ik.analog
summary(ik.analog)

## compare training set dissimilarities with normals
## and derive cut-offs
ik.dij <- dissim(ik.analog)
plot(ik.dij)
```

distance	<i>Flexibly calculate dissimilarity or distance measures</i>
----------	--

Description

Flexibly calculates distance or dissimilarity measures between a training set x and a fossil or test set y . If y is not supplied then the pairwise dissimilarities between samples in the training set, x , are calculated.

Usage

```
distance(x, ...)

## Default S3 method:
distance(x, y, method = c("euclidean", "SQeuclidean",
  "chord", "SQchord", "bray", "chi.square",
  "SQchi.square", "information", "chi.distance",
  "manhattan", "kendall", "gower", "alt.gower",
  "mixed"),
  fast = TRUE,
  weights = NULL, R = NULL, ...)
```

```
## S3 method for class 'join'
distance(x, ...)
```

Arguments

x	data frame or matrix containing the training set samples, or and object of class join .
y	data frame or matrix containing the fossil or test set samples.
method	character; which choice of dissimilarity coefficient to use. One of the listed options. See Details below.
fast	logical; should fast versions of the dissimilarities be calculated? See details below.
weights	numeric; vector of weights for each descriptor.
R	numeric; vector of ranges for each descriptor.
...	arguments passed to other methods

Details

A range of dissimilarity coefficients can be used to calculate dissimilarity between samples. The following are currently available:

euclidean	$d_{jk} = \sqrt{\sum_i (x_{ij} - x_{ik})^2}$
-----------	--

SQeuclidean	$d_{jk} = \sum_i (x_{ij} - x_{ik})^2$
chord	$d_{jk} = \sqrt{\sum_i (\sqrt{x_{ij}} - \sqrt{x_{ik}})^2}$
SQchord	$d_{jk} = \sum_i (\sqrt{x_{ij}} - \sqrt{x_{ik}})^2$
bray	$d_{jk} = \frac{\sum_i x_{ij} - x_{ik} }{\sum_i (x_{ij} + x_{ik})}$
chi.square	$d_{jk} = \sqrt{\sum_i \frac{(x_{ij} - x_{ik})^2}{x_{ij} + x_{ik}}}$
SQchi.square	$d_{jk} = \sum_i \frac{(x_{ij} - x_{ik})^2}{x_{ij} + x_{ik}}$
information	$d_{jk} = \sum_i (p_{ij} \log(\frac{2p_{ij}}{p_{ij} + p_{ik}}) + p_{ik} \log(\frac{2p_{ik}}{p_{ij} + p_{ik}}))$
chi.distance	$d_{jk} = \sqrt{\sum_i (x_{ij} - x_{ik})^2 / (x_{i+} / x_{++})}$
manhattan	$d_{jk} = \sum_i (x_{ij} - x_{ik})$
kendall	$d_{jk} = \sum_i MAX_i - minimum(x_{ij}, x_{ik})$
gower	$d_{jk} = \sum_i \frac{ p_{ij} - p_{ik} }{R_i}$
alt.gower	$d_{jk} = \sqrt{2 \sum_i \frac{ p_{ij} - p_{ik} }{R_i}}$
	where R_i is the range of proportions for descriptor (variable) i
mixed	$d_{jk} = \frac{\sum_{i=1}^p w_i s_{jki}}{\sum_{i=1}^p w_i}$
	where w_i is the weight for descriptor i and s_{jki} is the similarity between samples j and k for descriptor (variable) i .

Argument `fast` determines whether fast C versions of some of the dissimilarity coefficients are used. The fast versions make use of `dist` for methods "euclidean", "SQeuclidean", "chord", "SQchord", and `vegdist` for method == "bray". These fast versions are used only when `x` is supplied, not when `y` is also supplied. Future versions of `distance` will include fast C versions of all the dissimilarity coefficients and for cases where `y` is supplied.

Value

A matrix of dissimilarities where columns are the samples in `y` and the rows the samples in `x`. If `y` is not provided then a square, symmetric matrix of pairwise sample dissimilarities for the training set `x` is returned.

The dissimilarity coefficient used (`method`) is returned as attribute "method".

warning

For `method = "mixed"` it is essential that a factor in `x` and `y` have the same levels in the two data frames. Previous versions of analogue would work even if this was not the case, which will have generated incorrect dissimilarities for `method = "mixed"` for cases where factors for a given species had different levels in `x` to `y`.

`distance` now checks for matching levels for each species (column) recorded as a factor. If the factor for any individual species has different levels in `x` and `y`, an error will be issued.

Note

The dissimilarities are calculated in native R code. As such, other implementations (see See Also below) will be quicker. This is done for one main reason - it is hoped to allow a user defined function to be supplied as argument "method" to allow for user-extension of the available coefficients.

The other advantage of distance over other implementations, is the simplicity of calculating only the required pairwise sample dissimilarities between each fossil sample (y) and each training set sample (x). To do this in other implementations, you would need to merge the two sets of samples, calculate the full dissimilarity matrix and then subset it to achieve similar results.

Author(s)

Gavin L. Simpson

References

- Faith, D.P., Minchin, P.R. and Belbin, L. (1987) Compositional dissimilarity as a robust measure of ecological distance. *Vegetatio* **69**, 57–68.
- Gavin, D.G., Oswald, W.W., Wahl, E.R. and Williams, J.W. (2003) A statistical approach to evaluating distance metrics and analog assignments for pollen records. *Quaternary Research* **60**, 356–367.
- Kendall, D.G. (1970) A mathematical approach to seriation. *Philosophical Transactions of the Royal Society of London - Series B* **269**, 125–135.
- Legendre, P. and Legendre, L. (1998) *Numerical Ecology*, 2nd English Edition. Elsevier Science BV, The Netherlands.
- Overpeck, J.T., Webb III, T. and Prentice I.C. (1985) Quantitative interpretation of fossil pollen spectra: dissimilarity coefficients and the method of modern analogues. *Quaternary Research* **23**, 87–108.
- Prentice, I.C. (1980) Multidimensional scaling as a research tool in Quaternary palynology: a review of theory and methods. *Review of Palaeobiology and Palynology* **31**, 71–104.

See Also

[vegdist](#) in package **vegan**, [daisy](#) in package **cluster**, and [dist](#) provide comparable functionality for the case of missing y and are implemented in compiled code, so will be faster.

Examples

```
## simple example using dummy data
train <- data.frame(matrix(abs(runif(200)), ncol = 10))
rownames(train) <- LETTERS[1:20]
colnames(train) <- as.character(1:10)
fossil <- data.frame(matrix(abs(runif(100)), ncol = 10))
colnames(fossil) <- as.character(1:10)
rownames(fossil) <- letters[1:10]

## calculate distances/dissimilarities between train and fossil
## samples
test <- distance(train, fossil)

## using a different coefficient, chi-square distance
test <- distance(train, fossil, method = "chi.distance")

## calculate pairwise distances/dissimilarities for training
## set samples
```

```

test2 <- distance(train)

## Using distance on an object of class join
dists <- distance(join(train, fossil))
str(dists)

## calculate Gower's general coefficient for mixed data
## first, make a couple of variables factors
fossil[,4] <- factor(sample(rep(1:4, length = 10), 10))
train[,4] <- factor(sample(rep(1:4, length = 20), 20))
## now fit the mixed coefficient
test3 <- distance(train, fossil, "mixed")

## Example from page 260 of Legendre & Legendre (1998)
x1 <- t(c(2,2,NA,2,2,4,2,6))
x2 <- t(c(1,3,3,1,2,2,2,5))
Rj <- c(1,4,2,4,1,3,2,5) # supplied ranges

distance(x1, x2, method = "mixed", R = Rj)

## note this gives 1 - 0.66 (not 0.66 as the answer in
## Legendre & Legendre) as this is expressed as a
## distance whereas Legendre & Legendre describe the
## coefficient as similarity coefficient

```

fuse

Fused dissimilarities

Description

Combines dissimilarities from two or more dissimilarity objects into a single dissimilarity object so that both original dissimilarities contribute equally. Weighted combinations of the original objects can also be created.

Usage

```

fuse(..., weights = NULL)

## S3 method for class 'matrix'
fuse(..., weights = NULL)

## S3 method for class 'dist'
fuse(..., weights = NULL)

```

Arguments

... objects to fuse. Methods currently exist for objects of class "matrix" and "dist" objects. Method dispatch is performed on the first object specified in A minimum of two objects must be supplied.

weights numeric; vector of weights that sum to 1, one weight for each object supplied in

Details

Fuses, or combines, dissimilarity objects in a very flexible way to create a single dissimilarity object that incorporates the separate dissimilarities. In analogue matching, we may wish to combine information from two or more proxies, such as diatoms and cladocera, or from biological and chemical or physical data in the case of matching modern samples.

The function can also be used to fuse dissimilarity objects created from a single data set but using different dissimilarity coefficients. In this way one could create a new dissimilarity object combining dissimilarity based on abundance data and presence absence data into a single measure.

fuse uses the method of Melssen et al. (2006) to combine dissimilarities. The dissimilarities in each dissimilarity object are scaled so that the maximum dissimilarity in each object is 1. The scaled dissimilarity objects are then weighted according to the supplied weights. If no weights are supplied (the default) the dissimilarity objects are weighted equally; `weights = rep(1/N, N)`, where N is the number of dissimilarity objects fused.

$$D_{fused}(j, k) = \sum_{i=1}^N w_i D_{ijk}$$

where $D_{fused}(j, k)$ is the fused dissimilarity between samples j and k , w_i is the weight assigned to the i th dissimilarity object and D_{ijk} is the dissimilarity between j and k for the i th dissimilarity object.

Value

fuse returns an object of class "dist" with the attribute "method" set to "fuse".

This is the case even if the supplied objects are full dissimilarity matrices. If you want a full dissimilarity object, use `as.matrix.dist` on the returned object.

The returned object contains an extra attribute "weights", which records the weights used in the fusing.

Author(s)

Gavin L. Simpson

References

Melssen W., Wehrens R. and Buydens L. (2006) Supervised Kohonen networks for classification problems. *Chemometrics and intelligent laboratory systems* **83**, 99–113.

See Also

[dist](#), [vegdist](#), [distance](#).

Examples

```

train1 <- data.frame(matrix(abs(runif(100)), ncol = 10))
train2 <- data.frame(matrix(sample(c(0,1), 100, replace = TRUE),
                               ncol = 10))
rownames(train1) <- rownames(train2) <- LETTERS[1:10]
colnames(train1) <- colnames(train2) <- as.character(1:10)

d1 <- vegdist(train1, method = "bray")
d2 <- vegdist(train2, method = "jaccard")

dd <- fuse(d1, d2, weights = c(0.6, 0.4))
dd
str(dd)

```

getK

Extract and set the number of analogues

Description

An extractor function to access the number of analogues used in particular models. The stored value of k can be updated using `setK`.

Usage

```

getK(object, ...)

## S3 method for class 'mat'
getK(object, weighted = FALSE, ...)

## S3 method for class 'bootstrap.mat'
getK(object, which = c("bootstrap", "model"),
      prediction = FALSE, ...)

## S3 method for class 'predict.mat'
getK(object, which = c("model", "bootstrap"),
      ...)

setK(object, weighted = FALSE) <- value

## S3 replacement method for class 'mat'
setK(object, weighted = FALSE) <- value

```

Arguments

`object` an R object; currently only for objects of class `mat` and class `bootstrap.mat`.
`weighted` logical; extract/set number of analogues for a weighted or un-weighted model?

which	character; which k should be extracted, the one from the model or the one from the bootstrap results?
prediction	logical; should the extracted k be the one that is minimum for the test set (newdata) or the model (object).
...	further arguments to other methods.
value	integer; replacement value for k .

Details

getK is a generic accessor function, and setK<- is a generic replacement function.

Objects of class `bootstrap.mat` contain several different k 's. If no predictions are performed, there will be two k 's, one for the model and one from bootstrapping the model. Where predictions are performed **with** `newenv` supplied, in addition to the k 's above, there will be two k ' for the predictions, one for the model-based and one for the bootstrap-based predictions. To select k for the predictions, use `prediction = TRUE`. Argument `which` determines whether the model-based or the bootstrap-based k is returned.

Value

For getK, an integer value that is the number of analogues stored for use. The returned object has attributes "auto" and "weighted". "auto" refers to whether the extracted value of k was set automatically (TRUE) or by the user (FALSE). "weighted" states if the returned value is for a weighted analysis or an un-weighted analysis (FALSE).

For setK<-, the updated object.

Author(s)

Gavin L. Simpson

See Also

`mat`

Examples

```
## Imbrie and Kipp Sea Surface Temperature
data(ImbrieKipp)
data(SumSST)
data(V12.122)

## merge training set and core samples
dat <- join(ImbrieKipp, V12.122, verbose = TRUE)

## extract the merged data sets and convert to proportions
ImbrieKipp <- dat[[1]] / 100
ImbrieKippCore <- dat[[2]] / 100

## fit a MAT model
ik.mat <- mat(ImbrieKipp, SumSST, method = "chord")
```

```
## How many analogues gives lowest RMSE?
getK(ik.mat)
## note that this value was chosen automatically

## Now set k to be 10
setK(ik.mat) <- 10

## check
getK(ik.mat)
```

gradientDist

Positions of samples along a unit-length ordination gradient.

Description

Extracts information as to the locations of samples along an ordination gradient. `gradientDist()` standardises the entire gradient to the interval 0, ..., 1, to allow comparison between methods or data sets.

Usage

```
gradientDist(object, ...)

## Default S3 method:
gradientDist(object, order, na.rm = TRUE, ...)

## S3 method for class 'cca'
gradientDist(object, na.rm = TRUE, axis = 1L,
             scaling = 0, ...)

## S3 method for class 'prcurve'
gradientDist(object, na.rm = TRUE, ...)
```

Arguments

object	an R object of an appropriate type. For the default method, any R object that can be coerced to a vector.
order	numeric vector indicating the ordering of points along the gradient.
na.rm	logical; should missing values be removed?
axis	numeric, length 1; the ordination axis to take as the gradient.
scaling	Scaling to apply to the site scores. Default is to do no scaling. See scores.cca for details.
...	additional arguments passed to other methods. In the "cca" method, these are also passed to scores.cca .

Value

A numeric vector of positions along the gradient, scaled to the range 0, ..., 1.

Author(s)

Gavin L. Simpson

See Also

See [cca](#) and [prcurve](#) for functions that produce objects that `gradientDist()` can work with.

Examples

```
data(abernethy)

## Remove the Depth and Age variables
abernethy2 <- abernethy[, -(37:38)]

## Fit PCA
aber.pca <- rda(abernethy2)

## Distance along the first PCA axis
gradientDist(aber.pca)
```

hist.residLen	<i>Histogram plot for residual lengths</i>
---------------	--

Description

Base graphics histogram plot method for [residLen](#) objects.

Usage

```
## S3 method for class 'residLen'
hist(x, breaks = "Sturges", freq = TRUE,
     probs = c(0.9, 0.95, 0.99), ncol = 1, lcol = "red",
     llty = "dashed", xlab = NULL, ylab = NULL,
     main = "Residual distances", rug = TRUE, ...)
```

Arguments

x	Object of class "residLen", the result of a call to residLen .
breaks	How breakpoints for the histogram are determined. See <code>hist</code> for more details.
freq	logical; if TRUE, the histogram graphic is a representation of frequencies, the counts component of the result; if FALSE, probability densities, component density, are plotted (so that the histogram has a total area of one). Defaults to TRUE

<code>probs</code>	numeric; vector of probability quantiles to compute from the sets of residual distances.
<code>ncol</code>	numeric; number of columns for the plot layout. Choices are 1 or 2. Determines whether the histograms are plotted above or beside each other.
<code>lcol, llty</code>	colour and line-type for the quantiles.
<code>xlab, ylab</code>	Axis labels. If not supplied, suitable defaults are generated, depending on whether RDA or CCA was used as the underlying ordination model.
<code>main</code>	character; title for the plot.
<code>rug</code>	logical; should rug plots of the actual distances be drawn?
<code>...</code>	additional arguments passed to <code>hist</code> .

Value

A plot on the current device.

Returns a list with two components (`train` and `passive`), each of which is an object returned by `hist`.

Author(s)

Gavin L. Simpson

See Also

[residLen](#), [plot.residLen](#), [histogram.residLen](#), [densityplot.residLen](#).

Examples

```
## load the Imbrie and Kipp example data
data(ImbrieKipp, SumSST, V12.122)

## squared residual lengths for Core V12.122
rlens <- residLen(ImbrieKipp, SumSST, V12.122)
rlens

## plot a histogram of the residual distances
hist(rlens)
```

`histogram.residLen` *Lattice histogram plot for residual lengths*

Description

Lattice [histogram](#) method for [residLen](#) objects.

Usage

```
## S3 method for class 'residLen'  
histogram(x, ..., xlab = NULL, ylab = NULL,  
          type = c("percent", "count", "density"))
```

Arguments

x	Object of class "residLen", the result of a call to residLen .
xlab, ylab	Axis labels. If not supplied, suitable defaults are generated, depending on whether RDA or CCA was used as the underlying ordination model.
type	Character string indicating type of histogram to be drawn. "percent" and "count" give relative frequency and frequency histograms, and can be misleading when breakpoints are not equally spaced. "density" produces a density scale histogram. See histogram for further details.
...	Additional arguments passed to histogram .

Value

Returns an object of class "trellis". See [histogram](#) for details.

Author(s)

Gavin L. Simpson

See Also

[residLen](#), [plot.residLen](#), [hist.residLen](#), [densityplot.residLen](#).

Examples

```
## load the Imbrie and Kipp example data  
data(ImbrieKipp, SumSST, V12.122)  
  
## squared residual lengths for Core V12.122  
rlens <- residLen(ImbrieKipp, SumSST, V12.122)  
rlens  
  
## plot a histogram of the residual distances  
histogram(rlens)
```

ImbrieKipp

Imbrie and Kipp foraminifera training set

Description

The classic Imbrie and Kipp (1971) training set of counts on 27 species of foraminifera from 61 ocean sediment core surface samples and associated measures of summer and winter sea-surface temperatures and salinity at each location.

110 sediment cores samples from core V12-122 are also supplied in V12.122.

Usage

```
data(SumSST)
data(WinSST)
data(Salinity)
data(V12.122)
```

Format

ImbrieKipp is a data frame with 61 observations on the following 27 species:

- O.univ a numeric vector
- G.cglob a numeric vector
- G.ruber a numeric vector
- G.tenel a numeric vector
- G.saccu a numeric vector
- G.rubes a numeric vector
- G.pacL a numeric vector
- G.pacR a numeric vector
- G.bullo a numeric vector
- G.falco a numeric vector
- G.calid a numeric vector
- G.aequi a numeric vector
- G.gluti a numeric vector
- G.duter a numeric vector
- G.infla a numeric vector
- G.trnL a numeric vector
- G.trnR a numeric vector
- G.crasf a numeric vector
- G.scitu a numeric vector
- G.mentu a numeric vector

P.obliq a numeric vector
C.nitid a numeric vector
S.dehis a numeric vector
G.digit a numeric vector
Other a numeric vector
G.quin a numeric vector
G.hirsu a numeric vector

Summer and Winter sea-surface temperatures, and salinity values for the 61 sites in the Imbrie and Kipp training set (ImbrieKipp):

SumSST a numeric vector of summer sea-surface water temperatures. Values are in degrees C.

WinSST a numeric vector of winter sea-surface water temperatures. Values are in degrees C.

Salinity a numeric vector of sea water salinity values.

V12.122 is a data frame with 110 observations from core V12-122 on the following 28 species:

O.univ a numeric vector
G.cglob a numeric vector
G.ruber a numeric vector
G.tenel a numeric vector
G.saccu a numeric vector
G.rubes a numeric vector
G.pacL a numeric vector
G.pacR a numeric vector
G.bullo a numeric vector
G.falco a numeric vector
G.calid a numeric vector
G.aequi a numeric vector
G.gluti a numeric vector
G.duter a numeric vector
G.infla a numeric vector
G.trnL a numeric vector
G.trnR a numeric vector
G.crasf a numeric vector
G.scitu a numeric vector
G.mentu a numeric vector
P.obliq a numeric vector
C.nitid a numeric vector
S.dehis a numeric vector
G.digit a numeric vector
G.hexag a numeric vector
G.cglom a numeric vector
cfH.pel a numeric vector
Other a numeric vector

Source

Imbrie and Kipp (1971) TODO: Get the full reference.

These data were provided in electronic format by Prof. H. John B. Birks.

Examples

