

Package ‘verification’

January 2, 2012

Version 1.31

Date 2010-01-01

Title Forecast verification utilities.

Author NCAR - Research Application Program

Maintainer Matt Pocernich <pocernic@ucar.edu>

Depends methods, waveslim, fields, boot, CircStats

Description This package contains utilities for verification of discrete, continuous, probabilistic forecasts and forecast expressed as parametric distributions.

License GPL (>= 2)

LazyLoad yes

LazyData yes

Repository CRAN

Date/Publication 2010-01-11 07:42:51

R topics documented:

analysis.dat	2
attribute	2
brier	4
conditional.quantile	6
crps	7
crps.circ	9
crpsDecompostion	10
disc.dat	11
discrimination.plot	11
forecast.dat	13
fss	13
int.scale.verify	14

IS	16
leps	18
lines.verify	20
lookup	21
measurement.error	21
multi.cont	22
observation.error	24
performance.diagram	25
plot.int.scale	26
pop	27
precip.ensemble	27
prob.fracs.dat	28
probcont2disc	28
rerv	29
reliability.plot	30
roc.area	32
roc.plot	33
rps	36
table.stats	38
table.stats.boot	39
value	40
verify	42

Index **45**

analysis.dat	<i>Spatial Observation Dataset.</i>
--------------	-------------------------------------

Description

This spatial data set is used by the `int.scale.verify` function. This consists of the outcome of the weather forecast presented in `forecast.dat`. This data was provided by Barbara Casati and is from the Met Office.

attribute	<i>Attribute plot</i>
-----------	-----------------------

Description

An attribute plot illustrates the reliability, resolution and uncertainty of a forecast with respect to the observation. The frequency of binned forecast probabilities are plotted against proportions of binned observations. A perfect forecast would be indicated by a line plotted along the 1:1 line. Uncertainty is described as the vertical distance between this point and the 1:1 line. The relative frequency for each forecast value is displayed in parenthesis.

Usage

```
## Default S3 method:
attribute(x, obar.i, prob.y = NULL, obar = NULL, class =
"none", main = NULL, CI = FALSE, n.boot = 100, alpha = 0.05, tck = 0.01, freq = TRUE, pred = NULL, obs = NU
## S3 method for class 'prob.bin'
attribute(x, ...)
```

Arguments

x	A vector of forecast probabilities or a “prob.bin” class object produced by the verify function.
obar.i	A vector of observed relative frequency of forecast bins.
prob.y	Relative frequency of forecasts of forecast bins.
obar	Climatological or sample mean of observed events.
class	Class of object. If prob.bin, the function will use the data to estimate confidence intervals.
main	Plot title.
CI	Confidence Intervals. This is only an option if the data is accessible by using the verify command first. Calculated by bootstrapping the observations and prediction, then calculating PODy and PODn values.
n.boot	Number of bootstrap samples.
alpha	Confidence interval. By default = 0.05
tck	Tick width on confidence interval whiskers.
freq	Should the frequencies be plotted. Default = TRUE
pred	Required to create confidence intervals
obs	Required to create confidence intervals
thres	thresholds used to create bins for plotting confidence intervals.
bins	Should probabilities be binned or treated as unique predictions?
...	Graphical parameters

Note

Points and bins are plotted at the mid-point of bins. This can create distorted graphs if forecasts are created at irregular intervals.

Author(s)

Matt Pocerlich <pocernic@rap.ucar.edu>

References

Hsu, W. R., and A.H. Murphy, 1986: The attributes diagram: A geometrical framework for assessing the quality of probability forecasts. *Int. J. Forecasting* **2**, 285-293.

Wilks, D. S. (2005) *Statistical Methods in the Atmospheric Sciences* Chapter 7, San Diego: Academic Press.

See Also

[verify_reliability.plot](#)

Examples

```
## Data from Wilks, table 7.3 page 246.
y.i  <- c(0,0.05, seq(0.1, 1, 0.1))
obars.i <- c(0.006, 0.019, 0.059, 0.15, 0.277, 0.377, 0.511,
            0.587, 0.723, 0.779, 0.934, 0.933)
prob.y<- c(0.4112, 0.0671, 0.1833, 0.0986, 0.0616, 0.0366,
           0.0303, 0.0275, 0.245, 0.022, 0.017, 0.203)
obars<- 0.162

attribute(y.i, obars.i, prob.y, obars, main = "Sample Attribute Plot")

## Function will work with a ‘‘prob.bin’’ class objects as well.
## Note this is a random forecast.
obs<- round(runif(100))
pred<- runif(100)

A<- verify(obs, pred, frcst.type = "prob", obs.type = "binary")
attribute(A, main = "Alternative plot", xlab = "Alternate x label" )
## to add a line from another model
obs<- round(runif(100))
pred<- runif(100)

B<- verify(obs, pred, frcst.type = "prob", obs.type = "binary")
lines.attrib(B, col = "green")

## Same with confidence intervals
attribute(A, main = "Alternative plot", xlab = "Alternate x label", CI =
TRUE)

#### add lines to plot
data(pop)
pop.convert() ## internal function used to make binary observations for the pop figure.
### note the use of bins = FALSE
mod24 <- verify(d$obs_rain, d$p24_rain, bins = FALSE)

mod48 <- verify(d$obs_rain, d$p48_rain, bins = FALSE)
plot(mod24, freq = FALSE)

lines.attrib(mod48, col = "green", lwd = 2, type = "b")
```

Description

Calculates verification statistics for probabilistic forecasts of binary events.

Usage

```
brier(obs, pred, baseline, thresholds = seq(0,1,0.1), bins = TRUE, ... )
```

Arguments

obs	Vector of binary observations
pred	Vector of probabilistic predictions [0,1]
baseline	Vector of climatological (no - skill) forecasts. If this is null, a sample climatology will be calculated.
thresholds	Values used to bin the forecasts. By default the bins are {[0,0.1), [0.1, 0.2),}
bins	If TRUE, thresholds define bins into which the probabilistic forecasts are entered and assigned the midpoint as a forecast. Otherwise, each unique forecast is considered as a separate forecast. For example, set bins to FALSE when dealing with a finite number of probabilities generated by an ensemble forecast.
...	Optional arguments

Value

baseline.tf	Logical indicator of whether climatology was provided.
bs	Brier score
bs.baseline	Brier Score for climatology
ss	Skill score
bs.reliability	Reliability portion of Brier score.
bs.resolution	Resolution component of Brier score.
bs.uncert	Uncertainty component of Brier score.
y.i	Forecast bins – described as the center value of the bins.
obar.i	Observation bins – described as the center value of the bins.
prob.y	Proportion of time using each forecast
obar	Forecast based on climatology or average sample observations.
check	Reliability - resolution + uncertainty should equal brier score.

Note

This function is used within `verify`.

Author(s)

Matt Pocerlich <pocerlich@rap.ucar.edu>

References

Wilks, D. S. (1995) *Statistical Methods in the Atmospheric Sciences* Chapter 7, San Diego: Academic Press.

Examples

```
# probabilistic/ binary example

pred<- runif(100)
obs<- round(runif(100))
brier(obs, pred)
```

conditional.quantile *Conditional Quantile Plot*

Description

This function creates a conditional quantile plot as shown in Murphy, et al (1989) and Wilks (1995).

Usage

```
conditional.quantile(pred, obs, bins = NULL, thrs = c(10, 20), main,
... )
```

Arguments

pred	Forecasted value. ($[n,1]$ vector; $n = \#$ of forecasts)
obs	Observed value. ($[n,1]$ vector; $n = \#$ of observations)
bins	Bins for forecast and observed values. A minimum number of values are required to calculate meaningful statistics. So for variables that are continuous, such as temperature, it is frequently convenient to bin these values. ($[m,1]$ vector; $m = \#$ of bins)
thrs	The minimum number of values in a bin required to calculate the 25th and 75th quantiles and the 10th and 90th percentiles respectively. The median values will always be displayed. ($[2,1]$ vector)
main	Plot title
...	Plotting options.

Value

This function produces a conditional.quantile plot. The y axis represents the observed values, while the x axis represents the forecasted values. The histogram along the bottom axis illustrates the frequency of each forecast.

Note

In the example below, the median line extends beyond the range of the quartile or 10th and 90th percentile lines. This is because there are not enough points in each bin to calculate these quartile values. That is, there are fewer than the limits set in the `thrs` input.

