

Package ‘forecast’

June 18, 2017

Version 8.1

Title Forecasting Functions for Time Series and Linear Models

Description Methods and tools for displaying and analysing univariate time series forecasts including exponential smoothing via state space models and automatic ARIMA modelling.

Depends R (>= 3.0.2),

Imports stats, graphics, tseries, fracdiff, Rcpp (>= 0.11.0), nnet, colorspace, parallel, ggplot2 (>= 2.0.0), magrittr, lmtest, zoo, timeDate

Suggests testthat, knitr, rmarkdown, expsmooth, rticles

LinkingTo Rcpp (>= 0.11.0), RcppArmadillo (>= 0.2.35)

LazyData yes

ByteCompile TRUE

BugReports <https://github.com/robjhyndman/forecast/issues>

License GPL (>= 3)

URL <http://github.com/robjhyndman/forecast>

VignetteBuilder knitr

RoxygenNote 6.0.1.9000

NeedsCompilation yes

Author Rob Hyndman [aut, cre, cph],
Mitchell O'Hara-Wild [aut],
Christoph Bergmeir [aut],
Slava Razbash [aut],
Earo Wang [aut]

Maintainer Rob Hyndman <Rob.Hyndman@monash.edu>

Repository CRAN

Date/Publication 2017-06-17 23:07:45 UTC

R topics documented:

forecast-package	4
accuracy	4
Acf	6
arfima	8
Arima	9
arima.errors	11
arimaorder	12
auto.arima	13
autolayer	15
autolayer.mts	16
autoplot.acf	17
autoplot.decomposed.ts	19
autoplot.mforecast	20
baggedETS	22
bats	23
bizdays	25
bld.mbb.bootstrap	26
BoxCox	27
BoxCox.lambda	28
checkresiduals	29
croston	30
CV	32
CVar	33
dm.test	34
dshw	35
easter	37
ets	38
findfrequency	40
fitted.fracdiff	41
forecast	42
forecast.baggedETS	44
forecast.bats	46
forecast.ets	47
forecast.fracdiff	49
forecast.HoltWinters	51
forecast.lm	53
forecast.mlm	54
forecast.mts	56
forecast.nnetar	58
forecast.stl	60
forecast.StructTS	62
fourier	64
gas	65
getResponse	66
gghistogram	67
gglagplot	68

ggmonthplot	69
ggseasonplot	70
ggtsdisplay	71
gold	73
is.acf	73
is.constant	74
is.forecast	74
ma	75
meanf	76
monthdays	77
msts	78
na.interp	79
ndiffs	80
nnetar	82
plot.Arima	84
plot.bats	85
plot.ets	86
plot.forecast	87
residuals.tslm	90
rwf	91
seasadj	93
seasonal	94
seasonaldummy	95
ses	97
simulate.ets	99
sindexf	101
splinef	102
StatForecast	103
subset.ts	106
taylor	107
tbats	108
tbats.components	109
thetaf	110
tsclean	112
tsCV	113
tslm	114
tsoutliers	115
wineind	116
woolyrnq	116

 forecast-package

Forecasting Functions for Time Series and Linear Models

Description

Methods and tools for displaying and analysing univariate time series forecasts including exponential smoothing via state space models and automatic ARIMA modelling.

Details

Package: forecast
 Type: Package
 License: GPL3
 LazyLoad: yes

Author(s)

Rob J Hyndman

Maintainer: Rob.Hyndman@monash.edu

 accuracy

Accuracy measures for a forecast model

Description

Returns range of summary measures of the forecast accuracy. If x is provided, the function measures test set forecast accuracy based on $x-f$. If x is not provided, the function only produces training set accuracy measures of the forecasts based on $f["x"]-fitted(f)$. All measures are defined and discussed in Hyndman and Koehler (2006).

Usage

```
accuracy(f, ...)
```

```
## Default S3 method:
```

```
accuracy(f, x, test = NULL, d = NULL, D = NULL, ...)
```

Arguments

<code>f</code>	An object of class “forecast”, or a numerical vector containing forecasts. It will also work with <code>Arima</code> , <code>ets</code> and <code>lm</code> objects if <code>x</code> is omitted – in which case training set accuracy measures are returned.
<code>...</code>	Additional arguments depending on the specific method.
<code>x</code>	An optional numerical vector containing actual values of the same length as object, or a time series overlapping with the times of <code>f</code> .
<code>test</code>	Indicator of which elements of <code>x</code> and <code>f</code> to test. If <code>test</code> is <code>NULL</code> , all elements are used. Otherwise <code>test</code> is a numeric vector containing the indices of the elements to use in the test.
<code>d</code>	An integer indicating the number of lag-1 differences to be used for the denominator in MASE calculation. Default value is 1 for non-seasonal series and 0 for seasonal series.
<code>D</code>	An integer indicating the number of seasonal differences to be used for the denominator in MASE calculation. Default value is 0 for non-seasonal series and 1 for seasonal series.

Details

The measures calculated are:

- ME: Mean Error
- RMSE: Root Mean Squared Error
- MAE: Mean Absolute Error
- MPE: Mean Percentage Error
- MAPE: Mean Absolute Percentage Error
- MASE: Mean Absolute Scaled Error
- ACF1: Autocorrelation of errors at lag 1.

By default, the MASE calculation is scaled using MAE of training set naive forecasts for non-seasonal time series, training set seasonal naive forecasts for seasonal time series and training set mean forecasts for non-time series data. If `f` is a numerical vector rather than a forecast object, the MASE will not be returned as the training data will not be available.

See Hyndman and Koehler (2006) and Hyndman and Athanasopoulos (2014, Section 2.5) for further details.

Value

Matrix giving forecast accuracy measures.

Author(s)

Rob J Hyndman

References

Hyndman, R.J. and Koehler, A.B. (2006) "Another look at measures of forecast accuracy". *International Journal of Forecasting*, **22**(4), 679-688. Hyndman, R.J. and Athanasopoulos, G. (2014) "Forecasting: principles and practice", OTexts. Section 2.5 "Evaluating forecast accuracy". <http://www.otexts.org/fpp/2/5>.

Examples

```
fit1 <- rwf(EuStockMarkets[1:200,1],h=100)
fit2 <- meanf(EuStockMarkets[1:200,1],h=100)
accuracy(fit1)
accuracy(fit2)
accuracy(fit1,EuStockMarkets[201:300,1])
accuracy(fit2,EuStockMarkets[201:300,1])
plot(fit1)
lines(EuStockMarkets[1:300,1])
```

Acf

(Partial) Autocorrelation and Cross-Correlation Function Estimation

Description

The function `Acf` computes (and by default plots) an estimate of the autocorrelation function of a (possibly multivariate) time series. Function `Pacf` computes (and by default plots) an estimate of the partial autocorrelation function of a (possibly multivariate) time series. Function `Ccf` computes the cross-correlation or cross-covariance of two univariate series.

Usage

```
Acf(x, lag.max = NULL, type = c("correlation", "covariance", "partial"),
    plot = TRUE, na.action = na.contiguous, demean = TRUE, ...)
```

```
Pacf(x, lag.max = NULL, plot = TRUE, na.action = na.contiguous,
     demean = TRUE, ...)
```

```
Ccf(x, y, lag.max = NULL, type = c("correlation", "covariance"),
    plot = TRUE, na.action = na.contiguous, ...)
```

```
taperedacf(x, lag.max = NULL, type = c("correlation", "partial"),
           plot = TRUE, calc.ci = TRUE, level = 95, nsim = 100, ...)
```

```
taperedpacf(x, ...)
```

Arguments

x	a univariate or multivariate (not Ccf) numeric time series object or a numeric vector or matrix.
lag.max	maximum lag at which to calculate the acf. Default is $10 \cdot \log_{10}(N/m)$ where N is the number of observations and m the number of series. Will be automatically limited to one less than the number of observations in the series.
type	character string giving the type of acf to be computed. Allowed values are “correlation” (the default), “covariance” or “partial”.
plot	logical. If TRUE (the default) the resulting acf, pacf or ccf is plotted.
na.action	function to handle missing values. Default is <code>na.contiguous</code> . Useful alternatives are <code>na.pass</code> and <code>na.interp</code> .
demean	Should covariances be about the sample means?
...	Additional arguments passed to the plotting function.
y	a univariate numeric time series object or a numeric vector.
calc.ci	If TRUE, confidence intervals for the ACF/PACF estimates are calculated.
level	Percentage level used for the confidence intervals.
nsim	The number of bootstrap samples used in estimating the confidence intervals.

Details

The functions improve the `acf`, `pacf` and `ccf` functions. The main differences are that `acf` does not plot a spike at lag 0 when `type=="correlation"` (which is redundant) and the horizontal axes show lags in time units rather than seasonal units.

The tapered versions implement the ACF and PACF estimates and plots described in Hyndman (2015), based on the banded and tapered estimates of autocovariance proposed by McMurry and Politis (2010).

Value

The `acf`, `Pacf` and `Ccf` functions return objects of class "acf" as described in `acf` from the stats package. The `taperedacf` and `taperedpacf` functions return objects of class "mpacf".

Author(s)

Rob J Hyndman

References

Hyndman, R.J. (2015). Discussion of “High-dimensional autocovariance matrices and optimal linear prediction”. *Electronic Journal of Statistics*, 9, 792-796.

McMurry, T. L., & Politis, D. N. (2010). Banded and tapered estimates for autocovariance matrices and the linear process bootstrap. *Journal of Time Series Analysis*, 31(6), 471-482.

See Also

[acf](#), [pacf](#), [ccf](#), [tsdisplay](#)

Examples

```
Acf(wineind)
Pacf(wineind)
## Not run:
taperedacf(wineind, nsim=50)
taperedpacf(wineind, nsim=50)

## End(Not run)
```

 arfima

Fit a fractionally differenced ARFIMA model

Description

An ARFIMA(p,d,q) model is selected and estimated automatically using the Hyndman-Khandakar (2008) algorithm to select p and q and the Haslett and Raftery (1989) algorithm to estimate the parameters including d.

Usage

```
arfima(y, drange = c(0, 0.5), estim = c("mle", "ls"), model = NULL,
       lambda = NULL, biasadj = FALSE, x = y, ...)
```

Arguments

y	a univariate time series (numeric vector).
drange	Allowable values of d to be considered. Default of $c(0, 0.5)$ ensures a stationary model is returned.
estim	If <code>estim=="ls"</code> , then the ARMA parameters are calculated using the Haslett-Raftery algorithm. If <code>estim=="mle"</code> , then the ARMA parameters are calculated using full MLE via the arima function.
model	Output from a previous call to <code>arfima</code> . If model is passed, this same model is fitted to y without re-estimating any parameters.
lambda	Box-Cox transformation parameter. Ignored if NULL. Otherwise, data transformed before model is estimated.
biasadj	Use adjusted back-transformed mean for Box-Cox transformations. If TRUE, point forecasts and fitted values are mean forecast. Otherwise, these points can be considered the median of the forecast densities.
x	Deprecated. Included for backwards compatibility.
...	Other arguments passed to auto.arima when selecting p and q.

Details

This function combines `fracdiff` and `auto.arima` to automatically select and estimate an ARFIMA model. The fractional differencing parameter is chosen first assuming an ARFIMA(2,d,0) model. Then the data are fractionally differenced using the estimated `d` and an ARMA model is selected for the resulting time series using `auto.arima`. Finally, the full ARFIMA(p,d,q) model is re-estimated using `fracdiff`. If `estim="mle"`, the ARMA coefficients are refined using `arima`.

Value

A list object of S3 class "fracdiff", which is described in the `fracdiff` documentation. A few additional objects are added to the list including `x` (the original time series), and the residuals and fitted values.

Author(s)

Rob J Hyndman and Farah Yasmeeen

References

J. Haslett and A. E. Raftery (1989) Space-time Modelling with Long-memory Dependence: Assessing Ireland's Wind Power Resource (with discussion); *Applied Statistics* **38**, 1-50.

Hyndman, R.J. and Khandakar, Y. (2008) "Automatic time series forecasting: The forecast package for R", *Journal of Statistical Software*, **26**(3).

See Also

`fracdiff`, `auto.arima`, `forecast.fracdiff`.

Examples

```
library(fracdiff)
x <- fracdiff.sim( 100, ma=-.4, d=.3)$series
fit <- arfima(x)
tsdisplay(residuals(fit))
```

Arima

Fit ARIMA model to univariate time series

Description

Largely a wrapper for the `arima` function in the stats package. The main difference is that this function allows a drift term. It is also possible to take an ARIMA model from a previous call to `Arima` and re-apply it to the data `y`.

Usage

```
Arima(y, order = c(0, 0, 0), seasonal = c(0, 0, 0), xreg = NULL,
      include.mean = TRUE, include.drift = FALSE, include.constant,
      lambda = model$lambda, biasadj = FALSE, method = c("CSS-ML", "ML",
      "CSS"), model = NULL, x = y, ...)
```

Arguments

<code>y</code>	a univariate time series of class <code>ts</code> .
<code>order</code>	A specification of the non-seasonal part of the ARIMA model: the three components (p, d, q) are the AR order, the degree of differencing, and the MA order.
<code>seasonal</code>	A specification of the seasonal part of the ARIMA model, plus the period (which defaults to <code>frequency(y)</code>). This should be a list with components <code>order</code> and <code>period</code> , but a specification of just a numeric vector of length 3 will be turned into a suitable list with the specification as the <code>order</code> .
<code>xreg</code>	Optionally, a vector or matrix of external regressors, which must have the same number of rows as <code>y</code> .
<code>include.mean</code>	Should the ARIMA model include a mean term? The default is <code>TRUE</code> for undifferenced series, <code>FALSE</code> for differenced ones (where a mean would not affect the fit nor predictions).
<code>include.drift</code>	Should the ARIMA model include a linear drift term? (i.e., a linear regression with ARIMA errors is fitted.) The default is <code>FALSE</code> .
<code>include.constant</code>	If <code>TRUE</code> , then <code>include.mean</code> is set to be <code>TRUE</code> for undifferenced series and <code>include.drift</code> is set to be <code>TRUE</code> for differenced series. Note that if there is more than one difference taken, no constant is included regardless of the value of this argument. This is deliberate as otherwise quadratic and higher order polynomial trends would be induced.
<code>lambda</code>	Box-Cox transformation parameter. Ignored if <code>NULL</code> . Otherwise, data transformed before model is estimated.
<code>biasadj</code>	Use adjusted back-transformed mean for Box-Cox transformations. If <code>TRUE</code> , point forecasts and fitted values are mean forecast. Otherwise, these points can be considered the median of the forecast densities.
<code>method</code>	Fitting method: maximum likelihood or minimize conditional sum-of-squares. The default (unless there are missing values) is to use conditional-sum-of-squares to find starting values, then maximum likelihood.
<code>model</code>	Output from a previous call to <code>Arima</code> . If <code>model</code> is passed, this same model is fitted to <code>y</code> without re-estimating any parameters.
<code>x</code>	Deprecated. Included for backwards compatibility.
<code>...</code>	Additional arguments to be passed to <code>arima</code> .

Details

See the `arima` function in the stats package.

Value

See the [arima](#) function in the stats package. The additional objects returned are

x	The time series data
xreg	The regressors used in fitting (when relevant).

Author(s)

Rob J Hyndman

See Also

[auto.arima](#), [forecast.Arima](#).

Examples

```
library(ggplot2)
WWWusage %>%
  Arima(order=c(3,1,0)) %>%
  forecast(h=20) %>%
  autoplot

# Fit model to first few years of AirPassengers data
air.model <- Arima(window(AirPassengers,end=1956+11/12),order=c(0,1,1),
                  seasonal=list(order=c(0,1,1),period=12),lambda=0)
plot(forecast(air.model,h=48))
lines(AirPassengers)

# Apply fitted model to later data
air.model2 <- Arima(window(AirPassengers,start=1957),model=air.model)

# Forecast accuracy measures on the log scale.
# in-sample one-step forecasts.
accuracy(air.model)
# out-of-sample one-step forecasts.
accuracy(air.model2)
# out-of-sample multi-step forecasts
accuracy(forecast(air.model,h=48,lambda=NULL),
        log(window(AirPassengers,start=1957)))
```

arima.errors

Errors from a regression model with ARIMA errors

Description

Returns time series of the regression residuals from a fitted ARIMA model.

Usage

```
arima.errors(object)
```

Arguments

object An object containing a time series model of class Arima.

Details

This is a deprecated function which is identical to `residuals.Arima(object, type="regression")`. Regression residuals are equal to the original data minus the effect of any regression variables. If there are no regression variables, the errors will be identical to the original series (possibly adjusted to have zero mean).

Value

A ts object

Author(s)

Rob J Hyndman

See Also

[residuals.Arima](#).

arimaorder

Return the order of an ARIMA or ARFIMA model

Description

Returns the order of a univariate ARIMA or ARFIMA model.

Usage

```
arimaorder(object)
```

Arguments

object An object of class "Arima", "dQuote" or "fracdiff". Usually the result of a call to `arima`, `Arima`, `auto.arima`, `ar`, `arfima` or `fracdiff`.

Value

A numerical vector giving the values p , d and q of the ARIMA or ARFIMA model. For a seasonal ARIMA model, the returned vector contains the values p , d , q , P , D , Q and m , where m is the period of seasonality.

Author(s)

Rob J Hyndman

See Also[ar](#), [auto.arima](#), [Arima](#), [arima](#), [arfima](#).**Examples**

```
WWWusage %>% auto.arima %>% arimaorder
```

 auto.arima

Fit best ARIMA model to univariate time series

Description

Returns best ARIMA model according to either AIC, AICc or BIC value. The function conducts a search over possible model within the order constraints provided.

Usage

```
auto.arima(y, d = NA, D = NA, max.p = 5, max.q = 5, max.P = 2,
  max.Q = 2, max.order = 5, max.d = 2, max.D = 1, start.p = 2,
  start.q = 2, start.P = 1, start.Q = 1, stationary = FALSE,
  seasonal = TRUE, ic = c("aicc", "aic", "bic"), stepwise = TRUE,
  trace = FALSE, approximation = (length(x) > 150 | frequency(x) > 12),
  truncate = NULL, xreg = NULL, test = c("kpss", "adf", "pp"),
  seasonal.test = c("ocsb", "ch"), allowdrift = TRUE, allowmean = TRUE,
  lambda = NULL, biasadj = FALSE, parallel = FALSE, num.cores = 2,
  x = y, ...)
```

Arguments

y	a univariate time series
d	Order of first-differencing. If missing, will choose a value based on KPSS test.
D	Order of seasonal-differencing. If missing, will choose a value based on OCSB test.
max.p	Maximum value of p
max.q	Maximum value of q
max.P	Maximum value of P
max.Q	Maximum value of Q
max.order	Maximum value of p+q+P+Q if model selection is not stepwise.
max.d	Maximum number of non-seasonal differences
max.D	Maximum number of seasonal differences

<code>start.p</code>	Starting value of p in stepwise procedure.
<code>start.q</code>	Starting value of q in stepwise procedure.
<code>start.P</code>	Starting value of P in stepwise procedure.
<code>start.Q</code>	Starting value of Q in stepwise procedure.
<code>stationary</code>	If TRUE, restricts search to stationary models.
<code>seasonal</code>	If FALSE, restricts search to non-seasonal models.
<code>ic</code>	Information criterion to be used in model selection.
<code>stepwise</code>	If TRUE, will do stepwise selection (faster). Otherwise, it searches over all models. Non-stepwise selection can be very slow, especially for seasonal models.
<code>trace</code>	If TRUE, the list of ARIMA models considered will be reported.
<code>approximation</code>	If TRUE, estimation is via conditional sums of squares and the information criteria used for model selection are approximated. The final model is still computed using maximum likelihood estimation. Approximation should be used for long time series or a high seasonal period to avoid excessive computation times.
<code>truncate</code>	An integer value indicating how many observations to use in model selection. The last truncate values of the series are used to select a model when truncate is not NULL and approximation=TRUE. All observations are used if either truncate=NULL or approximation=FALSE.
<code>xreg</code>	Optionally, a vector or matrix of external regressors, which must have the same number of rows as y.
<code>test</code>	Type of unit root test to use. See ndiffs for details.
<code>seasonal.test</code>	This determines which seasonal unit root test is used. See nsdiffs for details.
<code>allowdrift</code>	If TRUE, models with drift terms are considered.
<code>allowmean</code>	If TRUE, models with a non-zero mean are considered.
<code>lambda</code>	Box-Cox transformation parameter. Ignored if NULL. Otherwise, data transformed before model is estimated.
<code>biasadj</code>	Use adjusted back-transformed mean for Box-Cox transformations. If TRUE, point forecasts and fitted values are mean forecast. Otherwise, these points can be considered the median of the forecast densities.
<code>parallel</code>	If TRUE and <code>stepwise = FALSE</code> , then the specification search is done in parallel. This can give a significant speedup on multicore machines.
<code>num.cores</code>	Allows the user to specify the amount of parallel processes to be used if <code>parallel = TRUE</code> and <code>stepwise = FALSE</code> . If NULL, then the number of logical cores is automatically detected and all available cores are used.
<code>x</code>	Deprecated. Included for backwards compatibility.
<code>...</code>	Additional arguments to be passed to arima .

Details

The default arguments are designed for rapid estimation of models for many time series. If you are analysing just one time series, and can afford to take some more time, it is recommended that you set `stepwise=FALSE` and `approximation=FALSE`.

