

Package ‘hmix’

July 22, 2025

Type Package

Title Hidden Markov Model for Predicting Time Sequences with Mixture Sampling

Version 1.0.2

Maintainer Giancarlo Vercellino <giancarlo.vercellino@gmail.com>

Description An algorithm for time series analysis that leverages hidden Markov models, cluster analysis, and mixture distributions to segment data, detect patterns and predict future sequences.

License GPL-3

Encoding UTF-8

LazyData true

RoxygenNote 7.2.3

Imports normalp (>= 0.7.2), glogis (>= 1.0-2), gld (>= 2.6.6), edfun (>= 0.2.0), purrr (>= 1.0.1), HMM (>= 1.0.1), mc2d (>= 0.2.0), cubature (>= 2.1.0), dplyr (>= 1.1.2)

URL https://rpubs.com/giancarlo_vercellino/hmix

Suggests testthat (>= 3.0.0)

Config/testthat/edition 3

Depends R (>= 2.10)

NeedsCompilation no

Author Giancarlo Vercellino [aut, cre, cph]

Repository CRAN

Date/Publication 2024-09-10 09:50:12 UTC

Contents

dummy_set	2
hmix	2
Index	4

dummy_set	<i>A simple data set with stock close prices</i>
-----------	--

Description

A data frame with the close prices for AMZN, NVDA and IBM

Usage

```
dummy_set
```

Format

An object of class `data.frame` with 1925 rows and 4 columns.

Source

Yahoo Finance

hmix	<i>hmix: an algorithm for time series analysis that leverages hidden Markov models, cluster analysis, and mixture distributions to segment data, detect patterns and predict future sequences.</i>
------	--

Description

An algorithm for time series analysis that leverages hidden Markov models, cluster analysis, and mixture distributions to segment data, detect patterns and predict future sequences.

hmix function segments the time series with k-means clustering, fits an HMM to model state transitions, and generates future predictions over a specified horizon. It evaluates model accuracy by calculating the Continuous Ranked Probability Score (CRPS) across multiple test points, producing error metrics that assess the model's predictive performance and robustness.

Usage

```
hmix(
  ts,
  horizon,
  centers = 10,
  n_hidden = 4,
  seed = 42,
  n_tests = 20,
  warmup = 0.5
)
```

Arguments

<code>ts</code>	A numeric vector representing the time series data.
<code>horizon</code>	Integer. The prediction horizon, specifying how many future points to forecast.
<code>centers</code>	Integer. Number of clusters for k-means clustering. Default: 10.
<code>n_hidden</code>	Integer. Number of hidden states in the Hidden Markov Model. Default: 4.
<code>seed</code>	Integer. Random seed for reproducibility. Default: 42.
<code>n_tests</code>	Integer. Number of testing points for back-testing. Default: 20.
<code>warmup</code>	Numeric. Proportion of the time series used as the warm-up period before testing. Default: 0.5.

Value

This function returns a list containing:

- `model`: The HMM model along with its estimated parameters.
 - `hmm_model`: The object includes classified observations, initial HMM and trained HMM.
 - `pred_funs`: Prediction functions for each point in horizon (`rfun`, `dfun`, `pfun`, `qfun`)
- `error_sets`: A list of error metrics calculated for each testing point (at this time, CRPS).

Author(s)

Maintainer: Giancarlo Vercellino <giancarlo.vercellino@gmail.com> [copyright holder]

See Also

Useful links:

- https://rpubs.com/giancarlo_vercellino/hmix

Examples

```
# Example usage of hmix function:
result <- hmix(dummy_set$AMZN, horizon = 10, centers = 5, n_hidden = 3, n_tests = 2)
print(result$model)
print(result$error_sets)

# Random sampling for each point in predicted horizon
result$model$pred_funs$t1$rfun(10)

# ICDF for each point in horizon
result$model$pred_funs$t5$qfun(c(0, 1))

# PDF for each point in horizon
result$model$pred_funs$t8$dfun(tail(ts))

# CDF for each point in horizon
result$model$pred_funs$t10$pfun(tail(ts))
```

Index

* **datasets**

dummy_set, [2](#)

dummy_set, [2](#)

hmix, [2](#)

hmix-package (hmix), [2](#)