

Package ‘tsfknn’

May 8, 2026

Type Package

Title Time Series Forecasting Using Nearest Neighbors

Version 0.6.0

Description Allows forecasting time series using nearest neighbors regression Francisco Martinez, Maria P. Frias, Maria D. Perez-Godoy and Antonio J. Rivera (2019) <[doi:10.1007/s10462-017-9593-z](https://doi.org/10.1007/s10462-017-9593-z)>. When the forecasting horizon is higher than 1, two multi-step ahead forecasting strategies can be used. The model built is autoregressive, that is, it is only based on the observations of the time series. The nearest neighbors used in a prediction can be consulted and plotted.

Maintainer Francisco Martinez <fmartin@ujaen.es>

License GPL-2

Encoding UTF-8

RoxygenNote 7.2.3

Depends R (>= 3.6.0)

Suggests knitr, rmarkdown, testthat (>= 3.0.0)

Imports ggplot2 (>= 3.1.1), graphics, Rcpp, stats, utils

VignetteBuilder knitr

URL <https://github.com/francisco martinezdelrio/tsfknn>

BugReports <https://github.com/francisco martinezdelrio/tsfknn/issues>

LinkingTo Rcpp

Config/testthat/edition 3

NeedsCompilation yes

Author Francisco Martinez [aut, cre]

Repository CRAN

Date/Publication 2023-12-20 16:20:02 UTC

Contents

autoplot.knnForecast	2
knn_examples	3
knn_forecasting	3
nearest_neighbors	5
n_training_examples	5
plot.knnForecastRO	6
predict.knnForecast	6
rolling_origin	7
Index	9

autoplot.knnForecast *Create a ggplot object from a knnForecast object*

Description

It uses a knnForecast object to create a ggplot object that plots a time series and its forecast using KNN regression.

Usage

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

Arguments

object An object of class knnForecast.
 ... additional parameter, see details.

Details

Commonly used parameters are:

- `highlight`. A character string indicating what elements should be highlighted. Possible values are "none", "points" and "neighbors". The default value is "none".
- `faceting`. Logical. This applies only if the `highlight` parameter is set to "neighbors". It indicates whether the different nearest neighbors should be seen in different plots (TRUE, the default value) or in one plot.

Value

The ggplot object representing a plotting with the forecast.

Examples

```
pred <- knn_forecasting(USAccDeaths, h = 12, lags = 1:12, k = 2)
library(ggplot2)
autoplot(pred)
autoplot(pred, highlight = "neighbors")
```

`knn_examples`*Examples of the model associated with a prediction*

Description

It allows to see the examples of the model associated to a `knnForecast` object.

Usage

```
knn_examples(forecast)
```

Arguments

`forecast` A `knnForecast` object.

Value

A matrix including the features and targets of the examples associated with the model of a `knnForecast` object.

Examples

```
pred <- knn_forecasting(ts(1:8), h = 1, lags = 1:2, k = 2)
knn_examples(pred)
```

`knn_forecasting`*Time series forecasting using KNN regression*

Description

It applies KNN regression to forecast the future values of a time series. The lags used as autoregressive variables are set with the `lags` parameter. If the user does not set the number of nearest neighbors or the lags, these values are selected automatically.

Usage

```
knn_forecasting(
  timeS,
  h,
  lags = NULL,
  k = c(3, 5, 7),
  msas = c("recursive", "MIMO"),
  cf = c("mean", "median", "weighted"),
  transform = c("additive", "multiplicative", "none")
)
```

Arguments

timeS	A numeric vector or time series of class <code>ts</code> .
h	A positive integer. Number of values to forecast.
lags	An integer vector in increasing order expressing the lags used as autoregressive variables.
k	A positive integer. The <code>k</code> parameter in KNN regression. A vector of <code>k</code> values can also be used. In that case, the forecast is the average of the forecasts produced by the different models with the different <code>k</code> parameters.
msas	A string indicating the Multiple-Step Ahead Strategy used when more than one value is predicted. It can be "recursive" or "MIMO" (the default).
cf	A string. It indicates the combination function used to aggregate the targets associated with the nearest neighbors. It can be "median", "weighted" or "mean" (the default).
transform	A character value indicating whether the training samples are transformed. If the time series has a trend it is recommended. By default is "multiplicative" (multiplicative transformation). It is also possible a multiplicative transformation or no transformation.

Value

An object of class "knnForecast". The function `summary` can be used to obtain or print a summary of the results.

