

Package ‘tsintermittent’

May 8, 2026

Type Package

Title Intermittent Time Series Forecasting

Version 1.10

Date 2022-07-15

Description

Time series methods for intermittent demand forecasting. Includes Croston's method and its variants (Moving Average, SBA), and the TSB method. Users can obtain optimal parameters on a variety of loss functions, or use fixed ones (Kourentzes (2014) <[doi:10.1016/j.ijpe.2014.06.007](https://doi.org/10.1016/j.ijpe.2014.06.007)>). Intermittent time series classification methods and iMAPA that uses multiple temporal aggregation levels are also provided (Petropoulos & Kourentzes (2015) <[doi:10.1057/jors.2014.62](https://doi.org/10.1057/jors.2014.62)>).

LazyData yes

License GPL (>= 2)

Imports MAPA, parallel

URL <https://kourentzes.com/forecasting/2014/06/23/intermittent-demand-forecasting-package-for-r/>

NeedsCompilation no

Author Nikolaos Kourentzes [cre, aut] (ORCID:
<<https://orcid.org/0000-0003-0211-5218>>),
Fotios Petropoulos [ctb] (ORCID:
<<https://orcid.org/0000-0003-3039-4955>>)

Maintainer Nikolaos Kourentzes <nikolaos@kourentzes.com>

Repository CRAN

Date/Publication 2022-07-18 08:10:05 UTC

Contents

crost	2
crost.decomp	3
crost.ma	4
data.frc	6
idclass	7

imapa	8
sexsm	10
simID	11
ts.data1	12
ts.data2	12
tsb	13

Index	15
--------------	-----------

crost	<i>Croston's method and variants</i>
-------	--------------------------------------

Description

Croston's method and variants for intermittent demand series with fixed or optimised parameters.

Usage

```
crost(data,h=10,w=NULL,init=c("mean","naive"),nop=c(2,1),
      type=c("croston","sba","sbj"),cost=c("mar","msr","mae","mse"),
      init.opt=c(TRUE,FALSE),outplot=c(FALSE,TRUE),opt.on=c(FALSE,TRUE),
      na.rm=c(FALSE,TRUE))
```

Arguments

data	Intermittent demand time series.
h	Forecast horizon.
w	Smoothing parameters. If <code>w == NULL</code> then parameters are optimised. If <code>w</code> is a single parameter then the same is used for smoothing both the demand and the intervals. If two parameters are provided then the second is used to smooth the intervals.
init	Initial values for demand and intervals. This can be: 1. <code>c(z,x)</code> - Vector of two scalars, where first is initial demand and second is initial interval; 2. "naive" - Initial demand is first non-zero demand and initial interval is first interval; 3. "mean" - Same as "naive", but initial interval is the mean of all in sample intervals.
nop	Specifies the number of model parameters. Used only if they are optimised. 1. 1 - Demand and interval parameters are the same; 2. 2 - Different demand and interval parameters.
type	Croston's method variant: 1. "croston" Croston's method; 2. "sba" Syntetos-Boylan approximation; 3. "sbj" Shale-Boylan-Johnston.
cost	Cost function used for optimisation: 1. "mar" - Mean Absolute Rate; 2. "msr" - Mean Squared Rate; 3. "mae" - Mean Absolute Error; 4. "mse" - Mean Squared Error.
init.opt	If <code>init.opt==TRUE</code> then initial values are optimised.
outplot	If <code>TRUE</code> a plot of the forecast is provided.

opt.on	This is meant to use only by the optimisation function. When opt.on is TRUE then no checks on inputs are performed.
na.rm	A logical value indicating whether NA values should be remove using the method.

Value

model	Type of model fitted.
frc.in	In-sample demand rate.
frc.out	Out-of-sample demand rate.
weights	Smoothing parameters for demand and interval.
initial	Initialisation values for demand and interval smoothing.
component	List of c.in and c.out containing the non-zero demand and interval vectors for in- and out-of-sample respectively. Third element is the coefficient used to scale demand rate for sba and sbj.

Author(s)

Nikolaos Kourentzes

References

Optimisation of the methods described in: N. Kourentzes, 2014, On intermittent demand model optimisation and selection, International Journal of Production Economics, 156: 180-190. doi: [10.1016/j.ijpe.2014.06.007](https://doi.org/10.1016/j.ijpe.2014.06.007).

<https://kourentzes.com/forecasting/2014/06/11/on-intermittent-demand-model-optimisation-and-selection>

See Also

[tsb](#), [sexsm](#), [crost.ma](#).

