

Package ‘fable.bayesRecon’

May 28, 2026

Date 2026-05-05

Title Bayesian Reconciliation in the 'fable' Framework

Version 0.1.0

Description Implements the 'bayesRecon' probabilistic reconciliation methods within the 'fable' framework for hierarchical time series forecasting. Bayesian reconciliation (bayesRecon) methods are accessed via the 'reconcile' verb, following 'fable' conventions. For methodological background, see Corani et al. (2021) <[doi:10.1007/978-3-030-67664-3_13](https://doi.org/10.1007/978-3-030-67664-3_13)>, Zambon et al. (2024a) <[doi:10.1007/s11222-023-10343-y](https://doi.org/10.1007/s11222-023-10343-y)>, Zambon et al. (2024b) <<https://proceedings.mlr.press/v244/zambon24a.html>>, and Carrara et al. (2025) <[doi:10.48550/arXiv.2506.19554](https://doi.org/10.48550/arXiv.2506.19554)>.

License LGPL (>= 3)

URL <https://github.com/dazzimonti/fable.bayesRecon>

BugReports <https://github.com/dazzimonti/fable.bayesRecon/issues>

Encoding UTF-8

Imports stats, bayesRecon, fabletools (>= 0.7.0), purrr, vctrs, distributional, dplyr, tsibble, rlang

Suggests testthat (>= 3.0.0), cli, fable, tsibbledata, tibble, ggplot2, scales, knitr, rmarkdown

VignetteBuilder knitr

Config/testthat/edition 3

Depends R (>= 4.1)

Config/roxygen2/version 8.0.0

NeedsCompilation no

Author Dario Azzimonti [aut, cre, cph] (ORCID: <<https://orcid.org/0000-0001-5080-3061>>), Stefano Damato [aut] (ORCID: <<https://orcid.org/0009-0001-0225-3736>>), Lorenzo Zambon [aut] (ORCID: <<https://orcid.org/0000-0002-8939-993X>>), Chiara Carrara [aut] (ORCID: <<https://orcid.org/0009-0002-7349-6249>>), Giorgio Corani [aut] (ORCID: <<https://orcid.org/0000-0002-1541-8384>>)

Maintainer Dario Azzimonti <dario.azzimonti@gmail.com>

Repository CRAN

Date/Publication 2026-05-28 12:20:06 UTC

Contents

bayesRecon_BUIS	2
bayesRecon_MixCond	3
bayesRecon_t	5

Index	7
--------------	----------

bayesRecon_BUIS	<i>BUIS for probabilistic reconciliation of forecasts via conditioning</i>
-----------------	--

Description

Specifies Bottom-Up Importance Sampling (BUIS) reconciliation for use within `reconcile()`. The method uses the Bottom-Up Importance Sampling algorithm to draw samples from the reconciled forecast distribution, obtained via conditioning. Reconciliation is performed when `forecast()` is called on the resulting model. Marginal reconciled forecasts follow a sample distribution.

Usage

```
bayesRecon_BUIS(models, n_samples = 1000)
```

Arguments

<code>models</code>	A list of fitted models to reconcile.
<code>n_samples</code>	Number of samples to draw from the reconciled distribution.

Value

A model specification of class "bayesRecon_BUIS" (inheriting from "mdl_1st") to be passed to `reconcile`. The reconciliation is performed when `forecast` is called on the resulting mable (a model table). The forecast output is a fable object containing reconciled probabilistic forecasts represented as sample distributions (`dist_sample`).

References

Zambon, L., Azzimonti, D. & Corani, G. (2024). *Efficient probabilistic reconciliation of forecasts for real-valued and count time series*. *Statistics and Computing* 34 (1), 21. [doi:10.1007/s11222-02310343y](https://doi.org/10.1007/s11222-02310343y).

See Also

`fabletools::reconcile()`, `fabletools::aggregate_key()`, `bayesRecon_MixCond()`, `bayesRecon::reconc_BUIS()`

Examples

```
library(fable)
library(fabletools)
library(tsibble)
library(fable.bayesRecon)

# Small hierarchical example using tourism data
tourism_small <- tsibble::tourism |>
dplyr::filter(Region == "Melbourne") |>
fabletools::aggregate_key(Purpose, Trips = sum(Trips))

fit <- tourism_small |>
model(base = ETS(Trips ~ trend("A") + season("A"))) |>
reconcile(recon = bayesRecon_BUIS(base, n_samples = 300))

fc <- forecast(fit, h = "2 years")
fc
```

bayesRecon_MixCond *Probabilistic reconciliation of mixed hierarchies*

Description

`bayesRecon_MixCond` specifies Mixed-Conditioning (Mix-Cond) reconciliation for use within `reconcile()`. The method uses importance sampling to draw samples from the reconciled forecast distribution, obtained via conditioning, in the case of a mixed hierarchy.

`bayesRecon_TDcond` specifies Top-Down Conditioning reconciliation for use within `reconcile()`. The method uses a top-down conditioning algorithm: first, upper base forecasts are reconciled via conditioning using only the hierarchical constraints between the upper; then, the bottom distributions are updated via a probabilistic top-down procedure.

Reconciliation is performed when `forecast()` is called on the resulting model. Marginal reconciled forecasts follow a sample distribution.

Usage

```
bayesRecon_MixCond(models, n_samples = 1000, suppress_warnings = TRUE)
```

```
bayesRecon_TDcond(models, n_samples = 1000, suppress_warnings = TRUE)
```

Arguments

<code>models</code>	A list of fitted models to reconcile.
<code>n_samples</code>	Number of samples to draw from the reconciled distribution.
<code>suppress_warnings</code>	If TRUE, suppress warnings from reconciliation.