The number of seasonal differences is sometimes poorly chosen. If your data shows strong seasonality, try setting $D=1$ rather than relying on the automatic selection of D .

Non-stepwise selection can be slow, especially for seasonal data. The stepwise algorithm outlined in Hyndman and Khandakar (2008) is used except that the default method for selecting seasonal differences is now the OCSB test rather than the Canova-Hansen test. There are also some other minor variations to the algorithm described in Hyndman and Khandakar (2008).

Value

Same as for [Arima](#)

Author(s)

Rob J Hyndman

References

Hyndman, R.J. and Khandakar, Y. (2008) "Automatic time series forecasting: The forecast package for R", *Journal of Statistical Software*, **26**(3).

See Also

[Arima](#)

Examples

```
fit <- auto.arima(WWWusage)
plot(forecast(fit,h=20))
```

autolayer

Create a ggplot layer appropriate to a particular data type

Description

autolayer uses ggplot2 to draw a particular layer for an object of a particular class in a single command. This defines the S3 generic that other classes and packages can extend.

Usage

```
autolayer(object, ...)
```

Arguments

object an object, whose class will determine the behaviour of autoplot
... other arguments passed to specific methods

Value

a ggplot layer

See Also

[autoplot](#), [ggplot](#), [fortify](#)

autolayer.mts

Automatically create a ggplot for time series objects

Description

autoplot takes an object of type ts or mts and creates a ggplot object suitable for usage with stat_forecast.

Usage

```
## S3 method for class 'mts'
autolayer(object, colour = TRUE, series = NULL, ...)
```

```
## S3 method for class 'msts'
autolayer(object, series = NULL, ...)
```

```
## S3 method for class 'ts'
autolayer(object, colour = TRUE, series = NULL, ...)
```

```
## S3 method for class 'ts'
autoplot(object, series = NULL, ...)
```

```
## S3 method for class 'mts'
autoplot(object, colour = TRUE, facets = FALSE, ...)
```

```
## S3 method for class 'msts'
autoplot(object, series = NULL, ...)
```

```
## S3 method for class 'ts'
fortify(model, data, ...)
```

Arguments

object	Object of class “ts” or “mts”.
colour	If TRUE, the time series will be assigned a colour aesthetic
series	Identifies the timeseries with a colour, which integrates well with the functionality of geom_forecast .
...	Other plotting parameters to affect the plot.

facets	If TRUE, multiple time series will be faceted (and unless specified, colour is set to FALSE). If FALSE, each series will be assigned a colour.
model	Object of class “ts” to be converted to “data.frame”.
data	Not used (required for fortify method)

Details

`fortify.ts` takes a ts object and converts it into a data frame (for usage with `ggplot2`).

Value

None. Function produces a `ggplot` graph.

Author(s)

Mitchell O’Hara-Wild

See Also

[plot.ts](#), [fortify](#)

Examples

```
library(ggplot2)
autoplot(USAccDeaths)

lungDeaths <- cbind(mdeaths, fdeaths)
autoplot(lungDeaths)
autoplot(lungDeaths, facets=TRUE)
```

autoplot.acf	<i>ggplot (Partial) Autocorrelation and Cross-Correlation Function Estimation and Plotting</i>
--------------	--

Description

Produces a `ggplot` object of their equivalent `Acf`, `Pacf`, `Ccf`, `taperedacf` and `taperedpacf` functions.

Usage

```
## S3 method for class 'acf'
autoplot(object, ci = 0.95, ...)

ggAcf(x, lag.max = NULL, type = c("correlation", "covariance", "partial"),
      plot = TRUE, na.action = na.contiguous, demean = TRUE, ...)
```

```

ggPacf(x, lag.max = NULL, plot = TRUE, na.action = na.contiguous,
       demean = TRUE, ...)

ggCcf(x, y, lag.max = NULL, type = c("correlation", "covariance"),
      plot = TRUE, na.action = na.contiguous, ...)

## S3 method for class 'mpacf'
autoplot(object, ...)

ggtaperedacf(x, lag.max = NULL, type = c("correlation", "partial"),
            plot = TRUE, calc.ci = TRUE, level = 95, nsim = 100, ...)

ggtaperedpacf(x, ...)

```

Arguments

object	Object of class “acf”.
ci	coverage probability for confidence interval. Plotting of the confidence interval is suppressed if ci is zero or negative.
...	Other plotting parameters to affect the plot.
x	a univariate or multivariate (not Ccf) numeric time series object or a numeric vector or matrix.
lag.max	maximum lag at which to calculate the acf.
type	character string giving the type of acf to be computed. Allowed values are “correlation” (the default), “covariance” or “partial”.
plot	logical. If TRUE (the default) the resulting ACF, PACF or CCF is plotted.
na.action	function to handle missing values. Default is <code>na.contiguous</code> . Useful alternatives are <code>na.pass</code> and <code>na.interp</code> .
demean	Should covariances be about the sample means?
y	a univariate numeric time series object or a numeric vector.
calc.ci	If TRUE, confidence intervals for the ACF/PACF estimates are calculated.
level	Percentage level used for the confidence intervals.
nsim	The number of bootstrap samples used in estimating the confidence intervals.

Details

If `autoplot` is given an `acf` or `mpacf` object, then an appropriate `ggplot` object will be created.

`ggtaperedpacf`

Value

A `ggplot` object.

Author(s)

Mitchell O’Hara-Wild

See Also

[plot.acf](#), [Acf](#), [acf](#), [taperedacf](#)

Examples

```
library(ggplot2)
ggAcf(wineind)
wineind %>% Acf(plot=FALSE) %>% autoplot
## Not run:
wineind %>% taperedacf(plot=FALSE) %>% autoplot
ggtaperedacf(wineind)
ggtaperedpacf(wineind)
## End(Not run)
ggCcfc(mdeaths, fdeaths)
```

autoplot.decomposed.ts

Plot time series decomposition components using ggplot

Description

Produces a ggplot object of seasonally decomposed time series for objects of class “stl” (created with [stl](#)), class “seas” (created with [seas](#)), or class “decomposed.ts” (created with [decompose](#)).

Usage

```
## S3 method for class 'decomposed.ts'
autoplot(object, labels = NULL, range.bars = NULL,
  ...)

## S3 method for class 'stl'
autoplot(object, labels = NULL, range.bars = TRUE, ...)

## S3 method for class 'StructTS'
autoplot(object, labels = NULL, range.bars = TRUE, ...)

## S3 method for class 'seas'
autoplot(object, labels = NULL, range.bars = NULL, ...)
```

Arguments

object	Object of class “seas”, “stl”, or “decomposed.ts”.
labels	Labels to replace “seasonal”, “trend”, and “remainder”.
range.bars	Logical indicating if each plot should have a bar at its right side representing relative size. If NULL, automatic selection takes place.
...	Other plotting parameters to affect the plot.

Value

Returns an object of class ggplot.

Author(s)

Mitchell O'Hara-Wild

See Also

[seas](#), [stl](#), [decompose](#), [StructTS](#), [plot.stl](#).

Examples

```
library(ggplot2)
co2 %>% decompose %>% autoplot
nottem %>% stl(s.window='periodic') %>% autoplot

## Not run:
library(seasonal)
seas(USAccDeaths) %>% autoplot

## End(Not run)
```

autoplot.mforecast *Multivariate forecast plot*

Description

Plots historical data with multivariate forecasts and prediction intervals.

Usage

```
## S3 method for class 'mforecast'
autoplot(object, PI = TRUE, facets = TRUE,
  colour = FALSE, ...)

## S3 method for class 'mforecast'
autolayer(object, series = NULL, PI = TRUE, ...)

## S3 method for class 'mforecast'
plot(x, main = paste("Forecasts from", unique(x$method)),
  xlab = "time", ...)
```

Arguments

object	Multivariate forecast object of class mforecast. Used for ggplot graphics (S3 method consistency).
PI	If FALSE, confidence intervals will not be plotted, giving only the forecast line.
facets	If TRUE, multiple time series will be faceted. If FALSE, each series will be assigned a colour.
colour	If TRUE, the time series will be assigned a colour aesthetic
...	additional arguments to each individual plot.
series	Matches an unidentified forecast layer with a coloured object on the plot.
x	Multivariate forecast object of class mforecast.
main	Main title. Default is the forecast method. For autoplot, specify a vector of titles for each plot.
xlab	X-axis label. For autoplot, specify a vector of labels for each plot.

Details

autoplot will produce an equivalent plot as a ggplot object.

Author(s)

Mitchell O'Hara-Wild

References

Hyndman and Athanasopoulos (2014) *Forecasting: principles and practice*, OTexts: Melbourne, Australia. <http://www.otexts.org/fpp/>

See Also

[plot.forecast](#), [plot.ts](#)

Examples

```
library(ggplot2)

lungDeaths <- cbind(mdeaths, fdeaths)
fit <- tslm(lungDeaths ~ trend + season)
fcast <- forecast(fit, h=10)
plot(fcast)
autoplot(fcast)

carPower <- as.matrix(mtcars[,c("qsec", "hp")])
carmpg <- mtcars[, "mpg"]
fit <- lm(carPower ~ carmpg)
fcast <- forecast(fit, newdata=data.frame(carmpg=30))
plot(fcast, xlab="Year")
autoplot(fcast, xlab=rep("Year", 2))
```

 baggedETS

Forecasting using the bagged ETS method

Description

The bagged ETS forecasting method.

Usage

```
baggedETS(y, bootstrapped_series = bld.mbb.bootstrap(y, 100), ...)
```

Arguments

`y` A numeric vector or time series of class `ts`.
`bootstrapped_series` bootstrapped versions of `y`.
`...` Other arguments passed to [ets](#).

Details

This function implements the bagged ETS forecasting method described in Bergmeir et al. The [ets](#) function is applied to all bootstrapped series. Using the default parameters, the function [bld.mbb.bootstrap](#) is used to calculate the bootstrapped series with the Box-Cox and Loess-based decomposition (BLD) bootstrap. The function [forecast.baggedETS](#) can then be used to calculate forecasts.

Value

Returns an object of class "baggedETS".

The function `print` is used to obtain and print a summary of the results.

`models` A list containing the fitted ETS ensemble models.
`method` The name of the forecasting method as a character string
`y` The original time series.
`bootstrapped_series` The bootstrapped series.
`etsargs` The arguments passed through to [ets](#).
`fitted` Fitted values (one-step forecasts). The mean of the fitted values is calculated over the ensemble.
`residuals` Original values minus fitted values.

Author(s)

Christoph Bergmeir, Fotios Petropoulos

References

Bergmeir, C., R. J. Hyndman, and J. M. Benitez (2016). Bagging Exponential Smoothing Methods using STL Decomposition and Box-Cox Transformation. *International Journal of Forecasting* 32, 303-312.

Examples

```
fit <- baggedETS(WWWusage)
fcast <- forecast(fit)
plot(fcast)
```

bats	<i>BATS model (Exponential smoothing state space model with Box-Cox transformation, ARMA errors, Trend and Seasonal components)</i>
------	---

Description

Fits a BATS model applied to y , as described in De Livera, Hyndman & Snyder (2011). Parallel processing is used by default to speed up the computations.

Usage

```
bats(y, use.box.cox = NULL, use.trend = NULL, use.damped.trend = NULL,
     seasonal.periods = NULL, use.arma.errors = TRUE,
     use.parallel = length(y) > 1000, num.cores = 2, bc.lower = 0,
     bc.upper = 1, biasadj = FALSE, model = NULL, ...)
```

Arguments

<code>y</code>	The time series to be forecast. Can be numeric, msts or ts. Only univariate time series are supported.
<code>use.box.cox</code>	TRUE/FALSE indicates whether to use the Box-Cox transformation or not. If NULL then both are tried and the best fit is selected by AIC.
<code>use.trend</code>	TRUE/FALSE indicates whether to include a trend or not. If NULL then both are tried and the best fit is selected by AIC.
<code>use.damped.trend</code>	TRUE/FALSE indicates whether to include a damping parameter in the trend or not. If NULL then both are tried and the best fit is selected by AIC.
<code>seasonal.periods</code>	If <code>y</code> is a numeric then seasonal periods can be specified with this parameter.
<code>use.arma.errors</code>	TRUE/FALSE indicates whether to include ARMA errors or not. If TRUE the best fit is selected by AIC. If FALSE then the selection algorithm does not consider ARMA errors.
<code>use.parallel</code>	TRUE/FALSE indicates whether or not to use parallel processing.

<code>num.cores</code>	The number of parallel processes to be used if using parallel processing. If NULL then the number of logical cores is detected and all available cores are used.
<code>bc.lower</code>	The lower limit (inclusive) for the Box-Cox transformation.
<code>bc.upper</code>	The upper limit (inclusive) for the Box-Cox transformation.
<code>biasadj</code>	Use adjusted back-transformed mean for Box-Cox transformations. If TRUE, point forecasts and fitted values are mean forecast. Otherwise, these points can be considered the median of the forecast densities.
<code>model</code>	Output from a previous call to <code>bats</code> . If <code>model</code> is passed, this same model is fitted to <code>y</code> without re-estimating any parameters.
<code>...</code>	Additional arguments to be passed to <code>auto.arima</code> when choose an ARMA(p, q) model for the errors. (Note that <code>xreg</code> will be ignored, as will any arguments concerning seasonality and differencing, but arguments controlling the values of <code>p</code> and <code>q</code> will be used.)

Value

An object of class "bats". The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `bats` and associated functions. The fitted model is designated BATS(ω , p, q , ϕ , m_1, \dots, m_J) where ω is the Box-Cox parameter and ϕ is the damping parameter; the error is modelled as an ARMA(p, q) process and m_1, \dots, m_J list the seasonal periods used in the model.

Author(s)

Slava Razbash and Rob J Hyndman

References

De Livera, A.M., Hyndman, R.J., & Snyder, R. D. (2011), Forecasting time series with complex seasonal patterns using exponential smoothing, *Journal of the American Statistical Association*, **106**(496), 1513-1527.

Examples

```
## Not run:
fit <- bats(USAccDeaths)
plot(forecast(fit))

taylor.fit <- bats(taylor)
plot(forecast(taylor.fit))
## End(Not run)
```

bizdays	<i>Number of trading days in each season</i>
---------	--

Description

Returns number of trading days in each month or quarter of the observed time period in a major financial center.

Usage

```
bizdays(x, FinCenter = c("New York", "London", "NERC", "Tokyo", "Zurich"))
```

Arguments

x	Monthly or quarterly time series
FinCenter	Major financial center.

Details

Useful for trading days length adjustments. More on how to define "business days", please refer to [isBizday](#).

Value

Time series

Author(s)

Earo Wang

See Also

[monthdays](#)

Examples

```
x <- ts(rnorm(30), start = c(2013, 2), frequency = 12)
bizdays(x, FinCenter = "New York")
```

bld.mbb.bootstrap *Box-Cox and Loess-based decomposition bootstrap.*

Description

Generates bootstrapped versions of a time series using the Box-Cox and Loess-based decomposition bootstrap.

Usage

```
bld.mbb.bootstrap(x, num, block_size = if (frequency(x) > 1) 2 * frequency(x)
  else 8)
```

Arguments

x	Original time series.
num	Number of bootstrapped versions to generate.
block_size	Block size for the moving block bootstrap.

Details

The procedure is described in Bergmeir et al. Box-Cox decomposition is applied, together with STL or Loess (for non-seasonal time series), and the remainder is bootstrapped using a moving block bootstrap.

Value

A list with bootstrapped versions of the series. The first series in the list is the original series.

Author(s)

Christoph Bergmeir, Fotios Petropoulos

References

Bergmeir, C., R. J. Hyndman, and J. M. Benitez (2016). Bagging Exponential Smoothing Methods using STL Decomposition and Box-Cox Transformation. *International Journal of Forecasting* 32, 303-312.

See Also

[baggedETS](#).

Examples

```
bootstrapped_series <- bld.mbb.bootstrap(WWWusage, 100)
```

BoxCox	<i>Box Cox Transformation</i>
--------	-------------------------------

Description

BoxCox() returns a transformation of the input variable using a Box-Cox transformation. InvBoxCox() reverses the transformation.

Usage

BoxCox(x, lambda)

InvBoxCox(x, lambda, biasadj = FALSE, fvar = NULL)

Arguments

x	a numeric vector or time series of class ts.
lambda	transformation parameter.
biasadj	Use adjusted back-transformed mean for Box-Cox transformations. If TRUE, point forecasts and fitted values are mean forecast. Otherwise, these points can be considered the median of the forecast densities.
fvar	Optional parameter required if biasadj=TRUE. Can either be the forecast variance, or a list containing the interval level, and the corresponding upper and lower intervals.

Details

The Box-Cox transformation is given by

$$f_{\lambda}(x) = \frac{x^{\lambda} - 1}{\lambda}$$

if $\lambda \neq 0$. For $\lambda = 0$,

$$f_0(x) = \log(x)$$

Value

a numeric vector of the same length as x.

Author(s)

Rob J Hyndman & Mitchell O'Hara-Wild

References

Box, G. E. P. and Cox, D. R. (1964) An analysis of transformations. *JRSS B* **26** 211–246.

See Also[BoxCox.lambda](#)**Examples**

```
lambda <- BoxCox.lambda(lynx)
lynx.fit <- ar(BoxCox(lynx,lambda))
plot(forecast(lynx.fit,h=20,lambda=lambda))
```

`BoxCox.lambda`*Automatic selection of Box-Cox transformation parameter*

Description

If `method=="guerrero"`, Guerrero's (1993) method is used, where `lambda` minimizes the coefficient of variation for subseries of `x`.

Usage

```
BoxCox.lambda(x, method = c("guerrero", "loglik"), lower = -1, upper = 2)
```

Arguments

<code>x</code>	a numeric vector or time series of class <code>ts</code>
<code>method</code>	Choose method to be used in calculating <code>lambda</code> .
<code>lower</code>	Lower limit for possible <code>lambda</code> values.
<code>upper</code>	Upper limit for possible <code>lambda</code> values.

Details

If `method=="loglik"`, the value of `lambda` is chosen to maximize the profile log likelihood of a linear model fitted to `x`. For non-seasonal data, a linear time trend is fitted while for seasonal data, a linear time trend with seasonal dummy variables is used.

Value

a number indicating the Box-Cox transformation parameter.

Author(s)

Leanne Chhay and Rob J Hyndman

References

- Box, G. E. P. and Cox, D. R. (1964) An analysis of transformations. *JRSS B* **26** 211–246.
- Guerrero, V.M. (1993) Time-series analysis supported by power transformations. *Journal of Forecasting*, **12**, 37–48.

See Also

[BoxCox](#)

Examples

```
lambda <- BoxCox.lambda(AirPassengers,lower=0)
air.fit <- Arima(AirPassengers, order=c(0,1,1),
                seasonal=list(order=c(0,1,1),period=12), lambda=lambda)
plot(forecast(air.fit))
```

checkresiduals *Check that residuals from a time series model look like white noise*

Description

If `plot=TRUE`, produces a time plot of the residuals, the corresponding ACF, and a histogram. If the degrees of freedom for the model can be determined and `test` is not `FALSE`, the output from either a Ljung-Box test or Breusch-Godfrey test is printed.

Usage

```
checkresiduals(object, lag, df = NULL, test, plot = TRUE, ...)
```

Arguments

<code>object</code>	Either a time series model, a forecast object, or a time series (assumed to be residuals).
<code>lag</code>	Number of lags to use in the Ljung-Box or Breusch-Godfrey test. If missing, it is set to $\max(10, df+3)$ for non-seasonal data, and $\max(2m, df+3)$ for seasonal data, where <code>df</code> is the degrees of freedom of the model, and <code>m</code> is the seasonal period of the data.
<code>df</code>	Number of degrees of freedom for fitted model, required for the Ljung-Box or Breusch-Godfrey test. Ignored if the degrees of freedom can be extracted from <code>object</code> .
<code>test</code>	Test to use for serial correlation. By default, if <code>object</code> is of class <code>lm</code> , then <code>test="BG"</code> . Otherwise, <code>test="LB"</code> . Setting <code>test=FALSE</code> will prevent the test results being printed.
<code>plot</code>	Logical. If <code>TRUE</code> , will produce the plot.
<code>...</code>	Other arguments are passed to ggtsdisplay .

Value

None

Author(s)

Rob J Hyndman

See Also[ggtsdisplay](#), [Box.test](#), [bgtest](#)**Examples**

```
fit <- ets(WWWusage)
checkresiduals(fit)
```

`croston`*Forecasts for intermittent demand using Croston's method*

Description

Returns forecasts and other information for Croston's forecasts applied to `y`.

Usage

```
croston(y, h = 10, alpha = 0.1, x = y)
```

Arguments

<code>y</code>	a numeric vector or time series of class <code>ts</code>
<code>h</code>	Number of periods for forecasting.
<code>alpha</code>	Value of alpha. Default value is 0.1.
<code>x</code>	Deprecated. Included for backwards compatibility.

Details

Based on Croston's (1972) method for intermittent demand forecasting, also described in Shenstone and Hyndman (2005). Croston's method involves using simple exponential smoothing (SES) on the non-zero elements of the time series and a separate application of SES to the times between non-zero elements of the time series. The smoothing parameters of the two applications of SES are assumed to be equal and are denoted by `alpha`.