```
data(ImbrieKipp)
head(ImbrieKipp)

data(SumSST)
data(WinSST)
data(Salinity)

plot(cbind(SumSST, WinSST, Salinity))

data(V12.122)
head(V12.122)
```

join	<i>Merge species data sets on common columns (species)</i>
------	--

Description

Merges any number of species matrices on their common columns to create a new data set with number of columns equal to the number of unique columns across all data frames. Needed for analysis of fossil data sets with respect to training set samples.

Usage

```
join(..., verbose = FALSE, na.replace = TRUE, split = TRUE, value = 0,
      type = c("outer", "left", "inner"))

## S3 method for class 'join'
head(x, ...)

## S3 method for class 'join'
tail(x, ...)
```

Arguments

...	for join, data frames containing the data sets to be merged. For the head and tail methods, additional arguments to head and tail , in particular "n" to control the number of rows of each joined data set to display.
verbose	logical; if TRUE, the function prints out the dimensions of the data frames in "...", as well as those of the returned, merged data frame.

<code>na.replace</code>	logical; samples where a column in one data frame that have no matching column in the other will contain missing values (NA). If <code>na.replace</code> is TRUE, these missing values are replaced with zeros. This is standard practice in ecology and palaeoecology. If you want to replace with another value, then set <code>na.replace</code> to FALSE and do the replacement later.
<code>split</code>	logical; should the merged data sets samples be split back into individual data frames, but now with common columns (i.e. species)?
<code>value</code>	numeric; value to replace NA with if <code>na.replace</code> is TRUE.
<code>type</code>	logical; type of join to perform. "outer" returns the <i>union</i> of the variables in data frames to be merged, such that the resulting objects have columns for all variables found across all the data frames to be merged. "left" returns the left outer (or the left) join, such that the merged data frames contain the set of variables found in the first supplied data frame. "inner" returns the inner join, such that the merged data frame contain the intersection of the variables in the supplied data frames. See Details.
<code>x</code>	an object of class "join", usually the result of a call to join .

Details

When merging multiple data frames the set of variables in the merged data can be determined via a number of routes. `join` provides for two (currently) join types; the *outer* join and the *left outer* (or simply the *left*) join. Which type of join is performed is determined by the argument `type`.

The *outer* join returns the union of the set of variables found in the data frames to be merged. This means that the resulting data frame(s) contain columns for all the variable observed across all the data frames supplied for merging.

With the *left outer* join the resulting data frame(s) contain only the set of variables found in the first data frame provided.

The *inner* join returns the intersection of the set of variables found in the supplied data frames. The resulting data frame(s) contains the variables common to all supplied data frames.

Value

If `split = TRUE`, an object of class "join", a list of data frames, with as many components as the number of data frames originally merged.

Otherwise, an object of class `c("join", "data.frame")`, a data frame containing the merged data sets.

`head.join` and `tail.join` return a list, each component of which is the result of a call to [head](#) or [tail](#) on each data set compont of the joined object.

Author(s)

Gavin L. Simpson

See Also

[merge](#)

Examples

```

## load the example data
data(swapdiat, swappH, rlgh)

## merge training and test set on columns
dat <- join(swapdiat, rlgh, verbose = TRUE)

## extract the merged data sets and convert to proportions
swapdiat <- dat[[1]] / 100
rlgh <- dat[[2]] / 100

## merge training and test set using left join
head(join(swapdiat, rlgh, verbose = TRUE, type = "left"))

## load the example data
data(ImbrieKipp, SumSST, V12.122)

## merge training and test set on columns
dat <- join(ImbrieKipp, V12.122, verbose = TRUE)

## extract the merged data sets and convert to proportions
ImbrieKipp <- dat[[1]] / 100
V12.122 <- dat[[2]] / 100

## show just the first few lines of each data set
head(dat, n = 4)

## show just the last few lines of each data set
tail(dat, n = 4)

## merge training and test set using inner join
head(join(ImbrieKipp, V12.122, verbose = TRUE, type = "inner"))

## merge training and test set using outer join and replace
## NA with -99.9
head(join(ImbrieKipp, V12.122, verbose = TRUE, value = -99.9))

```

logitreg

Logistic regression models for assessing analogues/non-analogues

Description

Fits logistic regression models to each level of group to model the probability of two samples being analogues conditional upon the dissimilarity between the two samples.

Usage

```
logitreg(object, groups, k = 1, ...)
```

```
## Default S3 method:
logitreg(object, groups, k = 1, ...)

## S3 method for class 'analog'
logitreg(object, groups, k = 1, ...)

## S3 method for class 'logitreg'
summary(object, p = 0.9, ...)
```

Arguments

object	for <code>logitreg</code> ; a full dissimilarity matrix. For <code>summary.logitreg</code> an object of class "logitreg", the result of a call to <code>logitreg</code> .
groups	factor (or object that can be coerced to one) containing the group membership for each sample in object.
k	numeric; the k closest analogues to use in the model fitting.
p	probability at which to predict the dose needed.
...	arguments passed to other methods.

Details

Fits logistic regression models to each level of group to model the probability of two samples being analogues (i.e. in the same group) conditional upon the dissimilarity between the two samples.

This function can be seen as a way of directly modelling the probability that two sites are analogues, conditional upon dissimilarity, that can also be done less directly using [roc](#) and [bayesF](#).

Value

`logitreg` returns an object of class "logitreg"; a list whose components are objects returned by [glm](#). See [glm](#) for further details on the returned objects.

The components of this list take their names from `group`.

For `summary.logitreg` an object of class "summary.logitreg", a data frame with summary statistics of the model fits. The components of this data frame are:

In, Out	The number of analogue and non-analogue dissimilarities analysed in each group,
Est.(Dij), Std.Err	Coefficient and its standard error for dissimilarity from the logit model,
Z-value, p-value	Wald statistic and associated p-value for each logit model.
Dij(p=?), Std.Err(Dij)	The dissimilarity at which the posterior probability of two samples being analogues is equal to p , and its standard error. These are computed using dose.p .

Note

The function may generate warnings from function `glm.fit`. These should be investigated and not simply ignored.

If the message is concerns fitted probabilities being numerically 0 or 1, then check the fitted values of each of the models. These may well be numerically 0 or 1. Heed the warning in `glm` and read the reference cited therein which **may** indicate problems with the fitted models, such as (quasi-)complete separation.

Author(s)

Gavin L. Simpson

See Also

[roc](#), [bayesF](#), [glm](#).

Examples

```
## load the example data
data(swapdiat, swappH, rlgh)

## merge training and test set on columns
dat <- join(swapdiat, rlgh, verbose = TRUE)

## extract the merged data sets and convert to proportions
swapdiat <- dat[[1]] / 100
rlgh <- dat[[2]] / 100

## fit an analogue matching (AM) model using the squared chord distance
## measure - need to keep the training set dissimilarities
swap.ana <- analog(swapdiat, rlgh, method = "SQchord",
                  keep.train = TRUE)

## fit the ROC curve to the SWAP diatom data using the AM results
## Generate a grouping for the SWAP lakes
clust <- hclust(as.dist(swap.ana$train), method = "ward")
grps <- cutree(clust, 12)

## fit the logit models to the analog object
swap.lrm <- logitreg(swap.ana, grps)
swap.lrm

## summary statistics
summary(swap.lrm)

## plot the fitted logit curves
plot(swap.lrm, conf.type = "polygon")
```

 mat

Modern Analogue Technique transfer function models

Description

Modern Analogue Technique (MAT) transfer function models for palaeoecology. The fitted values are the, possibly weighted, averages of the environment for the k -closest modern analogues. MAT is a k -NN method.

Usage

```
mat(x, ...)

## Default S3 method:
mat(x, y,
     method = c("euclidean", "SQeuclidean", "chord", "SQchord",
                "bray", "chi.square", "SQchi.square",
                "information", "chi.distance", "manhattan",
                "kendall", "gower", "alt.gower", "mixed"),
     kmax, ...)

## S3 method for class 'formula'
mat(formula, data, subset, na.action,
     method = c("euclidean", "SQeuclidean", "chord", "SQchord",
                "bray", "chi.square", "SQchi.square",
                "information", "chi.distance", "manhattan",
                "kendall", "gower", "alt.gower", "mixed"),
     model = FALSE, ...)

## S3 method for class 'mat'
fitted(object, k, weighted = FALSE, ...)

## S3 method for class 'mat'
residuals(object, k, weighted = FALSE, ...)
```

Arguments

x	a data frame containing the training set data, usually species data.
y	a vector containing the response variable, usually environmental data to be predicted from x.
formula	a symbolic description of the model to be fit. The details of model specification are given below.
data	an optional data frame, list or environment (or object coercible by as.data.frame to a data frame) containing the variables in the model. If not found in data, the variables are taken from <code>environment(formula)</code> , typically the environment from which <code>wa</code> is called.

subset	an optional vector specifying a subset of observations to be used in the fitting process.
na.action	a function which indicates what should happen when the data contain NAs. The default is set by the na.action setting of <code>options</code> , and is <code>na.fail</code> if that is unset. The "factory-fresh" default is <code>na.omit</code> . Another possible value is NULL, no action. Value <code>na.exclude</code> can be useful.
method	a character string indicating the dissimilarity (distance) coefficient to be used to define modern analogues. See Details, below.
model	logical; If TRUE the model frame of the fit is returned.
kmax	numeric; limit the maximum number of analogues considered during fitting. By default, kmax is equal to $n - 1$, where n is the number of sites. For large data sets this is just wasteful as we wouldn't expect to be averaging over the entire training set. kmax can be used to restrict the upper limit on the number of analogues considered.
object	an object of class <code>mat</code> .
k	numeric; the k -closest analogue models' for which fitted values and residuals are returned. Overrides the default stored in the object.
weighted	logical; should weighted averages be used instead of simple averages?
...	arguments can be passed to <code>distance</code> to provide additional options required for some dissimilarities.

Details

The Modern Analogue Technique (MAT) is perhaps the simplest of the transfer function models used in palaeoecology. An estimate of the environment, x , for the response for a fossil sample, y , is the, possibly weighted, mean of that variable across the k -closest modern analogues selected from a modern training set of samples. If used, weights are the reciprocal of the dissimilarity between the fossil sample and each modern analogue.

A typical model has the form $\text{response} \sim \text{terms}$ where response is the (numeric) response data frame and terms is a series of terms which specifies a linear predictor for response. A typical form for terms is `.`, which is shorthand for "all variables" in data. If `.` is used, data must also be provided. If specific species (variables) are required then terms should take the form `spp1 + spp2 + spp3`.

Pairwise sample dissimilarity is defined by dissimilarity or distance coefficients. A variety of coefficients are supported — see `distance` for details of the supported coefficients.

k is chosen by the user. The simplest choice for k is to evaluate the RMSE of the difference between the predicted and observed values of the environmental variable of interest for the training set samples for a sequence of models with increasing k . The number of analogues chosen is the value of k that has lowest RMSE. However, it should be noted that this value is biased as the data used to build the model are also used to test the predictive power.

An alternative approach is to employ an optimisation data set on which to evaluate the size of k that provides the lowest RMSEP. This may be impractical with smaller sample sizes.

A third option is to bootstrap re-sample the training set many times. At each bootstrap sample, predictions for samples in the bootstrap test set can be made for $k = 1, \dots, n$, where n is the number

of samples in the training set. k can be chosen from the model with the lowest RMSEP. See function [bootstrap.mat](#) for further details on choosing k .

The output from [summary.mat](#) can be used to choose k in the first case above. For predictions on an optimisation or test set see [predict.mat](#). For bootstrap resampling of [mat](#) models, see [bootstrap.mat](#).

The fitted values are for the training set and are taken as the, possibly weighted, mean of the environmental variable in question across the k -closest analogues. The fitted value for each sample does **not** include a contribution from itself — it is the closest analogue, having zero dissimilarity. This spurious distance is ignored and analogues are ordered in terms of the non-zero distances to other samples in the training set, with the k -closest contributing to the fitted value.

Typical usages for [residuals.mat](#) are:

```
resid(object, k, weighted = FALSE, ...)
```

Value

`mat` returns an object of class `mat` with the following components:

<code>standard</code>	list; the model statistics based on simple averages of k -closest analogues. See below.
<code>weighted</code>	list; the model statistics based on weighted of k -closest analogues. See below.
<code>Dij</code>	matrix of pairwise sample dissimilarities for the training set x .
<code>orig.x</code>	the original training set data.
<code>orig.y</code>	the original environmental data or response, y .
<code>call</code>	the matched function call.
<code>method</code>	the dissimilarity coefficient used.

If `model = TRUE` then additional components "terms" and "model" are returned containing the [terms](#) object and model frame used.

`fitted.mat` returns a list with the following components:

- `estimatednumeric`; a vector of fitted values.
- `knumeric`; this is the k -closest analogue model with lowest apparent RMSE.
- `weightedlogical`; are the fitted values the weighted averages of the environment for the k -closest analogues. If `FALSE`, the fitted values are the average of the environment for the k -closest analogues.

Note

The object returned by `mat` contains lists "standard" and "weighted" both with the following elements:

- `esta` matrix of estimated values for the training set samples for models using k analogues, where $k = 1, \dots, n$. n is the number of samples in the training set. Rows contain the values for each model of size k , with columns containing the estimates for each training set sample.

- `residmatrix`; as for `"est"`, but containing the model residuals.
- `rmsepvector`; containing the leave-one-out root mean square error or prediction.
- `avg.biasvector`; contains the average bias (mean of residuals) for models using k analogues, where $k = 1, \dots, n$. n is the number of samples in the training set.
- `max.biasvector`; as for `"avg.bias"`, but containing the maximum bias statistics.
- `r.squaredvector`; as for `"avg.bias"`, but containing the R^2 statistics.

Author(s)

Gavin L. Simpson

References

Gavin, D.G., Oswald, W.W., Wahl, E.R. and Williams, J.W. (2003) A statistical approach to evaluating distance metrics and analog assignments for pollen records. *Quaternary Research* **60**, 356–367.

Overpeck, J.T., Webb III, T. and Prentice I.C. (1985) Quantitative interpretation of fossil pollen spectra: dissimilarity coefficients and the method of modern analogues. *Quaternary Research* **23**, 87–108.

Prell, W.L. (1985) The stability of low-latitude sea-surface temperatures: an evaluation of the CLIMAP reconstruction with emphasis on the positive SST anomalies, Report TR 025. U.S. Department of Energy, Washington, D.C.

Sawada, M., Viau, A.E., Vettoretti, G., Peltier, W.R. and Gajewski, K. (2004) Comparison of North-American pollen-based temperature and global lake-status with CCCma AGCM2 output at 6 ka. *Quaternary Science Reviews* **23**, 87–108.

See Also

[summary.mat](#), [bootstrap.mat](#) for bootstrap resampling of MAT models, [predict.mat](#) for making predictions from MAT models, [fitted.mat](#) and [resid.mat](#) for extraction of fitted values and residuals from MAT models respectively. [plot.mat](#) provides a `plot.lm`-like plotting tool for MAT models.

Examples

```
## Imbrie and Kipp Sea Surface Temperature
data(ImbrieKipp)
data(SumSST)
data(V12.122)

## merge training set and core samples
dat <- join(ImbrieKipp, V12.122, verbose = TRUE)

## extract the merged data sets and convert to proportions
ImbrieKipp <- dat[[1]] / 100
ImbrieKippCore <- dat[[2]] / 100

## fit the MAT model using the squared chord distance measure
ik.mat <- mat(ImbrieKipp, SumSST, method = "chord")
```

```
ik.mat

## model summary
summary(ik.mat)

## fitted values
fitted(ik.mat)

## model residuals
resid(ik.mat)

## draw summary plots of the model
par(mfrow = c(2,2))
plot(ik.mat)
par(mfrow = c(1,1))

## reconstruct for the RLGH core data
coreV12.mat <- predict(ik.mat, V12.122, k = 3)
coreV12.mat
summary(coreV12.mat)

## draw the reconstruction
reconPlot(coreV12.mat, use.labels = TRUE, display.error = "bars",
          xlab = "Depth", ylab = "SumSST")

## fit the MAT model using the squared chord distance measure
## and restrict the number of analogues we fit models for to 1:20
ik.mat2 <- mat(ImbrieKipp, SumSST, method = "chord", kmax = 20)
ik.mat2

## model summary
summary(ik.mat2)

## fitted values
fitted(ik.mat2)

## model residuals
resid(ik.mat2)

## draw summary plots of the model
par(mfrow = c(2,2))
plot(ik.mat2)
par(mfrow = c(1,1))

## reconstruct for the RLGH core data
coreV12.mat2 <- predict(ik.mat, V12.122, k = 3)
coreV12.mat2
summary(coreV12.mat2)

## draw the reconstruction
reconPlot(coreV12.mat2, use.labels = TRUE, display.error = "bars",
          xlab = "Depth", ylab = "SumSST")
```

mcarlo

*Monte Carlo simulation of dissimilarities***Description**

Permutations and Monte Carlo simulations to define critical values for dissimilarity coefficients for use in MAT reconstructions.

Usage

```
mcarlo(object, ...)

## Default S3 method:
mcarlo(object, nsamp = 10000,
        type = c("paired", "complete", "bootstrap", "permuted"),
        replace = FALSE,
        method = c("euclidean", "SQeuclidean", "chord", "SQchord",
                  "bray", "chi.square", "SQchi.square",
                  "information", "chi.distance", "manhattan",
                  "kendall", "gower", "alt.gower", "mixed"),
        is.dcmat = FALSE, diag = FALSE, ...)

## S3 method for class 'mat'
mcarlo(object, nsamp = 10000,
        type = c("paired", "complete", "bootstrap", "permuted"),
        replace = FALSE, diag = FALSE, ...)

## S3 method for class 'analog'
mcarlo(object, nsamp = 10000,
        type = c("paired", "complete", "bootstrap", "permuted"),
        replace = FALSE, diag = FALSE, ...)
```

Arguments

object	an R object. Currently only object's of class "mat", "analog" or matrix-like object of species data allowed.
nsamp	numeric; number of permutations or simulations to draw.
type	character; the type of permutation or simulation to perform. See Details, below.
replace	logical; should sampling be done with replacement?
method	character; for raw species matrices, the dissimilarity coefficient to use. This is predefined when fitting a MAT model with mat or analogue matching via analogue and is ignored in the "mcarlo" methods for classes "mat" and "analog".
is.dcmat	logical; is "object" a dissimilarity matrix. Not meant for general use; used internally by "mat" and "analogue" methods to instruct the "default" method that "object" is already a dissimilarity matrix, so there is no need to recalculate.

`diag` logical; should the dissimilarities include the diagonal (zero) values of the dissimilarity matrix. See Details.

... arguments passed to or from other methods.

Details

Only "type" "paired" and "bootstrap" are currently implemented.

`distance` produces square, symmetric dissimilarity matrices for training sets. The upper triangle of these matrices is a duplicate of the lower triangle, and as such is redundant. `mcarlo` works on the lower triangle of these dissimilarity matrices, representing all pairwise dissimilarity values for training set samples. The default is **not** to include the diagonal (zero) values of the dissimilarity matrix. If you feel that these diagonal (zero) values are part of the population of dissimilarities then use "`diag = TRUE`" to include them in the permutations.

Value

A vector of simulated dissimilarities of length "`nsamp`". The "`method`" used is stored in attribute "`method`".

Note

The performance of these permutation and simulation techniques still needs to be studied. This function is provided for pedagogic reasons. Although recommended by Sawada et al (2004), sampling with replacement ("`replace = TRUE`") and including diagonal (zero) values ("`diag = TRUE`") simulates too many zero distances. This is because the same training set sample can, on occasion be drawn twice leading to a zero distance. It is impossible to find in nature two samples that will be perfectly similar, and as such sampling **with** replacement **and** "`diag = TRUE`" seems undesirable at best.

Author(s)

Gavin L. Simpson

References

Sawada, M., Viau, A.E., Vettoretti, G., Peltier, W.R. and Gajewski, K. (2004) Comparison of North-American pollen-based temperature and global lake-status with CCCma AGCM2 output at 6 ka. *Quaternary Science Reviews* **23**, 87–108.

See Also

`mat` for fitting MAT models and `analog` for analogue matching. `roc` as an alternative method for determining critical values for dissimilarity measures when one has grouped data.

`plot.mcarlo` provides a plotting method to visualise the distribution of simulated dissimilarities.

Examples

```

## Imbrie and Kipp example
## load the example data
data(ImbrieKipp)
data(SumSST)
data(V12.122)

## merge training and test set on columns
dat <- join(ImbrieKipp, V12.122, verbose = TRUE)

## extract the merged data sets and convert to proportions
ImbrieKipp <- dat[[1]] / 100
V12.122 <- dat[[2]] / 100

## perform the modified method of Sawada (2004) - paired sampling,
## with replacement
ik.mcarlo <- mcarlo(ImbrieKipp, method = "chord", nsamp = 1000,
                   type = "paired", replace = FALSE)

ik.mcarlo

## plot the simulated distribution
layout(matrix(1:2, ncol = 1))
plot(ik.mcarlo)
layout(1)

```

minDC

Extract minimum dissimilarities

Description

Minimum dissimilarity is a useful indicator of reliability of reconstructions performed via MAT and other methods, and for analogue matching. Minimum dissimilarity for a sample is the smallest dissimilarity between it and the training set samples.

Usage

```

minDC(x, ...)

