Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani
Università degli Studi di Padova

FoReco: From foundations to frontiers

Daniele Girolimetto

Department of Statistical Sciences, University of Padova

danigiro.github.io
danigiro
daniele.girolimetto@unipd.it

Forecast reconciliation

Present and future of reconciliation packages in **R** (R Core Team, 2024)

- Post-forecasting process aimed to improve the quality of the base forecasts (however obtained) of a linearly constrained multiple time series by exploiting **cross-sectional** (e.g., spatial) and/or **temporal** constraints of the target forecasts

Target	Base forecasts		Reconciled forecasts
$\boldsymbol{C}\boldsymbol{y}_h = \boldsymbol{0}$	$\boldsymbol{C}\widehat{\boldsymbol{y}}_h \neq \boldsymbol{0}$	$\rightarrow$	$\boldsymbol{C}\widetilde{\boldsymbol{y}}_h = \boldsymbol{0}$

- **Forecasting examples**: Sales, Tourism, Energy demand, Healthcare, Supply chain ...
- **R packages**:

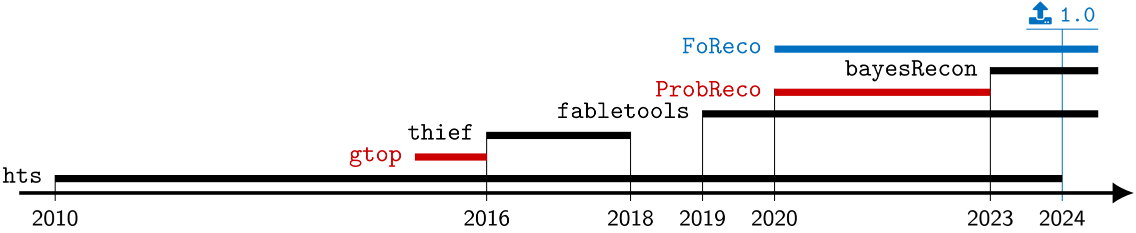

FoReco

- Robust set of tools implementing forecast reconciliation with a variety of approaches accounting for different constraints
- Additional resources and examples on GitHub and on the documentation page

Date	Version	Highlights
2020 - 10 - 01	`FoReco 0.1`	The journey begins! First release of FoReco
2021 - 05 - 21	`FoReco 0.2`	**Level-conditional forecast reconciliation**
2022 - 06 - 16	`FoReco 0.2.4`	Complete support to **multiple linearly constrained time series**, not just hierarchical/grouped → `lcmat()`
2023 - 05 - 16	`FoReco 0.2.6`	**Probabilistic forecast reconciliation**
2024 - 08 - 20	`FoReco 1.0`	Turning point! Package redesign with **unified notation**, **improved computation speed**, development of FoReco's **core engine**

Hierarchical, grouped and linearly constrained time series

Genuine hierarchical time series

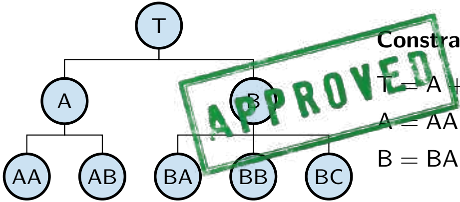

Constraints

$\mathrm{T} = \mathrm{A} + \mathrm{B}$

$\mathrm{A} = \mathrm{AA} + \mathrm{AB}$

$\mathrm{B} = \mathrm{BA} + \mathrm{BB} + \mathrm{BC}$

General linearly constrained time series

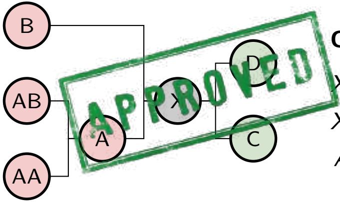

Constraints

$X = \mathrm{A} + \mathrm{B}$

$X = \mathrm{C} + \mathrm{D}$

$A = \mathrm{AA} + \mathrm{AB}$

Grouped time series: two or more genuine hierarchies sharing the same top and bottom variables

T	A	B
Z	AZ	BZ
Y	AY	BY

=

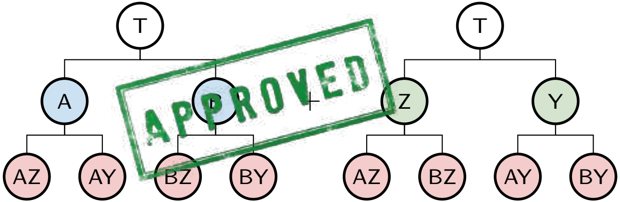

Constraints

$\mathrm{T} = \mathrm{A} + \mathrm{B}$
$\mathrm{A} = \mathrm{AZ} + \mathrm{AY}$
$\mathrm{B} = \mathrm{BZ} + \mathrm{BY}$

$\mathrm{T} = \mathrm{Z} + \mathrm{Y}$
$\mathrm{Z} = \mathrm{AZ} + \mathrm{BZ}$
$\mathrm{Y} = \mathrm{AY} + \mathrm{BY}$

Australian Tourism Demand

Today example: `data(vndata)`

Geographical division, g.d.