Note that prediction intervals are not computed as Croston's method has no underlying stochastic model. The separate forecasts for the non-zero demands, and for the times between non-zero demands do have prediction intervals based on ETS(A,N,N) models.

Value

An object of class "forecast" is a list containing at least the following elements:

model	A list containing information about the fitted model. The first element gives the model used for non-zero demands. The second element gives the model used for times between non-zero demands. Both elements are of class forecast.
method	The name of the forecasting method as a character string
mean	Point forecasts as a time series
x	The original time series (either object itself or the time series used to create the model stored as object).
residuals	Residuals from the fitted model. That is y minus fitted values.
fitted	Fitted values (one-step forecasts)

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `croston` and associated functions.

Author(s)

Rob J Hyndman

References

Croston, J. (1972) "Forecasting and stock control for intermittent demands", *Operational Research Quarterly*, **23**(3), 289-303.

Shenstone, L., and Hyndman, R.J. (2005) "Stochastic models underlying Croston's method for intermittent demand forecasting". *Journal of Forecasting*, **24**, 389-402.

See Also

[ses](#).

Examples

```
y <- rpois(20,lambda=.3)
fcast <- croston(y)
plot(fcast)
```

CV	<i>Cross-validation statistic</i>
----	-----------------------------------

Description

Computes the leave-one-out cross-validation statistic (also known as PRESS – prediction residual sum of squares), AIC, corrected AIC, BIC and adjusted R² values for a linear model.

Usage

```
CV(obj)
```

Arguments

obj output from `lm` or `tslm`

Value

Numerical vector containing CV, AIC, AICc, BIC and AdjR2 values.

Author(s)

Rob J Hyndman

See Also

[AIC](#)

Examples

```
y <- ts(rnorm(120,0,3) + 20*sin(2*pi*(1:120)/12), frequency=12)
fit1 <- tslm(y ~ trend + season)
fit2 <- tslm(y ~ season)
CV(fit1)
CV(fit2)
```

CVar *k-fold Cross-Validation applied to an autoregressive model*

Description

CVar computes the errors obtained by applying an autoregressive modelling function to subsets of the time series y using k -fold cross-validation as described in Bergmeir, Hyndman and Koo (2015).

Usage

```
CVar(y, k = 10, FUN = nnetar, cvtrace = FALSE, ...)
```

Arguments

<code>y</code>	Univariate time series
<code>k</code>	Number of folds to use for cross-validation.
<code>FUN</code>	Function to fit an autoregressive model. Currently, it only works with the nnetar function.
<code>cvtrace</code>	Provide progress information.
<code>...</code>	Other arguments are passed to FUN.

Value

A list containing information about the model and accuracy for each fold, plus other summary information computed across folds.

Author(s)

Gabriel Caceres and Rob J Hyndman

References

Bergmeir, C., Hyndman, R.J., Koo, B. (2015) A note on the validity of cross-validation for evaluating time series prediction. Monash working paper 10/15. <https://robjhyndman.com/publications/cv-time-series/>.

See Also

[CV](#), [tsCV](#).

Examples

```
modelcv <- CVar(lynx, k=5, lambda=0.15)
print(modelcv)
print(modelcv$fold1)
```

dm.test

Diebold-Mariano test for predictive accuracy

Description

The Diebold-Mariano test compares the forecast accuracy of two forecast methods.

Usage

```
dm.test(e1, e2, alternative = c("two.sided", "less", "greater"), h = 1,
        power = 2)
```

Arguments

e1	Forecast errors from method 1.
e2	Forecast errors from method 2.
alternative	a character string specifying the alternative hypothesis, must be one of "two.sided" (default), "greater" or "less". You can specify just the initial letter.
h	The forecast horizon used in calculating e1 and e2.
power	The power used in the loss function. Usually 1 or 2.

Details

This function implements the modified test proposed by Harvey, Leybourne and Newbold (1997). The null hypothesis is that the two methods have the same forecast accuracy. For `alternative="less"`, the alternative hypothesis is that method 2 is less accurate than method 1. For `alternative="greater"`, the alternative hypothesis is that method 2 is more accurate than method 1. For `alternative="two.sided"`, the alternative hypothesis is that method 1 and method 2 have different levels of accuracy.

Value

A list with class "htest" containing the following components:

statistic	the value of the DM-statistic.
parameter	the forecast horizon and loss function power used in the test.
alternative	a character string describing the alternative hypothesis.
p.value	the p-value for the test.
method	a character string with the value "Diebold-Mariano Test".
data.name	a character vector giving the names of the two error series.

Author(s)

George Athanasopoulos

References

Diebold, F.X. and Mariano, R.S. (1995) Comparing predictive accuracy. *Journal of Business and Economic Statistics*, **13**, 253-263.

Harvey, D., Leybourne, S., & Newbold, P. (1997). Testing the equality of prediction mean squared errors. *International Journal of forecasting*, **13**(2), 281-291.

Examples

```
# Test on in-sample one-step forecasts
f1 <- ets(WWWusage)
f2 <- auto.arima(WWWusage)
accuracy(f1)
accuracy(f2)
dm.test(residuals(f1),residuals(f2),h=1)

# Test on out-of-sample one-step forecasts
f1 <- ets(WWWusage[1:80])
f2 <- auto.arima(WWWusage[1:80])
f1.out <- ets(WWWusage[81:100],model=f1)
f2.out <- Arima(WWWusage[81:100],model=f2)
accuracy(f1.out)
accuracy(f2.out)
dm.test(residuals(f1.out),residuals(f2.out),h=1)
```

dshw

Double-Seasonal Holt-Winters Forecasting

Description

Returns forecasts using Taylor's (2003) Double-Seasonal Holt-Winters method.

Usage

```
dshw(y, period1 = NULL, period2 = NULL, h = 2 * max(period1, period2),
     alpha = NULL, beta = NULL, gamma = NULL, omega = NULL, phi = NULL,
     lambda = NULL, biasadj = FALSE, armethod = TRUE, model = NULL)
```

Arguments

y	Either an <code>msts</code> object with two seasonal periods or a numeric vector.
period1	Period of the shorter seasonal period. Only used if y is not an <code>msts</code> object.
period2	Period of the longer seasonal period. Only used if y is not an <code>msts</code> object.
h	Number of periods for forecasting.
alpha	Smoothing parameter for the level. If NULL, the parameter is estimated using least squares.

beta	Smoothing parameter for the slope. If NULL, the parameter is estimated using least squares.
gamma	Smoothing parameter for the first seasonal period. If NULL, the parameter is estimated using least squares.
omega	Smoothing parameter for the second seasonal period. If NULL, the parameter is estimated using least squares.
phi	Autoregressive parameter. If NULL, the parameter is estimated using least squares.
lambda	Box-Cox transformation parameter. Ignored if NULL. Otherwise, data transformed before model is estimated.
biasadj	Use adjusted back-transformed mean for Box-Cox transformations. If TRUE, point forecasts and fitted values are mean forecast. Otherwise, these points can be considered the median of the forecast densities. By default, the value is taken from what was used when fitting the model.
armethod	If TRUE, the forecasts are adjusted using an AR(1) model for the errors.
model	If it's specified, an existing model is applied to a new data set.

Details

Taylor's (2003) double-seasonal Holt-Winters method uses additive trend and multiplicative seasonality, where there are two seasonal components which are multiplied together. For example, with a series of half-hourly data, one would set `period1=48` for the daily period and `period2=336` for the weekly period. The smoothing parameter notation used here is different from that in Taylor (2003); instead it matches that used in Hyndman et al (2008) and that used for the `ets` function.

Value

An object of class "forecast" which is a list that includes the following elements:

model	A list containing information about the fitted model
method	The name of the forecasting method as a character string
mean	Point forecasts as a time series
x	The original time series.
residuals	Residuals from the fitted model. That is x minus fitted values.
fitted	Fitted values (one-step forecasts)

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `dshw`.

Author(s)

Rob J Hyndman

References

Taylor, J.W. (2003) Short-term electricity demand forecasting using double seasonal exponential smoothing. *Journal of the Operational Research Society*, **54**, 799-805.

Hyndman, R.J., Koehler, A.B., Ord, J.K., and Snyder, R.D. (2008) *Forecasting with exponential smoothing: the state space approach*, Springer-Verlag. <http://www.exponentialsmoothing.net>.

See Also

[HoltWinters, etc.](#)

Examples

```
## Not run:
fcast <- dshw(taylor)
plot(fcast)

t <- seq(0,5,by=1/20)
x <- exp(sin(2*pi*t) + cos(2*pi*t*4) + rnorm(length(t),0,.1))
fit <- dshw(x,20,5)
plot(fit)

## End(Not run)
```

easter

Easter holidays in each season

Description

Returns a vector of 0's and 1's or fractional results if Easter spans March and April in the observed time period. Easter is defined as the days from Good Friday to Easter Sunday inclusively, plus optionally Easter Monday if `easter.mon=TRUE`.

Usage

```
easter(x, easter.mon = FALSE)
```

Arguments

`x` Monthly or quarterly time series
`easter.mon` If TRUE, the length of Easter holidays includes Easter Monday.

Details

Useful for adjusting calendar effects.

Value

Time series

Author(s)

Earo Wang

Examples

```
easter(wineind, easter.mon = TRUE)
```

ets

Exponential smoothing state space model

Description

Returns ets model applied to *y*.

Usage

```
ets(y, model = "ZZZ", damped = NULL, alpha = NULL, beta = NULL,
    gamma = NULL, phi = NULL, additive.only = FALSE, lambda = NULL,
    biasadj = FALSE, lower = c(rep(1e-04, 3), 0.8), upper = c(rep(0.9999,
    3), 0.98), opt.crit = c("lik", "amse", "mse", "sigma", "mae"), nmse = 3,
    bounds = c("both", "usual", "admissible"), ic = c("aicc", "aic", "bic"),
    restrict = TRUE, allow.multiplicative.trend = FALSE,
    use.initial.values = FALSE, ...)
```

Arguments

<i>y</i>	a numeric vector or time series of class <i>ts</i>
<i>model</i>	Usually a three-character string identifying method using the framework terminology of Hyndman et al. (2002) and Hyndman et al. (2008). The first letter denotes the error type ("A", "M" or "Z"); the second letter denotes the trend type ("N", "A", "M" or "Z"); and the third letter denotes the season type ("N", "A", "M" or "Z"). In all cases, "N"=none, "A"=additive, "M"=multiplicative and "Z"=automatically selected. So, for example, "ANN" is simple exponential smoothing with additive errors, "MAM" is multiplicative Holt-Winters' method with multiplicative errors, and so on. It is also possible for the model to be of class "ets", and equal to the output from a previous call to <i>ets</i> . In this case, the same model is fitted to <i>y</i> without re-estimating any smoothing parameters. See also the <i>use.initial.values</i> argument.

damped	If TRUE, use a damped trend (either additive or multiplicative). If NULL, both damped and non-damped trends will be tried and the best model (according to the information criterion ic) returned.
alpha	Value of alpha. If NULL, it is estimated.
beta	Value of beta. If NULL, it is estimated.
gamma	Value of gamma. If NULL, it is estimated.
phi	Value of phi. If NULL, it is estimated.
additive.only	If TRUE, will only consider additive models. Default is FALSE.
lambda	Box-Cox transformation parameter. Ignored if NULL. Otherwise, data transformed before model is estimated. When lambda is specified, additive.only is set to TRUE.
biasadj	Use adjusted back-transformed mean for Box-Cox transformations. If TRUE, point forecasts and fitted values are mean forecast. Otherwise, these points can be considered the median of the forecast densities.
lower	Lower bounds for the parameters (alpha, beta, gamma, phi)
upper	Upper bounds for the parameters (alpha, beta, gamma, phi)
opt.crit	Optimization criterion. One of "mse" (Mean Square Error), "amse" (Average MSE over first nmse forecast horizons), "sigma" (Standard deviation of residuals), "mae" (Mean of absolute residuals), or "lik" (Log-likelihood, the default).
nmse	Number of steps for average multistep MSE ($1 \leq \text{nmse} \leq 30$).
bounds	Type of parameter space to impose: "usual" indicates all parameters must lie between specified lower and upper bounds; "admissible" indicates parameters must lie in the admissible space; "both" (default) takes the intersection of these regions.
ic	Information criterion to be used in model selection.
restrict	If TRUE (default), the models with infinite variance will not be allowed.
allow.multiplicative.trend	If TRUE, models with multiplicative trend are allowed when searching for a model. Otherwise, the model space excludes them. This argument is ignored if a multiplicative trend model is explicitly requested (e.g., using model="MMN").
use.initial.values	If TRUE and model is of class "ets", then the initial values in the model are also not re-estimated.
...	Other undocumented arguments.

Details

Based on the classification of methods as described in Hyndman et al (2008).

The methodology is fully automatic. The only required argument for ets is the time series. The model is chosen automatically if not specified. This methodology performed extremely well on the M3-competition data. (See Hyndman, et al, 2002, below.)

Value

An object of class "ets".

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `ets` and associated functions.

Author(s)

Rob J Hyndman

References

Hyndman, R.J., Koehler, A.B., Snyder, R.D., and Grose, S. (2002) "A state space framework for automatic forecasting using exponential smoothing methods", *International J. Forecasting*, **18**(3), 439–454.

Hyndman, R.J., Akram, Md., and Archibald, B. (2008) "The admissible parameter space for exponential smoothing models". *Annals of Statistical Mathematics*, **60**(2), 407–426.

Hyndman, R.J., Koehler, A.B., Ord, J.K., and Snyder, R.D. (2008) *Forecasting with exponential smoothing: the state space approach*, Springer-Verlag. <http://www.exponentialsMOOTHING.net>.

See Also

[HoltWinters](#), [rwf](#), [Arima](#).

Examples

```
fit <- ets(USAccDeaths)
plot(forecast(fit))
```

findfrequency

Find dominant frequency of a time series

Description

`findfrequency` returns the period of the dominant frequency of a time series. For seasonal data, it will return the seasonal period. For cyclic data, it will return the average cycle length.

Usage

```
findfrequency(x)
```

Arguments

`x` a numeric vector or time series of class `ts`

Details

The dominant frequency is determined from a spectral analysis of the time series. First, a linear trend is removed, then the spectral density function is estimated from the best fitting autoregressive model (based on the AIC). If there is a large (possibly local) maximum in the spectral density function at frequency f , then the function will return the period $1/f$ (rounded to the nearest integer). If no such dominant frequency can be found, the function will return 1.

Value

an integer value

Author(s)

Rob J Hyndman

Examples

```
findfrequency(USAccDeaths) # Monthly data
findfrequency(taylor) # Half-hourly data
findfrequency(lynx) # Annual data
```

fitted.fracdiff *h-step in-sample forecasts for time series models.*

Description

Returns h-step forecasts for the data used in fitting the model.

Usage

```
## S3 method for class 'fracdiff'
fitted(object, h = 1, ...)

## S3 method for class 'Arima'
fitted(object, h = 1, ...)

## S3 method for class 'ar'
fitted(object, ...)

## S3 method for class 'bats'
fitted(object, h = 1, ...)

## S3 method for class 'ets'
fitted(object, h = 1, ...)
```

```
## S3 method for class 'nnetar'  
fitted(object, h = 1, ...)  
  
## S3 method for class 'tbats'  
fitted(object, h = 1, ...)
```

Arguments

object	An object of class "Arima", "bats", "tbats", "ets" or "nnetar".
h	The number of steps to forecast ahead.
...	Other arguments.

Value

A time series of the h-step forecasts.

Author(s)

Rob J Hyndman & Mitchell O'Hara-Wild

See Also

[forecast.Arima](#), [forecast.bats](#), [forecast.tbats](#), [forecast.ets](#), [forecast.nnetar](#), [residuals.Arima](#), [residuals.bats](#), [residuals.tbats](#), [residuals.ets](#), [residuals.nnetar](#).

Examples

```
fit <- ets(WWWusage)  
plot(WWWusage)  
lines(fitted(fit), col='red')  
lines(fitted(fit, h=2), col='green')  
lines(fitted(fit, h=3), col='blue')  
legend("topleft", legend=paste("h =", 1:3), col=2:4, lty=1)
```

forecast

Forecasting time series

Description

forecast is a generic function for forecasting from time series or time series models. The function invokes particular *methods* which depend on the class of the first argument.

Usage

```

forecast(object, ...)

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

## S3 method for class 'ts'
forecast(object, h = ifelse(frequency(object) > 1, 2 *
  frequency(object), 10), level = c(80, 95), fan = FALSE, robust = FALSE,
  lambda = NULL, find.frequency = FALSE,
  allow.multiplicative.trend = FALSE, model = NULL, ...)

```

Arguments

<code>object</code>	a time series or time series model for which forecasts are required
<code>...</code>	Additional arguments affecting the forecasts produced. If <code>model=NULL</code> , <code>forecast.ts</code> passes these to <code>ets</code> or <code>stlf</code> depending on the frequency of the time series. If <code>model</code> is not <code>NULL</code> , the arguments are passed to the relevant modelling function.
<code>h</code>	Number of periods for forecasting
<code>level</code>	Confidence level for prediction intervals.
<code>fan</code>	If <code>TRUE</code> , <code>level</code> is set to <code>seq(51, 99, by=3)</code> . This is suitable for fan plots.
<code>robust</code>	If <code>TRUE</code> , the function is robust to missing values and outliers in <code>object</code> . This argument is only valid when <code>object</code> is of class <code>ts</code> .
<code>lambda</code>	Box-Cox transformation parameter.
<code>find.frequency</code>	If <code>TRUE</code> , the function determines the appropriate period, if the data is of unknown period.
<code>allow.multiplicative.trend</code>	If <code>TRUE</code> , then ETS models with multiplicative trends are allowed. Otherwise, only additive or no trend ETS models are permitted.
<code>model</code>	An object describing a time series model; e.g., one of of class <code>ets</code> , <code>Arima</code> , <code>bats</code> , <code>tbats</code> , or <code>nnetar</code> .

Details

For example, the function `forecast.Arima` makes forecasts based on the results produced by `arima`.

If `model=NULL`, the function `forecast.ts` makes forecasts using `ets` models (if the data are non-seasonal or the seasonal period is 12 or less) or `stlf` (if the seasonal period is 13 or more).

If `model` is not `NULL`, `forecast.ts` will apply the `model` to the `object` time series, and then generate forecasts accordingly.

Value

An object of class "forecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract various useful features of the value returned by `forecast$model`.

An object of class "forecast" is a list usually containing at least the following elements:

<code>model</code>	A list containing information about the fitted model
<code>method</code>	The name of the forecasting method as a character string
<code>mean</code>	Point forecasts as a time series
<code>lower</code>	Lower limits for prediction intervals
<code>upper</code>	Upper limits for prediction intervals
<code>level</code>	The confidence values associated with the prediction intervals
<code>x</code>	The original time series (either object itself or the time series used to create the model stored as object).
<code>residuals</code>	Residuals from the fitted model. For models with additive errors, the residuals will be <code>x</code> minus the fitted values.
<code>fitted</code>	Fitted values (one-step forecasts)

Author(s)

Rob J Hyndman

See Also

Other functions which return objects of class "forecast" are [forecast.ets](#), [forecast.Arima](#), [forecast.HoltWinters](#), [forecast.StructTS](#), [meanf](#), [rwf](#), [splinef](#), [thetaf](#), [croston](#), [ses](#), [holt](#), [hw](#).

Examples

```
WWWusage %>% forecast %>% plot
fit <- ets(window(WWWusage, end=60))
fc <- forecast(WWWusage, model=fit)
```

`forecast.baggedETS` *Forecasting using the bagged ETS method*

Description

Returns forecasts and other information for bagged ETS models.

Usage

```
## S3 method for class 'baggedETS'
forecast(object, h = ifelse(frequency(object$x) > 1, 2 *
  frequency(object$x), 10), ...)
```

Arguments

object	An object of class "baggedETS" resulting from a call to baggedETS .
h	Number of periods for forecasting.
...	Additional arguments passed to ets

Details

Intervals are calculated as min and max values over the point forecasts from the ETS models in the ensemble. I.e., the intervals are not prediction intervals, but give an indication of how different the forecasts within the ensemble are.

Value

An object of class "forecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts and prediction intervals.

An object of class "forecast" is a list containing at least the following elements:

model	A list containing information about the fitted model
method	The name of the forecasting method as a character string
mean	Point forecasts as a time series
lower	Lower limits for prediction intervals
upper	Upper limits for prediction intervals
level	The confidence values associated with the prediction intervals
x	The original time series (either object itself or the time series used to create the model stored as object).
xreg	The external regressors used in fitting (if given).
residuals	Residuals from the fitted model. That is x minus fitted values.
fitted	Fitted values (one-step forecasts)
...	Other arguments

Author(s)

Christoph Bergmeir, Fotios Petropoulos

References

Bergmeir, C., R. J. Hyndman, and J. M. Benitez (2016). Bagging Exponential Smoothing Methods using STL Decomposition and Box-Cox Transformation. *International Journal of Forecasting* 32, 303-312.

See Also

[baggedETS](#).

Examples

```
fit <- baggedETS(WWWusage)
fcast <- forecast(fit)
plot(fcast)
```

forecast.bats

Forecasting using BATS and TBATS models

Description

Forecasts h steps ahead with a BATS model. Prediction intervals are also produced.