Examples

```
crost(ts.data1, outplot=TRUE)
```

crost.decomp

Croston's decomposition

Description

Apply Croston's decomposition on a time series.

Usage

```
crost.decomp(data, init=c("naive", "mean"))
```

Arguments

data	Intermittent demand time series.
init	Initial values for intervals. This can be: 1. x - Numerical value; 2. "naive" - Initial interval is the first interval from start of time series; 3. "mean" - Initial interval is the mean of all in sample intervals.

Value

demand	Non-zero demand vector.
interval	Intervals vector.

Author(s)

Nikolaos Kourentzes

See Also

[crost](#), [crost.ma](#).

Examples

```
crost.decomp(ts.data1)
```

crost.ma

Moving average with Croston's method decomposition

Description

Moving average with Croston's method decomposition for intermittent demand series with fixed or optimised parameters.

Usage

```
crost.ma(data,h=10,w=NULL,nop=c(2,1),type=c("croston","sba","sbj"),
cost=c("mar","msr","mae","mse"),outplot=c(FALSE,TRUE),
na.rm=c(FALSE,TRUE))
```

Arguments

data	Intermittent demand time series.
h	Forecast horizon.
w	Moving average order. If w == NULL then moving average orders are optimised. If w is a single value then the same order is used for smoothing both the demand and the intervals. If two values are provided then the second is used to smooth the intervals.

nop	Specifies the number of model parameters. Used only if they are optimised. 1. 1 - Demand and interval moving average order are the same; 2. 2 - Different demand and interval orders.
type	Croston's method variant: 1. "croston" Croston's method; 2. "sba" Syntetos-Boylan approximation; 3. "sbj" Shale-Boylan-Johnston.
cost	Cost function used for optimisation: 1. "mar" - Mean Absolute Rate; 2. "msr" - Mean Squared Rate; 3. "mae" - Mean Absolute Error; 4. "mse" - Mean Squared Error.
outplot	If TRUE a plot of the forecast is provided.
na.rm	A logical value indicating whether NA values should be remove using the method.

Value

model	Type of model fitted.
frc.in	In-sample demand rate.
frc.out	Out-of-sample demand rate.
order	Moving averages orders for demand and interval.
component	List of c.in and c.out containing the non-zero demand and interval vectors for in- and out-of-sample respectively. Third element is the coefficient used to scale demand rate for sba and sbj.

Author(s)

Nikolaos Kourentzes

References

Optimisation cost functions described in: N. Kourentzes, 2014, On intermittent demand model optimisation and selection, International Journal of Production Economics, 156: 180-190. doi: [10.1016/j.ijpe.2014.06.007](https://doi.org/10.1016/j.ijpe.2014.06.007).

<https://kourentzes.com/forecasting/2014/06/11/on-intermittent-demand-model-optimisation-and-selection/>

See Also

[crost](#), [tsb](#), [sexsm](#).

Examples

```
crost.ma(ts.data1, outplot=TRUE)
```

data.frc *Wrapper to forecasts data.frames with a single call*

Description

Wrapper to forecasts data.frames with a single call.

Usage

```
data.frc(data.in,method=c("crost","crost.ma","tsb","sexsm","imapa","auto"),...)
```

Arguments

data.in	Data frame with time series. This can also be a matrix or array with each column being a different time series.
method	Which method to use for forecasting: "crost", "crost.ma", "tsb", "sexsm", "imapa", "auto". "auto" uses PKa classification to select automatically between Croston, SBA and SES.
...	Additional inputs to pass to forecasting functions. See individual function documentation for options.

Value

frc.out	Data frame containing forecasts for all time series.
out	List with detailed output per series. To access individual outputs of the list use: <code>sapply(out, get, x="element")</code> , where "element" could be for example "frc.in".

Author(s)

Nikolaos Kourentzes

References

By default methods are optimised using the cost functions introduced by: N. Kourentzes, 2014, On intermittent demand model optimisation and selection, International Journal of Production Economics, 156: 180-190. doi: [10.1016/j.ijpe.2014.06.007](https://doi.org/10.1016/j.ijpe.2014.06.007).

<https://kourentzes.com/forecasting/2014/06/11/on-intermittent-demand-model-optimisation-and-selection/>

The PK approximate classification is described in: F. Petropoulos and N. Kourentzes, 2015, Journal of Operational Research Society. <https://link.springer.com/article/10.1057/jors.2014.62>.