Value

A model specification of class "bayesRecon_MixCond" or "bayesRecon_TDcond" (inheriting from "mdl_1st") to be passed to `reconcile`. The reconciliation is performed when `forecast` is called on the resulting mable (a model table). The forecast output is a fable object containing reconciled probabilistic forecasts represented as sample distributions (`dist_sample`).

References

Zambon, L., Azzimonti, D., Rubattu, N., Corani, G. (2024). *Probabilistic reconciliation of mixed-type hierarchical time series*. Proceedings of the Fortieth Conference on Uncertainty in Artificial Intelligence, PMLR 244:4078-4095. <https://proceedings.mlr.press/v244/zambon24a.html>.

See Also

`fabletools::reconcile()`, `fabletools::aggregate_key()`, `bayesRecon_BUIS()`, `bayesRecon::reconc_MixCond()`, `bayesRecon::reconc_TDcond()`

Examples

```
library(tsibble)
library(dplyr)
library(tibble)
library(fable)
library(fabletools)

# Mixed hierarchy with integer-valued bottom series and one upper aggregate.
# Two low-rate Poisson count series at the bottom; their sum forms the
# continuous-looking upper aggregate.
set.seed(42)
n <- 60
idx <- tsibble::yearmonth("2019 Jan") + 0:(n - 1)

counts <- dplyr::bind_rows(
  tibble::tibble(Month = idx, Item = "A", Sales = rpois(n, lambda = 3)),
  tibble::tibble(Month = idx, Item = "B", Sales = rpois(n, lambda = 5)) |>
  tsibble::as_tsibble(index = Month, key = Item) |>
  fabletools::aggregate_key(Item, Sales = sum(Sales))

# Fit a base model on every level, then reconcile via MixCond:
# the upper (aggregated) forecast is treated as Gaussian, while the
# bottom-level forecast samples are treated as discrete via importance
# sampling.
fit <- counts |>
model(base = ETS(Sales)) |>
reconcile(mc = bayesRecon_MixCond(base))

# Alternative reconciliation via TDcond
fit_TDcond <- counts |>
model(base = ETS(Sales)) |>
```

```
reconcile(mc = bayesRecon_TDcond(base))
```

```
fc <- forecast(fit, h = 2)
fc
```

bayesRecon_t	<i>t-Rec: reconciliation via conditioning with uncertain covariance via multivariate t-distribution</i>
--------------	---

Description

Specifies t-Rec reconciliation for use within `reconcile()`.

Reconciles base forecasts by conditioning on the hierarchical constraints. The base forecasts are assumed to be jointly Gaussian, conditionally on the covariance matrix of the forecast errors. A Bayesian approach is adopted to account for the uncertainty of the covariance matrix. An Inverse-Wishart prior is specified on the covariance matrix, leading to a multivariate t-distribution for the base forecasts. The reconciliation via conditioning is in closed-form, yielding a multivariate t reconciled distribution.

Reconciliation is performed when `forecast()` is called on the resulting model. Marginal reconciled forecasts follow a Student-t distribution.

Usage

```
bayesRecon_t(models, ...)
```

Arguments

`models` A list of fitted models to reconcile.

`...` Additional arguments passed to other methods, including:

- `prior`: Optional list with entries `nu` and `Psi` specifying the parameters of the Inverse-Wishart prior distribution for the covariance matrix. If not provided, the prior is estimated from the data.
- `freq`: Optional frequency of the time series, used for estimating the naive covariance matrix via seasonal naive residuals. If not provided, the frequency is inferred from the data.
- `criterion`: Criterion for estimating the naive covariance matrix (default: "RSS").
- `l_shr`: Optional shrinkage parameter (between 0 and 1) for the covariance matrix of the residuals.

Value

A model specification of class "bayesRecon_t" (inheriting from "mdl_1st") to be passed to `reconcile`. The reconciliation is performed when `forecast` is called on the resulting mable (a model table). The forecast output is a fable object containing reconciled probabilistic forecasts represented as sample distributions (`dist_student_t`).

References

Carrara, C., Corani, G., Azzimonti, D., & Zambon, L. (2025). Modeling the uncertainty on the covariance matrix for probabilistic forecast reconciliation. arXiv preprint arXiv:2506.19554. <https://arxiv.org/abs/2506.19554>

See Also

`fabletools::reconcile()`, `fabletools::aggregate_key()`, `fabletools::min_trace()`, `bayesRecon::reconc_t()`

Examples

```
library(fable)
library(fabletools)
library(tsibble)
library(fable.bayesRecon)

# Small hierarchical example using tourism data
tourism_small <- tsibble::tourism |>
  dplyr::filter(Region == "Melbourne") |>
  fabletools::aggregate_key(Purpose, Trips = sum(Trips))

fit <- tourism_small |>
  model(base = ETS(Trips ~ trend("A") + season("A"))) |>
  reconcile(recon = bayesRecon_t(base))

fc <- forecast(fit, h = "2 years")
fc
```

Index

bayesRecon::reconc_BUIS(), 2
bayesRecon::reconc_MixCond(), 4
bayesRecon::reconc_t(), 6
bayesRecon::reconc_TDcond(), 4
bayesRecon_BUIS, 2
bayesRecon_BUIS(), 4
bayesRecon_MixCond, 3
bayesRecon_MixCond(), 2
bayesRecon_t, 5
bayesRecon_TDcond (bayesRecon_MixCond),
3

dist_sample, 2, 4
dist_student_t, 6

fabletools::aggregate_key(), 2, 4, 6
fabletools::min_trace(), 6
fabletools::reconcile(), 2, 4, 6
forecast, 2, 4, 6

reconcile, 2, 4, 6