Usage

```
## S3 method for class 'bats'
forecast(object, h, level = c(80, 95), fan = FALSE,
         biasadj = NULL, ...)

## S3 method for class 'tbats'
forecast(object, h, level = c(80, 95), fan = FALSE,
         biasadj = NULL, ...)
```

Arguments

object	An object of class "bats". Usually the result of a call to bats .
h	Number of periods for forecasting. Default value is twice the largest seasonal period (for seasonal data) or ten (for non-seasonal data).
level	Confidence level for prediction intervals.
fan	If TRUE, level is set to seq(51, 99, by=3). This is suitable for fan plots.
biasadj	Use adjusted back-transformed mean for Box-Cox transformations. If TRUE, point forecasts and fitted values are mean forecast. Otherwise, these points can be considered the median of the forecast densities.
...	Other arguments, currently ignored.

Value

An object of class "forecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `forecast.bats`.

An object of class "forecast" is a list containing at least the following elements:

model	A copy of the bats object
method	The name of the forecasting method as a character string
mean	Point forecasts as a time series
lower	Lower limits for prediction intervals
upper	Upper limits for prediction intervals
level	The confidence values associated with the prediction intervals
x	The original time series (either object itself or the time series used to create the model stored as object).
residuals	Residuals from the fitted model.
fitted	Fitted values (one-step forecasts)

Author(s)

Slava Razbash and Rob J Hyndman

References

De Livera, A.M., Hyndman, R.J., & Snyder, R. D. (2011), Forecasting time series with complex seasonal patterns using exponential smoothing, *Journal of the American Statistical Association*, **106**(496), 1513-1527.

See Also

[bats](#), [tbats](#), [forecast.ets](#).

Examples

```
## Not run:  
fit <- bats(USAccDeaths)  
plot(forecast(fit))  
  
taylor.fit <- bats(taylor)  
plot(forecast(taylor.fit))  
  
## End(Not run)
```

Description

Returns forecasts and other information for univariate ETS models.

Usage

```
## S3 method for class 'ets'
forecast(object, h = ifelse(object$m > 1, 2 * object$m, 10),
  level = c(80, 95), fan = FALSE, simulate = FALSE, bootstrap = FALSE,
  npaths = 5000, PI = TRUE, lambda = object$lambda, biasadj = NULL, ...)
```

Arguments

object	An object of class "ets". Usually the result of a call to <code>ets</code> .
h	Number of periods for forecasting
level	Confidence level for prediction intervals.
fan	If TRUE, level is set to <code>seq(51,99,by=3)</code> . This is suitable for fan plots.
simulate	If TRUE, prediction intervals are produced by simulation rather than using analytic formulae. Errors are assumed to be normally distributed.
bootstrap	If TRUE, then prediction intervals are produced by simulation using resampled errors (rather than normally distributed errors).
npaths	Number of sample paths used in computing simulated prediction intervals.
PI	If TRUE, prediction intervals are produced, otherwise only point forecasts are calculated. If PI is FALSE, then level, fan, simulate, bootstrap and npaths are all ignored.
lambda	Box-Cox transformation parameter. Ignored if NULL. Otherwise, forecasts back-transformed via an inverse Box-Cox transformation.
biasadj	Use adjusted back-transformed mean for Box-Cox transformations. If TRUE, point forecasts and fitted values are mean forecast. Otherwise, these points can be considered the median of the forecast densities. By default, the value is taken from what was used when fitting the model.
...	Other arguments.

Value

An object of class "forecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `forecast.ets`.

An object of class "forecast" is a list containing at least the following elements:

model	A list containing information about the fitted model
method	The name of the forecasting method as a character string
mean	Point forecasts as a time series
lower	Lower limits for prediction intervals
upper	Upper limits for prediction intervals
level	The confidence values associated with the prediction intervals

x	The original time series (either object itself or the time series used to create the model stored as object).
residuals	Residuals from the fitted model. For models with additive errors, the residuals are $x - \text{fitted values}$. For models with multiplicative errors, the residuals are equal to $x / (\text{fitted values}) - 1$.
fitted	Fitted values (one-step forecasts)

Author(s)

Rob J Hyndman

See Also

[ets](#), [ses](#), [holt](#), [hw](#).

Examples

```
fit <- ets(USAccDeaths)
plot(forecast(fit,h=48))
```

forecast.fracdiff *Forecasting using ARIMA or ARFIMA models*

Description

Returns forecasts and other information for univariate ARIMA models.

Usage

```
## S3 method for class 'fracdiff'
forecast(object, h = 10, level = c(80, 95),
         fan = FALSE, lambda = object$lambda, biasadj = NULL, ...)

## S3 method for class 'Arima'
forecast(object, h = ifelse(object$arma[5] > 1, 2 *
         object$arma[5], 10), level = c(80, 95), fan = FALSE, xreg = NULL,
         lambda = object$lambda, bootstrap = FALSE, npaths = 5000,
         biasadj = NULL, ...)

## S3 method for class 'ar'
forecast(object, h = 10, level = c(80, 95), fan = FALSE,
         lambda = NULL, bootstrap = FALSE, npaths = 5000, biasadj = FALSE, ...)
```

Arguments

object	An object of class "Arima", "ar" or "fracdiff". Usually the result of a call to <code>arima</code> , <code>auto.arima</code> , <code>ar</code> , <code>arfima</code> or <code>fracdiff</code> .
h	Number of periods for forecasting. If <code>xreg</code> is used, <code>h</code> is ignored and the number of forecast periods is set to the number of rows of <code>xreg</code> .
level	Confidence level for prediction intervals.
fan	If TRUE, level is set to <code>seq(51, 99, by=3)</code> . This is suitable for fan plots.
lambda	Box-Cox transformation parameter. Ignored if NULL. Otherwise, forecasts back-transformed via an inverse Box-Cox transformation.
biasadj	Use adjusted back-transformed mean for Box-Cox transformations. If TRUE, point forecasts and fitted values are mean forecast. Otherwise, these points can be considered the median of the forecast densities. By default, the value is taken from what was used when fitting the model.
...	Other arguments.
xreg	Future values of an regression variables (for class <code>Arima</code> objects only).
bootstrap	If TRUE, then prediction intervals computed using simulation with resampled errors.
npaths	Number of sample paths used in computing simulated prediction intervals when <code>bootstrap=TRUE</code> .

Details

For `Arima` or `ar` objects, the function calls `predict.Arima` or `predict.ar` and constructs an object of class "forecast" from the results. For `fracdiff` objects, the calculations are all done within `forecast.fracdiff` using the equations given by Peiris and Perera (1988).

Value

An object of class "forecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `forecast.Arima`.

An object of class "forecast" is a list containing at least the following elements:

model	A list containing information about the fitted model
method	The name of the forecasting method as a character string
mean	Point forecasts as a time series
lower	Lower limits for prediction intervals
upper	Upper limits for prediction intervals
level	The confidence values associated with the prediction intervals
x	The original time series (either object itself or the time series used to create the model stored as object).
residuals	Residuals from the fitted model. That is <code>x</code> minus fitted values.
fitted	Fitted values (one-step forecasts)

Author(s)

Rob J Hyndman

References

Peiris, M. & Perera, B. (1988), On prediction with fractionally differenced ARIMA models, *Journal of Time Series Analysis*, **9**(3), 215-220.

See Also

[predict.Arima](#), [predict.ar](#), [auto.arima](#), [Arima](#), [arima](#), [ar](#), [arfima](#).

Examples

```
fit <- Arima(WWWusage,c(3,1,0))
plot(forecast(fit))

library(fracdiff)
x <- fracdiff.sim( 100, ma=-.4, d=.3)$series
fit <- arfima(x)
plot(forecast(fit,h=30))
```

forecast.HoltWinters *Forecasting using Holt-Winters objects*

Description

Returns forecasts and other information for univariate Holt-Winters time series models.

Usage

```
## S3 method for class 'HoltWinters'
forecast(object, h = ifelse(frequency(object$x) > 1, 2 *
  frequency(object$x), 10), level = c(80, 95), fan = FALSE, lambda = NULL,
  biasadj = NULL, ...)
```

Arguments

object	An object of class "HoltWinters". Usually the result of a call to HoltWinters .
h	Number of periods for forecasting
level	Confidence level for prediction intervals.
fan	If TRUE, level is set to seq(51,99,by=3). This is suitable for fan plots.
lambda	Box-Cox transformation parameter. Ignored if NULL. Otherwise, forecasts back-transformed via an inverse Box-Cox transformation.

biasadj	Use adjusted back-transformed mean for Box-Cox transformations. If TRUE, point forecasts and fitted values are mean forecast. Otherwise, these points can be considered the median of the forecast densities.
...	Other arguments.

Details

This function calls `predict.HoltWinters` and constructs an object of class "forecast" from the results.

It is included for completeness, but the `ets` is recommended for use instead of `HoltWinters`.

Value

An object of class "forecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `forecast.HoltWinters`.

An object of class "forecast" is a list containing at least the following elements:

model	A list containing information about the fitted model
method	The name of the forecasting method as a character string
mean	Point forecasts as a time series
lower	Lower limits for prediction intervals
upper	Upper limits for prediction intervals
level	The confidence values associated with the prediction intervals
x	The original time series (either object itself or the time series used to create the model stored as object).
residuals	Residuals from the fitted model.
fitted	Fitted values (one-step forecasts)

Author(s)

Rob J Hyndman

See Also

`predict.HoltWinters`, `HoltWinters`.

Examples

```
fit <- HoltWinters(WWWusage, gamma=FALSE)
plot(forecast(fit))
```

forecast.lm	<i>Forecast a linear model with possible time series components</i>
-------------	---

Description

forecast.lm is used to predict linear models, especially those involving trend and seasonality components.

Usage

```
## S3 method for class 'lm'
forecast(object, newdata, h = 10, level = c(80, 95),
         fan = FALSE, lambda = object$lambda, biasadj = NULL, ts = TRUE, ...)
```

Arguments

object	Object of class "lm", usually the result of a call to lm or tslm .
newdata	An optional data frame in which to look for variables with which to predict. If omitted, it is assumed that the only variables are trend and season, and h forecasts are produced.
h	Number of periods for forecasting. Ignored if newdata present.
level	Confidence level for prediction intervals.
fan	If TRUE, level is set to seq(51,99,by=3). This is suitable for fan plots.
lambda	Box-Cox transformation parameter. Ignored if NULL. Otherwise, forecasts back-transformed via an inverse Box-Cox transformation.
biasadj	Use adjusted back-transformed mean for Box-Cox transformations. If TRUE, point forecasts and fitted values are mean forecast. Otherwise, these points can be considered the median of the forecast densities.
ts	If TRUE, the forecasts will be treated as time series provided the original data is a time series; the newdata will be interpreted as related to the subsequent time periods. If FALSE, any time series attributes of the original data will be ignored.
...	Other arguments passed to predict.lm() .

Details

forecast.lm is largely a wrapper for [predict.lm\(\)](#) except that it allows variables "trend" and "season" which are created on the fly from the time series characteristics of the data. Also, the output is reformatted into a forecast object.

Value

An object of class "forecast".

The function summary is used to obtain and print a summary of the results, while the function plot produces a plot of the forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `forecast.lm`.

An object of class "forecast" is a list containing at least the following elements:

<code>model</code>	A list containing information about the fitted model
<code>method</code>	The name of the forecasting method as a character string
<code>mean</code>	Point forecasts as a time series
<code>lower</code>	Lower limits for prediction intervals
<code>upper</code>	Upper limits for prediction intervals
<code>level</code>	The confidence values associated with the prediction intervals
<code>x</code>	The historical data for the response variable.
<code>residuals</code>	Residuals from the fitted model. That is <code>x</code> minus fitted values.
<code>fitted</code>	Fitted values

Author(s)

Rob J Hyndman

See Also

[tslm](#), [lm](#).

Examples

```
y <- ts(rnorm(120,0,3) + 1:120 + 20*sin(2*pi*(1:120)/12), frequency=12)
fit <- tslm(y ~ trend + season)
plot(forecast(fit, h=20))
```

forecast.mlm

Forecast a multiple linear model with possible time series components

Description

`forecast.mlm` is used to predict multiple linear models, especially those involving trend and seasonality components.

Usage

```
## S3 method for class 'mlm'
forecast(object, newdata, h = 10, level = c(80, 95),
         fan = FALSE, lambda = object$lambda, biasadj = NULL, ts = TRUE, ...)
```

Arguments

object	Object of class "mlm", usually the result of a call to <code>lm</code> or <code>tslm</code> .
newdata	An optional data frame in which to look for variables with which to predict. If omitted, it is assumed that the only variables are trend and season, and h forecasts are produced.
h	Number of periods for forecasting. Ignored if newdata present.
level	Confidence level for prediction intervals.
fan	If TRUE, level is set to <code>seq(51,99,by=3)</code> . This is suitable for fan plots.
lambda	Box-Cox transformation parameter. Ignored if NULL. Otherwise, forecasts back-transformed via an inverse Box-Cox transformation.
biasadj	Use adjusted back-transformed mean for Box-Cox transformations. If TRUE, point forecasts and fitted values are mean forecast. Otherwise, these points can be considered the median of the forecast densities.
ts	If TRUE, the forecasts will be treated as time series provided the original data is a time series; the newdata will be interpreted as related to the subsequent time periods. If FALSE, any time series attributes of the original data will be ignored.
...	Other arguments passed to <code>forecast.lm()</code> .

Details

`forecast.mlm` is largely a wrapper for `forecast.lm()` except that it allows forecasts to be generated on multiple series. Also, the output is reformatted into a `mforecast` object.

Value

An object of class "mforecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `forecast.lm`.

An object of class "mforecast" is a list containing at least the following elements:

model	A list containing information about the fitted model
method	The name of the forecasting method as a character string
mean	Point forecasts as a multivariate time series
lower	Lower limits for prediction intervals of each series
upper	Upper limits for prediction intervals of each series
level	The confidence values associated with the prediction intervals
x	The historical data for the response variable.
residuals	Residuals from the fitted model. That is x minus fitted values.
fitted	Fitted values

Author(s)

Mitchell O'Hara-Wild

See Also

[tslm](#), [forecast.lm](#), [lm](#).

Examples

```
lungDeaths <- cbind(mdeaths, fdeaths)
fit <- tslm(lungDeaths ~ trend + season)
fcast <- forecast(fit, h=10)

carPower <- as.matrix(mtcars[,c("qsec", "hp")])
carmpg <- mtcars[, "mpg"]
fit <- lm(carPower ~ carmpg)
fcast <- forecast(fit, newdata=data.frame(carmpg=30))
```

forecast.mts

Forecasting time series

Description

mforecast is a class of objects for forecasting from multivariate time series or multivariate time series models. The function invokes particular *methods* which depend on the class of the first argument.

Usage

```
## S3 method for class 'mts'
forecast(object, h = ifelse(frequency(object) > 1, 2 *
  frequency(object), 10), level = c(80, 95), fan = FALSE, robust = FALSE,
  lambda = NULL, find.frequency = FALSE,
  allow.multiplicative.trend = FALSE, ...)
```

Arguments

object	a multivariate time series or multivariate time series model for which forecasts are required
h	Number of periods for forecasting
level	Confidence level for prediction intervals.
fan	If TRUE, level is set to seq(51, 99, by=3). This is suitable for fan plots.
robust	If TRUE, the function is robust to missing values and outliers in object. This argument is only valid when object is of class mts.

lambda	Box-Cox transformation parameter.
find.frequency	If TRUE, the function determines the appropriate period, if the data is of unknown period.
allow.multiplicative.trend	If TRUE, then ETS models with multiplicative trends are allowed. Otherwise, only additive or no trend ETS models are permitted.
...	Additional arguments affecting the forecasts produced.

Details

For example, the function `forecast.mlm` makes multivariate forecasts based on the results produced by `tslm`.

Value

An object of class "mforecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the multivariate forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract various useful features of the value returned by `forecast$model`.

An object of class "mforecast" is a list usually containing at least the following elements:

model	A list containing information about the fitted model
method	The name of the forecasting method as a character string
mean	Point forecasts as a time series
lower	Lower limits for prediction intervals
upper	Upper limits for prediction intervals
level	The confidence values associated with the prediction intervals
x	The original time series (either object itself or the time series used to create the model stored as object).
residuals	Residuals from the fitted model. For models with additive errors, the residuals will be x minus the fitted values.
fitted	Fitted values (one-step forecasts)

Author(s)

Rob J Hyndman & Mitchell O'Hara-Wild

See Also

Other functions which return objects of class "mforecast" are `forecast.mlm`, `forecast.varest`.

forecast.nnetar *Forecasting using neural network models*

Description

Returns forecasts and other information for univariate neural network models.

Usage

```
## S3 method for class 'nnetar'
forecast(object, h = ifelse(object$m > 1, 2 * object$m, 10),
  PI = FALSE, level = c(80, 95), fan = FALSE, xreg = NULL,
  lambda = object$lambda, bootstrap = FALSE, npaths = 1000,
  innov = NULL, ...)
```

Arguments

object	An object of class "nnetar" resulting from a call to arima .
h	Number of periods for forecasting. If xreg is used, h is ignored and the number of forecast periods is set to the number of rows of xreg.
PI	If TRUE, prediction intervals are produced, otherwise only point forecasts are calculated. If PI is FALSE, then level, fan, bootstrap and npaths are all ignored.
level	Confidence level for prediction intervals.
fan	If TRUE, level is set to seq(51, 99, by=3). This is suitable for fan plots.
xreg	Future values of external regressor variables.
lambda	Box-Cox transformation parameter. Ignored if NULL. Otherwise, forecasts back-transformed via an inverse Box-Cox transformation.
bootstrap	If TRUE, then prediction intervals computed using simulations with resampled residuals rather than normally distributed errors. Ignored if innov is not NULL.
npaths	Number of sample paths used in computing simulated prediction intervals.
innov	Values to use as innovations for prediction intervals. Must be a matrix with h rows and npaths columns (vectors are coerced into a matrix). If present, bootstrap is ignored.
...	Additional arguments passed to simulate.nnetar

Details

Prediction intervals are calculated through simulations and can be slow. Note that if the network is too complex and overfits the data, the residuals can be arbitrarily small; if used for prediction interval calculations, they could lead to misleadingly small values.

Value

An object of class "forecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `forecast.nnetar`.

An object of class "forecast" is a list containing at least the following elements:

<code>model</code>	A list containing information about the fitted model
<code>method</code>	The name of the forecasting method as a character string
<code>mean</code>	Point forecasts as a time series
<code>lower</code>	Lower limits for prediction intervals
<code>upper</code>	Upper limits for prediction intervals
<code>level</code>	The confidence values associated with the prediction intervals
<code>x</code>	The original time series (either object itself or the time series used to create the model stored as object).
<code>xreg</code>	The external regressors used in fitting (if given).
<code>residuals</code>	Residuals from the fitted model. That is <code>x</code> minus fitted values.
<code>fitted</code>	Fitted values (one-step forecasts)
<code>...</code>	Other arguments

Author(s)

Rob J Hyndman

See Also

[nnetar](#).

Examples

```
fit <- nnetar(lynx)
fcast <- forecast(fit)
plot(fcast)
```

Description

Forecasts of STL objects are obtained by applying a non-seasonal forecasting method to the seasonally adjusted data and re-seasonalizing using the last year of the seasonal component.

Usage

```
## S3 method for class 'stl'
forecast(object, method = c("ets", "arima", "naive", "rwdrift"),
  etsmodel = "ZZN", forecastfunction = NULL,
  h = frequency(object$time.series) * 2, level = c(80, 95), fan = FALSE,
  lambda = NULL, biasadj = NULL, xreg = NULL, newxreg = NULL,
  allow.multiplicative.trend = FALSE, ...)

stlm(y, s.window = 7, robust = FALSE, method = c("ets", "arima"),
  modelfunction = NULL, model = NULL, etsmodel = "ZZN", lambda = NULL,
  biasadj = FALSE, xreg = NULL, allow.multiplicative.trend = FALSE,
  x = y, ...)

## S3 method for class 'stlm'
forecast(object, h = 2 * object$m, level = c(80, 95),
  fan = FALSE, lambda = object$lambda, biasadj = NULL, newxreg = NULL,
  allow.multiplicative.trend = FALSE, ...)

stlf(y, h = frequency(x) * 2, s.window = 7, t.window = NULL,
  robust = FALSE, lambda = NULL, biasadj = FALSE, x = y, ...)
```

Arguments

object	An object of class <code>stl</code> or <code>stlm</code> . Usually the result of a call to <code>stl</code> or <code>stlm</code> .
method	Method to use for forecasting the seasonally adjusted series.
etsmodel	The ets model specification passed to <code>ets</code> . By default it allows any non-seasonal model. If <code>method!="ets"</code> , this argument is ignored.
forecastfunction	An alternative way of specifying the function for forecasting the seasonally adjusted series. If <code>forecastfunction</code> is not <code>NULL</code> , then <code>method</code> is ignored. Otherwise <code>method</code> is used to specify the forecasting method to be used.
h	Number of periods for forecasting.
level	Confidence level for prediction intervals.
fan	If <code>TRUE</code> , level is set to <code>seq(51,99,by=3)</code> . This is suitable for fan plots.
lambda	Box-Cox transformation parameter. Ignored if <code>NULL</code> . Otherwise, data transformed before decomposition and back-transformed after forecasts are computed.

biasadj	Use adjusted back-transformed mean for Box-Cox transformations. If TRUE, point forecasts and fitted values are mean forecast. Otherwise, these points can be considered the median of the forecast densities.
xreg	Historical regressors to be used in <code>auto.arima()</code> when <code>method=="arima"</code> .
newxreg	Future regressors to be used in <code>forecast.Arima()</code> .
allow.multiplicative.trend	If TRUE, then ETS models with multiplicative trends are allowed. Otherwise, only additive or no trend ETS models are permitted.
...	Other arguments passed to <code>forecast.stl</code> , <code>modelfunction</code> or <code>forecastfunction</code> .
y	A univariate numeric time series of class <code>ts</code>
s.window	Either the character string "periodic" or the span (in lags) of the loess window for seasonal extraction.
robust	If TRUE, robust fitting will used in the loess procedure within <code>stl</code> .
modelfunction	An alternative way of specifying the function for modelling the seasonally adjusted series. If <code>modelfunction</code> is not NULL, then <code>method</code> is ignored. Otherwise <code>method</code> is used to specify the time series model to be used.
model	Output from a previous call to <code>stlm</code> . If a <code>stlm</code> model is passed, this same model is fitted to <code>y</code> without re-estimating any parameters.
x	Deprecated. Included for backwards compatibility.
t.window	A number to control the smoothness of the trend. See <code>stl</code> for details.