## Default S3 method:
minDC(x, ...)

## S3 method for class 'predict.mat'
minDC(x, ...)

## S3 method for class 'analog'
minDC(x, probs = c(0.01, 0.02, 0.05, 0.1), ...)

## S3 method for class 'wa'

```

```
minDC(x, y,
      method = c("euclidean", "SQeuclidean", "chord", "SQchord",
                  "bray", "chi.square", "SQchi.square", "information",
                  "chi.distance", "manhattan", "kendall", "gower",
                  "alt.gower", "mixed"),
      percent = FALSE, probs = c(0.01, 0.025, 0.05, 0.1), ...)
```

Arguments

x	an object of class "predict.mat", "analog" or a object with a component named "minDC".
probs	numeric; vector of probabilities with values in [0,1].
y	an optional matrix-like object containing fossil samples for which the minimum dissimilarities to training samples are to be calculated.
method	character; which choice of dissimilarity coefficient to use. One of the listed options. See distance .
percent	logical; Are the data percentages? If TRUE, the data (x and y) will be divided by 100 to convert them to the proportions expected by distance .
...	other arguments to be passed to other methods. Currently ignored.

Value

minDC returns an object of class "minDC".

An object of class minDC is a list with some or all of the following components:

minDC	numeric; vector of minimum dissimilarities.
method	character; the dissimilarity coefficient used.
quantiles	numeric; named vector of probability quantiles for distribution of dissimilarities of modern training set.

Note

The "default" method of minDC will attempt to extract the relevant component of the object in question. This may be useful until a specific minDC method is written for a given class.

Author(s)

Gavin L. Simpson

See Also

[predict.mat](#), and [plot.minDC](#) for a plotting method.

Examples

```
## Imbrie and Kipp example
## load the example data
data(ImbrieKipp)
data(SumSST)
data(V12.122)

## merge training and test set on columns
dat <- join(ImbrieKipp, V12.122, verbose = TRUE)

## extract the merged data sets and convert to proportions
ImbrieKipp <- dat[[1]] / 100
V12.122 <- dat[[2]] / 100

## fit the MAT model using the squared chord distance measure
ik.mat <- mat(ImbrieKipp, SumSST, method = "SQchord")
ik.mat

## reconstruct for the V12-122 core data
v12.mat <- predict(ik.mat, V12.122)

## extract the minimum DC values
v12.mdc <- minDC(v12.mat)
v12.mdc

## draw a plot of minimum DC by time
plot(v12.mdc, use.labels = TRUE, xlab = "Depth (cm.)")
```

optima

Weighted averaging optima and tolerance ranges

Description

Computes weighted average optima and tolerance ranges from species abundances and values of the environment.

Usage

```
optima(x, ...)

## Default S3 method:
optima(x, env, ...)

## Default S3 method:
tolerance(x, env, useN2 = TRUE, ...)
```

Arguments

x	Species data matrix or data frame.
env	Numeric; variable for which optima or tolerances are required.
useN2	logical; should Hill's N2 values be used to produce un-biased tolerances?
...	Arguments passed to other methods.

Value

Both functions return a named vector containing the WA optima or tolerances for the environmental gradient specified by env.

Note

Objects of class "optima" or "tolerance" can be coerced to data frames using methods for [as.data.frame](#).

Author(s)

Gavin L. Simpson

See Also

[wa](#)

Examples

```
## Load the Imbrie & Kipp data and
## summer sea-surface temperatures
data(ImbrieKipp)
data(SumSST)

## WA optima
(opt <- optima(ImbrieKipp, SumSST))

## WA tolerances
(tol <- tolerance(ImbrieKipp, SumSST, useN2 = TRUE))

## convert to data frame
as.data.frame(opt)
as.data.frame(tol)
```

panel.Loess	<i>Loess smooths to stratigraphic diagrams</i>
-------------	--

Description

A modified version of [panel.loess](#), for drawing Loess smooths on stratigraphic diagrams.

Usage

```
panel.Loess(x, y,
            span = 1/3, degree = 1,
            family = c("symmetric", "gaussian"),
            evaluation = 50,
            lwd = plot.line$lwd,
            lty = plot.line$lty,
            col,
            col.line = plot.line$col,
            type,
            ...)
```

Arguments

x, y	variables defining the contents of the panel.
span, degree, family, evaluation	arguments passed to loess.smooth , for which panel.Loess is a wrapper.
lwd, lty, col, col.line	graphical parameters.
type	for compatibility with panel.loess , but is ignored within the function.
...	graphical parameters can be supplied. Color can usually be specified by col, col.line, the latter overriding the first for lines.

Details

The standard panel function [panel.loess](#) treats the data as the x-axis acting as the time component. For stratigraphic plots where time flows along the y-axis, we want the smoother to be fitted with the x-axis data as the response and the time component (y-axis) as the predictor.

This modified version of [panel.loess](#) flips the two axes to produce the desired effect. Note also that it does not have argument `horizontal` as this is not required or supported by [Stratiplot](#). In other respects, panel.Loess is equivalent to the lattice panel function [panel.loess](#).

User should note that warnings can be generated by the fitting function if span is set too small for species with few observations. In such cases, the user is directed to the help page for [loess.smooth](#), but increasing span slightly can often stop the warnings.

Author(s)

Gavin L. Simpson, slightly modified from the Lattice function [panel.loess](#) by Deepayan Sarkar.

See Also

[loess.smooth](#), [panel.loess](#).

panel.Stratplot *Panel function for stratigraphic diagrams*

Description

A Lattice panel function for drawing individual panels on stratigraphic diagrams using a range of plot types commonly used within palaeoecology.

Usage

```
panel.Stratplot(x, y,
               type = "l",
               col,
               pch = plot.symbol$pch,
               cex = plot.symbol$cex,
               col.line = plot.line$col,
               col.symbol = plot.symbol$col,
               col.refline = ref.line$col,
               col.smooth = "red",
               col.poly = plot.line$col,
               lty = plot.line$lty,
               lwd = plot.line$lwd,
               lty.smooth = plot.line$lty,
               lwd.smooth = 2,
               fill = plot.symbol$fill,
               zones = NULL,
               col.zones = plot.line$col,
               lty.zones = plot.line$lty,
               lwd.zones = plot.line$lwd,
               gridh = -1, gridv = -1,
               ...)
```

Arguments

`x, y` variables defining the contents of the panel.

`type` character vector consisting of one or more of the following: "l", "p", "o", "b", "h", "g", "smooth", and "poly". It type has more than one element, the effects of each component are combined, though note that some elements will over-plot, and hence obscure, earlier elements.

For types "l", "p", "o", "b" and "g" the standard Lattice interpretation is observed. See [panel.xyplot](#) for further details. Note that type "b" is the same as type "o".

"g" adds a reference grid using `panel.grid` in the background.

	For "h", histogram-like bars are plotted from the y-axis , not from the x-axis with <code>plot</code> and <code>panel.loess</code> .
	For "smooth" a loess fit is added to each panel using <code>panel.Loess</code> .
	For "poly", a shaded polygon, or silhouette, is drawn for each panel.
<code>col, col.line, col.symbol, col.poly, col.refline, col.smooth, col.zones</code>	colour parameters. For all but <code>col.smooth</code> , default colours are obtained from <code>plot.symbol</code> and <code>plot.line</code> using <code>trellis.par.get</code> . <code>col.refline</code> controls the colour of the reference line drawn at value 0 on the x-axis, as well as the colour of the grid lines if drawn.
<code>pch, cex, lty, lwd, fill</code>	other graphical parameters, defaults for which are obtained from <code>plot.symbol</code> and <code>plot.line</code> using <code>trellis.par.get</code> .
<code>lty.smooth</code>	line type for the loess smoother. The default is obtained from <code>plot.line</code> using <code>trellis.par.get</code> .
<code>lwd.smooth</code>	The line width for the loess smoother.
<code>zones</code>	numeric; vector of zone boundary positions on scale of the depth/time (y-)axis.
<code>lty.zones, lwd.zones</code>	line type and width for the zone markers. The defaults are obtained from <code>plot.line</code> .
<code>gridh, gridv</code>	numeric arguments corresponding to <code>h</code> and <code>v</code> of <code>panel.grid</code> , which control the number of grid lines drawn.
<code>...</code>	extra arguments passed on to the underlying panel functions; <code>panel.points</code> , <code>panel.lines</code> , <code>panel.segments</code> , <code>panel.polygon</code> , <code>panel.Loess</code> and <code>panel.refline</code> .

Details

Creates stratigraphic scatter plots of `x` and `y`, with various modifications possible via the type argument.

Zones can be drawn on the panels by supplying the numeric vector of zone boundaries as argument `zones`. The panel function will then draw horizontal lines across the panels at the desired y-axis locations. Note that the panel function does **not** attempt to identify the zone boundaries automatically; these must be determined via a chronological (constrained) cluster analysis function or similar.

Note that all the arguments controlling the display can be supplied directly to a high-level call of the function `Stratplot`.

Author(s)

Gavin L. Simpson

See Also

`Stratplot`, `panel.Loess`, `panel.xyplot`.

Description

Fits a palaeoecological transfer function model using principal component regression, using an optional transformation of the matrix of predictor variables when these are species abundance data.

Usage

```
## Default S3 method:
pcr(x, y, ncomp, tranFun, ...)

## S3 method for class 'formula'
pcr(formula, data, subset, na.action, ..., model = FALSE)

Hellinger(x, apply = FALSE)

ChiSquare(x, apply = FALSE)

## S3 method for class 'pcr'
performance(object, ...)

## S3 method for class 'pcr'
residuals(object, comps = NULL, ...)

## S3 method for class 'pcr'
fitted(object, comps = NULL, ...)

## S3 method for class 'pcr'
coef(object, comps = NULL, ...)

## S3 method for class 'pcr'
screeplot(x, restrict = NULL,
          display = c("RMSE", "avgBias", "maxBias", "R2"),
          xlab = NULL, ylab = NULL, main = NULL, sub = NULL, ...)

## S3 method for class 'pcr'
eigenvals(x, ...)
```

Arguments

x	Matrix or data frame of predictor variables. Usually species composition or abundance data for transfer function models. For screeplot and eigenvals, an object of class "pcr".
y	Numeric vector; the response variable to be modelled.

ncomp	numeric; number of principal components to build models for. If not supplied the largest possible number of components is determined.
tranFun	function; a function or name of a function that performs a transformation of the predictor variables x . The function must be self-contained as no arguments are passed to the function when it is applied. See Details for more information.
formula	a model formula.
data	an optional data frame, list or environment (or object coercible by <code>as.data.frame</code> to a data frame) containing the variables specified on the RHS of the model formula. If not found in data, the variables are taken from <code>environment(formula)</code> , typically the environment from which <code>pcr</code> is called.
subset	an optional vector specifying a subset of observations to be used in the fitting process.
na.action	a function which indicates what should happen when the data contain NAs. The default is set by the <code>na.action</code> setting of options, and is <code>na.fail</code> if that is unset. The 'factory-fresh' default is <code>na.omit</code> . Another possible value is <code>NULL</code> , no action. Value <code>na.exclude</code> can be useful.
model	logical; If TRUE the model frame is returned?
apply	logical; should an existing transformation, using pre-computed meta-parameters, be applied?
object	an object of class "pcr".
comps	numeric; which components to return.
restrict	numeric; limit the number of components on the screeplot.
display	character; which model performance statistic should be drawn on the screeplot?
xlab, ylab, main, sub	character; labels for the plot.
...	Arguments passed to other methods.

Details

When applying cross-validation (CV) to transfer function models, any transformation of the predictors must be applied separately during each iteration of the CV procedure to the part of the data used in fitting the model. In the same way, any samples to be predicted from the model must use any meta-parameters derived from the training data only. For example, centring is applied to the training data only and the variable means used to centre the training data are used to centre the test samples. The variable means should not be computed on a combination of the training and test samples.

When using PCR, we might wish to apply a transformation to the species data predictor variables such that the PCA of those data preserves a dissimilarity coefficient other than the Euclidean distance. This transformation is applied to allow PCA to better describe patterns in the species data (Legendre & Gallagher 2001).

How this is handled in `pcr` is to take a user-supplied function that takes a single argument, the matrix of predictor variables. The function should return a matrix of the same dimension as the input. If any meta-parameters are required for subsequent use in prediction, these should be returned as attribute "parms", attached to the matrix.

Two example transformation functions are provided implementing the Hellinger and Chi Square transformations of Legendre & Gallagher (2001). Users can base their transformation functions on these. `ChiSquare()` illustrates how meta-parameters should be returned as the attribute "parms".

Value

Returns an object of class "pcr", a list with the following components:

fitted.values	matrix; the PCR estimates of the response. The columns contain fitted values using C components, where C is the Cth column of the matrix.
coefficients	matrix; regression coefficients for the PCR. Columns as per fitted above.
residuals	matrix; residuals, where the Cth column represents a PCR model using C components.
scores	
loadings	
Yloadings	
xMeans	numeric; means of the predictor variables in the training data.
yMean	numeric; mean of the response variable in the training data.
varExpl	numeric; variance explained by the PCR model. These are the squares of the singular values.
totvar	numeric; total variance in the training data
call	the matched call.
tranFun	transformation function used. NA if none supplied/used.
tranParms	list; meta parameters used to computed the transformed training data.
performance	data frame; cross-validation performance statistics for the model.
ncomp	numeric; number of principal components computed

Author(s)

Gavin L. Simpson

See Also

[wa](#)

Examples

```
## Load the Imbrie & Kipp data and
## summer sea-surface temperatures
data(ImbrieKipp)
data(SumSST)

## normal interface and apply Hellinger transformation
mod <- pcr(ImbrieKipp, SumSST, tranFun = Hellinger)
mod

## formula interface, but as above
mod2 <- pcr(SumSST ~ ., data = ImbrieKipp, tranFun = Hellinger)
mod2

## Several standard methods are available
```

```
fitted(mod, comps = 1:4)
resid(mod, comps = 1:4)
coef(mod, comps = 1:4)

## Eigenvalues can be extracted
eigenvals(mod)

## screeplot method
screeplot(mod)
```

performance

Transfer function model performance statistics

Description

A simple extractor function to access the model performance statistics of transfer function models.

Usage

```
performance(object, ...)

## S3 method for class 'wa'
performance(object, ...)

## S3 method for class 'predict.wa'
performance(object, ...)

## S3 method for class 'bootstrap.wa'
performance(object, ...)
```

Arguments

`object` A transfer function object.
`...` Arguments passed to other methods. Currently ignored.

Value

A named vector containing the extracted model performance statistics.

Author(s)

Gavin L. Simpson

See Also

[wa](#), [predict.wa](#), [bootstrap.wa](#).

Examples

```
## continue the example from ?wa
example(wa)

## the model performance statistics
performance(mod)
```

plot.dissimilarities *Plots the distribution of extracted dissimilarities*

Description

Produces a plot of the distribution of the extracted dissimilarities and a reference normal distribution with comparable mean and sd.

Usage

```
## S3 method for class 'dissimilarities'
plot(x, prob = 0.05,
     legend = TRUE, n.rnorm = 1e+05, col = "black",
     col.ref = "red", lty = "solid", lty.quant = "dotted",
     xlab = NULL, ylab = NULL, main = NULL, sub = NULL, ...)
```

Arguments

x	an object of class "dissimilarities".
prob	numeric; density probability defining the threshold for close modern analogues.
legend	logical; draw a legend on the plotted figure?
n.rnorm	numeric; number of random normal deviates for reference line.
col, col.ref	colours for the dissimilarity and reference density functions drawn.
lty, lty.quant	line types for the dissimilarity and reference density functions drawn.
xlab, ylab	character; x- and y-axis labels.
main, sub	character; main and subtitle for the plot.
...	graphical arguments passed to other graphics functions.

Value

A plot on the currently active device.

Author(s)

Gavin L. Simpson

See Also

[dissimilarities](#)

Examples

```
## Imbrie and Kipp example
## load the example data
data(ImbrieKipp)
data(SumSST)
data(V12.122)

## merge training and test set on columns
dat <- join(ImbrieKipp, V12.122, verbose = TRUE)

## extract the merged data sets and convert to proportions
ImbrieKipp <- dat[[1]] / 100
V12.122 <- dat[[2]] / 100

## analog matching between SWAPImbrie & Kipp and V12.122 core
ik.analog <- analog(ImbrieKipp, V12.122, method = "chord")
ik.analog
summary(ik.analog)

## compare training set dissimilarities with normals
## and derive cut-offs
ik.dij <- dissim(ik.analog)
plot(ik.dij)
```

plot.logitreg

Produces plots of analogue logistic regression models

Description

Draws the fitted logistic regression function describing the posterior probability that two sites are analogues conditional upon the dissimilarity between the two samples. Confidence intervals are also computed and displayed if requested.

Usage

```
## S3 method for class 'logitreg'
plot(x, group = "all", npred = 100,
     conf.int = 0.9, conf.type = c("none", "polygon", "lines"),
     xlab = expression(D[ij]), ylab = "Pr (A+ | d)",
     rug = TRUE, ticksize = 0.02,
     col = "red", ref.col = "lightgrey",
     lwd = 2, conf.lwd = 1, conf.lty = "dashed",
     shade = "lightgrey", ...)
```

Arguments

x object to plot; an object of class "logitreg", usually the result of a call to [logitreg](#).

group	The group to plot the logit model for. Can be one of the group labels or "Combined" to draw the individual logit models. Alternatively, and the default, is to use "all", which divides the plotting region into the required number of plotting regions and draws all the fitted curves.
npred	number of points at which the fitted curves are evaluated for plotting purposes.
conf.int	numeric; the confidence interval required.
conf.type	character; how should the confidence interval be drawn. Default is not to draw the confidence interval.
xlab, ylab	character; the x and y axis labels.
rug	logical; should rug plots be drawn?
ticksize	The size of the tick marks used in rug plots.
col	The colour in which to draw the line representing the fitted values.
ref.col	The colour of the reference lines drawn at 0 and 1.
lwd	The line width in which to draw the line representing the fitted values.
conf.lwd, conf.lty	Line width and line type for the confidence interval. Only used if conf.type = "lines" is specified.
shade	The colour for the fill and border of the confidence interval if conf.type = "polygon" is specified.
...	arguments passed on to plot.

Value

A plot on the current device.

Author(s)

Gavin L. Simpson

See Also

[logitreg](#) for an example, [roc](#)

plot.mat

Plot diagnostics for a mat object

Description

Five plots (selectable by which) are currently available: a plot of estimated against observed values, a plot of residuals against estimated values, and screeplots of the apparent RMSE, average bias and maximum bias for MAT models of size k , where $k = 1, \dots, n$. By default, the first three and '5' are provided.

Usage