Author(s)

Matt Pocernich <pocernic@rap.ucar.edu>

References

Murphy, A. H., B.G. Brown and Y. Chen. (1989) "Diagnostic Verification of Temperature Forecasts," *Weather and Forecasting*; December, 1989.

Examples

```
set.seed(10)
m<- seq(10, 25, length = 1000)
frcst <- round(rnorm(1000, mean = m, sd = 2) )
obs<- round(rnorm(1000, mean = m, sd = 2 ))
bins <- seq(0, 30,1)
thrs<- c( 10, 20) # number of obs needed for a statistic to be printed #1,4 quartile, 2,3 quartiles

conditional.quantile(frcst, obs, bins, thrs, main = "Sample Conditional Quantile Plot")
#### Or plots a 'cont.cont' class object.

obs<- rnorm(100)
pred<- rnorm(100)
baseline <- rnorm(100, sd = 0.5)

A<- verify(obs, pred, baseline = baseline, frcst.type = "cont", obs.type = "cont")
plot(A)
```

crps

Continuous Ranked Probability Score

Description

Calculates the crps for a forecast made in terms of a normal probability distribution and an observation expressed in terms of a continuous variable.

Usage

```
crps(obs, pred, ...)
```

Arguments

obs	A vector of observations.
pred	A vector or matrix of the mean and standard deviation of a normal distribution. If the vector has a length of 2, it is assumed that these values represent the mean and standard deviation of the normal distribution that will be used for all forecasts.
...	Optional arguments

Value

crps	Continuous ranked probability scores
CRPS	Mean of crps
ign	Ignorance score
IGN	Mean of the ignorance score

Note

This function is used within `verify`.

Author(s)

Matt Pocernich <pocernic@rap.ucar.edu>

References

Gneiting, T., Westveld, A., Raftery, A. and Goldman, T, 2004: *Calibrated Probabilistic Forecasting Using Ensemble Model Output Statistics and Minimum CRPS Estimation*. Technical Report no. 449, Department of Statistics, University of Washington. [Available online at <http://www.stat.washington.edu/www/research/reports/>]

Examples

```
# probabilistic/ binary example

x <- runif(100) ## simulated observation.
crps(x, c(0,1))

## simulated forecast in which mean and sd differs for each forecast.
frcs<- data.frame( runif(100, -2, 2), runif(100, 1, 3 ) )
crps(x, frcs)
```

`crps.circ`*CRPS Statistics for circular data*

Description

Computes the circular CRPS when the verifying direction is 'x' and the probabilistic forecast is a von Mises distribution on the circle with mean 'mu' and concentration parameter 'kappa'.

Usage

```
crps.circ(x, mu, kappa )
```

Arguments

x	A single direction in angles ($0 \leq x < 2\pi$)
mu	A single direction in angles ($0 \leq \mu < 2\pi$)
kappa	A single positive number ($\text{kappa} \geq 0$)

Value

`crps.circ` Circular crps

Author(s)

Nick Johnson<nickaj@gmail.com>

References

Grimm, E.P., T. Gneiting, V.J. Berrocal and N.A. Johnson (2006). The continuous ranked probability score for circular variables and its application to mesoscale forecast ensemble verification. , **132**, pp. 1-17, *Q.J.R. Meteorolo. Soc.*.

Examples

```
x           <- pi/3
mu          <- pi/3.2
kappa      <- 10000
ES.VM.Comp2.Max.Kappa <- 1500

crps.circ( x = pi/3, mu = pi/3.2, kappa = 10000)
```

crpsDecomposition *Decomposition of Continuous Ranked Probability Score*

Description

The CRPS measures the distance between the predicted and the observed cumulative density functions (CDFs) of scalar variables. Furthermore, the `crpsDecomposition` function provides the reliability and resolution terms obtained by the CRPS decomposition proposed by Hersbach. The alpha, beta matrices and Heavisides vectors of outliers calculated in the CRPS decomposition are also returned. To speed up calculation time, these matrices/vectors can then be used to recalculate the CRPS's in a bootstrap by using the `crpsFromAlphaBeta` function.

Usage

```
crpsDecomposition(obs, eps)
crpsFromAlphaBeta(alpha, beta, heaviside0, heavisideN)
```

Arguments

<code>obs</code>	Vector of observations
<code>eps</code>	Matrix of ensemble forecast. Each column represent a member.
<code>alpha</code>	Matrix of alpha (returned from <code>crpsDecomposition</code>)
<code>beta</code>	Vector of beta (returned from <code>crpsDecomposition</code>)
<code>heaviside0</code>	Vector of Heaviside for outlier $i=0$ (returned from <code>crpsDecomposition</code>)
<code>heavisideN</code>	Vector of Heaviside for outlier $i=N$ (returned from <code>crpsDecomposition</code>)

Value

<code>CRPS</code>	CRPS score
<code>CRPSpot</code>	The potential CRPS (Resolution - Uncertainty)
<code>Reli</code>	The Reliability term of the CRPS
<code>alpha</code>	Matrix (Nobservation rows x Nmember +1 columns) of alpha used in the CRPS decomposition.
<code>beta</code>	Matrix (Nobservation rows x Nmember +1 columns) of beta used in the CRPS decomposition.
<code>heaviside0</code>	Vector (Nobservation length) of Heaviside for outlier $i=0$ used in the CRPS decomposition
<code>heavisideN</code>	Vector (Nobservation length) of Heaviside for outlier $i=N$ used in the CRPS decomposition

Author(s)

Ronald Frenette <Ronald.Frenette@ec.gc.ca>

References

G. Candille, P. L. Houtekamer, and G. Pellerin: Verification of an Ensemble Prediction System against Observations, *Monthly Weather Review*, **135**, pp. 2688-2699

Hershbach, Hans, 2000. Decomposition of the Continuous Ranked Probability Score for Ensemble Prediction Systems. *Weather and Forecasting*, **15**, (5) pp. 559-570.

Examples

```
data(precip.ensemble)

#Observations are in the column
obs<-precip.ensemble[,3]

#Forecast values of ensemble are in the column 4 to 54
eps<-precip.ensemble[,4:54]

#CRPS calculation
c<-crpsDecomposition(obs,eps)

#CRPS with alpha and beta
#Resampling indices
nObs<-length(obs)
i<-sample(seq(nObs),nObs,replace=TRUE)
crps2<-crpsFromAlphaBeta(c$alpha[i,],c$beta[i,],c$heaviside0[i],c$heavisideN[i])
```

disc.dat

Discrimination plot dataset.

Description

This dataset is used to illustrate the `discrimination.plot` function.

Usage

```
data(disc.dat)
```

discrimination.plot

Discrimination plot

Description

This function creates a plot of discrimination plots (overlay histograms). In the context of verification, this is often used to compare the distribution of event and no-event forecasts. This may be useful in comparing any set of observations. By default, boxplots of groups appear as upper marginal plots. These may be suppressed.

Usage

```
discrimination.plot(group.id, value, breaks = 11, main =
"Discrimination Plot", xlim = NULL, ylim = NULL, legend =
FALSE, leg.txt = paste("Model", unique(group.id) ), marginal = TRUE, cols =
seq(2, length(unique(group.id)) + 1), xlab = "Forecast", ... )
```

Arguments

group.id	A vector identifying groups. A histogram is created for each unique value.
value	A vector of values corresponding to the group.id vector used to create the histograms
breaks	Number of breaks in the x-axis of the histogram. The range of values is taken to be the range of prediction values.
main	Title for plot.
xlim	Range of histogram - x axis - main plot coordinates.
ylim	Range of histogram - y axis - main plot coordinates.
legend	Should there be a legend? Default = FALSE
leg.txt	Legend text. If FALSE or if a marginal plot is created, no legend is added.
cols	A vector showing the colors to be used in the histograms and in the marginal boxplots
marginal	Should a boxplots be placed in the top margin? Defaults to TRUE
xlab	Label of the x-axis on the main plot.
...	Additional plotting options.

Author(s)

Matt Pocernich <pocernic@rap.ucar.edu>

Examples

```
# A sample forecast.

data(disc.dat)
discrimination.plot(disc.dat$group.id, disc.dat$frcst, main = "Default Plot")

discrimination.plot(disc.dat$group.id, disc.dat$frcst, main = "New Labels", cex = 1.2, leg.txt = c("Low", "Med", "H"))

discrimination.plot(disc.dat$group.id, disc.dat$frcst, main = "Without Marginal Plots ", marginal = FALSE)
```

forecast.dat	<i>Forecast forecast dataset.</i>
--------------	-----------------------------------

Description

This spatial data set is an example used by the `int.scale.verify` function. It consists of the Nimrod forecasting system and was provided by Barbara Casati.

fss	<i>Fractional Skill Score</i>
-----	-------------------------------

Description

Calculates the fractional skill score for spatial forecasts and spatial observations.

