Package ‘wired’

April 13, 2026

Type Package

Title Weighted Adaptive Prediction with Structured Dependence

Version 1.0.1

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

Description Builds a joint probabilistic forecast across series and horizons using adaptive copulas (Gaussian/t) with shrinkage-repaired correlations. At the low level it calls a probabilistic mixer per series and horizon, which backtests several simple predictors, predicts next-window Continuous Ranked Probability Score (CRPS), and converts those scores into softmax weights to form a calibrated mixture (r/q/p/dfun). The mixer blends eight simple predictors: a naive predictor that wraps the last move in a PERT distribution; an arima predictor using auto.arima for one-step forecasts; an Exponentially Weighted Moving Average (EWMA) gaussian predictor with mean/variance under a Gaussian; a historical bootstrap predictor that resamples past horizon-aligned moves; a drift residual bootstrap predictor combining linear trend with bootstrapped residuals; a volatility-scaled naive predictor centering on the last move and scaling by recent volatility; a robust median mad predictor using median/MAD with Laplace or Normal shape; and a shrunk quantile predictor that fits a few quantile regressions over time and interpolates to a full predictive. The function then couples the per-series mixtures on a common transform (additive/multiplicative/log-multiplicative), simulates coherent draws, and returns both transformed- and level-scale samplers and summaries.

License GPL-3

RoxygenNote 7.3.3

Imports mc2d (>= 0.2.1), forecast (>= 8.24.0), quantreg (>= 6.1), MASS (>= 7.3-65), imputeTS (>= 3.4)

Encoding UTF-8

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

Suggests knitr, testthat (>= 3.0.0)

Config/testthat/edition 3

Depends R (>= 4.1.0)

NeedsCompilation no

Author Giancarlo Vercellino [aut, cre, cph]

Repository CRAN

Date/Publication 2026-04-13 06:20:02 UTC

Contents

`wired` *wired: Weighted Adaptive Prediction with Structured Dependence*

Description

Builds a joint probabilistic forecast across series and horizons using adaptive copulas (Gaussian/t) with shrinkage-repaired correlations. At the low level it calls a probabilistic mixer per series and horizon, which backtests several simple predictors, predicts next-window Continuous Ranked Probability Score (CRPS), and converts those scores into softmax weights to form a calibrated mixture (r/q/p/dfun). The mixer blends eight simple predictors: a naive predictor that wraps the last move in a PERT distribution; an arima predictor using auto.arima for one-step forecasts; an Exponentially Weighted Moving Average (EWMA) gaussian predictor with mean/variance under a Gaussian; a historical bootstrap predictor that resamples past horizon-aligned moves; a drift residual bootstrap predictor combining linear trend with bootstrapped residuals; a volatility-scaled naive predictor centering on the last move and scaling by recent volatility; a robust median mad predictor using median/MAD with Laplace or Normal shape; and a shrunk quantile predictor that fits a few quantile regressions over time and interpolates to a full predictive. The function then couples the per-series mixtures on a common transform (additive/multiplicative/log-multiplicative), simulates coherent draws, and returns both transformed- and level-scale samplers and summaries.

Usage

```
wired(
  ts_set,
  future,
  dates = NULL,
  mode = c("additive", "multiplicative", "log_multiplicative"),
  n_testing = 30,
  dep_metric = c("kendall", "spearman", "pearson"),
  corr_adapt = c("static", "ewma", "rolling", "regime"),
  ewma_lambda = 0.15,
  roll_window = 60,
  shrink_alpha = 0.05,
  copula = c("gaussian", "t"),
  t_df = 7,
  stress_fun = c("mean_abs", "rms"),
  calm_q = 0.5,
  stress_q = 0.85,
  stress_smooth = 5,
  stress_blend_k = 8,
  seed = 123,
  u_eps = 1e-06,
```

```
  ...
)
```

Arguments

`ts_set` A matrix, or data frame of numeric time series.