Details

`stlm` takes a time series `y`, applies an STL decomposition, and models the seasonally adjusted data using the model passed as `modelfunction` or specified using `method`. It returns an object that includes the original STL decomposition and a time series model fitted to the seasonally adjusted data. This object can be passed to the `forecast.stlm` for forecasting.

`forecast.stlm` forecasts the seasonally adjusted data, then re-seasonalizes the results by adding back the last year of the estimated seasonal component.

`stlf` combines `stlm` and `forecast.stlm`. It takes a `ts` argument, applies an STL decomposition, models the seasonally adjusted data, reseasonalizes, and returns the forecasts. However, it allows more general forecasting methods to be specified via `forecastfunction`.

`forecast.stl` is similar to `stlf` except that it takes the STL decomposition as the first argument, instead of the time series.

Note that the prediction intervals ignore the uncertainty associated with the seasonal component. They are computed using the prediction intervals from the seasonally adjusted series, which are then reseasonalized using the last year of the seasonal component. The uncertainty in the seasonal component is ignored.

The time series model for the seasonally adjusted data can be specified in `stlm` using either `method` or `modelfunction`. The `method` argument provides a shorthand way of specifying `modelfunction` for a few special cases. More generally, `modelfunction` can be any function with first argument a `ts` object, that returns an object that can be passed to `forecast`. For example, `forecastfunction=ar` uses the `ar` function for modelling the seasonally adjusted series.

The forecasting method for the seasonally adjusted data can be specified in `stlf` and `forecast.stl` using either method or `forecastfunction`. The `method` argument provides a shorthand way of specifying `forecastfunction` for a few special cases. More generally, `forecastfunction` can be any function with first argument a `ts` object, and other `h` and `level`, which returns an object of class `forecast`. For example, `forecastfunction=thetaf` uses the `thetaf` function for forecasting the seasonally adjusted series.

Value

`stlm` returns an object of class `stlm`. The other functions return objects of class `forecast`.

There are many methods for working with `forecast` objects including `summary` to obtain and print a summary of the results, while `plot` produces a plot of the forecasts and prediction intervals. The generic accessor functions `fitted.values` and `residuals` extract useful features.

Author(s)

Rob J Hyndman

See Also

[stl](#), [forecast.ets](#), [forecast.Arima](#).

Examples

```
tsmod <- stlm(USAccDeaths, modelfunction=ar)
plot(forecast(tsmod, h=36))

decomp <- stl(USAccDeaths,s.window="periodic")
plot(forecast(decomp))

plot(stlf(AirPassengers, lambda=0))
```

forecast.StructTS *Forecasting using Structural Time Series models*

Description

Returns forecasts and other information for univariate structural time series models.

Usage

```
## S3 method for class 'StructTS'
forecast(object, h = ifelse(object$coef["epsilon"] > 1e-10,
  2 * object$xtsp[3], 10), level = c(80, 95), fan = FALSE, lambda = NULL,
  biasadj = NULL, ...)
```

Arguments

object	An object of class "StructTS". Usually the result of a call to StructTS .
h	Number of periods for forecasting
level	Confidence level for prediction intervals.
fan	If TRUE, level is set to seq(51,99,by=3). This is suitable for fan plots.
lambda	Box-Cox transformation parameter. Ignored if NULL. Otherwise, forecasts back-transformed via an inverse Box-Cox transformation.
biasadj	Use adjusted back-transformed mean for Box-Cox transformations. If TRUE, point forecasts and fitted values are mean forecast. Otherwise, these points can be considered the median of the forecast densities.
...	Other arguments.

Details

This function calls `predict.StructTS` and constructs an object of class "forecast" from the results.

Value

An object of class "forecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `forecast.StructTS`.

An object of class "forecast" is a list containing at least the following elements:

model	A list containing information about the fitted model
method	The name of the forecasting method as a character string
mean	Point forecasts as a time series
lower	Lower limits for prediction intervals
upper	Upper limits for prediction intervals
level	The confidence values associated with the prediction intervals
x	The original time series (either object itself or the time series used to create the model stored as object).
residuals	Residuals from the fitted model. That is x minus fitted values.
fitted	Fitted values (one-step forecasts)

Author(s)

Rob J Hyndman

See Also

[StructTS](#).

Examples

```
fit <- StructTS(WWWusage,"level")
plot(forecast(fit))
```

fourier

Fourier terms for modelling seasonality

Description

fourier returns a matrix containing terms from a Fourier series, up to order K, suitable for use in [Arima](#), [auto.arima](#), or [tslm](#).

Usage

```
fourier(x, K, h = NULL)
```

```
fourierf(x, K, h)
```

Arguments

x	Seasonal time series: a ts or a msts object
K	Maximum order(s) of Fourier terms
h	Number of periods ahead to forecast (optional)

Details

fourierf is deprecated, instead use the h argument in fourier.

The period of the Fourier terms is determined from the time series characteristics of x. When h is missing, the length of x also determines the number of rows for the matrix returned by fourier. Otherwise, the value of h determines the number of rows for the matrix returned by fourier, typically used for forecasting. The values within x are not used.

When x is a ts object, the value of K should be an integer and specifies the number of sine and cosine terms to return. Thus, the matrix returned has 2*K columns.

When x is a msts object, then K should be a vector of integers specifying the number of sine and cosine terms for each of the seasonal periods. Then the matrix returned will have 2*sum(K) columns.

Value

Numerical matrix.

Author(s)

Rob J Hyndman

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

```
library(ggplot2)

# Using Fourier series for a "ts" object
# K is chosen to minimize the AICc
deaths.model <- auto.arima(USAccDeaths, xreg=fourier(USAccDeaths,K=5), seasonal=FALSE)
deaths.fcast <- forecast(deaths.model, xreg=fourier(USAccDeaths, K=5, h=36))
autoplot(deaths.fcast) + xlab("Year")

# Using Fourier series for a "msts" object
taylor.lm <- tslm(taylor ~ fourier(taylor, K = c(3, 3)))
taylor.fcast <- forecast(taylor.lm,
  data.frame(fourier(taylor, K = c(3, 3), h = 270)))
autoplot(taylor.fcast)
```

gas

Australian monthly gas production

Description

Australian monthly gas production: 1956–1995.

Usage

gas

Format

Time series data

Source

Australian Bureau of Statistics.

Examples

```
plot(gas)
seasonplot(gas)
tsdisplay(gas)
```

`getResponse`*Get response variable from time series model.*

Description

`getResponse` is a generic function for extracting the historical data from a time series model (including Arima, ets, ar, fracdiff), a linear model of class `lm`, or a forecast object. The function invokes particular *methods* which depend on the class of the first argument.

Usage

```
getResponse(object, ...)  
  
## Default S3 method:  
getResponse(object, ...)  
  
## S3 method for class 'lm'  
getResponse(object, ...)  
  
## S3 method for class 'Arima'  
getResponse(object, ...)  
  
## S3 method for class 'fracdiff'  
getResponse(object, ...)  
  
## S3 method for class 'ar'  
getResponse(object, ...)  
  
## S3 method for class 'tbats'  
getResponse(object, ...)  
  
## S3 method for class 'bats'  
getResponse(object, ...)  
  
## S3 method for class 'mforecast'  
getResponse(object, ...)
```

Arguments

<code>object</code>	a time series model or forecast object.
<code>...</code>	Additional arguments that are ignored.

Value

A numerical vector or a time series object of class `ts`.

Author(s)

Rob J Hyndman

`gghistogram`*Histogram with optional normal and kernel density functions*

Description

Plots a histogram and density estimates using ggplot.

Usage

```
gghistogram(x, add.normal = FALSE, add.kde = FALSE, add.rug = TRUE, bins,  
            boundary = 0)
```

Arguments

<code>x</code>	a numerical vector.
<code>add.normal</code>	Add a normal density function for comparison
<code>add.kde</code>	Add a kernel density estimate for comparison
<code>add.rug</code>	Add a rug plot on the horizontal axis
<code>bins</code>	The number of bins to use for the histogram. Selected by default using the Friedman-Diaconis rule given by nclass.FD
<code>boundary</code>	A boundary between two bins.
<code>...</code>	Not used (for consistency with <code>lag.plot</code>)

Value

None.

Author(s)

Rob J Hyndman

See Also[hist](#), [geom_histogram](#)**Examples**

```
gghistogram(lynx, add.kde=TRUE)
```

gglagplot

*Time series lag ggplots***Description**

Plots a lag plot using ggplot.

Usage

```
gglagplot(x, lags = ifelse(frequency(x) > 9, 16, 9), set.lags = 1:lags,
  diag = TRUE, diag.col = "gray", do.lines = TRUE, colour = TRUE,
  continuous = frequency(x) > 12, labels = FALSE, seasonal = TRUE, ...)
```

```
gglagchull(x, lags = ifelse(frequency(x) > 1, min(12, frequency(x)), 4),
  set.lags = 1:lags, diag = TRUE, diag.col = "gray", ...)
```

Arguments

<code>x</code>	a time series object (type <code>ts</code>).
<code>lags</code>	number of lag plots desired, see arg <code>set.lags</code> .
<code>set.lags</code>	vector of positive integers specifying which lags to use.
<code>diag</code>	logical indicating if the <code>x=y</code> diagonal should be drawn.
<code>diag.col</code>	color to be used for the diagonal if(<code>diag</code>).
<code>do.lines</code>	if <code>TRUE</code> , lines will be drawn, otherwise points will be drawn.
<code>colour</code>	logical indicating if lines should be coloured.
<code>continuous</code>	Should the colour scheme for years be continuous or discrete?
<code>labels</code>	logical indicating if labels should be used.
<code>seasonal</code>	Should the line colour be based on seasonal characteristics (<code>TRUE</code>), or sequential (<code>FALSE</code>).
<code>...</code>	Not used (for consistency with <code>lag.plot</code>)

Details

“`gglagplot`” will plot time series against lagged versions of themselves. Helps visualising ‘auto-dependence’ even when auto-correlations vanish.

“`gglagchull`” will layer convex hulls of the lags, layered on a single plot. This helps visualise the change in ‘auto-dependence’ as lags increase.

Value

None.

Author(s)

Mitchell O’Hara-Wild

See Also[lag.plot](#)**Examples**

```
gglagplot(woolyrnq)
gglagplot(woolyrnq, seasonal=FALSE)

lungDeaths <- cbind(mdeaths, fdeaths)
gglagplot(lungDeaths, lags=2)
gglagchull(lungDeaths, lags=6)

gglagchull(woolyrnq)
```

`ggmonthplot`*Create a seasonal subseries ggplot*

Description

Plots a subseries plot using ggplot. Each season is plotted as a separate mini time series. The blue lines represent the mean of the observations within each season.

Usage

```
ggmonthplot(x, labels = NULL, times = time(x), phase = cycle(x), ...)
```

```
ggsubseriesplot(x, labels = NULL, times = time(x), phase = cycle(x), ...)
```

Arguments

<code>x</code>	a time series object (type <code>ts</code>).
<code>labels</code>	A vector of labels to use for each 'season'
<code>times</code>	A vector of times for each observation
<code>phase</code>	A vector of seasonal components
<code>...</code>	Not used (for consistency with <code>monthplot</code>)

Details

The `ggmonthplot` function is simply a wrapper for `ggsubseriesplot` as a convenience for users familiar with [monthplot](#).

Value

Returns an object of class `ggplot`.

Author(s)

Mitchell O'Hara-Wild

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

```
ggsubseriesplot(AirPassengers)
ggsubseriesplot(woolryrnq)
```

 ggseasonplot

Seasonal plot

Description

Plots a seasonal plot as described in Hyndman and Athanasopoulos (2014, chapter 2). This is like a time plot except that the data are plotted against the seasons in separate years.

Usage

```
ggseasonplot(x, season.labels = NULL, year.labels = FALSE,
             year.labels.left = FALSE, type = NULL, col = NULL, continuous = FALSE,
             polar = FALSE, labelgap = 0.04, ...)
```

```
seasonplot(x, s, season.labels = NULL, year.labels = FALSE,
           year.labels.left = FALSE, type = "o", main, xlab = NULL, ylab = "",
           col = 1, labelgap = 0.1, ...)
```

Arguments

<code>x</code>	a numeric vector or time series of class <code>ts</code> .
<code>season.labels</code>	Labels for each season in the "year"
<code>year.labels</code>	Logical flag indicating whether labels for each year of data should be plotted on the right.
<code>year.labels.left</code>	Logical flag indicating whether labels for each year of data should be plotted on the left.
<code>type</code>	plot type (as for plot). Not yet supported for <code>ggseasonplot</code> .
<code>col</code>	Colour
<code>continuous</code>	Should the colour scheme for years be continuous or discrete?
<code>polar</code>	Plot the graph on seasonal coordinates

labelgap	Distance between year labels and plotted lines
...	additional arguments to <code>plot</code> .
s	seasonal frequency of x
main	Main title.
xlab	X-axis label.
ylab	Y-axis label.

Value

None.

Author(s)

Rob J Hyndman & Mitchell O'Hara-Wild

References

Hyndman and Athanasopoulos (2014) *Forecasting: principles and practice*, OTexts: Melbourne, Australia. <http://www.otexts.org/fpp/>

See Also

[monthplot](#)

Examples

```
ggseasonplot(AirPassengers, col=rainbow(12), year.labels=TRUE)
ggseasonplot(AirPassengers, year.labels=TRUE, continuous=TRUE)

seasonplot(AirPassengers, col=rainbow(12), year.labels=TRUE)
```

ggtsdisplay

Time series display

Description

Plots a time series along with its acf and either its pacf, lagged scatterplot or spectrum.

Usage

```
ggtsdisplay(x, plot.type = c("partial", "histogram", "scatter", "spectrum"),
  points = TRUE, smooth = FALSE, lag.max, na.action = na.contiguous,
  theme = NULL, ...)

tsdisplay(x, plot.type = c("partial", "histogram", "scatter", "spectrum"),
  points = TRUE, ci.type = c("white", "ma"), lag.max,
  na.action = na.contiguous, main = NULL, xlab = "", ylab = "",
  pch = 1, cex = 0.5, ...)
```

Arguments

<code>x</code>	a numeric vector or time series of class <code>ts</code> .
<code>plot.type</code>	type of plot to include in lower right corner.
<code>points</code>	logical flag indicating whether to show the individual points or not in the time plot.
<code>smooth</code>	logical flag indicating whether to show a smooth loess curve superimposed on the time plot.
<code>lag.max</code>	the maximum lag to plot for the <code>acf</code> and <code>pacf</code> . A suitable value is selected by default if the argument is missing.
<code>na.action</code>	function to handle missing values in <code>acf</code> , <code>pacf</code> and spectrum calculations. The default is <code>na.contiguous</code> . Useful alternatives are <code>na.pass</code> and <code>na.interp</code> .
<code>theme</code>	Adds a ggplot element to each plot, typically a theme.
<code>...</code>	additional arguments to <code>acf</code> .
<code>ci.type</code>	type of confidence limits for ACF that is passed to <code>acf</code> . Should the confidence limits assume a white noise input or for lag k an $MA(k - 1)$ input?
<code>main</code>	Main title.
<code>xlab</code>	X-axis label.
<code>ylab</code>	Y-axis label.
<code>pch</code>	Plotting character.
<code>cex</code>	Character size.

Details

`ggtsdisplay` will produce the equivalent plot using `ggplot` graphics.

Value

None.

Author(s)

Rob J Hyndman

References

Hyndman and Athanasopoulos (2014) *Forecasting: principles and practice*, OTexts: Melbourne, Australia. <http://www.otexts.org/fpp/>

See Also

[plot.ts](#), [Acf](#), [spec.ar](#)

Examples

```
library(ggplot2)
ggtsdisplay(USAccDeaths, plot.type="scatter", theme=theme_bw())

tsdisplay(diff(WWWusage))
ggtsdisplay(USAccDeaths, plot.type="scatter")
```

gold

Daily morning gold prices

Description

Daily morning gold prices in US dollars. 1 January 1985 – 31 March 1989.

Usage

```
gold
```

Format

Time series data

Source

Time Series Data Library. <http://data.is/TSDLdemo>

Examples

```
tsdisplay(gold)
```

is.acf

Is an object a particular model type?

Description

Returns true if the model object is of a particular type

Usage

```
is.acf(x)
is.Arima(x)
is.baggedETS(x)
is.bats(x)
is.ets(x)
is.nnetar(x)
is.nnetarmodels(x)
is.stlm(x)
```

Arguments

x object to be tested

is.constant	<i>Is an object constant?</i>
-------------	-------------------------------

Description

Returns true if the object's numerical values do not vary.

Usage

```
is.constant(x)
```

Arguments

x object to be tested

is.forecast	<i>Is an object a particular forecast type?</i>
-------------	---

Description

Returns true if the forecast object is of a particular type

Usage

```
is.forecast(x)
```

```
is.mforecast(x)
```

```
is.splineforecast(x)
```

Arguments

x object to be tested

ma	<i>Moving-average smoothing</i>
----	---------------------------------

Description

ma computes a simple moving average smoother of a given time series.

Usage

```
ma(x, order, centre = TRUE)
```

Arguments

x Univariate time series

order Order of moving average smoother

centre If TRUE, then the moving average is centred for even orders.

Details

The moving average smoother averages the nearest order periods of each observation. As neighbouring observations of a time series are likely to be similar in value, averaging eliminates some of the randomness in the data, leaving a smooth trend-cycle component.

$$\hat{T}_t = \frac{1}{m} \sum_{j=-k}^k y_{t+j}$$

where $k = \frac{m-1}{2}$

When an even order is specified, the observations averaged will include one more observation from the future than the past (k is rounded up). If `centre` is TRUE, the value from two moving averages (where k is rounded up and down respectively) are averaged, centering the moving average.

Value

Numerical time series object containing the simple moving average smoothed values.

Author(s)

Rob J Hyndman

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

```
plot(wineind)
sm <- ma(wineind,order=12)
lines(sm,col="red")
```

 meanf

Mean Forecast

Description

Returns forecasts and prediction intervals for an iid model applied to y .

Usage

```
meanf(y, h = 10, level = c(80, 95), fan = FALSE, lambda = NULL,
      biasadj = FALSE, x = y)
```

Arguments

y	a numeric vector or time series of class <code>ts</code>
h	Number of periods for forecasting
<code>level</code>	Confidence levels for prediction intervals.
<code>fan</code>	If TRUE, level is set to <code>seq(51,99,by=3)</code> . This is suitable for fan plots.
<code>lambda</code>	Box-Cox transformation parameter. Ignored if NULL. Otherwise, forecasts back-transformed via an inverse Box-Cox transformation.
<code>biasadj</code>	Use adjusted back-transformed mean for Box-Cox transformations. If TRUE, point forecasts and fitted values are mean forecast. Otherwise, these points can be considered the median of the forecast densities.
x	Deprecated. Included for backwards compatibility.

Details

The iid model is

$$Y_t = \mu + Z_t$$

where Z_t is a normal iid error. Forecasts are given by

$$Y_n(h) = \mu$$

where μ is estimated by the sample mean.

Value

An object of class "forecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `meanf`.

An object of class "forecast" is a list containing at least the following elements:

<code>model</code>	A list containing information about the fitted model
<code>method</code>	The name of the forecasting method as a character string
<code>mean</code>	Point forecasts as a time series
<code>lower</code>	Lower limits for prediction intervals
<code>upper</code>	Upper limits for prediction intervals
<code>level</code>	The confidence values associated with the prediction intervals
<code>x</code>	The original time series (either object itself or the time series used to create the model stored as object).
<code>residuals</code>	Residuals from the fitted model. That is <code>x</code> minus fitted values.
<code>fitted</code>	Fitted values (one-step forecasts)

Author(s)

Rob J Hyndman

See Also

[rwf](#)

Examples

```
nile.fcast <- meanf(Nile, h=10)
plot(nile.fcast)
```

monthdays

Number of days in each season

Description

Returns number of days in each month or quarter of the observed time period.

