

Package: ModalForecast (via r-universe)

May 13, 2026

Title Parametric Modal ARIMA Models using the SKD Family

Version 0.1.0

Description Implements parametric modal ARIMA models utilizing the Skewed Distribution (SKD) family. Current distributions supported are the Skew-Normal, Skewed Student-t, and Skewed Laplace. The conditional mode is parameterized and optimized via maximum likelihood using analytical gradients. Includes comprehensive residual diagnostics, robustness options (heavy tails, asymmetry), robust parametric bootstrap prediction intervals, and classical asymptotic inference via the Fisher Information matrix.

URL <https://github.com/chedgala/ModalForecast>

BugReports <https://github.com/chedgala/ModalForecast/issues>

Depends R (>= 3.5.0)

License GPL-3 | file LICENSE

Encoding UTF-8

LazyData true

RoxygenNote 7.3.3

Imports stats, graphics, forecast, ggplot2, gridExtra, scales, grid

Suggests rmarkdown, testthat (>= 3.0.0), knitr

Config/testthat/edition 3

Repository <https://chedgala.r-universe.dev>

Date/Publication 2026-04-28 19:16:03 UTC

RemoteUrl <https://github.com/chedgala/modalforecast>

RemoteRef HEAD

RemoteSha debde16d4dc44b9359c325419fb94c985c335296

Contents

auto.modal.arima	2
diagnostics	3
envelope	5
fit_modal_arima	5
forecast.modal_arima	7

Index	9
--------------	----------

auto.modal.arima	<i>Automatic selection of Parametric Modal ARIMA model</i>
------------------	--

Description

Automatic selection of Parametric Modal ARIMA model

Usage

```
auto.modal.arima(
  y,
  d = NA,
  max.p = 5,
  max.q = 5,
  ic = c("aic", "bic"),
  dist = c("normal", "t", "laplace")
)
```

Arguments

y	numeric vector or time series of observations
d	Integer, degree of differencing. If NA, it's determined automatically.
max.p	Maximum AR order
max.q	Maximum MA order
ic	Information criterion to be used in model selection ("aic", "bic")
dist	Character string specifying the error distribution. "normal" (default) for Skew-Normal, "t" for Skewed Student-t, "laplace" for Skewed Laplace.

Details

This function performs a grid search over AR and MA orders to find the optimal Modal ARIMA model based on the selected information criterion (AIC or BIC).

Value

An object of class `modal_arima`.

Note

GitHub repository: <https://github.com/chedgala/ModalForecast>

References

Galarza, C. E., Lachos, V. H., Cabral, C. R. B., and Castro, L. M. (2017). Robust quantile regression using a generalized class of skewed distributions. *Stat*, 6(1), 113-130.

See Also

[fit_modal_arma](#)

Examples

```
library(forecast)

# 1. Load Empirical Data (Lynx)
data(lynx)
y <- log10(lynx)

# 2. Find the best SKD Error Distribution (Normal vs T vs Laplace)
fit_n <- fit_modal_arma(y, order = c(2, 0, 0), dist = "normal")
fit_t <- fit_modal_arma(y, order = c(2, 0, 0), dist = "t")
fit_l <- fit_modal_arma(y, order = c(2, 0, 0), dist = "laplace")
c(Normal = AIC(fit_n), Student = AIC(fit_t), Laplace = AIC(fit_l))

# 3. Auto Model Selection globally on the winning distribution (Skew-Normal)
fit_auto <- auto.modal.arma(y, d=0, max.p=5, max.q=5, dist="normal")

# 4. Summary & Inferences
summary(fit_auto)

# 5. Run residual diagnostics and Envelopes
diagnostics(fit_auto)
envelope(fit_auto, B=100)

# 6. Produce forecasts with multiple prediction bands (alphas)
pred <- forecast(fit_auto, h=10, level = c(80, 95, 99))

# 7. Native integration with 'forecast' ecosystem
autoplot(pred)
accuracy(pred)
```

Description

Provides visual and statistical diagnostics for the residuals of a fitted modal ARIMA model. Produces a comprehensive diagnostic panel including time series plot with fitted modes, ACF/PACF of residuals, QQ-plot for normality, histogram of residuals, and Ljung-Box p-values, all implemented using ggplot2.

Usage

```
diagnostics(object, ...)

## S3 method for class 'modal_arima'
diagnostics(object, ...)
```

Arguments

```
object      An object of class modal_arima.
...         Additional arguments (unused).
```

Value

A list of ggplot objects (invisibly) and draws the panel.

Examples

```
library(forecast)

# 1. Load Empirical Data (Lynx)
data(lynx)
y <- log10(lynx)

# 2. Find the best SKD Error Distribution (Normal vs T vs Laplace)
fit_n <- fit_modal_arima(y, order = c(2, 0, 0), dist = "normal")
fit_t <- fit_modal_arima(y, order = c(2, 0, 0), dist = "t")
fit_l <- fit_modal_arima(y, order = c(2, 0, 0), dist = "laplace")
c(Normal = AIC(fit_n), Student = AIC(fit_t), Laplace = AIC(fit_l))

# 3. Auto Model Selection globally on the winning distribution (Skew-Normal)
fit_auto <- auto.modal.arima(y, d=0, max.p=5, max.q=5, dist="normal")

# 4. Summary & Inferences
summary(fit_auto)

# 5. Run residual diagnostics and Envelopes
diagnostics(fit_auto)
envelope(fit_auto, B=100)

# 6. Produce forecasts with multiple prediction bands (alphas)
pred <- forecast(fit_auto, h=10, level = c(80, 95, 99))

# 7. Native integration with 'forecast' ecosystem
autoplot(pred)
```

```
accuracy(pred)
```

 envelope

Simulation Envelope Diagnostics for Modal ARIMA Models

Description

Constructs simulation envelopes based on the theoretical distance distributions from the SKD family (Galarza et al., 2017).