Usage

```
fss(obs, pred, w = 0, FUN = mean, ...)
```

Arguments

obs	A matrix of binomial observed values.
pred	A matrix of binomial forecasted values
w	Box width. When $w = 0$, each pixel is considered alone. $w = 2$ creates a box with a length of 5 units.
FUN	Function to be applied to each subgrid. By default, the mean of each box is used to calculate the fraction of each subgrid.
...	Optional arguments

Value

Return the fraction skill score.

Note

This function contain several loops and consequently is not particularly fast.

Author(s)

Matt Pocerlich <pocernic@rap.ucar.edu>

References

Roberts, N.M and H.W. Lean: 2008: Scale-Selective Verification of Rainfall Accumulations from High-Resolution Forecasts of Convective Events. *Monthly Weather Review* **136**, 78-97.

Examples

```

# The following example replicates Figure 4 in Roberts and Lean (2008).
####      examples

LAG <- c(1,3,11,21)
box.radius <- seq(0,24,2)

OUT <- matrix(NA, nrow = length(box.radius), ncol = length(LAG) )

for(w in 1:4){

FRCS <- OBS <- matrix(0, nrow = 100, ncol = 100)

obs.id      <- 50
OBS[,obs.id] <- 1

FRCS[, obs.id + LAG[w]] <- 1

for(i in 1:length(box.radius)){

OUT[i, w] <- fss(obs = OBS, pred = FRCS, w = box.radius[i] )

} ### close i
} ### close w

b <- mean(OBS) ## base rate

fss.uniform <- 0.5 + b/2
fss.random  <- b

matplot(OUT, ylim = c(0,1), ylab = "FSS", xlab = "grid squares",
type = "b", lty = 1, axes = FALSE , lwd = 2)

abline(h = c(fss.uniform, fss.random), lty = 2)
axis(2)
box()
axis(1, at = 1:length(box.radius), lab = 2*box.radius + 1)
grid()

legend("bottomright", legend = LAG, col = 1:4, pch = as.character(1:4),
title = "Spatial Lag", inset = 0.02, lwd = 2 )

```

Description

For a spatial forecast, evaluates the forecast skill as a function of precipitation rate intensity and spatial scale of the error.

Usage

```
int.scale.verify(fracs, obs, thres = quantile(fracs, p = seq(0,0.9,0.1)), ... )
```

Arguments

fracs	Forecast matrix. Must be of 2^n dimensions.
obs	Observation matrix. Must be of 2^n dimensions.
thres	A vector of thresholds to be considered. By default, the percentiles 0, 90 are used.
...	Optional arguments may be passed to the image plot

Value

SSu1	Skill score as matrix. The rownames are the thresholds, the colnames are n where 2^n is the spatial scale of the skill score decomposition.
MSE	A matrix with the mean squared error of the forecast
l.fracs	Number of rows in forecast. Used in plotting routine.
thres	Thresholds used in model

Note

This function creates an image plot of the intensity plot of the skill scores as a function of spatial scale and threshold. The top row is equivalent to the bias of the forecast.

Author(s)

Barabara Casati <barbara.casati (at) ec.gc.ca>

References

B.Casati, D.B. Stephenson, G. Ross. *A new intensity-scale approach for the verification of spatial precipitation forecasts*. Meteorological Application (RMS), in press.

See Also

<http://www.met.rdg.ac.uk/~swr00bc/>

Examples

```
## simulated example
n<- 5
set.seed(10)
forecast1 <- matrix( log(rlnorm(n = (2^n *2^n) )) , nrow = 2^n)
obs1      <- matrix(log( rlnorm(n = (2^n *2^n) )) , nrow = 2^n)
int.scale.verify(forecast1, obs1, main = "Test Case")

## real example.  Data source referenced below.

data(analysis.dat)
data(forecast.dat)

require(waveslim)
require(fields)

A<- int.scale.verify(forecast.dat, analysis.dat,
  thres = c(0, 2^seq(-5,6)),  main = "NIMROD example" )

plot(A)
```

IS

Alternative Intensity Scale Function

Description

Beta version of intensity scale function created by Barbara Casati.

Usage

```
IS(fracs, obs, thres)
```

Arguments

fracs	Forecast matrix. Must be of 2^n dimensions.
obs	Observation matrix. Must be of 2^n dimensions.
thres	A vector of thresholds to be considered. By default, the percentiles 0, 90 are used.

Value

SSul = SSul, MSEul = MSEul, l.fracs = dim(fracs)[1], thres = thres, Bias = Bu, BR = BRu

SSul Skill score as matrix. The rownames are the thresholds, the colnames are n where 2^n is the spatial scale of the skill score decomposition.

MSEul	A matrix with the mean squared error of the forecast
l.fracs	Number of rows in forecast. Used in plotting routine.
thres	Thresholds used in model
Bias	Bias
BR	BR

Note

THIS IS A DRAFT FORM OF THIS FUNCTION. IT MAY CHANGE AND REPLACE `int.scale.verify`

Author(s)

Barbara Casati <barbara.casati@ec.gc.ca>

References

Casati et al (2004), A new intensity-scale approach for the verification of spatial precipitation forecasts, Meteorol. Appl, vol 11, 141-154 pp.

See Also

[int.scale.verify](#) and [plot.int.scale](#)

Examples

```
#####
# files.dat: read, create and write
#####

IS.NIMROD.case <- IS(forecast.dat, analysis.dat,c(0, 2^seq(-5,6)))
NIMROD.SSul <- IS.NIMROD.case$SSul
colnames(NIMROD.SSul) <- paste(c("0", "1/32", "1/16", "1/8", "1/4", "1/2", "1",
    "2", "4", "8", "16", "32", "64"), "mm/h")
rownames(NIMROD.SSul) <- paste(5*2^seq(0,8), "km")
# write.table(NIMROD.SSul, file="NIMROD.SSul.dat")

NIMROD.MSEul <- IS.NIMROD.case$MSEul
colnames(NIMROD.MSEul) <- paste(c("0", "1/32", "1/16", "1/8", "1/4", "1/2", "1",
    "2", "4", "8", "16", "32", "64"), "mm/h")
rownames(NIMROD.MSEul) <- paste(5*2^seq(0,8), "km")

#####
# colorbars for the images
#####

Nimrod.colorbar <- function(){
  colors = c(0,8,8,8,8,5,5,4,4,4,2,2,2)
  xlimbar = c(7,8)
  ylimbar = c(46.5,59.5)
  barlabels = c("0", "1/32", "1/16", "1/8", "1/4", "1/2", "1", "2", "4", "8", "16",
    "32", "64", "128")
}
```

```

ycoord = seq(ylimbar[1],ylimbar[2],length=length(colors)+1)
for(i in seq(1,length(colors))){
  polygon(x=c(xlimbar[1],xlimbar[1],xlimbar[2],xlimbar[2],xlimbar[1]),
  y=c(ycoord[i],ycoord[i+1],ycoord[i+1],ycoord[i],ycoord[i]),
  col = colors[i])}
axis(4,at=ycoord,labels=barlabels,las=TRUE)
mtext("mm/h",line=1,at=c(7.5,61),cex=1.5)
}

colorbar <- function(xlimbar,ylimbar,colors,barlabels){
ycoord = seq(ylimbar[1],ylimbar[2],length=(length(colors)+1))
for(i in seq(1,length(colors))){
  polygon(x=c(xlimbar[1],xlimbar[1],xlimbar[2],xlimbar[2],xlimbar[1]),
  y=c(ycoord[i],ycoord[i+1],ycoord[i+1],ycoord[i],ycoord[i]),
  col = colors[i])}
axis(4,at=ycoord,labels=barlabels,las=TRUE)
}

# images
#
par(oma=c(3,3,3,3), mfrow = c(1,1) )
image(seq(-12,8,length=256), seq(46.5,59.5,length=256), analysis.dat,
  xlim = c(-12,8), ylim = c(46.5,59.5), zlim = c(0,128),
  xlab = "longitude", ylab = "latitude",main="Nimrod analysis 29/05/99 15:00",
  col=c(0,8,8,8,8,5,5,4,4,4,2,2,2),breaks=c(0,2^seq(-5,7,1)))
world(xlim=c(-12,8),ylim=c(46.5,59.5), add = TRUE, lwd = 3)
Nimrod.colorbar()

#
par(oma=c(3,3,3,3))
image(seq(-12,8,length=256), seq(46.5,59.5,length=256), forecast.dat,
  xlim = c(-12,8), ylim = c(46.5,59.5), zlim = c(0,128),
  xlab = "longitude", ylab = "latitude",main="Nimrod forecast T+3h",
  col=c(0,8,8,8,8,5,5,4,4,4,2,2,2),breaks=c(0,2^seq(-5,7,1)))
world(xlim=c(-12,8),ylim=c(46.5,59.5), add = TRUE, lwd = 3)
Nimrod.colorbar()

#
par(oma=c(3,3,3,3))
image(seq(-6,6,1),seq(1,9),t(NIMROD.SSul),xlim=c(-6.5,8.5),zlim=c(-4,1),
  xlab="threshold (mm/h)",ylab="scale (km)",axes=FALSE,
  main="Intensity-Scale Skill Score", col=c(4,4,4,4,5,5,5,5,7,7))
axis(1,at = seq(-6,6,1), labels = c("0", "1/32", "1/16", "1/8", "1/4", "1/2",
"1", "2", "4", "8", "16", "32", "64"))
axis(2,at = seq(1,9),labels=5*2^seq(0,8,1))
rect(xleft=-6.5, ybottom=0.5, xright=6.5, ytop=9.5)

colorbar(xlimbar=c(7.5,8.5),ylimbar=c(0.5,9.5),colors=c(4,4,4,4,5,5,5,5,7,7),barlabels=seq(-4,1,0.5))