```
## S3 method for class 'mat'
plot(x,
      which = c(1:3, 5),
      weighted = FALSE,
      k,
      caption = c("Inferred vs Observed", "Residuals vs Fitted",
                  "Leave-one-out errors", "Average bias",
                  "Maximum bias"),
      max.bias = TRUE,
      n.bias = 10,
      restrict = 20,
      sub.caption = NULL,
      main = "",
      ask = prod(par("mfcol")) < length(which) &&
            dev.interactive(),
      ...,
      panel = if (add.smooth) panel.smooth else points,
      add.smooth = getOption("add.smooth"))
```

Arguments

x	an object of class "mat".
which	which aspects of the "mat" object to plot if a subset of the plots is required, specify a subset of the numbers 1:5.
weighted	logical; should the analysis use weighted mean of env data of analogues as fitted/estimated values?
k	numeric; the number of analogues to use. If missing k is chosen automatically as the k that achieves lowest RMSE.
caption	captions to appear above the plots.
max.bias	logical, should max bias lines be added to residuals.
n.bias	numeric, number of sections to calculate maximum bias for.
restrict	logical; restrict comparison of k-closest model to $k \leq \text{restrict}$.
sub.caption	common title-above figures if there are multiple; used as 'sub' (s.'title') otherwise. If NULL, as by default, a possibly shortened version of <code>deparse(x\$call)</code> is used.
main	title to each plot-in addition to the above caption.
ask	logical; if TRUE, the user is <i>asked</i> before each plot, see <code>par(ask=.)</code> .
...	graphical arguments passed to other graphics functions.
panel	panel function. The useful alternative to <code>points</code> , <code>panel.smooth</code> , can be chosen by <code>add.smooth = TRUE</code> .
add.smooth	logical indicating if a smoother should be added to fitted & residuals plots; see also <code>panel</code> above.

Details

This plotting function is modelled closely on [plot.lm](#) and many of the conventions and defaults for that function are replicated here.

`sub.caption` - by default the function call - is shown as a subtitle (under the x-axis title) on each plot when plots are on separate pages, or as a subtitle in the outer margin (if any) when there are multiple plots per page.

Value

One or more plots, drawn on the current device.

Author(s)

Gavin L. Simpson. Code borrows heavily from [plot.lm](#).

See Also

[mat](#)

Examples

```
## Imbrie and Kipp example
## load the example data
data(ImbrieKipp)
data(SumSST)
data(V12.122)

## merge training and test set on columns
dat <- join(ImbrieKipp, V12.122, verbose = TRUE)

## extract the merged data sets and convert to proportions
ImbrieKipp <- dat[[1]] / 100
V12.122 <- dat[[2]] / 100

## MAT
ik.mat <- mat(ImbrieKipp, SumSST, method = "chord")
ik.mat
summary(ik.mat)

## summary plot of MAT model
layout(matrix(1:4, ncol = 2, byrow = TRUE))
plot(ik.mat)
layout(1)
```

plot.mcarlo

*Plot Monte Carlo simulated dissimilarity distributions***Description**

A `plot.lm`-like plotting function for objects of class "mcarlo" to visualise the simulated distribution of dissimilarities.

Usage

```
## S3 method for class 'mcarlo'
plot(x,
      which = c(1:2),
      alpha = 0.05,
      caption = c("Distribution of dissimilarities",
                  expression(paste("Simulated probability Pr (Dissim < ",
                                   alpha, ")"))),
      col.poly = "lightgrey",
      border.poly = "lightgrey",
      ask = prod(par("mfcol")) < length(which) &&
              dev.interactive(),
      ...)
```

Arguments

x	an object of class "mcarlo", usually the result of a call to <code>mcarlo</code> .
which	numeric; which of the plots should be produced?
alpha	numeric; the Monte Carlo significance level to be marked on the cumulative frequency plot.
caption	character, length 2; captions to appear above the plots.
col.poly, border.poly	character; the colour to draw the region and border of the polygon enclosing the Monte Carlo significance on the cumulative frequency plot.
ask	logical; should the function wait for user confirmation to draw each plot? If missing, the function makes a reasonable attempt to guess the current situation and act accordingly.
...	additional graphical parameters to be passed to the plotting functions. Currently ignored.

Details

The "Distribution of dissimilarities" plot produces a histogram and kernel density estimate of the distribution of simulated dissimilarity values.

The "Simulated probability" plot shows a cumulative probability function of the simulated dissimilarity values, and highlights the proportion of the curve that is less than alpha.

Value

One or more plots on the current device.

Author(s)

Gavin L. Simpson

References

Sawada, M., Viau, A.E., Vettoretti, G., Peltier, W.R. and Gajewski, K. (2004) Comparison of North-American pollen-based temperature and global lake-status with CCCma AGCM2 output at 6 ka. *Quaternary Science Reviews* **23**, 87–108.

See Also

[mcarlo](#)

Examples

```
## Imbrie and Kipp example
## load the example data
data(ImbrieKipp)
data(SumSST)
data(V12.122)

## merge training and test set on columns
dat <- join(ImbrieKipp, V12.122, verbose = TRUE)

## extract the merged data sets and convert to proportions
ImbrieKipp <- dat[[1]] / 100
V12.122 <- dat[[2]] / 100

## perform the modified method of Sawada (2004) - paired sampling,
## with replacement
ik.mcarlo <- mcarlo(ImbrieKipp, method = "chord", nsamp = 1000,
                   type = "paired", replace = FALSE)
ik.mcarlo

## plot the simulated distribution
layout(matrix(1:2, ncol = 1))
plot(ik.mcarlo)
layout(1)
```

Description

Minimum dissimilarity is a useful indicator of reliability of reconstructions performed via MAT and other methods, and for analogue matching. Minimum dissimilarity for a sample is the smallest dissimilarity between it and the training set samples.

Usage

```
## S3 method for class 'minDC'
plot(x, depths, use.labels = FALSE,
      quantiles = TRUE, rev.x = TRUE, type = "l",
      xlim, ylim, xlab = "", ylab = "Dissimilarity",
      main = "", sub = NULL,
      col.quantile = "red", lty.quantile = "dotted",
      ...)
```

Arguments

x	an object of class "minDC".
depths	numeric; a vector of depths for which predicted values exist or will be generated. Can be missing, in which case, if use.labels = TRUE, the function will attempt to derive suitable values for you. See Details below.
use.labels	logical; should reconPlot attempt to derive values for argument depths from the names of the predicted values? Only use if depths is missing. See Details below.
quantiles	logical; should the probability quantiles be drawn on the plot?
rev.x	logical; should the depth/age axis be reversed (drawn from high to low)?
type	type of line drawn. See par and argument "type".
xlab, ylab	character; the x- and y-axis labels respectively.
main, sub	character; main title and subtitle for the plot.
xlim, ylim	numeric, length 2; the x- and y-limits for the plotted axes. If not provided, the function will calculate appropriate values to cover the range of plotted values and any quantile lines (if requested via "quantiles").
col.quantile	colour in which to draw the quantile lines.
lty.quantile	line type in which to draw the quantile lines.
...	arguments to be passed to methods, such as graphical parameters (see par). Currently ignored.

Details

Conventionally, these plots are drawn on a depth or an age scale. Argument depths is used to provide the depth or age axis, against which the predicted values are plotted.

If depths is not provided, then the function will try to derive the appropriate values from the labels of the predictions if use.labels = TRUE. You must provide depths or set use.labels = TRUE otherwise an error will result. The derived labels will be coerced to numerics. If your labels are coercible, then you'll either get nonsense on the plot or an error from R. If so, provide suitable values for depths.

Value

A plot on the currently active device.

Author(s)

Gavin L. Simpson

See Also

[minDC](#)

Examples

```
## Imbrie and Kipp example
## load the example data
data(ImbrieKipp)
data(SumSST)
data(V12.122)

## merge training and test set on columns
dat <- join(ImbrieKipp, V12.122, verbose = TRUE)

## extract the merged data sets and convert to proportions
ImbrieKipp <- dat[[1]] / 100
V12.122 <- dat[[2]] / 100

## fit the MAT model using the chord distance measure
(ik.mat <- mat(ImbrieKipp, SumSST, method = "chord"))

## reconstruct for the RLGH core data
v12.mat <- predict(ik.mat, V12.122)

## extract the minimum DC values
v12.mdc <- minDC(v12.mat)
v12.mdc

## draw a plot of minimum DC by time
plot(v12.mdc, use.labels = TRUE, xlab = "Depth (cm.)")
```

plot.pcurve

Plot a fitted principal curve in PCA space

Description

Projects the principal curve into PCA space and draws it and the underlying data in a biplot.

Usage

```
## S3 method for class 'prcurve'  
plot(x, data, axes = 1:2, seg = TRUE, col.seg = "forestgreen",  
      col.curve = "red", lwd.curve = 2, ...)
```

Arguments

x	an object of class "prcurve".
data	The data the principal curve was fitted to.
axes	numeric vector of length 2; this is passed to the choices argument of the scores function.
seg	logical; should segments be drawn between the observed points to the location on the principal curve on to which they project.
col.seg	The colour to draw the segments in.
col.curve	The colour to draw the principal curve in.
lwd.curve	The line thickness used to draw the principal curve.
...	additional arguments passed on to points when drawing the observations in PCA space.

Value

A plot on the currently active device. The function does not return anything.

Author(s)

Gavin L. Simpson

See Also

[prcurve](#); [rda](#) for the code used to perform the PCA.

Examples

```
## Load the Abernethy Forest data  
data(abernethy)  
  
## Remove the Depth and Age variables  
abernethy2 <- abernethy[, -(37:38)]  
  
## Fit the principal curve using varying complexity of smoothers  
## for each species  
aber.pc2 <- prcurve(abernethy2, method = "ca", trace = TRUE,  
                    vary = TRUE, penalty = 1.4)  
  
## Plot the curve  
plot(aber.pc2, abernethy2)
```

plot.residLen *Plot method for residual lengths*

Description

Base graphics plot method for [residLen](#) objects.

Usage

```
## S3 method for class 'residLen'  
plot(x, probs = c(0.9, 0.95, 0.99), ncol = 1,  
      lcol = "red", llty = "dashed", xlab = NULL, ylab = NULL,  
      main = "Residual distances", rug = TRUE, ...)
```

Arguments

x	Object of class "residLen", the result of a call to residLen .
probs	numeric; vector of probability quantiles to compute from the sets of residual distances.
ncol	numeric; number of columns for the plot layout. Choices are 1 or 2. Determines whether the histograms are plotted above or beside each other.
lcol, llty	colour and line-type for the quantiles.
xlab, ylab	Axis labels. If not supplied, suitable defaults are generated, depending on whether RDA or CCA was used as the underlying ordination model.
main	character; title for the plot.
rug	logical; should rug plots of the actual distances be drawn?
...	additional arguments passed to plot.

Value

A plot on the current device.

Returns, invisibly, a list with two components (train and passive), each an object of the type returned by [density](#).

Author(s)

Gavin L. Simpson

See Also

[residLen](#), [plot.residLen](#), [histogram.residLen](#), [densityplot.residLen](#).

Examples

```
## load the Imbrie and Kipp example data
data(ImbrieKipp, SumSST, V12.122)

## squared residual lengths for Core V12.122
rlens <- residLen(ImbrieKipp, SumSST, V12.122)
rlens

## plot a histogram of the residual distances
plot(rlens)
```

plot.roc

Plot ROC curves and associated diagnostics

Description

Produces up to four plots (selectable by "which") from the results of a call to `roc`, including the ROC curve itself.

Usage

```
## S3 method for class 'roc'
plot(x,
      which = c(1:3,5),
      group = "Combined",
      prior = NULL,
      show.stats = TRUE,
      abline.col = "grey",
      abline.lty = "dashed",
      inGroup.col = "red",
      outGroup.col = "blue",
      lty = "solid",
      caption = c("ROC curve", "Dissimilarity profiles",
                  "TPF - FPF vs Dissimilarity",
                  "Likelihood ratios"),
      legend = "topright",
      ask = prod(par("mfcol")) < length(which) && dev.interactive(),
      ...)
```

Arguments

<code>x</code>	an object of class "roc".
<code>which</code>	numeric vector; which aspects of "roc" object to plot if a subset of the plots is required, specify a subset of the numbers 1:5.
<code>group</code>	character vector of length 1 giving the name of the group to plot.

prior	numeric vector of length 2 (e.g. <code>c(0.5, 0.5)</code>) specifying the prior probabilities of analogue and no-analogue. Used to generate posterior probability of analogue using Bayes factors in plot 5 (<code>which = 5</code>).
show.stats	logical; should concise summary statistics of the ROC analysis be displayed on the plot?
abline.col	character string or numeric value; the colour used to draw vertical lines at the optimal dissimilarity on the plots.
abline.lty	Line type for indicator of optimal ROC dissimilarity threshold. See <code>par</code> for the allowed line types.
inGroup.col	character string or numeric value; the colour used to draw the density curve for the dissimilarities between sites in the same group.
outGroup.col	character string or numeric value; the colour used to draw the density curve for the dissimilarities between sites not in the same group.
lty	vector of at most length 2 (elements past the second in longer vectors are ignored) line types. The first element of <code>lty</code> will be used where a single line is drawn on a plot. Where two lines are drawn (for analogue and non-analogue cases), the first element pertains to the analogue group and the second element to the non-analogue group. See <code>par</code> for the allowed line types.
caption	vector of character strings, containing the captions to appear above each plot.
legend	character; position of legends drawn on plots. See Details section in <code>legend</code> for keywords that can be specified.
ask	logical; if TRUE, the user is <i>asked</i> before each plot, see <code>par(ask=.)</code> .
...	graphical arguments passed to other graphics functions.

Details

This plotting function is modelled closely on `plot.lm` and many of the conventions and defaults for that function are replicated here.

First, some definitions:

TPF True Positive Fraction, also known as *sensitivity*.

TNF True Negative Fraction, also known as *specificity*.

FPF False Positive Fraction; the complement of TNF, calculated as $1 - \text{TNF}$. This is often referred to a $1 - \text{specificity}$. A false positive is also known as a type I error.

FNF False Negative Fraction; the complement of TPF, calculated as $1 - \text{TPF}$. A false negative is also known as a type II error.

AUC The Area Under the ROC Curve.

The "ROC curve" plot (`which = 1`), draws the ROC curve itself as a plot of the False Positive Fraction against the True Positive Fraction. A diagonal 1:1 line represents no ability for the dissimilarity coefficient to differentiate between groups. The AUC statistic may also be displayed (see argument "show.stats" above).

The "Dissimilarity profile" plot (`which = 2`), draws the density functions of the dissimilarity values (d) for the correctly assigned samples and the incorrectly assigned samples. A dissimilarity

coefficient that is able to well distinguish the sample groupings will have density functions for the correctly and incorrectly assigned samples that have little overlap. Conversely, a poorly discriminating dissimilarity coefficient will have density profiles for the two assignments that overlap considerably. The point where the two curves cross is the optimal dissimilarity or critical value, d' . This represents the point where the difference between TPF and FPF is maximal. The value of d at the point where the difference between TPF and FPF is maximal will not necessarily be the same as the value of d' where the density profiles cross. This is because the ROC curve has been estimated at discrete points d , which may not include exactly the optimal d' , but which should be close to this value if the ROC curve is not sampled on too coarse an interval.

The "TPF - FPF vs Dissimilarity" plot, draws the difference between the ROC curve and the 1:1 line. The point where the ROC curve is farthest from the 1:1 line is the point at which the ROC curve has maximal slope. This is the optimal value for d , as discussed above.

The "Likelihood ratios" plot, draws two definitions of the slope of the ROC curve as the likelihood functions LR(+), and LR(-). LR(+), is the likelihood ratio of a positive test result, that the value of d assigns the sample to the group it belongs to. LR(-) is the likelihood ratio of a negative test result, that the value of d assigns the sample to the wrong group.

LR(+) is defined as $LR(+) = TPF/FPF$ (or sensitivity / (1 - specificity)), and LR(-) is defined as $LR(-) = FPF/TNF$ (or (1 - sensitivity) / specificity), in Henderson (1993).

The "probability of analogue" plot, draws the posterior probability of analogue given a dissimilarity. This is the LR(+) likelihood ratio values multiplied by the prior odds of analogue, for given values of the dissimilarity, and is then converted to a probability.

Value

One or more plots, drawn on the current device.

Author(s)

Gavin L. Simpson. Code borrows heavily from [plot.lm](#).

References

Brown, C.D., and Davis, H.T. (2006) Receiver operating characteristics curves and related decision measures: A tutorial. *Chemometrics and Intelligent Laboratory Systems* **80**, 24–38.

Gavin, D.G., Oswald, W.W., Wahl, E.R. and Williams, J.W. (2003) A statistical approach to evaluating distance metrics and analog assignments for pollen records. *Quaternary Research* **60**, 356–367.

Henderson, A.R. (1993) Assessing test accuracy and its clinical consequences: a primer for receiver operating characteristic curve analysis. *Annals of Clinical Biochemistry* **30**, 834–846.

See Also

[roc](#) for a complete example

plot.wa

*Plot diagnostics for a weighted averaging model***Description**

Two plots (selectable by which) are currently available: a plot of estimated against observed values, a plot of residuals against estimated values.

Usage

```
## S3 method for class 'wa'
plot(x,
      which = 1:2,
      caption = c("Inferred vs Observed", "Residuals vs Fitted"),
      max.bias = TRUE,
      n.bias = 10,
      sub.caption = NULL,
      main = "",
      ask = prod(par("mfcol")) < length(which) &&
            dev.interactive(),
      ...,
      panel = if (add.smooth) panel.smooth else points,
      add.smooth = getOption("add.smooth"))
```

Arguments

x	an object of class "wa".
which	which aspects of the "wa" object to plot if a subset of the plots is required, specify a subset of the numbers 1:2.
caption	captions to appear above the plots.
max.bias	logical, should max bias lines be added to residuals.
n.bias	numeric, number of sections to calculate maximum bias for.
sub.caption	common title-above figures if there are multiple; used as 'sub' (s.'title') otherwise. If NULL, as by default, a possibly shortened version of deparse(x\$call) is used.
main	title to each plot-in addition to the above caption.
ask	logical; if TRUE, the user is <i>asked</i> before each plot, see par(ask=.).
...	graphical arguments passed to other graphics functions.
panel	panel function. The useful alternative to points, panel.smooth, can be chosen by add.smooth = TRUE.
add.smooth	logical indicating if a smoother should be added to fitted & residuals plots; see also panel above.

Details

This plotting function is modelled closely on [plot.lm](#) and many of the conventions and defaults for that function are replicated here.

`sub.caption` - by default the function call - is shown as a subtitle (under the x-axis title) on each plot when plots are on separate pages, or as a subtitle in the outer margin (if any) when there are multiple plots per page.

Value

One or more plots, drawn on the current device.

Author(s)

Gavin L. Simpson. Code borrows heavily from [plot.lm](#).

See Also

[mat](#)

Examples

```
## continue the RLGH example from ?wa
example(wa)

## diagnostics for the WA model
par(mfrow = c(1,2))
plot(mod)
par(mfrow = c(1,1))
```

Pollen

North American Modern Pollen Database

Description

A database of model pollen samples from a network of sites from North America and Greenland, compiled by Whitmore et al. (2005). Associated climatic and vegetation data are also record.

Usage

```
data(Pollen)
```

```
data(Climate)
```

```
data(Biome)
```

```
data(Location)
```

Format

For **Pollen**, a data frame of 4833 samples and 134 pollen taxa.

For **Biome**, a data frame of 4833 samples on, currently, a single variable:

Fedorova Factor; Vegetation type (**Biome**) See Whitmore et al. (2005), Figs 3 & 4. Reclassified biomes from Fedorova et al (1994).

For **Location**, a data frame of the latitude and longitude locational data for 4833 samples.

Latitude Latitude of the sampling location in decimal degrees.

Longitude Longitude of the sampling location in decimal degrees.

For **Climate**, a data frame with 4833 observations on the following 32 variables.

t[*jan:dec*] numeric vectors; Mean monthly temperatures for the indicated month. Degrees C.

p[*jan:dec*] numeric vectors; Mean total monthly precipitation (mm) for the indicated month.

tave numeric; annual average temperature (Degrees C)

tmax numeric; maximum temperature (in Degrees C) observed over the period of record.

tmin numeric; minimum temperature (in Degrees C) observed over the period of record.

gdd0 numeric; Growing degree days computed using a base of 0 degrees C.

gdd5 numeric; Growing degree days computed using a base of 5 degrees C.

mtco numeric; mean temperature of the coldest month.

mtwa numeric; mean temperature of the warmest month.

annp numeric; mean annual total precipitation (mm).

Details

These datasets were extracted from Version 1.7 of the North American Modern Pollen Database.

All pollen species were included, however, only the Vegetation type (**Biome**) field of the AVHRR data and selected Climatic variables were extracted. Requests for additional variables to be included in the versions of the data included in the package should be sent to the package maintainer.

Warning

Note that the data for the pollen species are a mixture of types. The **DataForm** variable contains information on the type of data included for each site. The codes are:

RC raw counts

RP raw counts expressed as percentages

DC digitised counts

DP digitised counts expressed as percentages

PM counts in permille

This value is not known for all samples.

Source

The database is archived electronically at two sites:

<http://www.lpc.uottawa.ca/data/modern/index.html>

<http://www.geography.wisc.edu/faculty/williams/web/data.htm>

References

Whitmore, J., Gajewski, K., Sawada, M., Williams, J.W., Shuman, B., Bartlein, P.J., Minckley, T., Viau, A.E., Webb III, T., Anderson, P.M., and Brubaker L.B., 2005. North American and Greenland modern pollen data for multi-scale paleoecological and paleoclimatic applications. *Quaternary Science Reviews* **24**:1828–1848.

Williams, J.W., Shuman, B., Bartlein, P.J., Whitmore, J., Gajewski, K., Sawada, M., Minckley, T., Shafer, S., Viau, A.E., Webb, III, T., Anderson, P.M., Brubaker, L.B., Whitlock, C. and Davis, O.K., 2006. *An Atlas of Pollen-Vegetation-Climate Relationships for the United States and Canada*. American Association of Stratigraphic Palynologists Foundation, Dallas, TX, 293p.

Williams, J.W. and Shuman, B., 2008. Obtaining accurate and precise environmental reconstructions from the modern analog technique and North American surface pollen dataset. *Quaternary Science Reviews*, **27**: 669–687.

Fedorova, I.T., Volkova, Y.A., Varlyguin, E., 1994. *World vegetation cover. Digital raster data on a 30-minute cartesian orthonormal geodetic (lat/long) 1080x2160 grid*. In: Global Ecosystems Database Version 2.0. USDOC/NOAA National Geophysical Data Center, Boulder, CO.

Examples