Usage

```
envelope(object, ...)

## S3 method for class 'modal_arima'
envelope(object, B = 100, ...)
```

Arguments

object	An object of class <code>modal_arima</code> .
...	Additional arguments (unused).
B	Number of Monte Carlo replications for envelope construction. Default is 100.

Value

A ggplot object (invisibly) and draws the envelope plot.

 fit_modal_arima

Fit a Parametric Modal ARIMA Model using the SKD Family

Description

Fits a Modal ARIMA model where the conditional mode follows an ARIMA recursion and the innovations follow a member of the SKD (Skewed Distribution) family. Supports the Skew-Normal, Skewed Student-t, and Skewed Laplace distributions.

Usage

```
fit_modal_arima(y, order = c(1, 0, 0), dist = c("normal", "t", "laplace"))
```

Arguments

y	numeric vector or time series of observations.
order	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.
dist	Character string specifying the error distribution from the SKD family. One of "normal" (default) for the Skew-Normal, "t" for the Skewed Student-t (adds degrees-of-freedom parameter nu), or "laplace" for the Skewed Laplace distribution.

Value

An object of class `modal_arma` containing:

y	The original time series.
order	The ARIMA order (p,d,q).
coefficients	Named vector of estimated parameters.
loglik	The maximized log-likelihood.
hessian	The Hessian matrix at the optimum.
convergence	Convergence code from <code>optim</code> .
dist	The distribution used ("normal", "t", or "laplace").

Note

GitHub repository: <https://github.com/chedgala/ModalForecast>

References

Galarza, C. E., Lachos, V. H., Cabral, C. R. B., and Castro, L. M. (2017). Robust quantile regression using a generalized class of skewed distributions. *Stat*, 6(1), 113-130.

See Also

[auto.modal.arma](#)

Examples

```
library(forecast)

# 1. Load Empirical Data (Lynx)
data(lynx)
y <- log10(lynx)

# 2. Find the best SKD Error Distribution (Normal vs T vs Laplace)
fit_n <- fit_modal_arma(y, order = c(2, 0, 0), dist = "normal")
fit_t <- fit_modal_arma(y, order = c(2, 0, 0), dist = "t")
fit_l <- fit_modal_arma(y, order = c(2, 0, 0), dist = "laplace")
c(Normal = AIC(fit_n), Student = AIC(fit_t), Laplace = AIC(fit_l))
```

```

# 3. Auto Model Selection globally on the winning distribution (Skew-Normal)
fit_auto <- auto.modal.arma(y, d=0, max.p=5, max.q=5, dist="normal")

# 4. Summary & Inferences
summary(fit_auto)

# 5. Run residual diagnostics and Envelopes
diagnostics(fit_auto)
envelope(fit_auto, B=100)

# 6. Produce forecasts with multiple prediction bands (alphas)
pred <- forecast(fit_auto, h=10, level = c(80, 95, 99))

# 7. Native integration with 'forecast' ecosystem
autoplot(pred)
accuracy(pred)

```

forecast.modal_arma *Forecast methodology for Modal ARIMA*

Description

Forecast methodology for Modal ARIMA

Usage

```

## S3 method for class 'modal_arma'
forecast(
  object,
  h = 10,
  level = c(80, 95),
  interval = c("asymptotic", "bootstrap"),
  npaths = 1000,
  ...
)

```

Arguments

object	A modal_arma object.
h	The forecast horizon.
level	Confidence level for prediction intervals.
interval	Method for computing prediction intervals ("asymptotic" or "bootstrap").
npaths	Number of simulated paths for bootstrap intervals. Defaults to 1000.
...	Additional arguments.

Value

A forecast object.

Examples

```
library(forecast)

# 1. Load Empirical Data (Lynx)
data(lynx)
y <- log10(lynx)

# 2. Find the best SKD Error Distribution (Normal vs T vs Laplace)
fit_n <- fit_modal_arma(y, order = c(2, 0, 0), dist = "normal")
fit_t <- fit_modal_arma(y, order = c(2, 0, 0), dist = "t")
fit_l <- fit_modal_arma(y, order = c(2, 0, 0), dist = "laplace")
c(Normal = AIC(fit_n), Student = AIC(fit_t), Laplace = AIC(fit_l))

# 3. Auto Model Selection globally on the winning distribution (Skew-Normal)
fit_auto <- auto.modal.arma(y, d=0, max.p=5, max.q=5, dist="normal")

# 4. Summary & Inferences
summary(fit_auto)

# 5. Run residual diagnostics and Envelopes
diagnostics(fit_auto)
envelope(fit_auto, B=100)

# 6. Produce forecasts with multiple prediction bands (alphas)
pred <- forecast(fit_auto, h=10, level = c(80, 95, 99))

# 7. Native integration with 'forecast' ecosystem
autoplot(pred)
accuracy(pred)
```

Index

`auto.modal.arima`, [2](#), [6](#)

`diagnostics`, [3](#)

`envelope`, [5](#)

`fit_modal_arima`, [3](#), [5](#)

`forecast.modal_arima`, [7](#)