```

Description

Calculates the linear error in probability spaces. This is the mean absolute difference between the forecast cumulative distribution value (cdf) and the observation. This function creates the empirical cdf function for the observations using the sample population. Linear interpolation is used to estimate the cdf values between observation values. Therefore; this may produce awkward results with small datasets.

Usage

```
leps(x, pred, plot = TRUE, ... )
```

Arguments

x	A vector of observations or a verification object with “cont.cont” properties.
pred	A vector of predictions.
plot	Logical to generate a plot or not.
...	Additional plotting options.

Value

If assigned to an object, the following values are reported.

leps.0	Negatively oriented score on the [0,1] scale, where 0 is a perfect score.
leps.1	Positively oriented score proposed by Potts.

Author(s)

Matt Pocerlich <pocernic@rap.ucar.edu>

References

DeQue, Michel. (2003) “Continuous Variables” **Chapter 5**, *Forecast Verification: A Practitioner’s Guide in Atmospheric Science*.

Potts, J. M., Folland, C.K., Jolliffe, I.T. and Sexton, D. (1996) “Revised ‘LEPS’ scores fore assessing climate model simulations and long-range forecasts.” *J. Climate*, **9**, pp. 34-54.

Examples

```
obs <- rnorm(100, mean = 1, sd = sqrt(50))
pred<- rnorm(100, mean = 10, sd = sqrt(500))

leps(obs, pred, main = "Sample Plot")
## values approximated

OBS <- c(2.7, 2.9, 3.2, 3.3, 3.4, 3.4, 3.5, 3.8, 4, 4.2, 4.4, 4.4, 4.6,
5.8, 6.4)
PRED <- c(2.05, 3.6, 3.05, 4.5, 3.5, 3.0, 3.9, 3.2, 2.4, 5.3, 2.5, 2.8,
3.2, 2.8, 7.5)
```

```
a <- leps(OBS, PRED)
a
```

lines.verify

Add lines to ROC or attribute diagrams

Description

Add lines to attribute and verification diagrams from verify.prob.bin objects.

Usage

```
## S3 method for class 'roc'
lines(x, binormal = FALSE, ...)
## S3 method for class 'attrib'
lines(x, ...)
```

Arguments

x	An object created by the verify function with the prob.bin class
binormal	Logical value indicating whether the lines to be added to a ROC plot are empirical or a binormal fit.
...	Optional arguments for the lines function. These include color, line weight (ltw) and line stype (lty)

Note

This will soon be replaced the a lines command constructed using S4 class properites.

Author(s)

Matt Pocernich <pocernic@rap.ucar.edu>

See Also

[verify](#)

lookup	<i>Lookup table for crps.circ.</i>
--------	------------------------------------

Description

Saved lookup data to expedite integration in the crps.circ function.

See Also

[crps.circ](#)

measurement.error	<i>Skill score with measurement error.</i>
-------------------	--

Description

Skill score that incorporates measurement error. This function allows the user to incorporate measurement error in an observation in a skill score.

Usage

```
measurement.error( obs, frcs = NULL, theta = 0.5, CI =
  FALSE, t = 1, u = 0, h = NULL, ...)
```

Arguments

obs	Information about a forecast and observation can be done in one of two ways. First, the results of a contingency table can be entered as a vector containing c(n11, n10, n01, n00), where n11 are the number of correctly predicted events and n01 is the number of incorrectly predicted non-events. Actual forecasts and observations can be used. In this case, obs is a vector of binary outcomes [0,1].
frcs	If obs is entered as a contingency table, this argument is null. If obs is a vector of outcomes, this column is a vector of probabilistic forecasts.
theta	Loss value (cost) of making a incorrect forecast by a non-event. Defaults to 0.5.
CI	Calculate confidence intervals for skill score.
t	Probability of forecasting an event, when an event occurs. A perfect value is 1.
u	Probability of forecasting that no event will occur, when an event occurs. A perfect value is 0.
h	Threshold for converting a probabilistic forecast into a binary forecast. By default, this value is NULL and the theta is used as this threshold.
...	Optional arguments.

Value

z	Error code
k	Skill score
G	Likelihood ratio statistic
p	p-value for the null hypothesis that the forecast contains skill.
theta	Loss value. Loss associated with an incorrect forecast of a non-event.
ciLO	Lower confidence interval
ciHI	Upper confidence interval

Author(s)

Matt Pocernich <pocernic@rap.ucar.edu> (R - code)

W.M Briggs <wib2004(at)med.cornell.edu> (Method questions)

References

W.M. Briggs, 2004. *Incorporating Cost in the Skill Score* Technical Report, wm-briggs.com/public/skillocost.pdf.

W.M. Briggs and D. Ruppert, 2004. *Assessing the skill of yes/no forecasts. Submitting to Biometrics*.

J.P. Finley, 1884. Tornado forecasts. *Amer. Meteor. J.* 85-88. (Tornado data used in example.)

Examples

```
DAT<- data.frame( obs = round(runif(50)), frcs = runif(50))

A<- measurement.error(DAT$obs, DAT$frcs, CI = TRUE)
A
### Finley Data

measurement.error(c(28, 23, 72, 2680)) ## assuming perfect observation,
```

multi.cont

Multiple Contingency Table Statistics

Description

Provides a variety of statistics for a data summarized in a contingency table. This will work for a 2 by 2 table, but is more useful for tables of greater dimensions.

Usage

```
multi.cont(DAT, baseline = NULL)
```

Arguments

DAT	A contingency table in the form of a matrix. It is assumed that columns represent observation, rows represent forecasts.
baseline	A vector indicating the baseline probabilities of each category. By default, the baseline or naive forecasts is based on teh

Value

pc	Percent correct - events along the diagonal.
bias	Bias
ts	Threat score a.k.a. Critical success index (CSI)
hss	Heidke Skill Score
pss	Peirce Skill Score
gs	Gerrity Score
pc2	Percent correct by category (vector)
h	Hit Rate by category (vector)
false.alarm.ratio	False alarm ratio by category (vector)

Note

Some verification statistics for a contingency table assume that the forecasts and observations are ordered, while others do not. An example of an ordered or ordinal forecast is "low, medium and high". An example of an unordered or nominal forecast is "snow, rain, hail, and none." If the forecasts are ordered, it is possible to account for forecasts which are close to the the observed value. For example, the Gerrity score takes this closeness into account. The Pierce Skill Score does not.

For ordered forecast, it is assumed that the columns and rows of the input matrix are ordered sequentially.

When multiple values are returned, as in the case of pc2, h, f and false.alarm.ratio, these values are conditioned on that category having occurred. For example, in the example included in Jolliffe, given that a below average temperature was observed, the forecast had a bias of 2.3 and had a 0.47 chance of being detected.