Usage

```
monthdays(x)
```

Arguments

x time series

Details

Useful for month length adjustments

Value

Time series

Author(s)

Rob J Hyndman

See Also

[bizdays](#)

Examples

```
par(mfrow=c(2,1))
plot(ldeaths,xlab="Year",ylab="pounds",
     main="Monthly deaths from lung disease (UK)")
ldeaths.adj <- ldeaths/monthdays(ldeaths)*365.25/12
plot(ldeaths.adj,xlab="Year",ylab="pounds",
     main="Adjusted monthly deaths from lung disease (UK)")
```

msts

Multi-Seasonal Time Series

Description

msts is an S3 class for multi seasonal time series objects, intended to be used for models that support multiple seasonal periods. The msts class inherits from the ts class and has an additional "msts" attribute which contains the vector of seasonal periods. All methods that work on a ts class, should also work on a msts class.

Usage

```
msts(data, seasonal.periods, ts.frequency = floor(max(seasonal.periods)), ...)
```

Arguments

data	A numeric vector, ts object, matrix or data frame. It is intended that the time series data is univariate, otherwise treated the same as ts().
seasonal.periods	A vector of the seasonal periods of the msts.
ts.frequency	The seasonal period that should be used as frequency of the underlying ts object. The default value is <code>max(seasonal.periods)</code> .
...	Arguments to be passed to the underlying call to <code>ts()</code> . For example <code>start=c(1987,5)</code> .

Value

An object of class `c("msts", "ts")`. If there is only one seasonal period (i.e., `length(seasonal.periods)==1`), then the object is of class `"ts"`.

Author(s)

Slava Razbash and Rob J Hyndman

Examples

```
x <- msts(taylor, seasonal.periods=c(48,336), start=2000+22/52)
y <- msts(USAccDeaths, seasonal.periods=12, start=1949)
```

na.interp

Interpolate missing values in a time series

Description

Uses linear interpolation for non-seasonal series and a periodic stl decomposition with seasonal series to replace missing values.

Usage

```
na.interp(x, lambda = NULL)
```

Arguments

x	time series
lambda	a numeric value suggesting Box-cox transformation

Details

A more general and flexible approach is available using `na.approx` in the `zoo` package.

Value

Time series

Author(s)

Rob J Hyndman

See Also

[na.interp](#), [tsoutliers](#)

Examples

```
data(gold)
plot(na.interp(gold))
```

ndiffs

Number of differences required for a stationary series

Description

Functions to estimate the number of differences required to make a given time series stationary. `ndiffs` estimates the number of first differences and `nsdiffs` estimates the number of seasonal differences.

Usage

```
ndiffs(x, alpha = 0.05, test = c("kpss", "adf", "pp"), max.d = 2)
```

```
nsdiffs(x, m = frequency(x), test = c("ocsb", "ch"), max.D = 1)
```

Arguments

<code>x</code>	A univariate time series
<code>alpha</code>	Level of the test
<code>test</code>	Type of unit root test to use
<code>max.d</code>	Maximum number of non-seasonal differences allowed
<code>m</code>	Length of seasonal period
<code>max.D</code>	Maximum number of seasonal differences allowed

Details

ndiffs uses a unit root test to determine the number of differences required for time series x to be made stationary. If `test="kpss"`, the KPSS test is used with the null hypothesis that x has a stationary root against a unit-root alternative. Then the test returns the least number of differences required to pass the test at the level α . If `test="adf"`, the Augmented Dickey-Fuller test is used and if `test="pp"` the Phillips-Perron test is used. In both of these cases, the null hypothesis is that x has a unit root against a stationary root alternative. Then the test returns the least number of differences required to fail the test at the level α .

nsdiffs uses seasonal unit root tests to determine the number of seasonal differences required for time series x to be made stationary (possibly with some lag-one differencing as well). If `test="ch"`, the Canova-Hansen (1995) test is used (with null hypothesis of deterministic seasonality) and if `test="ocsb"`, the Osborn-Chui-Smith-Birchenhall (1988) test is used (with null hypothesis that a seasonal unit root exists).

Value

An integer.

Author(s)

Rob J Hyndman and Slava Razbash

References

Canova F and Hansen BE (1995) "Are Seasonal Patterns Constant over Time? A Test for Seasonal Stability", *Journal of Business and Economic Statistics* **13**(3):237-252.

Dickey DA and Fuller WA (1979), "Distribution of the Estimators for Autoregressive Time Series with a Unit Root", *Journal of the American Statistical Association* **74**:427-431.

Kwiatkowski D, Phillips PCB, Schmidt P and Shin Y (1992) "Testing the Null Hypothesis of Stationarity against the Alternative of a Unit Root", *Journal of Econometrics* **54**:159-178.

Osborn DR, Chui APL, Smith J, and Birchenhall CR (1988) "Seasonality and the order of integration for consumption", *Oxford Bulletin of Economics and Statistics* **50**(4):361-377.

Osborn, D.R. (1990) "A survey of seasonality in UK macroeconomic variables", *International Journal of Forecasting*, **6**:327-336.

Said E and Dickey DA (1984), "Testing for Unit Roots in Autoregressive Moving Average Models of Unknown Order", *Biometrika* **71**:599-607.

See Also

[auto.arima](#)

Examples

```
ndiffs(WWWusage)
ndiffs(diff(log(AirPassengers), 12))

nsdiffs(log(AirPassengers))
```

nnetar

Neural Network Time Series Forecasts

Description

Feed-forward neural networks with a single hidden layer and lagged inputs for forecasting univariate time series.

Usage

```
nnetar(y, p, P = 1, size, repeats = 20, xreg = NULL, lambda = NULL,
       model = NULL, subset = NULL, scale.inputs = TRUE, x = y, ...)
```

Arguments

y	A numeric vector or time series of class <code>ts</code> .
p	Embedding dimension for non-seasonal time series. Number of non-seasonal lags used as inputs. For non-seasonal time series, the default is the optimal number of lags (according to the AIC) for a linear AR(p) model. For seasonal time series, the same method is used but applied to seasonally adjusted data (from an <code>stl</code> decomposition).
P	Number of seasonal lags used as inputs.
size	Number of nodes in the hidden layer. Default is half of the number of input nodes (including external regressors, if given) plus 1.
repeats	Number of networks to fit with different random starting weights. These are then averaged when producing forecasts.
xreg	Optionally, a vector or matrix of external regressors, which must have the same number of rows as <code>y</code> . Must be numeric.
lambda	Box-Cox transformation parameter.
model	Output from a previous call to <code>nnetar</code> . If <code>model</code> is passed, this same model is fitted to <code>y</code> without re-estimating any parameters.
subset	Optional vector specifying a subset of observations to be used in the fit. Can be an integer index vector or a logical vector the same length as <code>y</code> . All observations are used by default.
scale.inputs	If <code>TRUE</code> , inputs are scaled by subtracting the column means and dividing by their respective standard deviations. If <code>lambda</code> is not <code>NULL</code> , scaling is applied after Box-Cox transformation.
x	Deprecated. Included for backwards compatibility.
...	Other arguments passed to <code>nnet</code> for <code>nnetar</code> .

Details

A feed-forward neural network is fitted with lagged values of y as inputs and a single hidden layer with size nodes. The inputs are for lags 1 to p , and lags m to mP where $m = \text{frequency}(y)$. If there are missing values in y or $xreg$, the corresponding rows (and any others which depend on them as lags) are omitted from the fit. A total of `repeats` networks are fitted, each with random starting weights. These are then averaged when computing forecasts. The network is trained for one-step forecasting. Multi-step forecasts are computed recursively.

For non-seasonal data, the fitted model is denoted as an NNAR(p,k) model, where k is the number of hidden nodes. This is analogous to an AR(p) model but with nonlinear functions. For seasonal data, the fitted model is called an NNAR(p,P,k)[m] model, which is analogous to an ARIMA($p,0,0$)($P,0,0$)[m] model but with nonlinear functions.

Value

Returns an object of class "nnetar".

The function `summary` is used to obtain and print a summary of the results.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `nnetar`.

<code>model</code>	A list containing information about the fitted model
<code>method</code>	The name of the forecasting method as a character string
<code>x</code>	The original time series.
<code>xreg</code>	The external regressors used in fitting (if given).
<code>residuals</code>	Residuals from the fitted model. That is x minus fitted values.
<code>fitted</code>	Fitted values (one-step forecasts)
<code>...</code>	Other arguments

Author(s)

Rob J Hyndman

Examples

```
fit <- nnetar(lynx)
fcast <- forecast(fit)
plot(fcast)

## Arguments can be passed to nnet()
fit <- nnetar(lynx, decay=0.5, maxit=150)
plot(forecast(fit))
lines(lynx)

## Fit model to first 100 years of lynx data
fit <- nnetar(window(lynx,end=1920), decay=0.5, maxit=150)
plot(forecast(fit,h=14))
lines(lynx)
```

```
## Apply fitted model to later data, including all optional arguments
fit2 <- nnetar(window(lynx,start=1921), model=fit)
```

plot.Arima

Plot characteristic roots from ARIMA model

Description

Produces a plot of the inverse AR and MA roots of an ARIMA model. Inverse roots outside the unit circle are shown in red.

Usage

```
## S3 method for class 'Arima'
plot(x, type = c("both", "ar", "ma"), main, xlab = "Real",
     ylab = "Imaginary", ...)
```

```
## S3 method for class 'ar'
plot(x, main, xlab = "Real", ylab = "Imaginary", ...)
```

```
## S3 method for class 'Arima'
autoplot(object, type = c("both", "ar", "ma"), ...)
```

```
## S3 method for class 'ar'
autoplot(object, ...)
```

Arguments

x	Object of class “Arima” or “ar”.
type	Determines if both AR and MA roots are plotted, of if just one set is plotted.
main	Main title. Default is "Inverse AR roots" or "Inverse MA roots".
xlab	X-axis label.
ylab	Y-axis label.
...	Other plotting parameters passed to par .
object	Object of class “Arima” or “ar”. Used for ggplot graphics (S3 method consistency).

Details

autoplot will produce an equivalent plot as a ggplot object.

Value

None. Function produces a plot

Author(s)

Rob J Hyndman & Mitchell O'Hara-Wild

See Also

[Arima](#), [ar](#)

Examples

```
library(ggplot2)

fit <- Arima(WWWusage, order=c(3,1,0))
plot(fit)
autoplot(fit)

fit <- Arima(woolynrq, order=c(2,0,0), seasonal=c(2,1,1))
plot(fit)
autoplot(fit)

plot(ar.ols(gold[1:61]))
autoplot(ar.ols(gold[1:61]))
```

plot.bats

Plot components from BATS model

Description

Produces a plot of the level, slope and seasonal components from a BATS or TBATS model.

Usage

```
## S3 method for class 'bats'
plot(x, main = "Decomposition by BATS model", ...)

## S3 method for class 'tbats'
autoplot(object, range.bars = FALSE, ...)

## S3 method for class 'bats'
autoplot(object, range.bars = FALSE, ...)

## S3 method for class 'tbats'
plot(x, main = "Decomposition by TBATS model", ...)
```

Arguments

x	Object of class “bats/tbats”.
main	Main title for plot.
...	Other plotting parameters passed to par .
object	Object of class “bats/tbats”.
range.bars	Logical indicating if each plot should have a bar at its right side representing relative size. If NULL, automatic selection takes place.

Value

None. Function produces a plot

Author(s)

Rob J Hyndman

See Also

[bats,tbats](#)

Examples

```
## Not run:  
fit <- tbats(USAccDeaths)  
plot(fit)  
autoplot(fit, range.bars = TRUE)  
## End(Not run)
```

plot.ets

Plot components from ETS model

Description

Produces a plot of the level, slope and seasonal components from an ETS model.

Usage

```
## S3 method for class 'ets'  
plot(x, ...)  
  
## S3 method for class 'ets'  
autoplot(object, range.bars = NULL, ...)
```

Arguments

x	Object of class “ets”.
...	Other plotting parameters to affect the plot.
object	Object of class “ets”. Used for ggplot graphics (S3 method consistency).
range.bars	Logical indicating if each plot should have a bar at its right side representing relative size. If NULL, automatic selection takes place.

Details

autoplot will produce an equivalent plot as a ggplot object.

Value

None. Function produces a plot

Author(s)

Rob J Hyndman & Mitchell O’Hara-Wild

See Also

[ets](#)

Examples

```
fit <- ets(USAccDeaths)
plot(fit)
plot(fit, plot.type="single", ylab="", col=1:3)

library(ggplot2)
autoplot(fit)
```

plot.forecast

Forecast plot

Description

Plots historical data with forecasts and prediction intervals.

Usage

```
## S3 method for class 'forecast'
plot(x, include, PI = TRUE, showgap = TRUE,
     shaded = TRUE, shadebars = (length(x$mean) < 5), shadecols = NULL,
     col = 1, fcol = 4, pi.col = 1, pi.lty = 2, ylim = NULL,
     main = NULL, xlab = "", ylab = "", type = "l", flty = 1, flwd = 2,
     ...)

## S3 method for class 'forecast'
autoplot(object, include, PI = TRUE,
         shadecols = c("#596DD5", "#D5DBFF"), fcol = "#0000AA", flwd = 0.5, ...)

## S3 method for class 'splineforecast'
autoplot(object, PI = TRUE, ...)

## S3 method for class 'forecast'
autolayer(object, series = NULL, PI = TRUE,
          showgap = TRUE, ...)

## S3 method for class 'splineforecast'
plot(x, fitcol = 2, type = "o", pch = 19, ...)
```

Arguments

x	Forecast object produced by forecast .
include	number of values from time series to include in plot. Default is all values.
PI	Logical flag indicating whether to plot prediction intervals.
showgap	If showgap=FALSE, the gap between the historical observations and the forecasts is removed.
shaded	Logical flag indicating whether prediction intervals should be shaded (TRUE) or lines (FALSE)
shadebars	Logical flag indicating if prediction intervals should be plotted as shaded bars (if TRUE) or a shaded polygon (if FALSE). Ignored if shaded=FALSE. Bars are plotted by default if there are fewer than five forecast horizons.
shadecols	Colors for shaded prediction intervals. To get default colors used prior to v3.26, set shadecols="oldstyle".
col	Colour for the data line.
fcol	Colour for the forecast line.
pi.col	If shaded=FALSE and PI=TRUE, the prediction intervals are plotted in this colour.
pi.lty	If shaded=FALSE and PI=TRUE, the prediction intervals are plotted using this line type.
ylim	Limits on y-axis.
main	Main title.
xlab	X-axis label.

ylab	Y-axis label.
type	1-character string giving the type of plot desired. As for plot.default .
flty	Line type for the forecast line.
flwd	Line width for the forecast line.
...	Other plotting parameters to affect the plot.
object	Forecast object produced by forecast . Used for ggplot graphics (S3 method consistency).
series	Matches an unidentified forecast layer with a coloured object on the plot.
fitcol	Line colour for fitted values.
pch	Plotting character (if type=="p" or type=="o").

Details

autoplot will produce a ggplot object.
plot.splineforecast autoplot.splineforecast

Value

None.

Author(s)

Rob J Hyndman & Mitchell O'Hara-Wild

References

Hyndman and Athanasopoulos (2014) *Forecasting: principles and practice*, OTexts: Melbourne, Australia. <http://www.otexts.org/fpp/>

See Also

[plot.ts](#)

Examples

```
library(ggplot2)

wine.fit <- hw(wineind,h=48)
plot(wine.fit)
autoplot(wine.fit)

fit <- tslm(wineind ~ fourier(wineind,4))
fcast <- forecast(fit, newdata=data.frame(fourier(wineind,4,20)))
autoplot(fcast)

fcast <- splinef(airmiles,h=5)
plot(fcast)
autoplot(fcast)
```

residuals.tslm *Residuals for various time series models*

Description

Returns time series of residuals from a fitted model.

Usage

```
## S3 method for class 'tslm'
residuals(object, ...)

## S3 method for class 'forecast'
residuals(object, type = c("innovation", "response"), ...)

## S3 method for class 'ar'
residuals(object, type = c("innovation", "response"), ...)

## S3 method for class 'Arima'
residuals(object, type = c("innovation", "response",
  "regression"), h = 1, ...)

## S3 method for class 'bats'
residuals(object, type = c("innovation", "response"), h = 1,
  ...)

## S3 method for class 'tbats'
residuals(object, type = c("innovation", "response"), h = 1,
  ...)

## S3 method for class 'ets'
residuals(object, type = c("innovation", "response"), h = 1,
  ...)

## S3 method for class 'fracdiff'
residuals(object, type = c("innovation", "response"), ...)

## S3 method for class 'nnetar'
residuals(object, type = c("innovation", "response"),
  h = 1, ...)

## S3 method for class 'stlm'
residuals(object, type = c("innovation", "response"), ...)
```

Arguments

object An object containing a time series model of class `ar`, `Arima`, `bats`, `ets`, `fracdiff`, `nnetar` or `stlm`. If `object` is of class `forecast`, then the function will return

	object\$residuals if it exists, otherwise it returns the differences between the observations and their fitted values.
...	Other arguments not used.
type	Type of residual.
h	If type='response', then the fitted values are computed for h-step forecasts.

Details

Innovation residuals correspond to the white noise process that drives the evolution of the time series model. Response residuals are the difference between the observations and the fitted values (equivalent to h-step forecasts). For functions with no h argument, h=1. For homoscedastic models, the innovation residuals and the response residuals for h=1 are identical. Regression residuals are available for regression models with ARIMA errors, and are equal to the original data minus the effect of the regression variables. If there are no regression variables, the errors will be identical to the original series (possibly adjusted to have zero mean). `arima.errors` is a deprecated function which is identical to `residuals.Arima(object, type="regression")`.

Value

A ts object

Author(s)

Rob J Hyndman

See Also

[fitted.Arima](#), [checkresiduals](#).

Examples

```
fit <- Arima(lynx,order=c(4,0,0), lambda=0.5)

plot(residuals(fit))
plot(residuals(fit, type='response'))
```

Description

`rwf()` returns forecasts and prediction intervals for a random walk with drift model applied to `y`. This is equivalent to an ARIMA(0,1,0) model with an optional drift coefficient. `naive()` is simply a wrapper to `rwf()` for simplicity. `snaive()` returns forecasts and prediction intervals from an ARIMA(0,0,0)(0,1,0)_m model where `m` is the seasonal period.

Usage

```
rwf(y, h = 10, drift = FALSE, level = c(80, 95), fan = FALSE,
    lambda = NULL, biasadj = FALSE, x = y)
```

```
naive(y, h = 10, level = c(80, 95), fan = FALSE, lambda = NULL,
    biasadj = FALSE, x = y)
```

```
snaive(y, h = 2 * frequency(x), level = c(80, 95), fan = FALSE,
    lambda = NULL, biasadj = FALSE, x = y)
```

Arguments

y	a numeric vector or time series of class ts
h	Number of periods for forecasting
drift	Logical flag. If TRUE, fits a random walk with drift model.
level	Confidence levels for prediction intervals.
fan	If TRUE, level is set to seq(51,99,by=3). This is suitable for fan plots.
lambda	Box-Cox transformation parameter. Ignored if NULL. Otherwise, forecasts back-transformed via an inverse Box-Cox transformation.
biasadj	Use adjusted back-transformed mean for Box-Cox transformations. If TRUE, point forecasts and fitted values are mean forecast. Otherwise, these points can be considered the median of the forecast densities.
x	Deprecated. Included for backwards compatibility.

Details

The random walk with drift model is

$$Y_t = c + Y_{t-1} + Z_t$$

where Z_t is a normal iid error. Forecasts are given by

$$Y_n(h) = ch + Y_n$$

. If there is no drift (as in naive), the drift parameter $c=0$. Forecast standard errors allow for uncertainty in estimating the drift parameter.

The seasonal naive model is

$$Y_t = Y_{t-m} + Z_t$$

where Z_t is a normal iid error.

Value

An object of class "forecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by naive or snaive.

An object of class "forecast" is a list containing at least the following elements:

model	A list containing information about the fitted model
method	The name of the forecasting method as a character string
mean	Point forecasts as a time series
lower	Lower limits for prediction intervals
upper	Upper limits for prediction intervals
level	The confidence values associated with the prediction intervals
x	The original time series (either object itself or the time series used to create the model stored as object).
residuals	Residuals from the fitted model. That is x minus fitted values.
fitted	Fitted values (one-step forecasts)

Author(s)

Rob J Hyndman

See Also

[Arima](#)

Examples

```
gold.fcast <- rwf(gold[1:60], h=50)
plot(gold.fcast)
```

```
plot(naive(gold,h=50),include=200)
```

```
plot(snaive(wineind))
```

seasadj

Seasonal adjustment

Description

Returns seasonally adjusted data constructed by removing the seasonal component.

Usage

```
seasadj(object, ...)  
  
## S3 method for class 'stl'  
seasadj(object, ...)  
  
## S3 method for class 'decomposed.ts'  
seasadj(object, ...)  
  
## S3 method for class 'tbats'  
seasadj(object, ...)  
  
## S3 method for class 'seas'  
seasadj(object, ...)
```

Arguments

object	Object created by decompose , stl or tbats .
...	Other arguments not currently used.

Value

Univariate time series.

Author(s)

Rob J Hyndman

See Also

[stl](#), [decompose](#), [tbats](#).