`future` Integer scalar: forecast horizon used both for marginal models and for the dependence transform lag.

`dates` Vector of date values for the plot. Default: NULL.

`mode` Transformation to be applied to the time series: one of ‘"additive"‘, ‘"multiplicative"‘, ‘"log_multiplicative"‘.

`n_testing` Integer; number of expanding-window evaluation points. Default: 30.

`dep_metric` Dependence estimator for the correlation prototype: ‘"kendall"‘, ‘"spearman"‘ (rank-based; mapped to Gaussian/t correlation), or ‘"pearson"‘ (linear correlation).

`corr_adapt` Time-adaptation mode for correlation: - ‘"static"‘: single correlation from all aligned history, - ‘"ewma"‘: exponentially weighted correlation (fast-reacting), - ‘"rolling"‘: correlation from the last ‘roll_window‘ rows, - ‘"regime"‘: blend calm vs stress correlations using a stress score.

`ewma_lambda` Numeric in (0,1); higher values react faster in ‘"ewma"‘. Effective memory is about 1/lambda.

`roll_window` Integer; rolling window size for ‘"rolling"‘ and as a fallback in ‘"regime"‘. It is truncated to available rows if necessary.

`shrink_alpha` Numeric in (0,1); shrink correlation toward identity to stabilize inversion and PD repair.

`copula` Copula family: ‘"gaussian"‘ or ‘"t"‘. The t-copula introduces symmetric tail dependence controlled by ‘t_df‘.

`t_df` Degrees of freedom for the t-copula; must be > 2. Lower values increase tail dependence.

`stress_fun` Stress score used by ‘"regime"‘: ‘"mean_abs"‘ = mean absolute transformed return per row; ‘"rms"‘ = root-mean-square per row.

`calm_q, stress_q` Numeric quantiles in (0,1) with ‘calm_q < stress_q‘. Rows with stress lower than ‘calm_q‘ form the calm set; rows with stress greater than ‘stress_q‘ form the stress set. If either set is too small the method falls back to a rolling correlation.

`stress_smooth` Integer (greater than 1); length of a trailing moving average applied to the stress score to reduce noise.

`stress_blend_k` Positive scalar controlling logistic sharpness when blending calm/stress correlations at the latest stress value. Larger ‘k‘, sharper switching.

`seed` Integer RNG seed used both for copula draws and mixture components. For strict reproducibility across runs/platforms, keep packages and R versions fixed.

`u_eps` Small positive number used to clip uniform copula draws away from 0 and 1 to avoid quantile extremes or infinite transforms.

`...` Additional arguments forwarded to internal functions.

Value

A list with:

res_by_h Named list `h1..hH` (one per horizon) of per-horizon fits and helpers.

rfun_* Joint draw helpers: `rfun_trafo(n)` and `rfun_level(n)` return 3-D arrays $H \times n \times p$ (transformed vs level scale), and `rfun_both(n)` returns `list(trafo=..., level=...)` with the same shapes.

plot Recorded base R plot object.

meta Wrapper-level settings and controls (e.g., `future`, `mode`, `n_testing`, dependence/correlation and copula parameters, and regime-stress controls).

Author(s)

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

See Also

Useful links:

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

Examples

```
set.seed(1)
n <- 300
ts_set <- data.frame(
  A = 100 + cumsum(rnorm(n, 0, 1)),
  B =  80 + cumsum(rnorm(n, 0, 1))
)

fitH <- wired(
  ts_set    = ts_set,
  future    = 2,
  mode      = "additive",
  n_testing = 2,

  dep_metric = "spearman",
  corr_adapt = "rolling",
  roll_window = 40,
  copula = "gaussian",
  seed = 123,

  n_crps_mc = 30,
  q_grid_size = 10
)

draws_level <- fitH$rfun_level(5)
print(dim(draws_level))

both <- fitH$rfun_both(5)
```

Index