```
data(Pollen)
```

```
data(Climate)
```

```
data(Biome)
```

```
data(Location)
```

prcurve

Fits a principal curve to m-dimensional data

Description

A principal curve is a non-parametric generalisation of the principal component and is a curve that passes through the *middle* of a cloud of data points for a certain definition of ‘middle’.

Usage

```
prcurve(X, method = c("ca", "pca", "random", "user"), start = NULL,
        smoother = smoothSpline, complexity, vary = FALSE,
        maxComp, finalCV = FALSE, axis = 1, rank = FALSE,
        stretch = 2, maxit = 10, trace = FALSE, thresh = 0.001,
        plotit = FALSE, ...)

initCurve(X, method = c("ca", "pca", "random", "user"), rank = FALSE,
          axis = 1, start)

smoothSpline(lambda, x, choose = TRUE, complexity, ...,
             penalty = 1, cv = FALSE, keep.data = FALSE,
             control.spar = list(low = 0))
```

Arguments

<code>X</code>	a matrix-like object containing the variables to which the principal curve is to be fitted.
<code>method</code>	character; method to use when initialising the principal curve. "ca" fits a correspondence analysis to <code>X</code> and uses the <code>axis</code> -th axis scores as the initial curve. "pca" does the same but fits a principal components analysis to <code>X</code> . "random" produces a random ordering as the initial curve.
<code>start</code>	numeric vector specifying the initial curve when <code>method = "user"</code> . Must be of length <code>nrow(X)</code> .
<code>smoother</code>	function; the choice of smoother used to fit the principal curve. Currently, the only option is <code>smoothSpline</code> which is a wrapper to smooth.spline .
<code>complexity</code>	numeric; the complexity of the fitted smooth functions. The function passed as argument <code>smoother</code> should arrange for this argument to be passed on to relevant aspect of the underlying smoother. In the case of <code>smoothSpline</code> , <code>complexity</code> is the <code>df</code> argument of smooth.spline .
<code>vary</code>	logical; should the complexity of the smoother fitted to each variable in <code>X</code> be allowed to vary (i.e. to allow a more or less smooth function for a particular variable. If <code>FALSE</code> the median complexity over all m variables is chosen as the fixed complexity for all m smooths.
<code>maxComp</code>	numeric; the upper limit on the allowed complexity.
<code>finalCV</code>	logical; should a final fit of the smooth function be performed using cross validation?
<code>axis</code>	numeric; the ordinaion axis to use as the initial curve.
<code>rank</code>	logical; should rank position on the gradient be used? Not yet implemented.
<code>stretch</code>	numeric; a factor by which the curve can be extrapolated when points are projected. Default is 2 (times the last segment length).
<code>maxit</code>	numeric; the maximum number of iterations.
<code>trace</code>	logical; print progress on the iterations be printed to the console?

thresh	numeric; convergence threshold on shortest distances to the curve. The algorithm is considered to have converged when the latest iteration produces a total residual distance to the curve that is within thresh of the value obtained during the previous iteration.
plotit	logical; should the fitting process be plotted? If TRUE, then the fitted principal curve and observations in X are plotted in principal component space.
...	arguments passed on to lower functions. In the case of prcurve, these additional arguments are passed solely on to the function smoother. In smoothSpline, ... is passed on to the underlying function smooth.spline and users should read that function's help page for further details.
lambda	the current projection function; the position that each sample projects to on the current principal curve. This is the predictor variable or covariate in the smooth function.
x	numeric vector; a column from X used as the response variable in the smooth function. The principal curve algorithm fits a separate scatterplot smoother (or similar smoother) to each variable in X in turn as the response.
choose	logical; should the underlying smoother function be allowed to choose the degree of smooth complexity for each variable in X?
penalty, cv, keep.data, control.spar	arguments to smooth.spline .

Value

An object of class "prcurve" with the following components:

s	a matrix corresponding to X, giving their projections onto the curve.
tag	an index, such that s[tag,] is smooth.
lambda	for each point, its arc-length from the beginning of the curve.
dist	the sum-of-squared distances from the points to their projections.
converged	logical; did the algorithm converge?
iter	numeric; the number of iterations performed.
totalDist	numeric; total sum-of-squared distances.
complexity	numeric vector; the complexity of the smoother fitted to each variable in X.
call	the matched call.

Note

The fitting function uses function [get.lam](#) in package **princurve** to find the projection of the data on to the fitted curve.

Author(s)

Gavin L. Simpson

Examples

```

data(abernethy)

## Plot the most common taxa
StratipLOT(Age ~ . - Depth, data =
  chooseTaxa(abernethy, max.abun = 15, n.occ = 10),
  type = c("g", "poly"), sort = "wa")

## Remove the Depth and Age variables
abernethy2 <- abernethy[, -(37:38)]

## Fit PCA and CA
aber.pca <- rda(abernethy2)
aber.ca <- cca(abernethy2)

## Fit the principal curve using the median complexity over
## all species
aber.pc <- prcurve(abernethy2, method = "ca", trace = TRUE,
  vary = FALSE, penalty = 1.4)

## Fit the principal curve using varying complexity of smoothers
## for each species
aber.pc2 <- prcurve(abernethy2, method = "ca", trace = TRUE,
  vary = TRUE, penalty = 1.4)

```

predict.mat

Predict method for Modern Analogue Technique models

Description

Predicted values based on a MAT model object.

Usage

```

## S3 method for class 'mat'
predict(object, newdata, k, weighted = FALSE,
  bootstrap = FALSE, n.boot = 1000,
  probs = c(0.01, 0.025, 0.05, 0.1), ...)

```

Arguments

object	an object of mat .
newdata	data frame; required only if predictions for some new data are required. Mst have the same number of columns, in same order, as x in mat . See example below or join for how to do this. If newdata not provided, the fitted values are returned.
k	number of analogues to use. If missing, k is chosen automatically as the k that achieves lowest RMSE.

weighted	logical; should the analysis use the weighted mean of environmental data of analogues as predicted values?
bootstrap	logical; should bootstrap derived estimates and sample specific errors be calculated-ignored if newdata is missing.
n.boot	numeric; the number of bootstrap samples to take.
probs	numeric; vector of probabilities with values in [0,1].
...	arguments passed to of from other methods.

Details

This function returns one or more of three sets of results depending on the supplied arguments:

Fitted values: the fitted values of the `mat` model are returned if `newdata` is missing.

Predicted values: the predicted values for some new samples are returned if `newdata` is supplied. Summary model statistics and estimated values for the training set are also returned.

Bootstrap predictions and standard errors: if `newdata` is supplied and `bootstrap = TRUE`, the predicted values for `newdata` plus bootstrap estimates and standard errors for the new samples and the training set are returned.

The latter is simply a wrapper for `bootstrap(model, newdata, ...)` - see `bootstrap.mat`.

Value

A object of class `predict.mat` is returned if `newdata` is supplied, otherwise an object of `fitted.mat` is returned. If `bootstrap = FALSE` then not all components will be returned.

observed	vector of observed environmental values.
model	a list containing the model or non-bootstrapped estimates for the training set. With the following components: <ul style="list-style-type: none"> • <code>estimated</code> estimated values for "y", the environment. • <code>residuals</code> model residuals. • <code>r.squared</code> Model R^2 between observed and estimated values of "y". • <code>avg.bias</code> Average bias of the model residuals. • <code>max.bias</code> Maximum bias of the model residuals. • <code>rmsep</code> Model error (RMSEP). • <code>k</code> numeric; indicating the size of model used in estimates and predictions.
bootstrap	a list containing the bootstrap estimates for the training set. With the following components: <ul style="list-style-type: none"> • <code>estimated</code> Bootstrap estimates for "y". • <code>residuals</code> Bootstrap residuals for "y". • <code>r.squared</code> Bootstrap derived R^2 between observed and estimated values of "y". • <code>avg.bias</code> Average bias of the bootstrap derived model residuals. • <code>max.bias</code> Maximum bias of the bootstrap derived model residuals. • <code>rmsep</code> Bootstrap derived RMSEP for the model.

	<ul style="list-style-type: none"> • s1Bootstrap derived S1 error component for the model. • s2Bootstrap derived S2 error component for the model. • knumeric; indicating the size of model used in estimates and predictions.
sample.errors	a list containing the bootstrap-derived sample specific errors for the training set. With the following components: <ul style="list-style-type: none"> • rmsepBootstrap derived RMSEP for the training set samples. • s1Bootstrap derived S1 error component for training set samples. • s2Bootstrap derived S2 error component for training set samples.
weighted	logical; whether the weighted mean was used instead of the mean of the environment for k -closest analogues.
auto	logical; whether "k" was chosen automatically or user-selected.
n.boot	numeric; the number of bootstrap samples taken.
predictions	a list containing the model and bootstrap-derived estimates for the new data, with the following components: <ul style="list-style-type: none"> • observedthe observed values for the new samples — only if newenv is provided. • modela list containing the model or non-bootstrapped estimates for the new samples. A list with the same components as model, above. • bootstrapa list containing the bootstrap estimates for the new samples, with some or all of the same components as bootstrap, above. • sample.errorsa list containing the bootstrap-derived sample specific errors for the new samples, with some or all of the same components as sample.errors, above.
method	the dissimilarity measure used to fit the underlying MAT models.
quantiles	probability quantiles of the pairwise dissimilarities computed from the training set.
minDC	smallest distances between each sample in newdata and the training set samples.
Di _j	dissimilarities between newdata and training set samples.

Author(s)

Gavin L. Simpson

References

Birks, H.J.B., Line, J.M., Juggins, S., Stevenson, A.C. and ter Braak, C.J.F. (1990). Diatoms and pH reconstruction. *Philosophical Transactions of the Royal Society of London; Series B*, **327**; 263–278.

See Also

[mat](#), [bootstrap.mat](#)

Examples

```
## Imbrie and Kipp example
## load the example data
data(ImbrieKipp)
data(SumSST)
data(V12.122)

## merge training and test set on columns
dat <- join(ImbrieKipp, V12.122, verbose = TRUE)

## extract the merged data sets and convert to proportions
ImbrieKipp <- dat[[1]] / 100
V12.122 <- dat[[2]] / 100

## fit the MAT model using the chord distance measure
(ik.mat <- mat(ImbrieKipp, SumSST, method = "chord"))

## predict for V12.122 data
predict(ik.mat, V12.122)
```

predict.wa

Predict from a weighted average model

Description

Model predictions and cross-validation predictions for weighted averaging transfer function models.

Usage

```
## S3 method for class 'wa'
predict(object, newdata,
        CV = c("none", "L00", "bootstrap", "nfold"),
        verbose = FALSE, n.boot = 100, nfold = 5, ...)
```

Arguments

object	an object of class "wa", usually the result of a call to wa
newdata	An optional data frame in which to look for variables with which to predict. If omitted, the fitted values are used.
CV	Should cross-validation be performed? Leave-one-out ("L00"), bootstrap ("bootstrap") and <i>k</i> -fold ("nfold") CV are currently available.
verbose	Should CV progress be printed to the console?
n.boot	The number of bootstrap samples or <i>k</i> -fold steps.
nfold	Number of subsets in <i>k</i> -fold CV.
...	further arguments passed to or from other methods.

Details

Not all CV methods produce the same output. CV = "bootstrap" and CV = "nfold" produce sample specific errors.

Value

An object of class "predict.wa", a list with the following components:

pred	A list with components pred and rmsep containing the predicted values and the sample specific errors if available.
performance	A list with model performance statistics.
model.pred	A list with components pred and rmsep containing the predicted values for the training set samples and the sample specific errors if available.
call	the matched function call.
CV.method	The CV method used.

Author(s)

Gavin L. Simpson and Jari Oksanen (*k*-fold CV)

References

Birks, H.J.B., Line, J.M., Juggins, S., Stevenson, A.C. and ter Braak, C.J.F. (1990). Diatoms and pH reconstruction. *Philosophical Transactions of the Royal Society of London; Series B*, **327**; 263–278.

See Also

[wa](#), [predict.mat](#), [performance](#), [reconPlot](#).

Examples

```
## Not run:
## Imbrie and Kipp
data(ImbrieKipp)
data(SumSST)
ik.wa <- wa(SumSST ~ ., data = ImbrieKipp, tol.dw = TRUE,
           min.tol = 2, small.tol = "min")
ik.wa

## load V12.122 core data
data(V12.122)
V12.122 <- V12.122 / 100

## predict summer sea-surface temperature for V12.122 core
v12.pred <- predict(ik.wa, V12.122, CV = "bootstrap", n.boot = 100)

## draw the fitted reconstruction
reconPlot(v12.pred, use.labels = TRUE, display = "bars")

## extract the model performance stats
```

```
performance(v12.pred)

## End(Not run)
```

reconPlot

Stratigraphic plots of palaeoenvironmental reconstructions

Description

Draws a palaeoenvironmental reconstruction of predicted environmental values for sub-fossil assemblages.

Usage

```
reconPlot(x, ...)
```

Default S3 method:

```
reconPlot(x, depths, errors,
          display.error = c("none", "bars", "lines"),
          rev.x = TRUE,
          col.error = "grey", lty.error = "dashed",
          type = "l",
          xlim, ylim,
          xlab = "", ylab = "", main = "",
          ...)
```

S3 method for class 'predict.mat'

```
reconPlot(x, depths, use.labels = FALSE,
          predictions = c("model", "bootstrap"),
          display.error = c("none", "bars", "lines"),
          sample.specific = TRUE, ...)
```

S3 method for class 'predict.wa'

```
reconPlot(x, depths, use.labels = FALSE,
          display.error = c("none", "bars", "lines"),
          sample.specific = TRUE, ...)
```

Arguments

x	An R object.
depths	numeric; a vector of depths for which predicted values exist or will be generated. Can be missing, in which case, if <code>use.labels = TRUE</code> , the function will attempt to derive suitable values for you. See Details below.
errors	numeric; a vector of errors for plotting error bars or lines.

<code>display.error</code>	character; hown should error bars be drawn on the plot? One of "none", "bars", or "lines". If "bars", error bars are drawn for each sample. If "lines", lines are drawn enclosing the region prediction +/- RMSEP.
<code>rev.x</code>	logical; should the depth/age axis be reversed (drawn from high to low)?
<code>col.error, lty.error</code>	the colour and type of line drawn. See par and arguments "col" and "lty".
<code>type</code>	type of line drawn. See par and argument "type".
<code>xlab, ylab</code>	character; the x- and y-axis labels respectively.
<code>main</code>	character; main title for the plot.
<code>xlim, ylim</code>	numeric, length 2; the x- and y-limits for the plotted axes. If not provided, the function will calculate appropriate values to cover the range of plotted values and any error bars (if requested via "error.bars").
<code>use.labels</code>	logical; should reconPlot attempt to derive values for argument depths from the names of the predicted values? Only use if depths is missing. See Details below.
<code>predictions</code>	character; one of "model" or "bootstrap". Which type of predicted values should be plotted? The model predictions ("model") or the bootstrap-derived predictions ("bootstrap").
<code>sample.specific</code>	logical; should sample specific errors be used? Only for predictions = "bootstrap".
<code>...</code>	arguments passed to other graphics functions.

Details

Conventionally, these plots are drawn on a depth or an age scale. Argument `depths` is used to provide the depth or age axis, against which the predicted values are plotted.

If `depths` is not provided, then the function will try to derive the appropriate values from the labels of the predictions if `use.labels = TRUE`. You must provide `depths` or set `use.labels = TRUE` otherwise an error will result. The derived labels will be coerced to numerics. If your labels are not coercible, then you'll either get nonsense on the plot or an error from R. If so, provide suitable values for `depths`.

Value

A plot on the currently active device.

Author(s)

Gavin L. Simpson

See Also

[mat](#), and [predict.mat](#) for MAT transfer functions and [wa](#) and [predict.wa](#) for WA models.

Examples

```
## Imbrie and Kipp example
## load the example data
data(ImbrieKipp)
data(SumSST)
data(V12.122)

## merge training and test set on columns
dat <- join(ImbrieKipp, V12.122, verbose = TRUE)

## extract the merged data sets and convert to proportions
ImbrieKipp <- dat[[1]] / 100
V12.122 <- dat[[2]] / 100

## Fit a MAT model
(ik.mat <- mat(ImbrieKipp, SumSST, method = "chord"))

## Reconstruct pH for the RLGH core
v12.pH <- predict(ik.mat, V12.122)

## draw the reconstruction
reconPlot(v12.pH, use.labels = TRUE, display.error = "bars",
          xlab = "Depth", ylab = "Summer Seas-surface Temperature")
```

residLen

Squared residual length diagnostics

Description

The squared residual length between the fitted values of a constrained ordination and the original species data is one diagnostic for transfer function models.

Usage

```
residLen(X, env, passive, method = c("cca", "rda"))

fittedY(ord, newdata, colsum)

sqrLLinear(Y, fitted)

sqrLUnimodal(Y, colsum, fitted)
```

Arguments

X	data frame; the training set species data.
env	vector; the training set environmental data.
passive	data frame; the passive samples species data.

method	the ordination technique to use. One of "rda" or "cca", with the latter the default.
ord	an ordination object, the result of a call to cca or rda .
newdata	Species data matrix for passive samples. Must have same columns as data used to fit ord.
colsum	column (species) sums for training set data used to fit ord.
Y	Original species data matrix, the response for which squared residual lengths are to be computed.
fitted	The fitted values of the response derived from the constrained ordination model.

Details

The squared residual lengths are computed for the training set samples and the passive samples separately. Passive samples that are poorly fitted in the transfer function model will have large squared residual distances between the observed species data and the fitted values from the constrained ordination.

`residLen` is the main user-interface function and can be called with either the training data and passive samples.

`fittedY` returns the fitted approximation of the passive sample response data (i.e. species data). `sqrlLinear` and `sqrlUnimodal` return the squared residual distances between the observed species data and the fitted values from the constrained ordination model.

Value

`fittedY` returns a matrix of fitted species abundances for passive samples.

`sqrlLinear` and `sqrlUnimodal` return a vector of residual distances.

`residLen` returns an object of class "residLen" with the attribute "method" set to "method". This object is a list with the following components:

train, passive	The squared residual lengths for the training set and the passive samples.
ordination	The fitted ordination.
call	The matched call.

Author(s)

Gavin L. Simpson

References

Ter Braak C.J.F. and Smilauer P. (2002) CANOCO Reference manual and CanoDraw for Windows User's guide: Software for Canonical Ordination (version 4.5). Microcomputer Power (Ithaca, NY, USA), 500pp.

See Also

[cca](#) and [predict.cca](#) for some of the underlying computations.

Examples

```
## load the Imbrie and Kipp example data
data(ImbrieKipp, SumSST, V12.122)

## squared residual lengths for Core V12.122
rlens <- residLen(ImbrieKipp, SumSST, V12.122)
rlens

## as before but using linear RDA
residLen(ImbrieKipp, SumSST, V12.122, method = "rda")
```

rlgh

Round Loch of Glenhead Diatoms

Description

Diatom sub-fossil assemblage counts from the Round Loch of Glenhead.

Usage

```
data(rlgh)
```

Format

A data frame with 101 observations on 139 diatom species.

Details

The samples are taken from various depths down a sediment core retrieved from the Round Loch of Glenhead, Galloway, Scotland.

References

Jones, V.J., Stevenson, A.C. and Battarbee, R.W. (1989) Acidification of lakes in Galloway, south west Scotland — A diatom and pollen study of the post-glacial history of the Round Loch of Glenhead. *Journal of Ecology* **77**(1), 1–23.

Examples

```
data(rlgh)
```

RMSEP	<i>Root mean square error of prediction</i>
-------	---

Description

Calculates or extracts the RMSEP from transfer function models.

Usage

```
RMSEP(object, ...)