Author(s)

Matt Pocerlich <pocernic@rap.ucar.edu>

References

- Gerrity, J.P. Jr (1992). A note on Gandin and Murphy's equitable skill score. *Mon. Weather Rev.*, **120**, 2707-2712.
- Jolliffe, I.T. and D.B. Stephenson (2003). Forecast verification: a practitioner's guide in atmospheric science. John Wiley and Sons. See chapter 4 concerning categorical events, written by R. E. Livezey.

See Also

binary.table

Examples

```
DAT<- matrix(c(7,4,4,14,9,8,14,16,24), nrow = 3) # from p. 80 - Jolliffe
multi.cont(DAT)

DAT<- matrix(c(3,8,7,8,13,14,4,18,25), ncol = 3) ## Jolliffe JJA
multi.cont(DAT)

DAT<- matrix(c(50,47,54,91,2364,205,71,170,3288), ncol = 3) # Wilks p. 245
multi.cont(DAT)

DAT<- matrix(c(28, 23, 72, 2680 ), ncol = 2) ## Finley
multi.cont(DAT)
## Finnish clouds
DAT<- matrix(c(65, 10, 21, 29,17,48, 18, 10, 128), nrow = 3, ncol = 3, byrow = TRUE)
multi.cont(DAT)
### alternatively, the verify function and summary can be used.

mod <- verify(DAT, frcst.type = "cat", obs.type = "cat")
summary(mod)
```

observation.error *Observation Error*

Description

Quantifies observation error through use of a “Gold Standard” of observations.

Usage

```
observation.error(obs, gold.standard = NULL, ...)
```

Arguments

obs	Observation made by method to be quantified. This information can be entered two ways. If obs is a vector of length 4, it is assumed that it contains the values c(n11, n10, n01, n00), where n11 are the number of correctly predicted events and n01 is the number of incorrectly predicted non-events.
gold.standard	The gold standard. This is considered a higher quality observation (coded {0, 1}).
...	Optional arguments.

Value

t	Probability of forecasting an event, when an event occurs. A perfect value is 1.
u	Probability of forecasting that no event will occur, when an event occurs. A perfect value is 0.

Note

This function is used to calculate values for the [measurement.error](#) function.

Author(s)

Matt Pocerlich <pocerlich@rap.ucar.edu>

See Also

[measurement.error](#)

Examples

```
obs <- round(runif(100))
gold <- round(runif(100) )
observation.error(obs, gold)

## Pirep gold standard

observation.error(c(43,10,17,4) )
```

performance.diagram *Performance Diagram*

Description

Creates plot displaying multiple skill scores on a single plot

Usage

```
performance.diagram(...)
```

Arguments

... Optional plotting parameters.

Note

Currently this function just produces the base plot. Points summarizing model performance can be added using the [points](#) function.

Author(s)

Matt Pocerlich <pocernic@rap.ucar.edu>

References

Roebber, P.J. (2009). Visualizing Multiple Measures of Forecast Quality, *Weather and Forecasting*, **24**, pp - 601 - 608.

Examples

```
performance.diagram(main = "Sample Plot")
RB1 <- matrix(c(95, 55, 42, 141), ncol = 2)
## at point
pts <- table.stats(RB1)
boot.pts <- table.stats.boot(RB1, R = 1000 )
## add confidence intervals
segments(x0 = 1 - pts$FAR, y0 = boot.pts["up", "pod"], x1 = 1 - pts$FAR, y1 = boot.pts["dw", "pod"], col = 2, lwd = 2)

segments(x0 = 1 - boot.pts["up", "far"], y0 = pts$POD, x1 = 1 - boot.pts["dw", "far"], y1 = pts$POD, col = 2, lwd = 2)
points(1 - pts$FAR, pts$POD, col = 2, cex = 2)
```

plot.int.scale

Plot Intensity Scale Objects.

Description

Plots objects from the [int.scale.verify](#) using the image function from the base package and image.plot from the fields package..

Usage

```
## S3 method for class 'int.scale'
plot(x, y = NULL, plot.mse = FALSE, main = NULL, ...)
```

Arguments

x	A object from the <code>int.scale.verify</code> function that has the class <code>int.scale</code> .
y	NULL
plot.mse	Should a plot be created of the mean squared errors? By default the skill scores are plotted.
main	Plot title
...	Optional arguments

Author(s)

Matt Pocerlich <pocernic@rap.ucar.edu>

See Also

[int.scale.verify](#), [image](#), [image.plot](#)

Examples

```
data(analysis.dat)
data(forecast.dat)

A<- int.scale.verify(forecast.dat, analysis.dat, thres = c(0, 2^seq(-5,6)))
plot(A)
plot(A, plot.mse = TRUE)
plot(A, main = "Test case")
```

pop	<i>Probability of precipitation (pop) data.</i>
-----	---

Description

These datasets are used to illustrate several functions including [value](#) and [roc.plot](#).

This forecasts are summaries of 24-hour probability of precipitation forecasts were made by the Finnish Meteorological Institute (FMI) during 2003, for daily precipitation in the city of Tampere in south central Finland. Light precipitation is considered rainfall greater than .2 mm. Rainfall accumulation is considered values greater than 4.4 mm. Rows of data either missing forecasts or observations have been removed.

This data has been kindly provided by Dr. Pertti Nurmi of the Finnish Meteorological Institute. <http://www.bom.gov.au/bmrc/wefor/staff/eee/verif/POP3/POP3.html>

Usage

```
data(pop)
```

precip.ensemble	<i>An ensemble of precipitation forecasts</i>
-----------------	---

Description

This is an example of an ensemble of precipitation forecasts. The data set contains forecast for 517 days (3 monsoon seasons) at lead times of 1 to 10 days. Observations and forecasts are in millimeters.

prob.fracs.dat	<i>Probabilistic Forecast Dataset.</i>
----------------	--

Description

This data set is used as an example of data used by the `roc.plot` function. The first column contains a probabilistic forecast for aviation icing. The second column contains a logical variable indicating whether or not icing was observed.

Usage

```
data(prob.fracs.dat)
```

References

PROBABILITY FORECASTS OF IN-FLIGHT ICING CONDITIONS Barbara G. Brown, Ben C. Bernstein, Frank McDonough and Tiffany A. O. Bernstein, 8th Conference on Aviation, Range, and Aerospace Meteorology, Dallas TX, 10-15 January 1999.

probcont2disc	<i>Converts continuous probability values into binned discrete probability forecasts.</i>
---------------	---

Description

Converts continuous probability values into binned discrete probability forecasts. This is useful in calculated Brier type scores for values with continuous probabilities. Each probability is assigned the value of the midpoint.

Usage

```
probcont2disc(x, bins = seq(0,1,0.1) )
```

Arguments

x	A vector of probabilities
bins	Bins. Defaults to 0 to 1 by 0.1 .

Value

A vector of discrete probabilities. E

Note

This function is used within `brier`.

Author(s)

Matt Pocernich <pocernic@rap.ucar.edu>

Examples

```
# probabilistic/ binary example

set.seed(1)
x <- runif(10) ## simulated probabilities.

probcont2disc(x)
probcont2disc(x, bins = seq(0,1,0.25) )

## probcont2disc(4, bins = seq(0,1,0.3)) ## gets error
```

rcrv *Reduced centered random variable*

Description

The RCRV provides information on the reliability of an ensemble system in terms of the bias and the dispersion. A perfectly reliable system has no bias and a dispersion equal to 1. The observational error is taken into account

Usage

```
rcrv(obs, epsMean, epsVariance, obsError)
```

Arguments

obs	A vector of observations
epsMean	A vector of the means of the ensemble
epsVariance	A vector of the variances of the ensemble
obsError	Observational error

Value

bias	The weighted bias between the ensemble and the observation. A value equal to 0 indicates no bias. A positive (negative) value indicates a positive (negative) bias
disp	The dispersion of the ensemble. A value equal to 1 indicates no dispersion. A value greater (smaller) than 1 indicates underdispersion (overdispersion)
y	Vector of y. Mean of y equals bias and standard deviation of y equals dispersion
obsError	Observational error (passed to function)

Author(s)

Ronald Frenette <Ronald.Frenette@ec.gc.ca>

References

G. Candille, C. P. L. Houtekamer, and G. Pellerin: Verification of an Ensemble Prediction System against Observations, *Monthly Weather Review*, **135**, pp. 2688-2699

Examples

```
data(precip.ensemble)
#Observations are in the column
obs<-precip.ensemble[,3]

#Forecast values of ensemble are in the column 4 to 54
eps<-precip.ensemble[,4:54]

#Means and variances of the ensemble
mean<-apply(eps,1,mean)
var<-apply(eps,1,var)

#observation error of 0.5mm
sig0 <- 0.5

rcrv(obs,mean,var,sig0)
```

reliability.plot *Reliability Plot*

Description

A reliability plot is a simple form of an attribute diagram that depicts the performance of a probabilistic forecast for a binary event. In this diagram, the forecast probability is plotted against the observed relative frequency. Ideally, this value should be near to each other and so points falling on the 1:1 line are desirable. For added information, if one or two forecasts are being verified, sharpness diagrams are presented in the corners of the plot. Ideally, these histograms should be relatively flat, indicating that each bin of probabilities is use an appropriate amount of times.