<https://kourentzes.com/forecasting/2014/05/13/forecast-combinations-for-intermittent-demand/>

See Also

[crost](#), [crost.ma](#), [tsb](#), [sexsm](#), [imapa](#), [idclass](#).

Examples

```
data.frc(simID(10,30),method="crost",type="sba",h=5)$frc.out
```

idclass *Time series categorisation for intermittent demand*

Description

Time series categorisation for intermittent demand data.

Usage

```
idclass(data,type=c("PKa","SBC","KHa","KH","PK"),a.in=NULL,
        outplot=c("summary","detail","none"),plot.focus=NULL)
```

Arguments

data	Time series dataset. Each column is a series. Alternatively this can be a single series.
type	Type of categorisation: 1. "SBC" - Syntetos Boylan Croston; 2. "KH" - Kostenko Hyndman (exact*); 3. "KHa" - Kostenko Hyndman (approximate); 4. "PK" - Petropoulos Kourentzes (exact*); 5. "PKa" - Petropoulos Kourentzes (approximate). *These are computationally expensive, as SBA is optimised for each time series.
a.in	Vector of SBA demand interval smoothing parameters. This must be same length as number of series. This is used for categorisations "KH" and "PK". If a.in == NULL then the parameters are calculated internally using MAR as a cost function.
outplot	Plot results of categorisation: 1. "summary" - simplified plot that reports number of series in each class and cut-off points; 2. "detail" - scatterplot between average interdemand interval (p) and squared coefficient of variation of non-zero demand (CV ²). Series that are categorised for SBA or SES are plotted in shaded areas; 3. "none" - do not produce plot.
plot.focus	Only relevant to outplot == "detail". Can be used to specify the maximum p and CV ² to plot, so that the scatterplot can be focused on the separation area between the categories. Use vector of two elements. First one is max p and second one is max CV ² . Example: plot.focus=c(1.5,1.5). If NULL then maximums are defined from the dataset.

Value

idx.croston	Index of series that are categorised under Croston.
idx.sba	Index of series that are categorised under SBA.
idx.ses	Index of series that are categorised under SES. Provided only for "PK" and "PKa" types.
cv2	Coefficient of variation squared of non-zero demand.
p	Inter-demand interval.
summary	Summary of number of series under each category.

Author(s)

Nikolaos Kourentzes

References

Classification schemes described in: F. Petropoulos and N. Kourentzes, 2015, Journal of Operational Research Society. <https://link.springer.com/article/10.1057/jors.2014.62>. <https://kourentzes.com/forecasting/2014/05/13/forecast-combinations-for-intermittent-demand/>

Optimisation of the methods described in: N. Kourentzes, 2014, On intermittent demand model optimisation and selection, International Journal of Production Economics, 156: 180-190. doi: [10.1016/j.ijpe.2014.06.007](https://doi.org/10.1016/j.ijpe.2014.06.007).

<https://kourentzes.com/forecasting/2014/06/11/on-intermittent-demand-model-optimisation-and-selection/>

See Also

[crost](#), [tsb](#), [imapa](#).

Examples

```
# Create/load some data. Each column is a time series
dataset <- simID(100,60,idi=1.15,cv2=0.3)
idclass(dataset)
```

imapa

MAPA for intermittent demand data

Description

MAPA for intermittent demand data with automatic model selection based on the PK classification.

Usage

```
imapa(data,h=10,w=NULL,minimumAL=1,maximumAL=NULL,comb=c("mean","median"),
      init.opt=c(TRUE,FALSE),para=c(0,1,2),outplot=c(0,1,2,3,4),model.fit=NULL,
      na.rm=c(FALSE,TRUE))
```

Arguments

data	Intermittent demand time series.
h	Forecast horizon.
w	Smoothing parameters. If w == NULL then parameters are optimised. If w is a single parameter then the same is used for smoothing both the demand and the intervals. If two parameters are provided then the second is used to smooth the intervals. SES is always optimised.
minimumAL	Lowest aggregation level to use. Default = 1.
maximumAL	Highest aggregation level to use. Default = maximum interval.

comb	Combination operator. One of "mean" or "median". Default is "mean".
init.opt	If init.opt==TRUE then Croston and SBA initial values are optimised.
paral	Use parallel processing. 0 = no; 1 = yes (requires initialised cluster); 2 = yes and initialise cluster. Default is 0.
outplot	Optional plot: 0 = No plot; 1 = Time series and combined forecast; 2 = Time series and all aggregation level forecasts; 3 = Summary model selection plot; 4 = Detailed model selection plot.
model.fit	Optional input with model types and parameters. This is the model.fit output from this function. If used it overrides other model settings.
na.rm	A logical value indicating whether NA values should be remove using the method.