## S3 method for class 'mat'
RMSEP(object, k, weighted = FALSE,
       ...)

## S3 method for class 'bootstrap.mat'
RMSEP(object, type = c("birks1990", "standard"),
       ...)

## S3 method for class 'bootstrap.wa'
RMSEP(object, type = c("birks1990", "standard"),
       ...)
```

Arguments

object	An R object.
k	numeric; the number of analogues to use in calculating the RMSEP. May be missing. If missing, k is extracted from the model using getK .
weighted	logical; Return the RMSEP for the weighted or unweighted model? The default is for an unweighted model.
type	The type of RMSEP to return/calculate. See Details, below.
...	Arguments passed to other methods.

Details

There are two forms of RMSEP in common usage. Within palaeoecology, the RMSEP of Birks et al. (1990) is most familiar:

$$\text{RMSEP} = \sqrt{s_1^2 + s_2^2}$$

where s_1 is the standard deviation of the out-of-bag (OOB) residuals and s_2 is the mean bias or the mean of the OOB residuals.

In the wider statistical literature, the following form of RMSEP is more commonly used:

$$\text{RMSEP} = \sqrt{\frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{n}}$$

where y_i are the observed values and \hat{y}_i the transfer function predictions/fitted values.

The first form of RMSEP is returned by default or if type = "birks1990" is supplied. The latter form is returned if type = "standard" is supplied.

The RMSEP for objects of class "mat" is a leave-one-out cross-validated RMSEP, and is calculated as for type = "standard".

Value

A numeric vector of length 1 that is the RMSEP of object.

Author(s)

Gavin L. Simpson

References

Birks, H.J.B., Line, J.M., Juggins, S., Stevenson, A.C. and ter Braak, C.J.F. (1990). Diatoms and pH reconstruction. *Philosophical Transactions of the Royal Society of London; Series B*, **327**; 263–278.

See Also

[mat](#), [bootstrap](#), [wa](#), [bootstrap.wa](#).

Examples

```
## Imbrie and Kipp example
## load the example data
data(ImbrieKipp)
data(SumSST)
data(V12.122)

## merge training and test set on columns
dat <- join(ImbrieKipp, V12.122, verbose = TRUE)

## extract the merged data sets and convert to proportions
ImbrieKipp <- dat[[1]] / 100
V12.122 <- dat[[2]] / 100

## fit the MAT model using the squared chord distance measure
(ik.mat <- mat(ImbrieKipp, SumSST, method = "chord"))

## Leave-one-out RMSEP for the MAT model
RMSEP(ik.mat)

## bootstrap training set
(ik.boot <- bootstrap(ik.mat, n.boot = 100))

## extract the Birks et al (1990) RMSEP
RMSEP(ik.boot)

## Calculate the alternative formulation
```

```
RMSEP(ik.boot, type = "standard")
```

```
roc ROC curve analysis
```

Description

Fits Receiver Operator Characteristic (ROC) curves to training set data. Used to determine the critical value of a dissimilarity coefficient that best discriminates between assemblage-types in palaeoecological data sets, whilst minimising the false positive error rate (FPF).

Usage

```
roc(object, groups, k = 1, ...)

## Default S3 method:
roc(object, groups, k = 1, thin = FALSE,
     max.len = 10000, ...)

## S3 method for class 'mat'
roc(object, groups, k = 1, ...)

## S3 method for class 'analog'
roc(object, groups, k = 1, ...)
```

Arguments

object	an R object.
groups	a vector of group memberships, one entry per sample in the training set data. Can be a factor, and will be coerced to one if supplied vector is not a factor.
k	numeric; the k closest analogues to use to calculate ROC curves.
thin	logical; should the points on the ROC curve be thinned? See Details, below.
max.len	numeric; length of analogue and non-analogue vectors. Used as limit to thin points on ROC curve to.
...	arguments passed to/from other methods.

Details

A ROC curve is generated from the within-group and between-group dissimilarities.

For each level of the grouping vector (groups) the dissimilarity between each group member and its k closest analogues within that group are compared with the k closest dissimilarities between the non-group member and group member samples.

If one is able to discriminate between members of different group on the basis of assemblage dissimilarity, then the dissimilarities between samples within a group will be small compared to the dissimilarities between group members and non group members.

thin is useful for large problems, where the number of analogue and non-analogue distances can conceivably be large and thus overflow the largest number R can work with. This option is also useful to speed up computations for large problems. If `thin == TRUE`, then the larger of the analogue or non-analogue distances is thinned to a maximum length of `max.len`. The smaller set of distances is scaled proportionally. In thinning, we approximate the distribution of distances by taking `max.len` (or a fraction of `max.len` for the smaller set of distances) equally-spaced probability quantiles of the distribution as a new set of distances.

Value

A list with two components; i, `statistics`, a summary of ROC statistics for each level of groups and a combined ROC analysis, and ii, `roc`, a list of ROC objects, one per level of groups. For the latter, each ROC object is a list, with the following components:

<code>TPF</code>	The true positive fraction.
<code>FPE</code>	The false positive error
<code>optimal</code>	The optimal dissimilarity value, assessed where the slope of the ROC curve is maximal.
<code>AUC</code>	The area under the ROC curve.
<code>se.fit</code>	Standard error of the AUC estimate.
<code>n.in</code>	numeric; the number of samples within the current group.
<code>n.out</code>	numeric; the number of samples not in the current group.
<code>p.value</code>	The p-value of a Wilcoxon rank sum test on the two sets of dissimilarities. This is also known as a Mann-Whitney test.
<code>roc.points</code>	The unique dissimilarities at which the ROC curve was evaluated
<code>max.roc</code>	numeric; the position along the ROC curve at which the slope of the ROC curve is maximal. This is the index of this point on the curve.
<code>prior</code>	numeric, length 2. Vector of observed prior probabilities of true analogue and true non-analogues in the group.
<code>analogue</code>	a list with components <code>yes</code> and <code>no</code> containing the dissimilarities for true analogue and true non-analogues in the group.

Author(s)

Gavin L. Simpson, based on code from Thomas Lumley to optimise the calculation of the ROC curve.

References

- Brown, C.D., and Davis, H.T. (2006) Receiver operating characteristics curves and related decision measures: A tutorial. *Chemometrics and Intelligent Laboratory Systems* **80**, 24–38.
- Gavin, D.G., Oswald, W.W., Wahl, E.R. and Williams, J.W. (2003) A statistical approach to evaluating distance metrics and analog assignments for pollen records. *Quaternary Research* **60**, 356–367.
- Henderson, A.R. (1993) Assessing test accuracy and its clinical consequences: a primer for receiver operating characteristic curve analysis. *Annals of Clinical Biochemistry* **30**, 834–846.

See Also

[mat](#) for fitting of MAT models. [bootstrap.mat](#) and [mcarlo](#) for alternative methods for selecting critical values of analogue-ness for dissimilarity coefficients.

Examples

```
## load the example data
data(swapdiat, swappH, rlgh)

## merge training and test set on columns
dat <- join(swapdiat, rlgh, verbose = TRUE)

## extract the merged data sets and convert to proportions
swapdiat <- dat[[1]] / 100
rlgh <- dat[[2]] / 100

## fit an analogue matching (AM) model using the squared chord distance
## measure - need to keep the training set dissimilarities
swap.ana <- analog(swapdiat, rlgh, method = "SQchord",
                  keep.train = TRUE)

## fit the ROC curve to the SWAP diatom data using the AM results
## Generate a grouping for the SWAP lakes
clust <- hclust(as.dist(swap.ana$train), method = "ward")
grps <- cutree(clust, 12)

## fit the ROC curve
swap.roc <- roc(swap.ana, groups = grps)
swap.roc

## draw the ROC curve
plot(swap.roc, 1)

## draw the four default diagnostic plots
layout(matrix(1:4, ncol = 2))
plot(swap.roc)
layout(1)
```

screepLOT

ScreepLOTS of model results

Description

Draws screepLOTS of performance statistics for models of varying complexity.

Usage

```
## S3 method for class 'mat'
```

```

screepplot(x, k, restrict = 20,
           display = c("rmsep", "avg.bias",
                      "max.bias", "r.squared"),
           weighted = FALSE, col = "red", xlab = NULL,
           ylab = NULL, main = NULL, sub = NULL, ...)

## S3 method for class 'bootstrap.mat'
screepplot(x, k, restrict = 20,
           display = c("rmsep", "avg.bias", "max.bias",
                      "r.squared"),
           legend = TRUE, loc.legend = "topright",
           col = c("red", "blue"),
           xlab = NULL, ylab = NULL,
           main = NULL, sub = NULL,
           ...,
           lty = c("solid", "dashed"))

```

Arguments

<code>x</code>	object of class <code>mat</code> and <code>bootstrap.mat</code> .
<code>k</code>	number of analogues to use. If missing 'k' is chosen automatically as the 'k' that achieves lowest RMSE.
<code>restrict</code>	logical; restrict comparison of k-closest model to $k \leq \text{restrict}$.
<code>display</code>	which aspect of x to plot? Partial match.
<code>weighted</code>	logical; should the analysis use weighted mean of env data of analogues as fitted/estimated values?
<code>xlab, ylab</code>	x- and y-axis labels respectively.
<code>main, sub</code>	main and subtitle for the plot.
<code>legend</code>	logical; should a legend be displayed on the figure?
<code>loc.legend</code>	character; a keyword for the location of the legend. See legend for details of allowed keywords.
<code>col</code>	Colours for lines drawn on the screepplot. Method for class "bootstrap.mat" takes a vector of two colours.
<code>lty</code>	vector detailing the line type to use in drawing the screepplot of the apparent and bootstrap statistics, respectively. Code currently assumes that <code>length(lty)</code> is 2.
<code>...</code>	arguments passed to other graphics functions.

Details

Screepplots are often used to graphically show the results of cross-validation or other estimate of model performance across a range of model complexity.

Four measures of model performance are currently available: i) root mean square error of prediction (RMSEP); ii) average bias — the mean of the model residuals; iii) maximum bias — the maximum average bias calculated for each of n sections of the gradient of the environmental variable; and v) model R^2 .

For the maximum bias statistic, the response (environmental) gradient is split into $n = 10$ sections.

For the [bootstrap](#) method, apparent and bootstrap versions of these statistics are available and plotted.

Note

Currently only models of class [mat](#) and [bootstrap.mat](#) are supported.

Author(s)

Gavin Simpson

See Also

[screepplot](#)

Examples

```
## Imbrie and Kipp example
## load the example data
data(ImbrieKipp)
data(SumSST)
data(V12.122)

## merge training and test set on columns
dat <- join(ImbrieKipp, V12.122, verbose = TRUE)

## extract the merged data sets and convert to proportions
ImbrieKipp <- dat[[1]] / 100
V12.122 <- dat[[2]] / 100

## fit the MAT model using the chord distance measure
(ik.mat <- mat(ImbrieKipp, SumSST, method = "chord"))

screepplot(ik.mat)
```

stdError

Standard error of MAT fitted and predicted values

Description

Computes the weighted standard deviation of the environment for the k -closest analogues for each sample. This measure is proposed as a measure of reconstruction uncertainty for MAT models.

Usage

```
stdError(object, ...)  
  
## S3 method for class 'mat'  
stdError(object, k, ...)  
  
## S3 method for class 'predict.mat'  
stdError(object, k, ...)
```

Arguments

object	Object for which the uncertainty measure is to be computed. Currently methods for <code>mat</code> and <code>predict.mat</code> .
k	numeric; how many analogues to take? If missing, the default, k is chosen using <code>getK</code> .
...	Additional arguments passed to other methods. Currently not used.

Value

A named numeric vector of weighted standard deviations of the environment for the k closest analogues used to compute the MAT predicted values.

The returned vector has attributes "k" and "auto", indicating the number of analogues used and whether this was determined from object or supplied by the user.

Author(s)

Gavin L. Simpson

See Also

`minDC`, `mat`, `predict.mat`.

Examples

```
## Imbrie and Kipp Sea Surface Temperature  
data(ImbrieKipp)  
data(SumSST)  
data(V12.122)  
  
## merge training set and core samples  
dat <- join(ImbrieKipp, V12.122, verbose = TRUE)  
  
## extract the merged data sets and convert to proportions  
ImbrieKipp <- dat[[1]] / 100  
ImbrieKippCore <- dat[[2]] / 100  
  
## fit the MAT model using the squared chord distance measure  
ik.mat <- mat(ImbrieKipp, SumSST, method = "SQchord")
```

```
## standard errors
stdError(ik.mat)

## reconstruct for the V12-122 core data
coreV12.mat <- predict(ik.mat, V12.122, k = 3)
## standard errors
stdError(coreV12.mat)
```

Stratiplot

Palaeoecological stratigraphic diagrams

Description

Draws palaeoecological stratigraphic diagrams of one or more variables as a function of depth/age, with the time dimension flowing from the bottom to the top of the y-axis, using the **Lattice** graphics package.

Usage

```
Stratiplot(x, ...)

## Default S3 method:
Stratiplot(x, y, type = "l", ylab = NULL, xlab = "",
           pages = 1, rev = TRUE, ylim, sort = c("none", "wa", "var"),
           svar = NULL, rev.sort = FALSE, strip = FALSE, topPad = 6,
           varTypes = "relative", absoluteSize = 0.5,
           zoneNames = NULL, drawLegend = TRUE, ...)

## S3 method for class 'formula'
Stratiplot(formula, data, subset, na.action, type = "l",
           ylab = NULL, xlab = "", pages = 1, ...)
```

Arguments

x	matrix-like object; the variables to be plotted.
y	numeric vector of depths/ages corresponding to rows in x. Length of y must be the same as nrow(x) or exactly equal to nrow(x) / ncol(x). See Details.
formula	an object of class "formula" (or one that can be coerced to that class): a symbolic description of the model to be fitted. The details of plot specification are given under 'Details'.
type	character; The type of plotting. Can be a vector. Note that not all Lattice 'type's are supported and some new types are allowed. See panel.Stratiplot for further details.
data	an optional data frame, list or environment (or object coercible by as.data.frame to a data frame) containing the variables to plot. If not found in data, the variables are taken from environment(formula), typically the environment from which Stratiplot is called.

subset	an optional vector specifying a subset of observations to be used in the fitting process.
na.action	a function which indicates what should happen when the data contain NAs. The default is set by the na.action setting of <code>options</code> , and is <code>na.fail</code> if that is unset. The ‘factory-fresh’ default is <code>na.omit</code> . Another possible value is <code>NULL</code> , no action. Value <code>na.exclude</code> can be useful.
ylab, xlab	the x- and y-axis labels.
pages	numeric; the number of pages to draw the plot over. May be useful for data sets with many species.
rev	logical; should the y-axis limits be reversed
ylim	user supplied limits for the y-axis (time/depth). If not supplied, suitable limits will be determined from the data. As such, in general use <code>ylim</code> need not be supplied. If you choose to supply your own <code>ylim</code> values, note the default for argument <code>rev</code> ; the default will reverse the values you supply to <code>ylim</code> .
sort	character; how should the variables (columns) of <code>x</code> be sorted on the plot. “wa” sorts by weighted averages of variable <code>svar</code> if not <code>NULL</code> or of <code>y</code> otherwise. The default when “wa” is specified is to order by wiehgted average of the depth/time axis – <code>y</code> . If “var”, then ordering is done as per the order of <code>svar</code> .
svar	vector; optional variable to sort columns of <code>x</code> by.
rev.sort	logical; should the sorting order be reversed.
strip	logical; Should panels have strips containing variable labels drawn on them? Default is <code>FALSE</code> , which labels each panel with a label resulting in a more familiar plot style.
topPad	numeric; additional padding for the top axis to accomodate long variable names. This is a temporary fudge until the actual space required can be automagically calculated from the variable names themselves. The currently gets most of the way there, but <code>topPad</code> is used to add some extra space if required.
varTypes	a character vector of length 1 or equal in length to the number of variables plotted. If length 1, the vector is expanded to the required length. Two values are allowed; i. “relative”, and ii. “absolute”. “relative” treats the indicated variable as a percentage type variable and the panel width is scaled relative to the percentage type variables plotted. “absolute” treats the indicated variable as an absolute variable whose panel width should be independent of the other panels. Use “absolute” for variables that are not species compositional data or other percentage type data.
absoluteSize	numeric, length 1. This controls the width of panels for variables marked as “absolute”, and is the proportion of the largest non-“absolute” panel.
zoneNames	character vector of labels, one per zone, with which to label the zone legend, if drawn (see argument <code>drawLegend</code>). See Details.
drawLegend	logical; should a legend for the zones
...	additional arguments passed to <code>panel.StratipLOT</code> and the underlying <code>xyplot</code> function.

Details

The function now includes preliminary code to handle both relative (proportional or percentage data) and absolute data types, and mixtures thereof. Mixtures can be specified by supplying a vector of types to `varTypes`, in the same order as the variables are drawn on the plot.

Plots can be specified symbolically using a formula. A typical model has the form $Y \sim \text{variables}$, where Y is either the core depths or sample ages/dates (to be plotted on the y-axis) and `variables` is a series of terms which specifies the variables to plot against Y . Terms should be specified with the form `var1 + var2 + var3` to plot only those variables. Other, standard, notation for formulae apply, such as model formulae used in `lm`.

Note that formula is **not** passed on to `xyplot`. Instead, the formula is parsed and evaluated within `StratipLOT` and an appropriate data structure formed to facilitate plotting via `xyplot`. As such, the special features of **Lattice** formulae cannot be used.

If zones are drawn on the stratigraphic plot, the `zoneNames` argument can be used to supply a set of names with which to label the zones using a legend. This legend is drawn on the right-hand side of the stratigraphic diagram if `drawLegend = TRUE` is supplied. The `zoneNames` must be supplied in stratigraphic order, as that is the order in which they are drawn in the legend. Whether this ordering is reversed or not will depend on the value of argument `rev`. It is up to the user to provide the labels in the correct order. Zones are specified by the zone boundaries (excluding the core sequence top and bottom), and as a result 1 more label is required than the number of zone boundaries supplied. If no `zoneNames` is not supplied, but a legend is requested, suitable names will be produced. If you do not wish to have any labelling at all, then set `zoneNames = ""` as this will get recycled to the correct length. See the Example section for an illustration of how this drawing zones works.

Value

Returns a Lattice plot object of class "trellis".

Note

The function currently doesn't know about ages/dates and will interpret these as 'depths' instead. This will be fixed in a future version.

Author(s)

Gavin L. Simpson.

See Also

[xyplot](#), [panel.StratipLOT](#), [panel.Loess](#).

Examples

```
data(V12.122)
Depths <- as.numeric(rownames(V12.122))

(plt <- StratipLOT(Depths ~ 0.univ + G.ruber + G.tene1 + G.pacR,
  data = V12.122, type = c("h", "l", "g", "smooth")))

## Order taxa by WA in depth --- ephasises change over time
```

```

(plt <- StratipLOT(Depths ~ 0.univ + G.ruber + G.tenel + G.pacR,
  data = V12.122, type = c("h"), sort = "wa"))

## Using the default interface
spp.want <- c("0.univ", "G.ruber", "G.tenel", "G.pacR")
(plt <- StratipLOT(V12.122[, spp.want], y = Depths,
  type = c("poly", "g")))

## Adding zones to a Stratigraphic plot
## Default labelling and draw zone legend
## Here we choose 4 arbitrary Depths as the zone boundaries
set.seed(123)
Zones <- sample(Depths, 4)
(plt <- StratipLOT(Depths ~ 0.univ + G.ruber + G.tenel + G.pacR,
  data = V12.122, type = c("poly", "g"),
  zones = Zones))

## As before, but supplying your own zone labels
zone.labs <- c("A", "B", "C", "D", "E")
(plt <- StratipLOT(Depths ~ 0.univ + G.ruber + G.tenel + G.pacR,
  data = V12.122, type = c("poly", "g"),
  zones = Zones, zoneNames = zone.labs))

## Suppress the drawing of the zone legend
(plt <- StratipLOT(Depths ~ 0.univ + G.ruber + G.tenel + G.pacR,
  data = V12.122, type = c("poly", "g"),
  zones = Zones, drawLegend = FALSE))