Examples

```
plot(AirPassengers)  
lines(seasadj(decompose(AirPassengers,"multiplicative")),col=4)
```

seasonal

Extract components from a time series decomposition

Description

Returns a univariate time series equal to either a seasonal component, trend-cycle component or remainder component from a time series decomposition.

Usage

```
seasonal(object)

trendcycle(object)

remainder(object)
```

Arguments

object Object created by [decompose](#), [stl](#) or [tbats](#).

Value

Univariate time series.

Author(s)

Rob J Hyndman

See Also

[stl](#), [decompose](#), [tbats](#), [seasadj](#).

Examples

```
plot(USAccDeaths)
fit <- stl(USAccDeaths, s.window="periodic")
lines(trendcycle(fit),col="red")

library(ggplot2)
autoplot(cbind(
  Data=USAccDeaths,
  Seasonal=seasonal(fit),
  Trend=trendcycle(fit),
  Remainder=remainder(fit)),
  facets=TRUE) +
  ylab("") + xlab("Year")
```

seasonaldummy

Seasonal dummy variables

Description

seasonaldummy returns a matrix of dummy variables suitable for use in [Arima](#), [auto.arima](#) or [tslm](#). The last season is omitted and used as the control.

Usage

```
seasonaldummy(x, h = NULL)
```

```
seasonaldummyf(x, h)
```

Arguments

x	Seasonal time series: a ts or a msts object
h	Number of periods ahead to forecast (optional)

Details

`seasonaldummyf` is deprecated, instead use the `h` argument in `seasonaldummy`.

The number of dummy variables is determined from the time series characteristics of `x`. When `h` is missing, the length of `x` also determines the number of rows for the matrix returned by `seasonaldummy`. the value of `h` determines the number of rows for the matrix returned by `seasonaldummy`, typically used for forecasting. The values within `x` are not used.

Value

Numerical matrix.

Author(s)

Rob J Hyndman

See Also

[fourier](#)

Examples

```
plot(ldeaths)

# Using seasonal dummy variables
month <- seasonaldummy(ldeaths)
deaths.lm <- tslm(ldeaths ~ month)
tsdisplay(residuals(deaths.lm))
ldeaths.fcast <- forecast(deaths.lm,
  data.frame(month=I(seasonaldummy(ldeaths,36))))
plot(ldeaths.fcast)

# A simpler approach to seasonal dummy variables
deaths.lm <- tslm(ldeaths ~ season)
ldeaths.fcast <- forecast(deaths.lm, h=36)
plot(ldeaths.fcast)
```


Description

Returns forecasts and other information for exponential smoothing forecasts applied to y .

Usage

```
ses(y, h = 10, level = c(80, 95), fan = FALSE, initial = c("optimal",
  "simple"), alpha = NULL, lambda = NULL, biasadj = FALSE, x = y, ...)
```

```
holt(y, h = 10, damped = FALSE, level = c(80, 95), fan = FALSE,
  initial = c("optimal", "simple"), exponential = FALSE, alpha = NULL,
  beta = NULL, phi = NULL, lambda = NULL, biasadj = FALSE, x = y, ...)
```

```
hw(y, h = 2 * frequency(x), seasonal = c("additive", "multiplicative"),
  damped = FALSE, level = c(80, 95), fan = FALSE, initial = c("optimal",
  "simple"), exponential = FALSE, alpha = NULL, beta = NULL,
  gamma = NULL, phi = NULL, lambda = NULL, biasadj = FALSE, x = y,
  ...)
```

Arguments

<code>y</code>	a numeric vector or time series of class <code>ts</code>
<code>h</code>	Number of periods for forecasting.
<code>level</code>	Confidence level for prediction intervals.
<code>fan</code>	If TRUE, level is set to <code>seq(51,99,by=3)</code> . This is suitable for fan plots.
<code>initial</code>	Method used for selecting initial state values. If <code>optimal</code> , the initial values are optimized along with the smoothing parameters using <code>ets</code> . If <code>simple</code> , the initial values are set to values obtained using simple calculations on the first few observations. See Hyndman & Athanasopoulos (2014) for details.
<code>alpha</code>	Value of smoothing parameter for the level. If NULL, it will be estimated.
<code>lambda</code>	Box-Cox transformation parameter. Ignored if NULL. Otherwise, data transformed before model is estimated. When <code>lambda=TRUE</code> , <code>additive.only</code> is set to FALSE.
<code>biasadj</code>	Use adjusted back-transformed mean for Box-Cox transformations. If TRUE, point forecasts and fitted values are mean forecast. Otherwise, these points can be considered the median of the forecast densities.
<code>x</code>	Deprecated. Included for backwards compatibility.
<code>...</code>	Other arguments passed to <code>forecast.ets</code> .
<code>damped</code>	If TRUE, use a damped trend.
<code>exponential</code>	If TRUE, an exponential trend is fitted. Otherwise, the trend is (locally) linear.

beta	Value of smoothing parameter for the trend. If NULL, it will be estimated.
phi	Value of damping parameter if damped=TRUE. If NULL, it will be estimated.
seasonal	Type of seasonality in hw model. "additive" or "multiplicative"
gamma	Value of smoothing parameter for the seasonal component. If NULL, it will be estimated.

Details

ses, holt and hw are simply convenient wrapper functions for `forecast(ets(...))`.

Value

An object of class "forecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `ets` and associated functions.

An object of class "forecast" is a list containing at least the following elements:

model	A list containing information about the fitted model
method	The name of the forecasting method as a character string
mean	Point forecasts as a time series
lower	Lower limits for prediction intervals
upper	Upper limits for prediction intervals
level	The confidence values associated with the prediction intervals
x	The original time series (either object itself or the time series used to create the model stored as object).
residuals	Residuals from the fitted model.
fitted	Fitted values (one-step forecasts)

Author(s)

Rob J Hyndman

References

Hyndman, R.J., Koehler, A.B., Ord, J.K., Snyder, R.D. (2008) *Forecasting with exponential smoothing: the state space approach*, Springer-Verlag: New York. <http://www.exponentialsmoothing.net>.

Hyndman, R.J., Athanasopoulos (2014) *Forecasting: principles and practice*, OTexts: Melbourne, Australia. <http://www.otexts.org/fpp>.

See Also

[ets](#), [HoltWinters](#), [rwf](#), [arima](#).

Examples

```
fcast <- holt(airmiles)
plot(fcast)
deaths.fcast <- hw(USAccDeaths,h=48)
plot(deaths.fcast)
```

simulate.ets

Simulation from a time series model

Description

Returns a time series based on the model object object.

Usage

```
## S3 method for class 'ets'
simulate(object, nsim = length(object$x), seed = NULL,
         future = TRUE, bootstrap = FALSE, innov = NULL, ...)

## S3 method for class 'Arima'
simulate(object, nsim = length(object$x), seed = NULL,
         xreg = NULL, future = TRUE, bootstrap = FALSE, innov = NULL,
         lambda = object$lambda, ...)

## S3 method for class 'ar'
simulate(object, nsim = object$n.used, seed = NULL,
         future = TRUE, bootstrap = FALSE, innov = NULL, ...)

## S3 method for class 'fracdiff'
simulate(object, nsim = object$n, seed = NULL,
         future = TRUE, bootstrap = FALSE, innov = NULL, ...)

## S3 method for class 'nnetar'
simulate(object, nsim = length(object$x), seed = NULL,
         xreg = NULL, future = TRUE, bootstrap = FALSE, innov = NULL,
         lambda = object$lambda, ...)
```

Arguments

object	An object of class "ets", "Arima", "ar" or "nnetar".
nsim	Number of periods for the simulated series. Ignored if either xreg or innov are not NULL.
seed	Either NULL or an integer that will be used in a call to <code>set.seed</code> before simulating the time series. The default, NULL, will not change the random generator state.

future	Produce sample paths that are future to and conditional on the data in object. Otherwise simulate unconditionally.
bootstrap	Do simulation using resampled errors rather than normally distributed errors or errors provided as innov.
innov	A vector of innovations to use as the error series. Ignored if bootstrap==TRUE. If not NULL, the value of nsim is set to length of innov.
...	Other arguments, not currently used.
xreg	New values of xreg to be used for forecasting. The value of nsim is set to the number of rows of xreg if it is not NULL.
lambda	Box-Cox parameter. If not NULL, the simulated series is transformed using an inverse Box-Cox transformation with parameter lambda.

Details

With `simulate.Arima()`, the object should be produced by [Arima](#) or [auto.arima](#), rather than [arima](#). By default, the error series is assumed normally distributed and generated using [rnorm](#). If `innov` is present, it is used instead. If `bootstrap=TRUE` and `innov=NULL`, the residuals are resampled instead.

When `future=TRUE`, the sample paths are conditional on the data. When `future=FALSE` and the model is stationary, the sample paths do not depend on the data at all. When `future=FALSE` and the model is non-stationary, the location of the sample paths is arbitrary, so they all start at the value of the first observation.

Value

An object of class "ts".

Author(s)

Rob J Hyndman

See Also

[ets](#), [Arima](#), [auto.arima](#), [ar](#), [arfima](#), [nnetar](#).

Examples

```
fit <- ets(USAccDeaths)
plot(USAccDeaths, xlim=c(1973,1982))
lines(simulate(fit, 36), col="red")
```

sindexf	<i>Forecast seasonal index</i>
---------	--------------------------------

Description

Returns vector containing the seasonal index for h future periods. If the seasonal index is non-periodic, it uses the last values of the index.

Usage

```
sindexf(object, h)
```

Arguments

object	Output from decompose or stl .
h	Number of periods ahead to forecast

Value

Time series

Author(s)

Rob J Hyndman

Examples

```
uk.stl <- stl(UKDriverDeaths, "periodic")
uk.sa <- seasadj(uk.stl)
uk.fcast <- holt(uk.sa, 36)
seasf <- sindexf(uk.stl, 36)
uk.fcast$mean <- uk.fcast$mean + seasf
uk.fcast$lower <- uk.fcast$lower + cbind(seasf, seasf)
uk.fcast$upper <- uk.fcast$upper + cbind(seasf, seasf)
uk.fcast$x <- UKDriverDeaths
plot(uk.fcast, main="Forecasts from Holt's method with seasonal adjustment")
```

splinef

*Cubic Spline Forecast***Description**

Returns local linear forecasts and prediction intervals using cubic smoothing splines.

Usage

```
splinef(y, h = 10, level = c(80, 95), fan = FALSE, lambda = NULL,
        biasadj = FALSE, method = c("gcv", "mle"), x = y)
```

Arguments

y	a numeric vector or time series of class <code>ts</code>
h	Number of periods for forecasting
level	Confidence level for prediction intervals.
fan	If TRUE, level is set to <code>seq(51,99,by=3)</code> . This is suitable for fan plots.
lambda	Box-Cox transformation parameter. Ignored if NULL. Otherwise, forecasts back-transformed via an inverse Box-Cox transformation.
biasadj	Use adjusted back-transformed mean for Box-Cox transformations. If TRUE, point forecasts and fitted values are mean forecast. Otherwise, these points can be considered the median of the forecast densities.
method	Method for selecting the smoothing parameter. If <code>method="gcv"</code> , the generalized cross-validation method from smooth.spline is used. If <code>method="mle"</code> , the maximum likelihood method from Hyndman et al (2002) is used.
x	Deprecated. Included for backwards compatibility.

Details

The cubic smoothing spline model is equivalent to an ARIMA(0,2,2) model but with a restricted parameter space. The advantage of the spline model over the full ARIMA model is that it provides a smooth historical trend as well as a linear forecast function. Hyndman, King, Pitrun, and Billah (2002) show that the forecast performance of the method is hardly affected by the restricted parameter space.

Value

An object of class "forecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `splinef`.

An object of class "forecast" containing the following elements:

model	A list containing information about the fitted model
method	The name of the forecasting method as a character string
mean	Point forecasts as a time series
lower	Lower limits for prediction intervals
upper	Upper limits for prediction intervals
level	The confidence values associated with the prediction intervals
x	The original time series (either object itself or the time series used to create the model stored as object).
onestepf	One-step forecasts from the fitted model.
fitted	Smooth estimates of the fitted trend using all data.
residuals	Residuals from the fitted model. That is x minus one-step forecasts.

Author(s)

Rob J Hyndman

References

Hyndman, King, Pitrun and Billah (2005) Local linear forecasts using cubic smoothing splines. *Australian and New Zealand Journal of Statistics*, **47**(1), 87-99. <https://robjhyndman.com/publications/splinefcast/>.

See Also

[smooth.spline](#), [arima](#), [holt](#).

Examples

```
fcast <- splinef(uspop,h=5)
plot(fcast)
summary(fcast)
```

StatForecast

Forecast plot

Description

Generates forecasts from `forecast.ts` and adds them to the plot. Forecasts can be modified via sending forecast specific arguments above.

Usage

```
StatForecast
```

```
GeomForecast
```

```
geom_forecast(mapping = NULL, data = NULL, stat = "forecast",
  position = "identity", na.rm = FALSE, show.legend = NA,
  inherit.aes = TRUE, PI = TRUE, showgap = TRUE, series = NULL, ...)
```

Arguments

mapping	Set of aesthetic mappings created by aes or aes_ . If specified and <code>inherit.aes = TRUE</code> (the default), it is combined with the default mapping at the top level of the plot. You must supply mapping if there is no plot mapping.
data	The data to be displayed in this layer. There are three options: If <code>NULL</code> , the default, the data is inherited from the plot data as specified in the call to ggplot . A <code>data.frame</code> , or other object, will override the plot data. All objects will be fortified to produce a data frame. See fortify for which variables will be created. A function will be called with a single argument, the plot data. The return value must be a <code>data.frame</code> , and will be used as the layer data.
stat	The stat object to use calculate the data.
position	Position adjustment, either as a string, or the result of a call to a position adjustment function.
na.rm	If <code>FALSE</code> (the default), removes missing values with a warning. If <code>TRUE</code> silently removes missing values.
show.legend	logical. Should this layer be included in the legends? <code>NA</code> , the default, includes if any aesthetics are mapped. <code>FALSE</code> never includes, and <code>TRUE</code> always includes.
inherit.aes	If <code>FALSE</code> , overrides the default aesthetics, rather than combining with them. This is most useful for helper functions that define both data and aesthetics and shouldn't inherit behaviour from the default plot specification, e.g. borders .
PI	If <code>FALSE</code> , confidence intervals will not be plotted, giving only the forecast line.
showgap	If <code>showgap=FALSE</code> , the gap between the historical observations and the forecasts is removed.
series	Matches an unidentified forecast layer with a coloured object on the plot.
...	Additional arguments for forecast.ts , other arguments are passed on to layer . These are often aesthetics, used to set an aesthetic to a fixed value, like <code>color = "red"</code> or <code>alpha = .5</code> . They may also be parameters to the paired geom/stat.

Format

An object of class `StatForecast` (inherits from `Stat`, `ggproto`, `gg`) of length 3.

Details

Multivariate forecasting is supported by having each time series on a different group.

You can also pass `geom_forecast` a forecast object to add it to the plot.

The aesthetics required for the forecasting to work includes forecast observations on the y axis, and the time of the observations on the x axis. Refer to the examples below. To automatically set up aesthetics, use `autoplot`.

Value

A layer for a ggplot graph.

Author(s)

Mitchell O'Hara-Wild

See Also

[forecast](#), [ggproto](#)

Examples

```
## Not run:
library(ggplot2)
autoplot(USAccDeaths) + geom_forecast()

lungDeaths <- cbind(mdeaths, fdeaths)
autoplot(lungDeaths) + geom_forecast()

# Using fortify.ts
p <- ggplot(aes(x=x, y=y), data=USAccDeaths)
p <- p + geom_line()
p + geom_forecast()

# Without fortify.ts
data <- data.frame(USAccDeaths=as.numeric(USAccDeaths), time=as.numeric(time(USAccDeaths)))
p <- ggplot(aes(x=time, y=USAccDeaths), data=data)
p <- p + geom_line()
p + geom_forecast()

p + geom_forecast(h=60)
p <- ggplot(aes(x=time, y=USAccDeaths), data=data)
p + geom_forecast(level=c(70,98))
p + geom_forecast(level=c(70,98),colour="lightblue")

#Add forecasts to multivariate series with colour groups
lungDeaths <- cbind(mdeaths, fdeaths)
autoplot(lungDeaths) + geom_forecast(forecast(mdeaths), series="mdeaths")

## End(Not run)
```

subset.ts

*Subsetting a time series***Description**

Various types of subsetting of a time series. Allows subsetting by index values (unlike [window](#)). Also allows extraction of the values of a specific season or subset of seasons in each year. For example, to extract all values for the month of May from a time series.

Usage

```
## S3 method for class 'ts'
subset(x, subset = NULL, month = NULL, quarter = NULL,
       season = NULL, start = NULL, end = NULL, ...)

## S3 method for class 'msts'
subset(x, subset = NULL, start = NULL, end = NULL, ...)
```

Arguments

x	a univariate time series to be subsetting
subset	optional logical expression indicating elements to keep; missing values are taken as false. subset must be the same length as x.
month	Numeric or character vector of months to retain. Partial matching on month names used.
quarter	Numeric or character vector of quarters to retain.
season	Numeric vector of seasons to retain.
start	Index of start of contiguous subset.
end	Index of end of contiguous subset.
...	Other arguments, unused.

Details

If character values for months are used, either upper or lower case may be used, and partial unambiguous names are acceptable. Possible character values for quarters are "Q1", "Q2", "Q3", and "Q4".

Value

If subset is used, a numeric vector is returned with no ts attributes. If start and/or end are used, a ts object is returned consisting of x[start:end], with the appropriate time series attributes retained. Otherwise, a ts object is returned with frequency equal to the length of month, quarter or season.

Author(s)

Rob J Hyndman

See Also

[subset](#), [window](#)

Examples

```
plot(subset(gas, month="November"))
subset(woolyrnq, quarter=3)
subset(USAccDeaths, start=49)
```

taylor

Half-hourly electricity demand

Description

Half-hourly electricity demand in England and Wales from Monday 5 June 2000 to Sunday 27 August 2000. Discussed in Taylor (2003), and kindly provided by James W Taylor. Units: Megawatts

Usage

```
taylor
```

Format

Time series data

Source

James W Taylor

References

Taylor, J.W. (2003) Short-term electricity demand forecasting using double seasonal exponential smoothing. *Journal of the Operational Research Society*, **54**, 799-805.

Examples

```
plot(taylor)
```

tbats	<i>TBATS model (Exponential smoothing state space model with Box-Cox transformation, ARMA errors, Trend and Seasonal components)</i>
-------	--

Description

Fits a TBATS model applied to y , as described in De Livera, Hyndman & Snyder (2011). Parallel processing is used by default to speed up the computations.

Usage

```
tbats(y, use.box.cox = NULL, use.trend = NULL, use.damped.trend = NULL,
      seasonal.periods = NULL, use.arma.errors = TRUE,
      use.parallel = length(y) > 1000, num.cores = 2, bc.lower = 0,
      bc.upper = 1, biasadj = FALSE, model = NULL, ...)
```

Arguments

<code>y</code>	The time series to be forecast. Can be numeric, msts or ts. Only univariate time series are supported.
<code>use.box.cox</code>	TRUE/FALSE indicates whether to use the Box-Cox transformation or not. If NULL then both are tried and the best fit is selected by AIC.
<code>use.trend</code>	TRUE/FALSE indicates whether to include a trend or not. If NULL then both are tried and the best fit is selected by AIC.
<code>use.damped.trend</code>	TRUE/FALSE indicates whether to include a damping parameter in the trend or not. If NULL then both are tried and the best fit is selected by AIC.
<code>seasonal.periods</code>	If <code>y</code> is numeric then seasonal periods can be specified with this parameter.
<code>use.arma.errors</code>	TRUE/FALSE indicates whether to include ARMA errors or not. If TRUE the best fit is selected by AIC. If FALSE then the selection algorithm does not consider ARMA errors.
<code>use.parallel</code>	TRUE/FALSE indicates whether or not to use parallel processing.
<code>num.cores</code>	The number of parallel processes to be used if using parallel processing. If NULL then the number of logical cores is detected and all available cores are used.
<code>bc.lower</code>	The lower limit (inclusive) for the Box-Cox transformation.
<code>bc.upper</code>	The upper limit (inclusive) for the Box-Cox transformation.
<code>biasadj</code>	Use adjusted back-transformed mean for Box-Cox transformations. If TRUE, point forecasts and fitted values are mean forecast. Otherwise, these points can be considered the median of the forecast densities.
<code>model</code>	Output from a previous call to <code>tbats</code> . If <code>model</code> is passed, this same model is fitted to <code>y</code> without re-estimating any parameters.

... Additional arguments to be passed to `auto.arima` when choose an ARMA(p, q) model for the errors. (Note that `xreg` will be ignored, as will any arguments concerning seasonality and differencing, but arguments controlling the values of p and q will be used.)

Value

An object with class `c("tbats", "bats")`. The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `bats` and associated functions. The fitted model is designated `TBATS(omega, p,q, phi, <m1,k1>,...,<mJ,kJ>)` where `omega` is the Box-Cox parameter and `phi` is the damping parameter; the error is modelled as an ARMA(p,q) process and `m1,...,mJ` list the seasonal periods used in the model and `k1,...,kJ` are the corresponding number of Fourier terms used for each seasonality.