An object of class `"knnForecast"` is a list containing at least the following components:

call	the matched call.
msas	the Multi-Step Ahead Strategy.
prediction	a time series with the forecast.
model	an object of class "knnModel" with the KNN model

Examples

```
pred <- knn_forecasting(USAccDeaths, h = 12, lags = 1:12, k = 2)
pred$prediction # To see a time series with the forecasts
plot(pred) # To see a plot with the forecast
```

nearest_neighbors	<i>Nearest neighbors associated with predictions</i>
-------------------	--

Description

It allows to check the new instances and their nearest neighbors used in a prediction associated with a "knnForecast" object.

Usage

```
nearest_neighbors(forecast)
```

Arguments

forecast A knnForecast object.

Value

A list including the new instances used in KNN regression and their nearest neighbors.

Examples

```
pred <- knn_forecasting(UKgas, h = 4, lags = 1:4, k = 2, msas = "MIMO")
nearest_neighbors(pred)
```

n_training_examples	<i>Number of training examples</i>
---------------------	------------------------------------

Description

It computes the number of training examples that would have a KNN model with the specified parameters.

Usage

```
n_training_examples(timeS, h, lags, msas = c("MIMO", "recursive"))
```

Arguments

timeS A numeric vector or time series of class ts.
h A positive integer. Number of values to forecast.
lags An integer vector in increasing order expressing the lags used as autoregressive variables.
msas A string indicating the Multiple-Step Ahead Strategy used when more than one value is predicted. It can be "recursive" or "MIMO" (the default).

Value

An integer.

Examples

```
n_training_examples(ts(1:10), h = 2, lags = 1:3, msas = "MIMO")
n_training_examples(ts(1:10), h = 2, lags = 1:3, msas = "recursive")
```

```
plot.knnForecastRO    Plot a prediction of a test set
```

Description

It uses a test set generated with the function [rolling_origin](#) and plots its forecast.

Usage

```
## S3 method for class 'knnForecastRO'
plot(x, h = NULL, ...)
```

Arguments

x	the object obtained from a call to rolling_origin .
h	an integer. The forecasting horizon. If NULL, the maximum forecasting horizon of all the test sets is used.
...	Other plotting parameters to affect the plot.

```
predict.knnForecast    Predict method for KNN models for time series forecasting.
```

Description

Predicted values based on a KNN model for time series forecasting.

Usage

```
## S3 method for class 'knnForecast'
predict(object, h, ...)
```

Arguments

object	a knnForecast object obtained by a call to the knn_forecasting function.
h	an integer. The forecasting horizon.
...	further arguments passed to or from other methods.

Details

If the models uses the MIMO strategy for multiple-step ahead prediction, the forecasting horizon is fixed to the model forecasting horizon.

Value

a `knnForecast` object with the prediction and information about the KNN model, see the documentation of `knn_forecasting` for the structure of `knnForecast` objects.

Examples

```
pred <- knn_forecasting(UKgas, h = 4, k = 1, msas = "recursive")
new_pred <- predict(pred, h = 6)
print(new_pred$prediction)
plot(new_pred) # To see a plot with the forecast
```

rolling_origin	<i>Assessing forecasting accuracy with rolling origin</i>
----------------	---

Description

It uses the model and the time series associated with the `knnForecast` object to asses the forecasting accuracy of the model using the last h values of the time series to build test sets applying a rolling origin evaluation.

Usage

```
rolling_origin(knnf, h = NULL, rolling = TRUE)
```

Arguments

<code>knnf</code>	A <code>knnForecast</code> object.
<code>h</code>	A positive integer. The forecast horizon. If <code>NULL</code> the prediction horizon of the <code>knnForecast</code> object is used.
<code>rolling</code>	A logical. If <code>TRUE</code> (the default), forecasting horizons from 1 to h are used. Otherwise, only horizon h is used.

Details

This function assesses the forecast accuracy of the model used by the `knnForecast` object. It uses h different test and training sets. The first test set consists of the last h values of the time series (the training set is formed by the previous values). The next test set consists of the last $h - 1$ values of the time series and so on (the last test set is formed by the last value of the time series).

Value

A list containing at least the following fields:

<code>test_sets</code>	a matrix containing the test sets used in the evaluation. Every row contains a different test set.
<code>predictions</code>	The predictions for the test sets.
<code>errors</code>	The errors for the test sets.
<code>global_accu</code>	Different measures of accuracy applied to all the errors.
<code>h_accu</code>	Different measures of accuracy applied to all the errors for every forecasting horizon.

Examples

```
pred <- knn_forecasting(UKgas, h = 4, lags = 1:4, k = 2)
ro <- rolling_origin(pred)
print(ro$global_accu)
```

Index

`autoplot.knnForecast`, 2

`knn_examples`, 3

`knn_forecasting`, 3, 6, 7

`n_training_examples`, 5

`nearest_neighbors`, 5

`plot.knnForecastR0`, 6

`predict.knnForecast`, 6

`rolling_origin`, 6, 7

`summary`, 4