## Add zones and draw a legend, but do not label the zones
(plt <- StratipLOT(Depths ~ 0.univ + G.ruber + G.tenel + G.pacR,
  data = V12.122, type = c("poly", "g"),
  zones = Zones, zoneNames = ""))

```

summary.analog

Summarise analogue matching results

Description

[summary](#) method for class "analog".

Usage

```

## S3 method for class 'analog'
summary(object, display = c("dist", "names", "quantiles"),
  k = 10, probs = c(0.01, 0.02, 0.05, 0.1, 0.2), ...)

```

Arguments

object an object of class "analog", usually as a result of a call to [analog](#).

display	character; one or more of the listed choices. Determines which aspects of the analog results are summarised.
k	number of analogues to use. If missing, k is chosen automatically as the k that achieves lowest RMSE.
probs	numeric; giving the probabilities of the distribution to return quantiles for. See quantile .
...	arguments passed to or from other methods.

Value

A list with one or more of the components listed below. Attributes "method", "train", "call" and "k" contain the dissimilarity coefficient used, whether the training set dissimilarities were saved, the matched function call and the number of close analogues to return respectively.

dists	a matrix of dissimilarities between training set samples and fossil samples. The number of rows is given by argument k. There is a column for each fossil sample.
names	a matrix of names of samples from the training set that are analogues for each fossil sample. The number of rows is given by argument k. There is a column for each fossil sample.
quantiles	numeric; the quantiles of the distribution of the pairwise dissimilarities for the training set for probabilities prob.

Author(s)

Gavin L. Simpson

See Also

[analog](#).

Examples

```
## Not run:
## continue the RLGH example from ?join
example(join)

## analog matching between SWAP and RLGH core
swap.analog <- analog(swapdiat, rlgh, method = "chord")
swap.analog
summary(swap.analog)

## End(Not run)
```

summary.bootstrap.mat *Summarise bootstrap resampling for MAT models*

Description

`summary` method for class "bootstrap.mat".

Usage

```
## S3 method for class 'bootstrap.mat'
summary(object, ...)
```

Arguments

`object` an object of class "bootstrap.mat", usually the result of a call to `bootstrap.mat`.
`...` arguments passed to or from other methods.

Value

A data frame with the following components:

<code>observed</code>	vector of observed environmental values.
<code>model</code>	a list containing the apparent or non-bootstrapped estimates for the training set. With the following components: <code>estimated</code> : estimated values for the response <code>residuals</code> : model residuals <code>r.squared</code> : Apparent R^2 between observed and estimated values of y <code>avg.bias</code> : Average bias of the model residuals <code>max.bias</code> : Maximum bias of the model residuals <code>rmse</code> : Apparent error (RMSE) for the model <code>k</code> : numeric; indicating the size of model used in estimates and predictions
<code>bootstrap</code>	a list containing the bootstrap estimates for the training set. With the following components: <code>estimated</code> : Bootstrap estimates for the response <code>residuals</code> : Bootstrap residuals for the response <code>r.squared</code> : Bootstrap derived R^2 between observed and estimated values of the response <code>avg.bias</code> : Average bias of the bootstrap derived model residuals <code>max.bias</code> : Maximum bias of the bootstrap derived model residuals <code>rmsep</code> : Bootstrap derived RMSEP for the model <code>s1</code> : Bootstrap derived S1 error component for the model <code>s2</code> : Bootstrap derived S2 error component for the model <code>k</code> : numeric; indicating the size of model used in estimates and predictions

sample.errors a list containing the bootstrap-derived sample specific errors for the training set. With the following components:

- rmsep: Bootstrap derived RMSEP for the training set samples
- s1: Bootstrap derived S1 error component for training set samples
- s2: Bootstrap derived S2 error component for training set samples

weighted logical; whether the weighted mean was used instead of the mean of the environment for k -closest analogues

auto logical; whether k was chosen automatically or user-selected

n.boot numeric; the number of bootstrap samples taken

call the matched call

call model type

predictions a list containing the apparent and bootstrap-derived estimates for the new data, with the following components:

- observed: the observed values for the new samples — only if newenv is provided
- model: a list containing the apparent or non-bootstrapped estimates for the new samples. A list with the same components as apparent, above
- bootstrap: a list containing the bootstrap estimates for the new samples, with some or all of the same components as bootstrap, above
- sample.errors: a list containing the bootstrap-derived sample specific errors for the new samples, with some or all of the same components as sample.errors, above

Author(s)

Gavin L. Simpson

See Also

[bootstrap.mat](#), [mat](#), [summary](#).

Examples

```
## Not run:
## continue the RLGH example from ?join
example(join)

## fit the MAT model using the squared chord distance measure
swap.mat <- mat(swapdiat, swappH, method = "SQchord")

## bootstrap training set
swap.boot <- bootstrap(swap.mat, k = 10, n.boot = 100)
swap.boot
summary(swap.boot)

## End(Not run)
```

summary.cma

*Summarise the extraction of close modern analogues***Description**

`summary` method for class "cma".

Usage

```
## S3 method for class 'cma'
summary(object, ...)
```

Arguments

`object` an object of class "cma", usually the result of a call to `cma`.
`...` arguments passed to or from other methods.

Value

An object of class "summary.cma" with the components of an object of class `cma`, plus:

`distances` a matrix of distances/dissimilarities. Individual columns contain the ordered close modern analogues for individual fossil samples. Rows of this matrix refer to the k^{th} closest analogue for each fossil sample. See notes below.

`samples` a matrix of sample names from the reference set that are close modern analogues for a fossil sample. Individual columns contain the ordered close modern analogues for individual fossil samples. Rows of this matrix refer to the k^{th} closest analogue for each fossil sample. See notes below.

Note

Currently, only objects of class `analog` are supported. The number of rows in the returned matrices is equal to the maximum number of close modern analogues identified for an individual fossil sample. If no close modern analogues exist for an individual fossil sample are identified, then the relevant column in "distances" will contain all missing values and in "samples" the string "none". Rows of individual columns will be padded with missing values if the number of close modern analogues for that sample is less than the maximum number of close modern analogues identified for a single sample.

Author(s)

Gavin L. Simpson

See Also

`cma`

Examples

```
## Not run:
## continue the RLGH example from ?join
example(join)

## analog matching between SWAP and RLGH core
swap.analog <- analog(swapdiat, rlgh, method = "chord")
swap.analog
summary(swap.analog)

## close modern analogues
swap.cma <- cma(swap.analog, cutoff = 0.6)
swap.cma
summary(swap.cma)

## End(Not run)
```

summary.mat

*Summarise Modern Analogue Technique models***Description**

`summary` method for class "mat".

Usage

```
## S3 method for class 'mat'
summary(object, k = 10,
        digits = min(2, getOption("digits") - 4), ...)
```

Arguments

object	an object of class "cma", usually the result of a call to <code>cma</code> .
k	numeric; maximum modern analogues to use to summarise model fits.
digits	numeric; the number of significant digits with which to format results.
...	arguments passed to or from other methods.

Value

A list with the components below. The number of analogues used, k is returned as attribute "k".

summ	a data.frame containing the model fits for training set samples. See notes below.
tbl	matrix of summary statistics for an un-weighted model.
tbl.W	matrix of summary statistics for a weighted model.
call	the matched function call
quantiles	the quantiles of the distribution of pairwise dissimilarities for the training set, for "probs = c(0.01, 0.02, 0.05, 0.1, 0.2)".

Note

The returned component "summ" contains the following:

Obs: the observed responses for the training set samples.

Est: the fitted values of the response for training set samples based on the average of k -closest analogues.

Resi: the residuals of the fitted model based on the average of k -closest analogues.

W.Est: the fitted values of the response for training set samples based on the weighted average of k -closest analogues.

W.Resi: the residuals of the fitted model based on the weighted average of k -closest analogues.

minDC: dissimilarity of closest analogue in training set for each training set sample.

minResi: smallest residual for an un-weighted model of size " k ".

k: size of model leading to minimal residual, "minResi".

minW.Resi: smallest residual for a weighted model of size " $k.W$ ".

k.W: size of model leading to minimal residual, "minW.Resi".

Author(s)

Gavin L. Simpson

See Also

[mat](#), [summary](#).

Examples

```
## Not run:
## continue the RLGH example from ?join
example(join)

## fit the MAT model using the squared chord distance measure
swap.mat <- mat(swapdiat, swappH, method = "SQchord")
swap.mat

## model summary
summary(swap.mat)

## model summary - evaluating models using k = 1, ..., 20
## analogues instead of the default, 10.
summary(swap.mat, k = 20)

## End(Not run)
```

summary.predict.mat *Summarise MAT model predictions*

Description

[summary](#) method for objects of class "predict.mat".

Usage

```
## S3 method for class 'predict.mat'  
summary(object, ...)
```

Arguments

object an object of class "predict.mat", usually the result of a call to [predict.mat](#).
... arguments passed to or from other methods.

Value

An object of class "summary.predict.mat", see [predict.mat](#) for more details.

Author(s)

Gavin L. Simpson

See Also

[predict.mat](#), [mat](#), [bootstrap.mat](#) and [summary](#).

Examples

```
## Not run:  
## continue the RLGH example from ?join  
example(join)  
  
## fit the MAT model using the squared chord distance measure  
swap.mat <- mat(swapdiat, swappH, method = "SQchord")  
  
## predict for RLGH data  
swap.pred <- predict(swap.mat, rlgh, bootstrap = FALSE)  
summary(swap.pred)  
  
## End(Not run)
```

`swapdiat`*SWAP sub-fossil diatom and pH training set*

Description

The Surface Waters Acidification Project (SWAP) Palaeolimnology Programme diatom surface sample training set.

Usage

```
data(swapdiat)
```

Format

A data frame of observations on 277 diatom taxa for the 167-lake SWAP diatom-pH training set.

Details

This data set contains the original 167-lake SWAP diatom-pH training set.

The variable names ([colnames](#)) are DIATCODE codes for individual taxa.

References

Stevenson, A.C., Juggins, S., Birks, H.J.B., Anderson, D.S., Anderson, N.J., Battarbee, R.W., Berge, F., Davis, R.B., Flower, R.J., Haworth, E.Y., Jones, V.J., Kingston, J.C., Kreiser, A.M., Line, J.M., Munro, M.A.R., and Renberg, I. (1995). *The Surface Waters Acidification Project Palaeolimnology programme: modern diatom/lake-water chemistry data-set*. ENSIS Publishing, London.

See Also

[swappH](#)

Examples

```
data(swapdiat)
```

`swappH`*SWAP sub-fossil diatom and pH training set*

Description

The Surface Waters Acidification Project (SWAP) Palaeolimnology Programme diatom-pH surface sample training set.

Usage

```
data(swappH)
```

Format

Numeric vector containing the pH of the 167 lakes from the SWAP diatom-pH training set. The values are the average of 4 quarterly samples.

References

Stevenson, A.C., Juggins, S., Birks, H.J.B., Anderson, D.S., Anderson, N.J., Battarbee, R.W., Berge, F., Davis, R.B., Flower, R.J., Haworth, E.Y., Jones, V.J., Kingston, J.C., Kreiser, A.M., Line, J.M., Munro, M.A.R., and Renberg, I. (1995). *The Surface Waters Acidification Project Palaeolimnology programme: modern diatom/lake-water chemistry data-set*. ENSIS Publishing, London.

See Also

[swapdiat](#)

Examples

```
data(swappH)
str(swappH)
```

`timetrack`*Timetracks of change in species composition*

Description

Project passive (e.g. sediment core) samples into an ordination of a set of training samples.

Usage

```

timetrack(X, passive, env, method = c("cca", "rda"),
          transform = "none", formula, scaling = 3,
          rank = "full", model = c("CCA", "CA"), ...)

## S3 method for class 'timetrack'
fitted(object, type = c("passive", "ordination"),
        model = NULL, ...)

## S3 method for class 'timetrack'
plot(x, choices = 1:2, pch = c(1,2),
      col = c("black","red"), ...)

```

Arguments

<code>X</code>	matrix-like object containing the training set or reference samples.
<code>passive</code>	matrix-like object containing the samples to be projected into the ordination of <code>X</code> . Usually a set of sediment core samples.
<code>env</code>	optional vector or matrix of environmental or constraining variables. If provided, a constrained ordination of <code>X</code> is performed.
<code>method</code>	character, resolving to an ordination function available in vegan . Currently only "cca", the default, and "rda" are supported.
<code>transform</code>	character; the name of the transformation to apply to both <code>X</code> and <code>passive</code> . The transformations are performed using <code>tran</code> and valid options are given by that function's <code>method</code> argument.
<code>formula</code>	a model formula; if provided, it defines the model formula for the ordination function and is supplied as argument <code>formula</code> to the ordination function.
<code>scaling</code>	numeric; the ordination scaling to apply. Useful options are likely to be 1 or 3 where the focus is on the samples.
<code>rank</code>	character; see argument of same name in function cca or rda .
<code>model</code>	character; see argument of same name in function cca or rda .
<code>object, x</code>	an object of class "timetrack".
<code>type</code>	character; which fitted values should be returned?
<code>choices</code>	numeric; the length-2 vector of ordination axes to plot.
<code>pch</code>	The length-2 vector of plotting characters. The first element is used for the ordination samples, the second for the passive samples.
<code>col</code>	The length-2 vector of plotting colours. The first element is used for the ordination samples, the second for the passive samples.
<code>...</code>	arguments passed to other methods. <code>timetrack</code> passes arguments on to <code>tran</code> and the ordination function given in <code>method</code> . <code>fitted</code> passes arguments on to other fitted methods as appropriate. <code>plot</code> passes arguments on to the underlying plotting functions.

Details

The timetrack is a way to visualise changes in species composition from sediment core samples within an underlying reference ordination or, usually, training set samples. This technique has been most often applied in situations where the underlying ordination is a constrained ordination and thence the timetrack of sediment core samples within the ordination reflects both the change in species composition and the indicative changes in the constraining variables.

The sediment core samples are projected passively into the underlying ordination. By projected passively, the locations of the core samples are predicted on the basis of the ordination species scores. A common set of species (columns) is required to passively place the sediment samples into the ordination. To achieve this, the left outer join of the species compositions of the training set and passive set is determined; the left outer join results in the passive data matrix having the same set of species (variables; columns) as the training set. Any training set species not in the passive set are added to the passive set with abundance 0. Any passive species not in the training set are removed from the passive set.

Value

The plot method results in a plot on the currently active device, whilst the fitted method returns the matrix of fitted locations on the set of ordination axes.

timetrack returns an object of class "timetrack", a list with the following components:

ordination	the ordination object, the result of the call to the function of the name method.
fitted.values	the matrix of fitted locations for the passive samples on the ordination axes.
method	the ordination function used.
formula	if supplied, the model formula used to define the ordination model.
scaling	the ordination scaling applied.
rank	The rank or the number of axes used in the approximation. The default is to use all axes (full rank) of the "model".
model	Show constrained ("CCA") or unconstrained ("CA") results.
labels	a list of names for the X, passive, and env arguments.
call	The matched function call.

Author(s)

Gavin L. Simpson

See Also

[cca](#) and [rda](#) for the underlying ordination functions.

Examples

```
## load the RLGH and SWAP data sets
data(rlgh, swapdiat)

## Fit the timetrack ordination
```

```

mod <- timetrack(swapdiat, rlgh, transform = "hellinger",
                method = "rda")
mod

## Plot the timetrack
plot(mod)

```

tran

Common data transformations and standardizations

Description

Provides common data transformations and standardizations useful for palaeoecological data. The function acts as a wrapper to function [decostand](#) in package [vegan](#) for several of the available options.

The formula method allows a convenient method for selecting or excluding subsets of variables before applying the chosen transformation.

Usage

```

## Default S3 method:
tran(x, method, a = 1, b = 0, p = 2, base = exp(1),
     na.rm = FALSE, na.value = 0, ...)

## S3 method for class 'formula'
tran(formula, data = NULL, subset = NULL,
     na.action = na.pass, ...)

```

Arguments

x	A matrix-like object.
method	transformation or standardization method to apply. See Details for available options.
a	Constant to multiply x by. method = "log" only. Can be a vector, in which case the vector of values to multiply each column of x by.
b	Constant to add to x before taking logs. method = "log" only. Can be a vector, in which case the vector of values to add to each column of x.
p	The power to use in the power transformation.
base	the base with respect to which logarithms are computed. See log for further details. The default is to compute natural logarithms.
na.rm	Should missing values be removed before some computations?
na.value	The value with which to replace missing values (NA).
...	Arguments passed to decostand , or other tran methods.
formula	A model formula describing the variables to be transformed. The formula should have only a right hand side, e.g. <code>~ foo + bar</code> .

data, subset, na.action

See `model.frame` for details on these arguments. `data` will generally be the object or environment within which the variables in the formula are searched for.

Details

The function offers following transformation and standardization methods for community data:

- `sqrt`: take the square roots of the observed values.
- `cubert`: take the cube root of the observed values.
- `rootroot`: take the fourth root of the observed values. This is also known as the root root transformation (Field et al 1982).
- `log`: take the logarithms of the observed values. The transformation applied can be modified by constants `a` and `b` and the base of the logarithms. The transformation applied is $x^* = \log_{\text{base}}(ax + b)$
- `reciprocal`: returns the multiplicative inverse or reciprocal, $1/x$, of the observed values.
- `freq`: divide by column (variable, species) maximum and multiply by the number of non-zero items, so that the average of non-zero entries is 1 (Oksanen 1983).
- `center`: centre all variables to zero mean.
- `range`: standardize values into range 0 . . . 1. If all values are constant, they will be transformed to 0.
- `percent`: convert observed count values to percentages.
- `proportion`: convert observed count values to proportions.
- `standardize`: scale `x` to zero mean and unit variance.
- `pa`: scale `x` to presence/absence scale (0/1).
- `missing`: replace missing values with `na.value`.
- `chi.square`: divide by row sums and square root of column sums, and adjust for square root of matrix total (Legendre & Gallagher 2001). When used with the Euclidean distance, the distances should be similar to the Chi-square distance used in correspondence analysis. However, the results from `cmdscale` would still differ, since CA is a weighted ordination method.
- `hellinger`: square root of observed values that have first been divided by row (site) sums (Legendre & Gallagher 2001).
- `wisconsin`: applies the Wisconsin double standardization, where columns (species, variables) are first standardized by maxima and then sites (rows) by site totals.
- `pcent2prop`: convert percentages to proportions.
- `prop2pcent`: convert proportions to percentages.
- `logRatio`: applies a log transformation (see `log` above) to the data, then centres the data by rows (by subtraction of the mean for row i from the observations in row i). Using this transformation subsequent to PCA results in Aitchison's Log Ratio Analysis (LRA), a means of dealing with closed compositional data such as common in palaeoecology (Aitchison, 1983).
- `power`: applies a power transformation.
- `rowCenter`: Centres `x` by rows through the subtraction of the corresponding row mean from the observations in the row.

Value

Returns the suitably transformed or standardized x . If x is a data frame, the returned value is likewise a data frame. The returned object also has an attribute "tran" giving the name of applied transformation or standardization "method".

Author(s)

Gavin L. Simpson. Much of the functionality of tran is provided by [decostand](#), written by Jari Oksanen.

References

- Aitchison, J. (1983) Principal components analysis of compositional data. *Biometrika* **70**(1); 57–65.
- Field, J.G., Clarke, K.R., & Warwick, R.M. (1982) A practical strategy for analysing multispecies distributions patterns. *Marine Ecology Progress Series* **8**; 37–52.
- Legendre, P. & Gallagher, E.D. (2001) Ecologically meaningful transformations for ordination of species data. *Oecologia* **129**; 271-280.
- Oksanen, J. (1983) Ordination of boreal heath-like vegetation with principal component analysis, correspondence analysis and multidimensional scaling. *Vegetatio* **52**; 181-189.

See Also

[decostand](#)

Examples

```
data(swapdiat)
## convert percentages to proportions
sptrans <- tran(swapdiat, "pcent2prop")

## apply Hellinger transformation
spHell <- tran(swapdiat, "hellinger")

## Dummy data to illustrate formula method
d <- data.frame(A = runif(10), B = runif(10), C = runif(10))
## simulate some missings
d[sample(10,3), 1] <- NA
## apply tran using formula
tran(~ . - B, data = d, na.action = na.pass,
     method = "missing", na.value = 0)
```

varExpl	<i>Variance explained by ordination axes</i>
---------	--

Description

Extracts information about the variance explained by ordination axes and expresses it in a variety of ways.

Usage

```
varExpl(object, ...)  
  