Value

frc.in	In-sample demand rate.
frc.out	Out-of-sample demand rate.
summary	An array containing information for each aggregation level: AL - Aggregation level; n - Number of observations of aggregated series; p - Average inter-demand interval; cv2 - Coefficient of variation squared of non-zero demand; model - Selected model, where 1 is Croston, 2 is SBA and 3 is SES; use - If == 0 then this aggregation level is ignored because it contains less than 4 observations.
model.fit	Parameters and initialisation values of fitted model in each aggregation level.

Note

Note on optimal model paramaters: This implementation of MAPA for intermittent demand data optimises smoothing parameters for all Croston's method, SBA and SES. Optimisation is based on: N. Kourentzes, 2014, On intermittent demand model optimisation and selection, International Journal of Production Economics, 156: 180-190. doi: [10.1016/j.ijpe.2014.06.007](https://doi.org/10.1016/j.ijpe.2014.06.007).

<https://kourentzes.com/forecasting/2014/06/11/on-intermittent-demand-model-optimisation-and-selection/>

Note on parallelisation: Option paral=2 incurs substantial overheads. For a single time series using no parallelisation seems to be as good. If imapa is to be applied on multiple series, then initialising the parallel cluster externally and using the option paral=1 is advised.

Author(s)

Nikolaos Kourentzes

References

Optimisation of the methods described in: F. Petropoulos and N. Kourentzes, 2015, Forecast Combinations for Intermittent Demand, Journal of Operational Research Society. <https://link.springer.com/article/10.1057/jors.2014.62>.

<https://kourentzes.com/forecasting/2014/05/13/forecast-combinations-for-intermittent-demand/>

See Also

[tsb](#), [sexsm](#), [crost](#), [idclass](#).

Examples

```
imapa(ts.data1, outplot=1)
```

sexsm

*Simple exponential smoothing***Description**

Simple exponential smoothing with fixed or optimised parameters.

Usage

```
sexsm(data, h=10, w=NULL, init=c("mean", "naive"),
      cost=c("mar", "msr", "mae", "mse"), init.opt=c(TRUE, FALSE),
      outplot=c(FALSE, TRUE), opt.on=c(FALSE, TRUE),
      na.rm=c(FALSE, TRUE))
```

Arguments

data	Intermittent demand time series.
h	Forecast horizon.
w	Smoothing parameter. If w == NULL then parameter is optimised.
init	Initial values for demand and intervals. This can be: 1. x - Numeric value for the initial level; 2. "naive" - Initial value is a naive forecast; 3. "mean" - Initial value is equal to the average of data.
cost	Cost function used for optimisation: 1. "mar" - Mean Absolute Rate; 2. "msr" - Mean Squared Rate; 3. "mae" - Mean Absolute Error; 4. "mse" - Mean Squared Error.
init.opt	If init.opt==TRUE then initial values are optimised.
outplot	If TRUE a plot of the forecast is provided.
opt.on	This is meant to use only by the optimisation function. When opt.on is TRUE then no checks on inputs are performed.
na.rm	A logical value indicating whether NA values should be remove using the method.

Value

model	Type of model fitted.
frc.in	In-sample demand.
frc.out	Out-of-sample demand.
alpha	Smoothing parameter.
initial	Initialisation value.

Author(s)

Nikolaos Kourentzes

References

Optimisation of the method described in: N. Kourentzes, 2014, On intermittent demand model optimisation and selection, *International Journal of Production Economics*, 156: 180-190. doi: [10.1016/j.ijpe.2014.06.007](https://doi.org/10.1016/j.ijpe.2014.06.007).

<https://kourentzes.com/forecasting/2014/06/11/on-intermittent-demand-model-optimisation-and-selection/>

See Also

[crost](#), [tsb](#), [crost.ma](#).

Examples

```
sexsm(ts.data1, outplot=TRUE)
```

simID

Simulator for Intermittent Demand Series

Description

Simulator of Intermittent Demand Series.

Usage

```
simID(n=1, obs=60, idi=2, cv2=0.5, level=NULL)
```

Arguments

n	Number of time series to be generated.
obs	Number of observation of each series.
idi	Average intermittent demand interval of each series.
cv2	Squared coefficient of variation of the non-zero demands.
level	Mean level of the non-zero demands. If NULL, then a random level in [10,100] is selected.