Usage

```
## Default S3 method:
reliability.plot(x, obar.i, prob.y, titl = NULL, legend.names = NULL, ... )
## S3 method for class 'verify'
reliability.plot(x, ...)
```

Arguments

x	Forecast probabilities (y_i) or a “prob.bin” class object from <code>verify</code> .
obar.i	Observed relative frequency \bar{o}_i .
prob.y	Relative frequency of forecasts
titl	Title
legend.names	Names of each model that will appear in the legend.
...	Graphical parameters.

Details

This function works either by entering vectors or on a `verify` class object.

Note

If a single `prob.bin` class object is used, a reliability plot along with a sharpness diagram is displayed. If two forecasts are provided in the form of a matrix of predictions, two sharpness diagrams are provided. If more forecasts are provided, the sharpness diagrams are not displayed.

Author(s)

Matt Pocerlich <pocerlich@rap.ucar.edu>

References

Wilks, D. S. (1995) *Statistical Methods in the Atmospheric Sciences* Chapter 7, San Diego: Academic Press.

Examples

```
## Data from Wilks, table 7.3 page 246.
y.i <- c(0,0.05, seq(0.1, 1, 0.1))
obar.i <- c(0.006, 0.019, 0.059, 0.15, 0.277, 0.377, 0.511, 0.587, 0.723, 0.779, 0.934, 0.933)
prob.y<- c(0.4112, 0.0671, 0.1833, 0.0986, 0.0616, 0.0366, 0.0303, 0.0275, 0.245, 0.022, 0.017, 0.203)
obar<- 0.162

reliability.plot(y.i, obar.i, prob.y, titl = "Test 1", legend.names =
c("Model A") )

## Function will work with a ‘‘prob.bin’’ class object as well.
## Note this is a very bad forecast.
obs<- round(runif(100))
pred<- runif(100)

A<- verify(obs, pred, frcst.type = "prob", obs.type = "binary")

reliability.plot(A, titl = "Alternative plot")
```

roc.area	<i>Area under curve (AUC) calculation for Response Operating Characteristic curve.</i>
----------	--

Description

This function calculates the area underneath a ROC curve following the process outlined in Mason and Graham (2002). The p-value produced is related to the Mann-Whitney U statistics. The p-value is calculated using the wilcox.test function which automatically handles ties and makes approximations for large values.

The p-value addresses the null hypothesis H_0 : The area under the ROC curve is 0.5 i.e. the forecast has no skill.

Usage

```
roc.area(obs, pred)
```

Arguments

obs	A binary observation (coded {0, 1 }).
pred	A probability prediction on the interval [0,1].

Value

A	Area under ROC curve, adjusted for ties in forecasts, if present
n.total	Total number of records
n.events	Number of events
n.noevents	Number of non-events
p.value	Unadjusted p-value

Note

This function is used internally in the roc.plot command to calculate areas.

Author(s)

Matt Pocerlich <pocerlich@rap.ucar.edu>

References

Mason, S.J. and N.E. Graham. (2002) "Areas beneath the relative operating characteristics (ROC) and relative operating levels (ROL) curves: Statistical significance and interpretation," *Q. J. R. Meteorol. Soc.* **30** (1982) 291-303.

See Also

[roc.plot](#), [verify](#), [wilcox.test](#)

Examples

```
# Data used from Mason and Graham (2002).
a<- c(1981, 1982, 1983, 1984, 1985, 1986, 1987, 1988, 1989, 1990,
      1991, 1992, 1993, 1994, 1995)
b<- c(0,0,0,1,1,1,0,1,1,0,0,0,0,1,1)
c<- c(.8, .8, 0, 1,1,.6, .4, .8, 0, 0, .2, 0, 0, 1,1)
d<- c(.928,.576, .008, .944, .832, .816, .136, .584, .032, .016, .28, .024, 0, .984, .952)

A<- data.frame(a,b,c, d)
names(A)<- c("year", "event", "p1", "p2")

## for model with ties
roc.area(A$event, A$p1)

## for model without ties
roc.area(A$event, A$p2)
```

roc.plot

Relative operating characteristic curve.

Description

This function creates Receiver Operating Characteristic (ROC) plots for one or more models. A ROC curve plots the false alarm rate against the hit rate for a probabilistic forecast for a range of thresholds. The area under the curve is viewed as a measure of a forecast's accuracy. A measure of 1 would indicate a perfect model. A measure of 0.5 would indicate a random forecast.

Usage

```
## Default S3 method:
roc.plot(x, pred, thresholds = NULL, binormal =
FALSE, legend = FALSE, leg.text = NULL, plot = "emp", CI = FALSE,
n.boot = 1000, alpha = 0.05, tck = 0.01, plot.thres = seq(0.1,
0.9, 0.1), show.thres = TRUE, main = "ROC Curve", xlab = "False Alarm Rate", ylab = "Hit Rate", extra = FA
## S3 method for class 'prob.bin'
roc.plot(x, ...)
```

Arguments

x A binary observation (coded {0, 1 }) or a verification object.

pred A probability prediction on the interval [0,1]. If multiple models are compared, this may be a matrix where each column represents a different prediction.

thresholds	Thresholds may be provided. These thresholds will be used to calculate the hit rate (H_h) and false alarm rate (F_f). If thresholds is NULL, all unique thresholds are used as a threshold. Alternatively, if the number of bins is specified, thresholds will be calculated using the specified numbers of quantiles.
binormal	If TRUE, in addition to the empirical ROC curve, the binormal ROC curve will be calculated. To get a plot draw, plot must be either “binorm” or “both”.
legend	Binomial. Defaults to FALSE indicating whether a legend should be displayed.
leg.text	Character vector for legend. If NULL, models are labeled “Model A”, “Model B”,...
plot	Either “emp” (default), “binorm” or “both” to determine which plot is shown. If set to NULL, a plot is not created
CI	Confidence Intervals. Calculated by bootstrapping the observations and prediction, then calculating POD _y and POD _n values.
n.boot	Number of bootstrap samples.
alpha	Confidence interval. By default = 0.05
tck	Tick width on confidence interval whiskers.
plot.thres	By default, displays the threshold levels on the ROC diagrams. To suppress these values, set it equal to NULL. If confidence intervals (CI) is set to TRUE, levels specified here will determine where confidence interval boxes are placed.
show.thres	Show thresholds for points indicated by plot.thres. Defaults to TRUE.
main	Title for plot.
xlab, ylab	Plot axes labels. Defaults to “Hit Rate” and “False Alarm Rate”, for the y and x axes respectively.
extra	Extra text describing binormal and empirical lines.
...	Additional plotting options.

Value

If assigned to an object, the following values are reported.

plot.data	The data used to generate the ROC plots. This is a array. Column headers are thresholds, empirical hit and false alarm rates, and binormal hit and false alarm rates. Each model is depicted on an array indexed by the third dimension.
roc.vol	The areas under the ROC curves. By default, this is printed on the plots. Areas and p-values are calculated with and without adjustments for ties along with the p-value for the area. These values are calculated using roc.area . The fifth column contains the area under the binormal curve, if binormal is selected.
A.boot	If confidence intervals are calculated, the area under the ROC curve are returned.

Note

Other packages in R provide functions to create ROC diagrams and different diagnostics. The **ROCR** package provides excellent functions to generate ROC diagrams with lines coded by threshold. Large datasets are handled by a sampling routine and the user may plot a number of threshold

dependent, contingency table scores. Arguably, this is a superior package with respect to ROC plotting.

There is not a minimum size required to create confidence limits or show thresholds. When there are few data points, it is possible to make some pretty unattractive graphs.