## S3 method for class 'cca'  
varExpl(object, axes = 1L, cumulative = FALSE,  
        pcent = FALSE, ...)  
  
## S3 method for class 'prcurve'  
varExpl(object, pcent = FALSE, ...)
```

Arguments

object	an R object of an appropriate type. Currently only for objects that inherit from classes "cca" or "prcurve".
axes	numeric vector indicating which axes to compute variance explained for.
cumulative	logical; should the variance be explained as a cumulative sum over the axes?
pcent	logical; should the variance explained be expressed as a percentage of the total variance.
...	additional arguments passed to other methods. Currently not used.

Value

A numeric vector variance explained by each axis.

Author(s)

Gavin L. Simpson

See Also

See [cca](#) and [prcurve](#) for functions that produce objects that `varExpl()` can work with.

Examples

```

data(abernethy)

## Remove the Depth and Age variables
abernethy2 <- abernethy[, -(37:38)]

## Fit PCA
aber.pca <- rda(abernethy2)

## Distance along the first PCA axis
varExpl(aber.pca)

```

wa

Weighted averaging transfer functions

Description

Implements the weighted averaging transfer function methodology. Tolerance down-weighting and inverse and classical deshrinking are supported.

Usage

```

wa(x, ...)

## Default S3 method:
wa(x, env,
  deshrink = c("inverse", "classical", "expanded", "none"),
  tol.dw = FALSE, useN2 = TRUE,
  na.tol = c("min", "mean", "max"),
  small.tol = c("min", "fraction", "absolute"),
  min.tol = NULL, f = 0.1, ...)

## S3 method for class 'formula'
wa(formula, data, subset, na.action,
  deshrink = c("inverse", "classical", "expanded", "none"),
  tol.dw = FALSE, useN2 = TRUE, na.tol = c("min", "mean", "max"),
  small.tol = c("min", "fraction", "absolute"), min.tol = NULL,
  f = 0.1, ..., model = FALSE)

## S3 method for class 'wa'
fitted(object, ...)

## S3 method for class 'wa'
residuals(object, ...)

## S3 method for class 'wa'

```

```
coef(object, ...)
```

```
waFit(x, y, tol.dw, useN2, deshrink, na.tol, small.tol,
      min.tol, f)
```

Arguments

x	The species training set data
env, y	The response vector
deshrink	Which deshrinking method to use? One of "inverse" or "classical", "expanded" or "none"
tol.dw	logical; should species with wider tolerances be given lower weight?
useN2	logical; should Hill's N2 values be used to produce un-biased tolerances?
na.tol	character; method to use to replace missing (NA) tolerances in WA computations. Missing values are replaced with the minimum, average or maximum tolerance observed that is not missing.
small.tol	character; method to replace small tolerances. See Details.
min.tol	numeric; threshold below which tolerances are treated as being 'small'. Default is not to replace small tolerances.
f	numeric, $0 < f < 1$; fraction of environmental gradient env to replace small tolerances with if <code>small.tol = "fraction"</code> is specified.
formula	a model formula
data	an optional data frame, list or environment (or object coercible by <code>as.data.frame</code> to a data frame) containing the variables specified on the RHS of the model formula. If not found in data, the variables are taken from <code>environment(formula)</code> , typically the environment from which <code>wa</code> is called.
subset	an optional vector specifying a subset of observations to be used in the fitting process.
na.action	a function which indicates what should happen when the data contain NAs. The default is set by the <code>na.action</code> setting of options, and is <code>na.fail</code> if that is unset. The 'factory-fresh' default is <code>na.omit</code> . Another possible value is <code>NULL</code> , no action. Value <code>na.exclude</code> can be useful.
model	logical. If TRUE the model frame is returned.
object	an Object of class "wa", the result of a call to <code>wa</code> .
...	arguments to other methods.

Details

A typical model has the form `response ~ terms` where `response` is the (numeric) response vector (the variable to be predicted) and `terms` is a series of terms which specifies a linear predictor for response. A terms specification of the form `first + second` indicates all the terms in `first` together with all the terms in `second` with duplicates removed. A specification of `.` is shorthand for all terms in `data` not already included in the model.

Species that have very small tolerances can dominate reconstructed values if tolerance down-weighting is used. In `wa`, small tolerances are defined as a tolerance that is $< \text{min.tol}$. The default is to not replace small tolerances, and the user needs to specify suitable values of `min.tol`. Function `tolerance` may be of use in computing tolerances before fitting the WA model.

Small tolerances can be adjusted in several ways:

`min` small tolerances are replaced by the smallest observed tolerance that is greater than, or equal to, `min.tol`. With this method, the replaced values will be no smaller than any other observed tolerance. This is the default in **analogue**.

`fraction` small tolerances are replaced by the fraction, `f`, of the observed environmental gradient in the training set, `env`.

`absolute` small tolerances are replaced by `min.tol`.

Function `waFit` is the workhorse implementing the actual WA computations. It performs no checks on the input data and returns a simple list containing the optima, tolerances, model tolerances, fitted values, coefficients and the numbers of samples and species. See `Value` below for details of each component.

Value

An object of class "wa", a list with the following components:

<code>wa.optima</code>	The WA optima for each species in the model.
<code>tolerances</code>	The actual tolerances calculated (these are weighted standard deviations).
<code>model.tol</code>	The tolerances used in the WA model computations. These will be similar to <code>tol</code> , but will not contain any NAs and any small tolerances will have been replaced with the appropriate value.
<code>fitted.values</code>	The fitted values of the response for each of the training set samples.
<code>residuals</code>	Model residuals.
<code>coefficients</code>	Deshrinking coefficients.
<code>rmse</code>	The RMSE of the model.
<code>r.squared</code>	The coefficient of determination of the observed and fitted values of the response.
<code>avg.bias, max.bias</code>	The average and maximum bias statistics.
<code>n.samp, n.spp</code>	The number of samples and species in the training set.
<code>deshrink</code>	The deshrinking regression method used.
<code>tol.dw</code>	logical; was tolerance down-weighting applied?
<code>call</code>	The matched function call.
<code>orig.x</code>	The training set species data.
<code>orig.env</code>	The response data for the training set.
<code>options.tol</code>	A list, containing the values of the arguments <code>useN2</code> , <code>na.tol</code> , <code>small.tol</code> , <code>min.tol</code> , and <code>f</code> .
<code>terms, model</code>	Model <code>terms</code> and <code>model.frame</code> components. Only returned by the <code>formula</code> method of <code>wa</code> .

Author(s)

Gavin L. Simpson and Jari Oksanen

See Also

[mat](#) for an alternative transfer function method.

Examples

```
data(ImbrieKipp)
data(SumSST)

## fit the WA model
mod <- wa(SumSST ~., data = ImbrieKipp)
mod

## extract the fitted values
fitted(mod)

## residuals for the training set
residuals(mod)

## deshrinking coefficients
coef(mod)

## diagnostics plots
par(mfrow = c(1,2))
plot(mod)
par(mfrow = c(1,1))

## tolerance DW
mod2 <- wa(SumSST ~ ., data = ImbrieKipp, tol.dw = TRUE,
          min.tol = 2, small.tol = "min")
mod2

## compare actual tolerances to working values
with(mod2, rbind(tolerances, model.tol))
```

weightedCor

Weighted correlation test of WA reconstruction

Description

Weighted correlation between WA optima from training set and axis 1 scores of constrained ordination fitted to fossil data with WA model predictions for fossil samples as constraints.

Usage

```
## Default S3 method:
weightedCor(x, env, fossil, method = c("rda", "cca"),
            test = TRUE, type = c("simulate", "permute"), sim = 999,
            verbose = TRUE, ...)

## S3 method for class 'weightedCor'
plot(x,
     type = c("bubble", "null"),
     weighted = TRUE,
     size = 0.25,
     xlab = paste(x$env, "WA Optima"),
     ylab = "Axis 1 Score",
     xlim,
     main = "",
     sub = NULL,
     border = "gray75",
     col = "gray75",
     obscol = "red",
     fg = "black", ...)
```

Arguments

x	training set covariates, a matrix-like object usually of species/proxy data. For the plot method, an object of class "weightedCor", the result of a call to weightedCor.
env	training set response, a vector usually of environmental data.
fossil	matrix of fossil/core species/proxy data for which a reconstruction is sought.
method	constrained ordination method. One of "rda" and "cca". Currently only "rda" is supported.
test	logical; should the observed correlation be tested?
type	the type of test to apply. One of "simulate" or "permute". The latter is currently not implemented. For the plot method, the type of plot to produce.
sim	numeric; number of simulations or permutations to perform as part of the test
verbose	logical; should the progress of the test be shown via a progress bar?
...	arguments passed to other methods. In the case of the plot method, additional graphical parameters can be supplied.
weighted	logical; should the null distribution plotted be of the weighted or normal correlation.
size	numeric; the size of the largest bubble in inches. See symbols and argument inches for details.
xlim,xlab,ylab,main,sub	graphical parameters with their usual meaning.
border, col	The border and fill colours for the histogram bars.
fg	The colour of the bubbles drawn on the bubble plot.
obscol	The colour of the indicator for the observed correlation.

Value

The plot method produces a plot on the current device. `weightedCor()` returns a list with the following components:

<code>wtdCorrel</code> , <code>Correl</code>	numeric; the observed weighted and standard correlation.
<code>data</code>	data frame; containing the training set WA Optima, axis 1 species scores, and mean abundance for each species.
<code>ord</code>	the fitted constrained ordination.
<code>model</code>	the fitted WA model.
<code>method</code>	the ordination method used.
<code>ndist</code>	the null distribution produced. NULL if argument <code>test</code> was FALSE.
<code>sim</code>	numeric; the number of simulations or permutations used to test the observed correlations.
<code>type</code>	the type of test performed.
<code>env</code>	the deparsed version of <code>env</code> argument. Used for plotting.
<code>call</code>	the matched function call.

Author(s)

Gavin L. Simpson

References

Telford R.J. and Birks, H.J.B. (2011) A novel method for assessing the statistical significance of quantitative reconstructions inferred from biotic assemblages. *Quaternary Science Reviews* **30**:1272-1278.

See Also

[wa](#) for details on fitting weighted average models.

Examples

```
data(ImbrieKipp, SumSST, V12.122)

Cor <- weightedCor(ImbrieKipp, env = SumSST,
                  fossil = V12.122, type = "simulate", sim = 49)

Cor

plot(Cor)
plot(Cor, type = "null")
```

Index

- *Topic **\textasciitildekwd1**
 - weightedCor, 121
- *Topic **\textasciitildekwd2**
 - weightedCor, 121
- *Topic **datasets**
 - abernethy, 4
 - ImbrieKipp, 36
 - Pollen, 76
 - rlgh, 90
 - swapdiat, 110
 - swappH, 111
- *Topic **hplot**
 - cma, 16
 - densityplot.residLen, 21
 - hist.residLen, 33
 - histogram.residLen, 34
 - panel.Loess, 54
 - panel.Stratiplot, 55
 - plot.dissimilarities, 61
 - plot.logitreg, 62
 - plot.mat, 63
 - plot.mcarlo, 66
 - plot.minDC, 67
 - plot.prcurve, 69
 - plot.residLen, 71
 - plot.roc, 72
 - plot.wa, 75
 - reconPlot, 86
 - screepLOT, 95
 - Stratiplot, 99
 - timetrack, 111
- *Topic **manip**
 - cma, 16
 - getK, 30
 - join, 38
 - minDC, 50
 - tran, 114
- *Topic **methods**
 - bayesF, 7
- bootstrap, 10
- bootstrap.wa, 13
- chooseTaxa, 15
- cma, 16
- crossval, 19
- densityplot.residLen, 21
- dissimilarities, 23
- distance, 25
- fuse, 28
- hist.residLen, 33
- histogram.residLen, 34
- logitreg, 40
- mat, 43
- minDC, 50
- optima, 52
- pcr, 57
- performance, 60
- plot.mat, 63
- plot.minDC, 67
- plot.residLen, 71
- plot.roc, 72
- plot.wa, 75
- predict.mat, 81
- predict.wa, 84
- residLen, 88
- RMSEP, 91
- roc, 93
- screepLOT, 95
- stdError, 97
- summary.analog, 102
- summary.bootstrap.mat, 104
- summary.cma, 106
- summary.mat, 107
- summary.predict.mat, 109
- timetrack, 111
- tran, 114
- wa, 118

- *Topic **models**
 - mat, 43

- roc, 93
- wa, 118
- *Topic **multivariate**
 - analog, 5
 - bootstrap, 10
 - bootstrap.wa, 13
 - bootstrapObject, 14
 - chooseTaxa, 15
 - dissimilarities, 23
 - distance, 25
 - fuse, 28
 - gradientDist, 32
 - join, 38
 - mat, 43
 - mcarlo, 48
 - plot.mcarlo, 66
 - prcurve, 78
 - residLen, 88
 - roc, 93
 - tran, 114
 - varExpl, 117
 - wa, 118
- *Topic **nonparametric**
 - prcurve, 78
- *Topic **package**
 - analogue-package, 3
- *Topic **regression**
 - wa, 118
- *Topic **smooth**
 - prcurve, 78
- *Topic **univar**
 - bayesF, 7
 - stdError, 97
- *Topic **utilities**
 - deshrink, 22
 - getK, 30
 - minDC, 50
 - performance, 60
 - RMSEP, 91
- *Topic **utility**
 - gradientDist, 32
 - varExpl, 117
- abernethy, 4
- analog, 5, 17–19, 24, 49, 102, 103, 106
- analogue, 48
- analogue (analogue-package), 3
- analogue-package, 3
- as.data.frame, 43, 53, 58, 99, 119
- as.data.frame.optima (optima), 52
- as.data.frame.tolerance (optima), 52
- as.matrix.dist, 29
- bayesF, 7, 41, 42
- Biome (Pollen), 76
- bootstrap, 10, 92, 97
- bootstrap.mat, 13, 14, 30, 31, 45, 46, 82, 83, 95, 97, 104, 105, 109
- bootstrap.wa, 13, 60, 92
- bootstrapObject, 11, 14
- boxplot, 19
- cca, 33, 89, 112, 113, 117
- ChiSquare (pcr), 57
- chooseTaxa, 15
- Climate (Pollen), 76
- cma, 6, 7, 16, 17, 106, 107
- cmdscale, 115
- coef.pcr (pcr), 57
- coef.wa (wa), 118
- colnames, 110
- crossval, 19
- daisy, 27
- decostand, 114, 116
- densityplot, 21
- densityplot.residLen, 21, 34, 35, 71
- deshrink, 22
- deshrinkPred (deshrink), 22
- dissim (dissimilarities), 23
- dissimilarities, 7, 23, 61
- dist, 26, 27, 29
- distance, 7, 25, 29, 44, 49, 51
- dose.p, 41
- eigenvals.pcr (pcr), 57
- fitted.bootstrap.mat (bootstrap), 10
- fitted.mat, 46, 82
- fitted.mat (mat), 43
- fitted.pcr (pcr), 57
- fitted.timetrack (timetrack), 111
- fitted.wa (wa), 118
- fittedY (residLen), 88
- fuse, 28
- get.lam, 80
- getK, 18, 30, 91, 98
- glm, 41, 42

- glm.fit, 42
- gradientDist, 32
- head, 38, 39
- head.join (join), 38
- Hellinger (pcr), 57
- hist.residLen, 22, 33, 35
- histogram, 34, 35
- histogram.residLen, 22, 34, 34, 71
- ImbrieKipp, 36
- initCurve (prcurve), 78
- join, 25, 38, 39, 81
- legend, 73, 96
- lm, 101
- Location (Pollen), 76
- loess.smooth, 54, 55
- log, 114
- logitreg, 40, 62, 63
- mat, 18, 30, 31, 43, 44, 45, 48, 49, 65, 76, 81–83, 87, 92, 95–98, 105, 108, 109, 121
- mcarlo, 48, 66, 67, 95
- merge, 39
- minDC, 50, 69, 98
- model.frame, 115, 120
- na.exclude, 44
- na.fail, 44
- na.omit, 44
- optima, 52
- options, 44, 100
- panel.grid, 56
- panel.lines, 56
- panel.Loess, 54, 56, 101
- panel.loess, 54–56
- panel.points, 56
- panel.polygon, 56
- panel.refline, 56
- panel.segments, 56
- panel.Stratiplot, 55, 99–101
- panel.xyplot, 55, 56
- par, 8, 18, 68, 73, 87
- pcr, 57
- performance, 60, 85
- performance.pcr (pcr), 57
- plot, 56
- plot.bayesF, 9
- plot.bayesF (bayesF), 7
- plot.cma (cma), 16
- plot.dissimilarities, 7, 24, 61
- plot.lm, 46, 65, 66, 73, 74, 76
- plot.logitreg, 62
- plot.mat, 46, 63
- plot.mcarlo, 49, 66
- plot.minDC, 51, 67
- plot.prcurve, 69
- plot.residLen, 22, 34, 35, 71, 71
- plot.roc, 72
- plot.timetrack (timetrack), 111
- plot.wa, 75
- plot.weightedCor (weightedCor), 121
- Pollen, 76
- prcurve, 33, 70, 78, 117
- predict.cca, 89
- predict.mat, 18, 45, 46, 51, 81, 85, 87, 98, 109
- predict.wa, 13, 60, 84, 87
- predWA (crossval), 19
- predWAT (crossval), 19
- print.analog (analog), 5
- print.bayesF (bayesF), 7
- print.bootstrap.mat (bootstrap), 10
- print.bootstrap.wa (bootstrap.wa), 13
- print.cma (cma), 16
- print.crossval (crossval), 19
- print.fitted.bootstrap.mat (bootstrap), 10
- print.fitted.mat (mat), 43
- print.logitreg (logitreg), 40
- print.mat (mat), 43
- print.mcarlo (mcarlo), 48
- print.minDC (minDC), 50
- print.optima (optima), 52
- print.pcr (pcr), 57
- print.performance (performance), 60
- print.predict.mat (predict.mat), 81
- print.predict.wa (predict.wa), 84
- print.residLen (residLen), 88
- print.residuals.bootstrap.mat (bootstrap), 10
- print.residuals.mat (mat), 43
- print.roc (roc), 93

- print.summary.analog (summary.analog), 102
- print.summary.bootstrap.mat (summary.bootstrap.mat), 104
- print.summary.cma (summary.cma), 106
- print.summary.logitreg (logitreg), 40
- print.summary.mat (summary.mat), 107
- print.summary.predict.mat (summary.predict.mat), 109
- print.summary.roc (roc), 93
- print.timetrack (timetrack), 111
- print.tolerance (optima), 52
- print.wa (wa), 118
- print.weightedCor (weightedCor), 121
- quantile, 17, 103
- rda, 70, 89, 112, 113
- reconPlot, 85, 86
- resid.bootstrap.mat (bootstrap), 10
- resid.mat, 46
- resid.mat (mat), 43
- residLen, 21, 22, 33–35, 71, 88
- residuals, 10, 12, 15
- residuals.bootstrap.mat (bootstrap), 10
- residuals.mat (mat), 43
- residuals.pcr (pcr), 57
- residuals.wa (wa), 118
- rlgh, 90
- RMSEP, 91
- roc, 9, 41, 42, 49, 63, 72, 74, 93
- Salinity (ImbrieKipp), 36
- scores, 70
- scores.cca, 32
- screepLOT, 95, 97
- screepLOT.pcr (pcr), 57
- setK<- (getK), 30
- smooth.spline, 79, 80
- smoothSpline (prcurve), 78
- sqlLinear (residLen), 88
- sqlUnimodal (residLen), 88
- stdError, 97
- Stratiplot, 54, 56, 99
- stripchart, 17–19
- summary, 102, 104–109
- summary.analog, 102
- summary.bootstrap.mat, 104
- summary.cma, 106
- summary.logitreg (logitreg), 40
- summary.mat, 45, 46, 107
- summary.predict.mat, 109
- summary.roc (roc), 93
- SumSST (ImbrieKipp), 36
- swapdiat, 110, 111
- swappH, 110, 111
- symbols, 122
- tail, 38, 39
- tail.join (join), 38
- terms, 45, 120
- timetrack, 111
- tolerance, 120
- tolerance.default (optima), 52
- tran, 114
- trellis.par.get, 56
- V12.122 (ImbrieKipp), 36
- varExpl, 117
- vegdist, 26, 27, 29
- wa, 20, 23, 53, 59, 60, 84, 85, 87, 92, 118, 123
- waFit (wa), 118
- weightedCor, 121
- WinSST (ImbrieKipp), 36
- xyplot, 100, 101