Value

series A data matrix containing all the generated series.

Author(s)

Fotios Petropoulos

References

This simulator assumes that non-zero demand arrivals follow a bernoulli distribution and the non-zero demands a negative binomial distribution. Petropoulos F., Makridakis S., Assimakopoulos V. & Nikolopoulos K. (2014) "'Horses for Courses' in demand forecasting", European Journal of Operational Research, Vol. 237, No. 1, pp. 152-163

See Also

[crost](#), [tsb](#), [idclass](#).

Examples

```
dataset <- t(simID(100,60,idi=1.15,cv2=0.3))
```

ts.data1	<i>Example intermittent demand series - 'ts.data1'</i>
----------	--

Description

Example intermittent demand series - 'ts.data1'.

Usage

```
ts.data1
```

Format

Time series data

Examples

```
plot(ts.data1)
```

ts.data2	<i>Example intermittent demand series - 'ts.data2'</i>
----------	--

Description

Example intermittent demand series - 'ts.data2'.

Usage

```
ts.data2
```

Format

Time series data

Examples

```
plot(ts.data2)
```

tsb	<i>TSB (Teunter-Syntetos-Babai) method</i>
-----	--

Description

TSB intermittent demand method with fixed or optimised parameters.

Usage

```
tsb(data,h=10,w=NULL,init=c("mean","naive"),
     cost=c("mar","msr","mae","mse"),
     init.opt=c(TRUE,FALSE),outplot=c(FALSE,TRUE),
     opt.on=c(FALSE,TRUE),na.rm=c(FALSE,TRUE))
```

Arguments

data	Intermittent demand time series.
h	Forecast horizon.
w	Smoothing parameters. If w == NULL then parameters are optimised. Otherwise first parameter is for demand and second for demand probability.
init	Initial values for demand and intervals. This can be: 1. c(z,x) - Vector of two scalars, where first is initial demand and second is initial interval; 2. "naive" - Initial demand is first non-zero demand and initial demand probability is again the first one; 3. "mean" - Same as "naive", but initial demand probability is the mean of all in sample probabilities.
cost	Cost function used for optimisation: 1. "mar" - Mean Absolute Rate; 2. "msr" - Mean Squared Rate; 3. "mae" - Mean Absolute Error; 4. "mse" - Mean Squared Error.
init.opt	If init.opt==TRUE then initial values are optimised.
outplot	If TRUE a plot of the forecast is provided.
opt.on	This is meant to use only by the optimisation function. When opt.on is TRUE then no checks on inputs are performed.
na.rm	A logical value indicating whether NA values should be remove using the method.

Value

model	Type of model fitted.
frc.in	In-sample demand rate.
frc.out	Out-of-sample demand rate.
weights	Smoothing parameters for demand and demand probability.
initial	Initialisation values for demand and demand probability smoothing.

Author(s)

Nikolaos Kourentzes

References

Optimisation of the method described in: N. Kourentzes, 2014, On intermittent demand model optimisation and selection, International Journal of Production Economics, 156: 180-190. doi: [10.1016/j.ijpe.2014.06.007](https://doi.org/10.1016/j.ijpe.2014.06.007).

<https://kourentzes.com/forecasting/2014/06/11/on-intermittent-demand-model-optimisation-and-selection/>

See Also

[crost](#), [sexsm](#), [crost.ma](#).

Examples

```
tsb(ts.data1, outplot=TRUE)
```

Index

- * **~croston ~crost ~crost.decomp**
crost.decomp, 3
- * **~croston ~crost.ma ~moving average**
crost.ma, 4
- * **~croston ~crost**
crost, 2
- * **~data.frc**
data.frc, 6
- * **~idclass ~crost ~croston ~tsb**
idclass, 7
- * **~imapa ~crost ~tsb ~ses ~sexsm**
imapa, 8
- * **~sexsm**
sexsm, 10
- * **~simID**
simID, 11
- * **~tsb**
tsb, 13
- * **datasets**
ts.data1, 12
ts.data2, 12

crost, 2, 4–6, 8, 9, 11, 12, 14
crost.decomp, 3
crost.ma, 3, 4, 4, 6, 11, 14

data.frc, 6

idclass, 6, 7, 9, 12
imapa, 6, 8, 8

sexsm, 3, 5, 6, 9, 10, 14
simID, 11

ts.data1, 12
ts.data2, 12
tsb, 3, 5, 6, 8, 9, 11, 12, 13