The roc.plot method can be used to summarize a "verify, prob.bin" class object created with the verify command. It is appropriate to use the roc plot for forecast which are not probabilities, but rather forecasts made on a continuous scale. The roc plot function can be used to summarize such forecasts but it is not possible to use the verify function to summarize such forecasts. An example is shown below.

Author(s)

Matt Pocerlich <pocerlich@rap.ucar.edu>

References

- Mason, I. (1982) "A model for assessment of weather forecasts," *Aust. Met. Mag* **30** (1982) 291-303.
- Mason, S.J. and N.E. Graham. (2002) "Areas beneath the relative operating characteristics (ROC) and relative operating levels (ROL) curves: Statistical significance and interpretation," *Q. J. R. Meteorol. Soc.* **128** pp. 2145-2166.
- Swets, John A. (1996) *Signal Detection Theory and ROC Analysis in Psychology and Diagnostics*, Lawrence Erlbaum Associates, Inc.

See Also

[pop](#) and [lines.roc](#)

Examples

```
# Data from Mason and Graham article.

a<- c(0,0,0,1,1,1,0,1,1,0,0,0,0,1,1)
b<- c(.8, .8, 0, 1,1,.6, .4, .8, 0, 0, .2, 0, 0, 1,1)
c<- c(.928,.576, .008, .944, .832, .816, .136, .584, .032, .016, .28, .024, 0, .984, .952)

A<- data.frame(a,b,c)
names(A)<- c("event", "p1", "p2")

## for model with ties
roc.plot(A$event, A$p1)

## for model without ties
roc.plot(A$event, A$p2)

### show binormal curve fit.

roc.plot(A$event, A$p2, binormal = TRUE)

# icing forecast
```

```

data(prob.fracs.dat)
A <- verify(prob.fracs.dat$obs, prob.fracs.dat$fracs/100)
roc.plot(A, main = "AWG Forecast")

# plotting a 'prob.bin' class object.
obs<- round(runif(100))
pred<- runif(100)

A<- verify(obs, pred, fracs.type = "prob", obs.type = "binary")

roc.plot(A, main = "Test 1", binormal = TRUE, plot = "both")

## show confidence intervals. MAY BE SLOW
roc.plot(A, threshold = seq(0.1,0.9, 0.1), main = "Test 1", CI = TRUE,
alpha = 0.1)

### example from forecast verification website.
data(pop)
pop.convert() ## internal function used to make binary observations for the pop figure.
### note the use of bins = FALSE !!
mod24 <- verify(d$obs_norain, d$p24_norain, bins = FALSE)

mod48 <- verify(d$obs_norain, d$p48_norain, bins = FALSE)

roc.plot(mod24, plot.thres = NULL)
lines.roc(mod48, col = 2, lwd = 2)
leg.txt <- c("24 hour forecast", "48 hour forecast")
legend( 0.6, 0.4, leg.txt, col = c(1,2), lwd = 2)

```

rps

Ranked Probability Score

Description

Calculates the ranked probability score (rps) and ranked probability skill score (rpss) for probabilistic forecasts of ordered events.

Usage

```
rps(obs, pred, baseline=NULL)
```

Arguments

obs	A vector of observed outcomes. These values correspond to columns of prediction probabilities.
pred	A matrix of probabilities for each outcome occurring. Each column represents a category of prediction.

baseline If NULL (default) the probability based on the sample data of each event to occur. Alternatively, a vector the same length of the as the number categories can be entered.

Value

rps Ranked probability scores
 rps Ranked probability skill score. Uses baseline or sample climatology as a references score.
 rps.clim Ranked probability score for baseline forecast.

Note

Perhaps the format of the data is best understood in the context of an example. Consider a probability of precipitation forecast of "none", "light" or "heavy". This could be [0.5, 0.3, 0.2]. If heavy rain occurred, the observed value would be 3, indicating event summarized in the third column occurred.

The RPS value is scaled to a [0,1] interval by dividing by (number of categories -1 . There is a discrepancy in the way this is explained in Wilks (2005) and the WWRP web page.

Author(s)

Matt Pocerlich <pocerlich@rap.ucar.edu>

References

WWRP/WGNE Joint Working Group on Verification - Forecast Verification - Issues, Methods and FAQ http://www.bom.gov.au/bmrc/wefor/staff/eee/verif/verif_web_page.html#RPS

Wilks, D. S. (2005) *Statistical Methods in the Atmospheric Sciences* Chapter 7, San Diego: Academic Press.

See Also

[crps](#)

Examples

```
### Example from Wilks, note without a baseline and only one
### forecast, the rps and ss are not too meaningfull.
```

```
rps( obs = c(1), pred = matrix(c(0.2, 0.5, 0.3), nrow = 1))
```

table.stats	<i>Verification statistics for a 2 by 2 Contingency Table</i>
-------------	---

Description

Provides a variety of statistics for a data summarized in a 2 by 2 contingency table.

Usage

```
table.stats(obs, pred, silent = FALSE)
```

Arguments

obs	Either a vector of contingency table counts, a vector of binary observations, or a 2 by 2 matrix in the form of a contingency table. (See note below.)
pred	Either null or a vector of binary forecasts.
silent	Should warning statements be suppressed.

Value

tab.out	Contingency table
TS	Threat score a.k.a. Critical success index (CSI)
POD	Hit Rate aka probability of detection
M	Miss rate
F	False Alarm RATE
FAR	False Alarm RATIO
HSS	Heidke Skill Score
PSS	Peirce Skill Score
KSS	Kuiper's Skill Score
PC	Percent correct - events along the diagonal.
BIAS	Bias
ETS	Equitable Threat Score
theta	Odds Ratio
log.theta	Log Odds Ratio
n.h	Degrees of freedom for log.theta
orss	Odds ratio skill score, aka Yules's Q
eds	Extreme Dependency Score
sed	Symmetric Extreme Dependency Score
sed.se	Standard Error for Symmetric Extreme Dependency Score

Note

Initially, table.stats was an internal function used by verify for binary events and multi.cont for categorical events. But occasionally, it is nice to use it directly.

Author(s)

Matt Pocerlich <pocerlich@rap.ucar.edu>

References

Jolliffe, I.T. and D.B. Stephenson (2003). Forecast verification: a practitioner's guide in atmospheric science. John Wiley and Sons. See chapter 3 concerning categorical events.

Stephenson, D.B. (2000). "Use of 'Odds Ratio for Diagnosing Forecast Skill." *Weather and Forecasting* **15** 221-232.

Hogan, R.J., O'Connor E.J. and Illingworth, 2009. "Verification of cloud-fraction forecasts." *Q.J.R. Meteorol. Soc.* **135**, 1494-1511.

See Also

verify and multi.cont

Examples

```
DAT<- matrix(c(28, 23, 72, 2680 ), ncol = 2) ## Finley
table.stats(DAT)
```

table.stats.boot	<i>Percentile bootstrap for 2 by 2 table</i>
------------------	--

Description

Performs a bootstrap on data from a 2 by 2 contingency table returning verification statistics. Potentially useful in creating error bars for performance diagrams.

Usage

```
table.stats.boot(CT, R = 100, alpha = 0.05)
```

Arguments

CT	Two by two contingency table. Columns summarize observed values. Rows summarize forecasted values.
R	Number of resamples
alpha	Confidence intervals.

Value

2 row matrix with upper and lower intervals for bias, pod, far, etc.

Author(s)

Matt Pocerlich <pocerlich@ucar.edu>

See Also

table.stats

Examples

```
### example from Roebber.
RB1 <- matrix(c(95, 55, 42, 141), ncol = 2)
table.stats.boot(RB1, R = 1000 )
```

value	<i>Forecast Value Function</i>
-------	--------------------------------

Description

Calculates the economic value of a forecast based on a cost/loss ratio.

Usage

```
value(obs, pred= NULL, baseline = NULL, cl = seq(0.05, 0.95, 0.05), plot = TRUE, all = FALSE, thresholds
```

Arguments

obs	A vector of binary observations or a contingency table summary of values in the form $c(n11, n01, n10, n00)$ where in nab $a = \text{obs}$, $b = \text{forecast}$.
pred	A vector of probabilistic predictions.
baseline	Baseline or naive forecast. Typically climatology.
cl	Cost loss ratio. The relative value of being unprepared and taking a loss to that of un-necessarily preparing. For example, $cl = 0.1$ indicates it would cost $\$1$ to prevent a $\$10$ loss. This defaults to the sequence 0.05 to 0.95 by 0.05.

plot	Should a plot be created? Default is TRUE
all	In the case of probabilistic forecasts, should value curves for each thresholds be displayed.
thresholds	Thresholds considered for a probabilistic forecast.
ylim, xlim	Plotting options.
...	Options to be passed into the plotting function.