Author(s)

Slava Razbash and Rob J Hyndman

References

De Livera, A.M., Hyndman, R.J., & Snyder, R. D. (2011), Forecasting time series with complex seasonal patterns using exponential smoothing, *Journal of the American Statistical Association*, **106**(496), 1513-1527.

See Also

[tbats.components](#).

Examples

```
## Not run:
fit <- tbats(USAccDeaths)
plot(forecast(fit))

taylor.fit <- tbats(taylor)
plot(forecast(taylor.fit))
## End(Not run)
```

tbats.components

Extract components of a TBATS model

Description

Extract the level, slope and seasonal components of a TBATS model.

Usage

```
tbats.components(x)
```

Arguments

x A tbats object created by [tbats](#).

Value

A multiple time series (mts) object.

Author(s)

Slava Razbash and Rob J Hyndman

References

De Livera, A.M., Hyndman, R.J., & Snyder, R. D. (2011), Forecasting time series with complex seasonal patterns using exponential smoothing, *Journal of the American Statistical Association*, **106**(496), 1513-1527.

See Also

[tbats](#).

Examples

```
## Not run:  
fit <- tbats(USAccDeaths, use.parallel=FALSE)  
components <- tbats.components(fit)  
plot(components)  
## End(Not run)
```

thetaf

Theta method forecast

Description

Returns forecasts and prediction intervals for a theta method forecast.

Usage

```
thetaf(y, h = ifelse(frequency(y) > 1, 2 * frequency(y), 10), level = c(80,  
95), fan = FALSE, x = y)
```

Arguments

y	a numeric vector or time series of class ts
h	Number of periods for forecasting
level	Confidence levels for prediction intervals.
fan	If TRUE, level is set to seq(51,99,by=3). This is suitable for fan plots.
x	Deprecated. Included for backwards compatibility.

Details

The theta method of Assimakopoulos and Nikolopoulos (2000) is equivalent to simple exponential smoothing with drift. This is demonstrated in Hyndman and Billah (2003).

The series is tested for seasonality using the test outlined in A&N. If deemed seasonal, the series is seasonally adjusted using a classical multiplicative decomposition before applying the theta method. The resulting forecasts are then reseasonalized.

Prediction intervals are computed using the underlying state space model.

More general theta methods are available in the [forecTheta](#) package.

Value

An object of class "forecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `rwf`.

An object of class "forecast" is a list containing at least the following elements:

model	A list containing information about the fitted model
method	The name of the forecasting method as a character string
mean	Point forecasts as a time series
lower	Lower limits for prediction intervals
upper	Upper limits for prediction intervals
level	The confidence values associated with the prediction intervals
x	The original time series (either object itself or the time series used to create the model stored as object).
residuals	Residuals from the fitted model. That is x minus fitted values.
fitted	Fitted values (one-step forecasts)

Author(s)

Rob J Hyndman

References

Assimakopoulos, V. and Nikolopoulos, K. (2000). The theta model: a decomposition approach to forecasting. *International Journal of Forecasting* **16**, 521-530.

Hyndman, R.J., and Billah, B. (2003) Unmasking the Theta method. *International J. Forecasting*, **19**, 287-290.

See Also

[arima](#), [meanf](#), [rwf](#), [ses](#)

Examples

```
nile.fcast <- thetaf(Nile)
plot(nile.fcast)
```

tsclean

Identify and replace outliers and missing values in a time series

Description

Uses supsmu for non-seasonal series and a periodic stl decomposition with seasonal series to identify outliers. To estimate missing values and outlier replacements, linear interpolation is used on the (possibly seasonally adjusted) series

Usage

```
tsclean(x, replace.missing = TRUE, lambda = NULL)
```

Arguments

x	time series
replace.missing	If TRUE, it not only replaces outliers, but also interpolates missing values
lambda	a numeric value giving the Box-Cox transformation parameter

Value

Time series

Author(s)

Rob J Hyndman

See Also

[na.interp](#), [tsoutliers](#), [supsmu](#)

Examples

```
cleangold <- tsclean(gold)
```

tsCV

Time series cross-validation

Description

tsCV computes the forecast errors obtained by applying forecastfunction to subsets of the time series y using a rolling forecast origin.

Usage

```
tsCV(y, forecastfunction, h = 1, ...)
```

Arguments

y	Univariate time series
forecastfunction	Function to return an object of class forecast. Its first argument must be a univariate time series, and it must have an argument h for the forecast horizon.
h	Forecast horizon
...	Other arguments are passed to forecastfunction.

Details

Let y contain the time series y_1, \dots, y_T . Then forecastfunction is applied successively to the time series y_1, \dots, y_t , for $t = 1, \dots, T - h$, making predictions $\hat{y}_{t+h|t}$. The errors are given by $e_{t+h} = y_{t+h} - \hat{y}_{t+h|t}$. These are returned as a vector, e_1, \dots, e_T . The first few errors may be missing as it may not be possible to apply forecastfunction to very short time series.

Value

Numerical time series object containing the forecast errors.

Author(s)

Rob J Hyndman

See Also

[CV](#), [CVar](#), [residuals.Arima](#).

Examples

```
#Fit an AR(2) model to each subset
far2 <- function(x, h){forecast(Arima(x, order=c(2,0,0)), h=h)}
e <- tsCV(lynx, far2, h=1)
```

 tslm

Fit a linear model with time series components

Description

tslm is used to fit linear models to time series including trend and seasonality components.

Usage

```
tslm(formula, data, subset, lambda = NULL, biasadj = FALSE, ...)
```

Arguments

formula	an object of class "formula" (or one that can be coerced to that class): a symbolic description of the model to be fitted.
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 environment(formula), typically the environment from which lm is called.
subset	an optional subset containing rows of data to keep. For best results, pass a logical vector of rows to keep. Also supports <code>subset()</code> functions.
lambda	Box-Cox transformation parameter. Ignored if NULL. Otherwise, data are transformed via a Box-Cox transformation.
biasadj	Use adjusted back-transformed mean for Box-Cox transformations. If TRUE, point forecasts and fitted values are mean forecast. Otherwise, these points can be considered the median of the forecast densities.
...	Other arguments passed to <code>lm()</code>

Details

tslm is largely a wrapper for `lm()` except that it allows variables "trend" and "season" which are created on the fly from the time series characteristics of the data. The variable "trend" is a simple time trend and "season" is a factor indicating the season (e.g., the month or the quarter depending on the frequency of the data).

Value

Returns an object of class "lm".

Author(s)

Mitchell O'Hara-Wild and Rob J Hyndman

See Also

[forecast.lm](#), [lm](#).

Examples

```
y <- ts(rnorm(120,0,3) + 1:120 + 20*sin(2*pi*(1:120)/12), frequency=12)
fit <- tslm(y ~ trend + season)
plot(forecast(fit, h=20))
```

tsoutliers

Identify and replace outliers in a time series

Description

Uses supsmu for non-seasonal series and a periodic stl decomposition with seasonal series to identify outliers and estimate their replacements.

Usage

```
tsoutliers(x, iterate = 2, lambda = NULL)
```

Arguments

x	time series
iterate	the number of iteration only for non-seasonal series
lambda	Allowing Box-cox transformation

Value

index	Indicating the index of outlier(s)
replacement	Suggested numeric values to replace identified outliers

Author(s)

Rob J Hyndman

See Also

[na.interp](#), [tsclean](#)

Examples

```
data(gold)
tsoutliers(gold)
```

wineind	<i>Australian total wine sales</i>
---------	------------------------------------

Description

Australian total wine sales by wine makers in bottles \leq 1 litre. Jan 1980 – Aug 1994.

Usage

```
wineind
```

Format

Time series data

Source

Time Series Data Library. <http://data.is/TSDLdemo>

Examples

```
tsdisplay(wineind)
```

woolyrnq	<i>Quarterly production of woollen yarn in Australia</i>
----------	--

Description

Quarterly production of woollen yarn in Australia: tonnes. Mar 1965 – Sep 1994.

Usage

```
woolyrnq
```

Format

Time series data

woolyrnq

117

Source

Time Series Data Library. <http://data.is/TSDLdemo>

Examples

```
tsdisplay(woolyrnq)
```

Index

- *Topic **datasets**
 - gas, [65](#)
 - gold, [73](#)
 - StatForecast, [103](#)
 - taylor, [107](#)
 - wineind, [116](#)
 - woolyrnq, [116](#)
- *Topic **hplot**
 - plot.Arima, [84](#)
 - plot.bats, [85](#)
 - plot.ets, [86](#)
- *Topic **htest**
 - dm.test, [34](#)
- *Topic **models**
 - CV, [32](#)
- *Topic **package**
 - forecast-package, [4](#)
- *Topic **stats**
 - forecast.lm, [53](#)
 - tslm, [114](#)
- *Topic **ts**
 - accuracy, [4](#)
 - Acf, [6](#)
 - arfima, [8](#)
 - Arima, [9](#)
 - arima.errors, [11](#)
 - arimaorder, [12](#)
 - auto.arima, [13](#)
 - autoplot.mforecast, [20](#)
 - baggedETS, [22](#)
 - bats, [23](#)
 - bizdays, [25](#)
 - bld.mbb.bootstrap, [26](#)
 - BoxCox, [27](#)
 - BoxCox.lambda, [28](#)
 - croston, [30](#)
 - CVar, [33](#)
 - dm.test, [34](#)
 - dshw, [35](#)
 - easter, [37](#)
 - ets, [38](#)
 - findfrequency, [40](#)
 - fitted.fracdiff, [41](#)
 - forecast, [42](#)
 - forecast.baggedETS, [44](#)
 - forecast.bats, [46](#)
 - forecast.ets, [47](#)
 - forecast.fracdiff, [49](#)
 - forecast.HoltWinters, [51](#)
 - forecast.nnetar, [58](#)
 - forecast.stl, [60](#)
 - forecast.StructTS, [62](#)
 - fourier, [64](#)
 - getResponse, [66](#)
 - ggseasonplot, [70](#)
 - ggtsdisplay, [71](#)
 - ma, [75](#)
 - meanf, [76](#)
 - monthdays, [77](#)
 - msts, [78](#)
 - na.interp, [79](#)
 - ndiffs, [80](#)
 - nnetar, [82](#)
 - plot.forecast, [87](#)
 - residuals.tslm, [90](#)
 - rwf, [91](#)
 - seasadj, [93](#)
 - seasonal, [94](#)
 - seasonaldummy, [95](#)
 - ses, [97](#)
 - simulate.ets, [99](#)
 - sindexf, [101](#)
 - splinef, [102](#)
 - subset.ts, [106](#)
 - tbats, [108](#)
 - tbats.components, [109](#)
 - thetaf, [110](#)
 - tsclean, [112](#)

- tsCV, 113
- tsoutliers, 115
- accuracy, 4
- Acf, 6, 19, 72
- acf, 7, 19, 72
- aes, 104
- aes_, 104
- AIC, 32
- ar, 12, 13, 50, 51, 61, 85, 100
- arfima, 8, 12, 13, 50, 51, 100
- Arima, 9, 12, 13, 15, 40, 51, 64, 85, 93, 95, 100
- arima, 8–14, 43, 50, 51, 58, 98, 100, 103, 112
- arima.errors, 11
- arimaorder, 12
- as.character.Arima (Arima), 9
- as.character.bats (bats), 23
- as.character.ets (ets), 38
- as.character.tbats (tbats), 108
- as.data.frame.forecast (forecast), 42
- as.data.frame.mforecast (forecast.mts), 56
- as.ts.forecast (forecast), 42
- auto.arima, 8, 9, 11–13, 13, 50, 51, 61, 64, 81, 95, 100
- autolayer, 15
- autolayer.forecast (plot.forecast), 87
- autolayer.mforecast (autoplot.mforecast), 20
- autolayer.mts (autolayer.mts), 16
- autolayer.mts, 16
- autolayer.ts (autolayer.mts), 16
- autoplot, 16
- autoplot.acf, 17
- autoplot.ar (plot.Arima), 84
- autoplot.Arima (plot.Arima), 84
- autoplot.bats (plot.bats), 85
- autoplot.decomposed.ts, 19
- autoplot.ets (plot.ets), 86
- autoplot.forecast (plot.forecast), 87
- autoplot.mforecast, 20
- autoplot.mpacf (autoplot.acf), 17
- autoplot.mts (autolayer.mts), 16
- autoplot.mts (autolayer.mts), 16
- autoplot.seas (autoplot.decomposed.ts), 19
- autoplot.splineforecast (plot.forecast), 87
- autoplot.stl (autoplot.decomposed.ts), 19
- autoplot.StructTS (autoplot.decomposed.ts), 19
- autoplot.tbats (plot.bats), 85
- autoplot.ts (autolayer.mts), 16
- baggedETS, 22, 26, 45, 46
- bats, 23, 46, 47, 86
- bgtest, 30
- bizdays, 25, 78
- bld.mbb.bootstrap, 22, 26
- borders, 104
- Box.test, 30
- BoxCox, 27, 29
- BoxCox.lambda, 28, 28
- Ccf (Acf), 6
- ccf, 7
- checkresiduals, 29, 91
- coef.ets (ets), 38
- croston, 30, 44
- CV, 32, 33, 113
- CVar, 33, 113
- decompose, 19, 20, 76, 94, 95, 101
- dm.test, 34
- dshw, 35
- easter, 37
- ets, 22, 36, 37, 38, 43, 45, 48, 49, 52, 60, 87, 97, 98, 100
- findfrequency, 40
- fitted.ar (fitted.fracdiff), 41
- fitted.Arima, 91
- fitted.Arima (fitted.fracdiff), 41
- fitted.bats (fitted.fracdiff), 41
- fitted.ets (fitted.fracdiff), 41
- fitted.fracdiff, 41
- fitted.nnetar (fitted.fracdiff), 41
- fitted.tbats (fitted.fracdiff), 41
- forecast, 42, 61, 62, 88, 89, 105
- forecast-package, 4
- forecast.ar (forecast.fracdiff), 49
- forecast.Arima, 11, 42–44, 61, 62
- forecast.Arima (forecast.fracdiff), 49
- forecast.baggedETS, 22, 44
- forecast.bats, 42, 46

- forecast.ets, [42](#), [44](#), [47](#), [47](#), [62](#)
- forecast.fracdiff, [9](#), [49](#), [50](#)
- forecast.HoltWinters, [44](#), [51](#)
- forecast.lm, [53](#), [55](#), [56](#), [115](#)
- forecast.mlm, [54](#), [57](#)
- forecast.mts, [56](#)
- forecast.nnetar, [42](#), [58](#)
- forecast.stl, [60](#)
- forecast.stlm (forecast.stl), [60](#)
- forecast.StructTS, [44](#), [62](#)
- forecast.tbats, [42](#)
- forecast.tbats (forecast.bats), [46](#)
- forecast.ts, [43](#), [104](#)
- forecTheta, [111](#)
- fortify, [16](#), [17](#), [104](#)
- fortify.ts (autolayer.mts), [16](#)
- fourier, [64](#), [96](#)
- fourierf (fourier), [64](#)
- fracdiff, [9](#), [12](#), [50](#)

- gas, [65](#)
- geom_forecast, [16](#)
- geom_forecast (StatForecast), [103](#)
- geom_histogram, [67](#)
- GeomForecast (StatForecast), [103](#)
- getResponse, [66](#)
- ggAcf (autoplot.acf), [17](#)
- ggCcf (autoplot.acf), [17](#)
- gghistogram, [67](#)
- gglagchull (gglagplot), [68](#)
- gglagplot, [68](#)
- ggmonthplot, [69](#)
- ggPacf (autoplot.acf), [17](#)
- ggplot, [16](#), [104](#)
- ggproto, [105](#)
- ggseasonplot, [70](#)
- ggsubseriesplot (ggmonthplot), [69](#)
- ggtaperedacf (autoplot.acf), [17](#)
- ggtaperedpacf (autoplot.acf), [17](#)
- ggtsdisplay, [29](#), [30](#), [71](#)
- gold, [73](#)

- hist, [67](#)
- holt, [44](#), [49](#), [103](#)
- holt (ses), [97](#)
- HoltWinters, [37](#), [40](#), [51](#), [52](#), [98](#)
- hw, [44](#), [49](#)
- hw (ses), [97](#)

- InvBoxCox (BoxCox), [27](#)
- is.acf, [73](#)
- is.Arima (is.acf), [73](#)
- is.baggedETS (is.acf), [73](#)
- is.bats (is.acf), [73](#)
- is.constant, [74](#)
- is.ets (is.acf), [73](#)
- is.forecast, [74](#)
- is.mforecast (is.forecast), [74](#)
- is.nnetar (is.acf), [73](#)
- is.nnetarmodels (is.acf), [73](#)
- is.splineforecast (is.forecast), [74](#)
- is.stlm (is.acf), [73](#)
- isBizday, [25](#)

- lag.plot, [69](#)
- layer, [104](#)
- lm, [32](#), [53–56](#), [114](#), [115](#)

- ma, [75](#)
- meanf, [44](#), [76](#), [112](#)
- mforecast (forecast.mts), [56](#)
- monthdays, [25](#), [77](#)
- monthplot, [69–71](#)
- msts, [35](#), [78](#)

- na.contiguous, [7](#), [18](#), [72](#)
- na.interp, [7](#), [18](#), [72](#), [79](#), [80](#), [112](#), [115](#)
- na.pass, [7](#), [18](#), [72](#)
- naive (rwf), [91](#)
- nclass.FD, [67](#)
- ndiffs, [14](#), [80](#)
- nnet, [82](#)
- nnetar, [33](#), [59](#), [82](#), [100](#)
- nsdiffs, [14](#)
- nsdiffs (ndiffs), [80](#)

- Pacf (Acf), [6](#)
- pacf, [7](#)
- par, [84](#), [86](#)
- plot, [70](#), [71](#)
- plot.acf, [19](#)
- plot.ar (plot.Arima), [84](#)
- plot.Arima, [84](#)
- plot.bats, [85](#)
- plot.default, [89](#)
- plot.ets, [86](#)
- plot.forecast, [21](#), [87](#)
- plot.mforecast (autoplot.mforecast), [20](#)

- plot.splineforecast (plot.forecast), 87
- plot.stl, 20
- plot.tbats (plot.bats), 85
- plot.ts, 17, 21, 72, 89
- predict.ar, 50, 51
- predict.Arima, 50, 51
- predict.HoltWinters, 52
- predict.lm, 53
- print.ARIMA (Arima), 9
- print.baggedETS (baggedETS), 22
- print.bats (bats), 23
- print.CVar (CVar), 33
- print.ets (ets), 38
- print.forecast (forecast), 42
- print.mforecast (forecast.mts), 56
- print.msts (msts), 78
- print.naive (rwf), 91
- print.nnetar (nnetar), 82
- print.nnetarmodels (nnetar), 82
- print.tbats (tbats), 108

- remainder (seasonal), 94
- residuals.ar (residuals.tslm), 90
- residuals.Arima, 12, 42, 113
- residuals.Arima (residuals.tslm), 90
- residuals.bats, 42
- residuals.bats (residuals.tslm), 90
- residuals.ets, 42
- residuals.ets (residuals.tslm), 90
- residuals.forecast (residuals.tslm), 90
- residuals.fracdiff (residuals.tslm), 90
- residuals.nnetar, 42
- residuals.nnetar (residuals.tslm), 90
- residuals.stlm (residuals.tslm), 90
- residuals.tbats, 42
- residuals.tbats (residuals.tslm), 90
- residuals.tslm, 90
- rnorm, 100
- rwf, 40, 44, 77, 91, 98, 112

- seas, 19, 20
- seasadj, 93, 95
- seasonal, 94
- seasonaldummy, 65, 95
- seasonaldummyf (seasonaldummy), 95
- seasonplot (ggseasonplot), 70
- ses, 31, 44, 49, 97, 112
- set.seed, 99
- simulate.ar (simulate.ets), 99
- simulate.Arima (simulate.ets), 99
- simulate.ets, 99
- simulate.fracdiff (simulate.ets), 99
- simulate.nnetar, 58
- simulate.nnetar (simulate.ets), 99
- sindexf, 101
- smooth.spline, 102, 103
- snaive (rwf), 91
- spec.ar, 72
- splinef, 44, 102
- StatForecast, 103
- stl, 19, 20, 60–62, 94, 95, 101
- stlf, 43
- stlf (forecast.stl), 60
- stlm (forecast.stl), 60
- StructTS, 20, 63
- subset, 107, 114
- subset.msts (subset.ts), 106
- subset.ts, 106
- summary.Arima (Arima), 9
- summary.ets (ets), 38
- summary.forecast (forecast), 42
- summary.mforecast (forecast.mts), 56
- supsmu, 112

- taperedacf, 19
- taperedacf (Acf), 6
- taperedpacf (Acf), 6
- taylor, 107
- tbats, 47, 86, 94, 95, 108, 110
- tbats.components, 109, 109
- thetaf, 44, 62, 110
- trendcycle (seasonal), 94
- tsclean, 112, 115
- tsCV, 33, 113
- tsdiag.ets (ets), 38
- tsdisplay, 7
- tsdisplay (ggtsdisplay), 71
- tslm, 32, 53–57, 64, 95, 114
- tsoutliers, 80, 112, 115

- window, 106, 107
- window.msts (msts), 78
- wineind, 116
- woolyrnq, 116