Value

If assigned to an object, the following values are reported.

vmax	Maximum value
V	Vector of values for each cl value
F	Conditional false alarm rate.
H	Conditional hit rate
cl	Vector of cost lost ratios.
s	Base rate

Author(s)

Matt Pocerlich <pocerlich@rap.ucar.edu>

References

Jolliffe, Ian and David B. Stephensen (2003) *Forecast Verification: A Practitioner's Guide in Atmospheric Science*, Chapter 8. Wiley

Examples

```
## value as a contingency table
## Finley tornado data
obs<- c(28, 72, 23, 2680)
value(obs)
aa <- value(obs)
aa$vmax # max value

## probabilistic forecast example
obs <- round(runif(100) )
pred <- runif(100)

value(obs, pred, main = "Sample Plot",
      thresholds = seq(0.02, 0.98, 0.02) )
#####
data(pop)
pop.convert()

value(obs = d$obs_rain, pred = d$p24_rain, all = TRUE)
```

 verify

Verification function

Description

Based on the type of inputs, this function calculates a range of verification statistics and skill scores. Additionally, it creates a verify class object that can be used in further analysis or with other methods such as plot and summary.

Usage

```
verify(obs, pred, baseline = NULL,
       frcst.type = "prob", obs.type = "binary", thresholds =
       seq(0,1,0.1), show = TRUE, bins = TRUE, ... )
```

Arguments

obs	The values with which the verifications are verified. May be a vector of length 4 if the forecast and predictions are binary data summarized in a contingency table. In this case, the value are entered in the order of c(n11, n01, n10, n00). If obs is a matrix, it is assumed to be a contingency table with observed values summarized in the columns and forecasted values summarized in the rows.
pred	Prediction of event. The prediction may be in the form of the a point prediction or the probability of a forecast. Let pred = NULL if obs is a contingency table.
baseline	In meteorology, climatology is the baseline that represents the no-skill forecast. In other fields this field would differ. This field is used to calculate certain skill scores. If left NULL, these statistics are calculated using sample climatology. If this is not NULL, the mean of these values is used as the baseline forecast. This interpretation is not appropriate for all applications. For example, if a baseline forecast is different for each forecast this will not work appropriately.
frcst.type	Forecast type. Either "prob", "binary", "norm.dist", "cat" or "cont". Defaults to "prob". "norm.dist" is used when the forecast is in the form of a normal distribution. See crps for more details.
obs.type	Observation type. Either "binary", "cat" or "cont". Defaults to "binary"
thresholds	Thresholds to be considered for point forecasts of continuous events.
show	Binary; if TRUE (the default), print warning message
bins	Binary; if TRUE (default), the probabilistic forecasts are placed in bins defined by the sequence defined in threshold and assigned the midpoint value.
...	Additional options.

Value

An object of the `verify` class. Depending on the type of data used, the following information may be returned. The following notation is used to describe which values are produced for which type of forecast/observations. (BB = binary/binary, PB = probabilistic/binary, CC = continuous/continuous, CTCT = categorical/categorical)

BS	Brier Score (PB)
BSS	Brier Skill Score(PB)
SS	Skill Score (BB)
hit.rate	Hit rate, aka POD _y , $\$h\$$ (PB, CTCT)
false.alarm.rate	False alarm rate, POD _n , $\$f\$$ (PB, CTCT)
TS	Threat Score or Critical Success Index (CSI)(BB, CTCT)
ETS	Equitable Threat Score (BB, CTCT)
BIAS	Bias (BB, CTCT)
PC	Percent correct or hit rate (BB, CTCT)
Cont.Table	Contingency Table (BB)
HSS	Heidke Skill Score(BB, CTCT)
KSS	Kuniper Skill Score (BB)
PSS	Pierce Skill Score (CTCT)
GS	Gerrity Score (CTCT)
ME	Mean error (CC)
MSE	Mean-squared error (CC)
MAE	Mean absolute error (CC)
theta	Odds Ratio (BB)
log.theta	Log Odds Ratio
n.h	Degrees of freedom for log.theta (BB)
orss	Odds ratio skill score, aka Yule's Q (BB)
eds	Extreme Dependency Score (BB)
sed	Symmetric Extreme Dependency Score (BB)
sed.se	Standard Error for Symmetric Extreme Dependency Score (BB)

Note

There are other packages in R and Bioconductor which are useful for verification tasks. This includes the **ROCR**, **ROC**, package and the "limma" package (in the Bioconductor repository.) Written by people in different fields, each provides tools for verification from different perspectives.

For the categorical forecast and verification, the Gerrity score only makes sense for forecast that have order, or are basically ordinal. It is assumed that the forecasts are listed in order. For example, if the rows of a contingency table were summarized as "medium, low, high", the Gerrity score will be incorrectly summarized.

Author(s)

Matt Pocerlich <pocerlich@ucar.edu>

References

Wilks, D. S. (1995) *Statistical Methods in the Atmospheric Sciences* Chapter 7, San Diego: Academic Press.

WMO Joint WWRP/WGNE Working Group on Verification Website

http://www.bom.gov.au/bmrc/wefor/staff/eee/verif/verif_web_page.html

See Also

table.stats

Examples

```
# binary/binary example
obs<- round(runif(100))
pred<- round(runif(100))

# binary/binary example
# Finley tornado data.

obs<- c(28, 72, 23, 2680)
A<- verify(obs, pred = NULL, frcst.type = "binary", obs.type = "binary")

summary(A)

# categorical/categorical example
# creates a simulated 5 category forecast and observation.
obs <- round(runif(100, 1,5) )
pred <- round(runif(100, 1,5) )

A<- verify(obs, pred, frcst.type = "cat", obs.type = "cat" )
summary(A)

# probabilistic/ binary example

pred<- runif(100)
A<- verify(obs, pred, frcst.type = "prob", obs.type = "binary")
summary(A)

# continuous/ continuous example
obs<- rnorm(100)
pred<- rnorm(100)
baseline <- rnorm(100, sd = 0.5)

A<- verify(obs, pred, baseline = baseline, frcst.type = "cont", obs.type = "cont")
summary(A)
```

Index

*Topic **datasets**

- analysis.dat, 2
- disc.dat, 11
- forecast.dat, 13
- lookup, 21
- pop, 27
- precip.ensemble, 27
- prob.fracs.dat, 28

*Topic **file**

- attribute, 2
- brier, 4
- conditional.quantile, 6
- crps, 7
- crps.circ, 9
- crpsDecompostion, 10
- discrimination.plot, 11
- fss, 13
- int.scale.verify, 14
- IS, 16
- leps, 18
- lines.verify, 20
- measurement.error, 21
- multi.cont, 22
- observation.error, 24
- performance.diagram, 25
- plot.int.scale, 26
- probcont2disc, 28
- reliability.plot, 30
- roc.area, 32
- roc.plot, 33
- rps, 36
- table.stats, 38
- value, 40
- verify, 42

analysis.dat, 2

attribute, 2

brier, 4

conditional.quantile, 6

crps, 7, 37

crps.circ, 9, 21

crpsDecomposition (crpsDecompostion), 10

crpsDecompostion, 10

crpsFromAlphaBeta (crpsDecompostion), 10

disc.dat, 11

discrimination.plot, 11, 11

forecast.dat, 2, 13

fss, 13

image, 27

image.plot, 27

int.scale.verify, 2, 13, 14, 17, 26, 27

IS, 16

leps, 18

lines.attrib (lines.verify), 20

lines.roc, 35

lines.roc (lines.verify), 20

lines.verify, 20

lookup, 21

measurement.error, 21, 25

multi.cont, 22

observation.error, 24

performance.diagram, 25

plot.int.scale, 17, 26

pop, 27, 35

precip.ensemble, 27

prob.fracs.dat, 28

probcont2disc, 28

rcrv, 29

reliability.plot, 4, 30

roc.area, 32, 34

roc.plot, 27, 33, 33

rps, [36](#)

table.stats, [38](#)

table.stats.boot, [39](#)

value, [27](#), [40](#)

verify, [4](#), [20](#), [33](#), [42](#)

wilcox.test, [33](#)