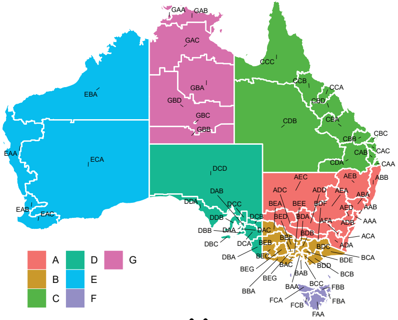

×

Purpose of travel, p.o.t.

Holiday, Visiting friends and relatives, Business, Other

- **Grouped ts** (geographical divisions × purpose of travel)

	AUS	States	Zones*	Regions	Tot
g.d.	1	7	21	76	105
p.o.t.	4	28	84	304	420
Tot	5	35	105	380	**525**

$n_a = 221$, $n_b = 304$, and $\mathbf{n} = \mathbf{525}$

- **Unique time series, no redundancy** (*6 Zones with only one Region are included in the Regions)
- **Temporal framework**, frequencies:
  - Monthly
  - Bi-Monthly
  - Quarterly
  - Four-Monthly
  - Semi-Annual
  - Annual

Cross-sectional framework

Hyndman et al. (2011); Panagiotelis et al. (2021); Girolimetto and Di Fonzo (2024b)

Zero-constrained

$$\boldsymbol{C}_{cs}\boldsymbol{y} = \boldsymbol{0}_{n_a}$$

$$\boldsymbol{C}_{cs} = \begin{bmatrix} \boldsymbol{I} & -\boldsymbol{A} \end{bmatrix}$$

not unique

$$\begin{bmatrix} 1 & -1 & -1 \end{bmatrix}$$

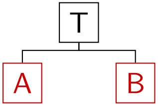

$\boldsymbol{A}$: $\boldsymbol{a} = \boldsymbol{A}\boldsymbol{b}$

Linear combination (or aggregation) matrix

$$\begin{bmatrix} 1 & 1 \end{bmatrix}$$

Structural

$$\boldsymbol{y} = \boldsymbol{S}_{cs}\boldsymbol{b}$$

$$\boldsymbol{S}_{cs} = \begin{bmatrix} \boldsymbol{A} \\ \boldsymbol{I}_{n_b} \end{bmatrix}$$

$$\begin{bmatrix} 1 \quad 1 \\ \boldsymbol{I}_2 \end{bmatrix}$$

General linearly constrained time series

- Genuine hierarchical time series
- Grouped time series

- Forecast reconciliation may be always expressed according to a **zero-constrained** representation
- A **structural-like** representation may be derived:

```
lcmat(cons_mat = C)
```

Optimal combination forecast reconciliation

Wickramasuriya et al. (2019); Panagiotelis et al. (2021)

1. Forecast **all series at all levels** of aggregation $\rightarrow$ base forecasts
2. Make the base forecasts **coherent** using least squares $\rightarrow$ reconciled forecasts

Two equivalent point forecast reconciliation formulae

Structural reconciliation approach

`approach = "strc"`

$$\widehat{\boldsymbol{y}} = \boldsymbol{S}\boldsymbol{\beta} + \boldsymbol{\varepsilon}$$

$$\Downarrow$$

$$\widetilde{\boldsymbol{y}} = \boldsymbol{S}\left(\boldsymbol{S}'\boldsymbol{W}^{-1}\boldsymbol{S}\right)^{-1}\boldsymbol{S}'\boldsymbol{W}^{-1}\widehat{\boldsymbol{y}} = \boldsymbol{S}\boldsymbol{G}\widehat{\boldsymbol{y}}$$

Projection reconciliation approach

`approach = "proj"`

$$\widehat{\boldsymbol{y}} = \boldsymbol{y} + \boldsymbol{\varepsilon}, \quad \text{s.t.} \quad \boldsymbol{C}\boldsymbol{y} = 0$$

$$\Downarrow$$

$$\widetilde{\boldsymbol{y}} = \left[\boldsymbol{I} - \boldsymbol{W}\boldsymbol{C}'\left(\boldsymbol{C}\boldsymbol{W}\boldsymbol{C}'\right)^{-1}\boldsymbol{C}\right]\widehat{\boldsymbol{y}} = \boldsymbol{M}\widehat{\boldsymbol{y}}$$

- The formulation of $\boldsymbol{W} = \mathrm{E}(\boldsymbol{\varepsilon}\boldsymbol{\varepsilon}')$ is conceptually complex; in practice, approximate forms are used, possibly using training/validation-set residuals

Projection vs structural approach: performance index

Cross-sectional wls reconciliation - *time* in seconds - median of 100 replications

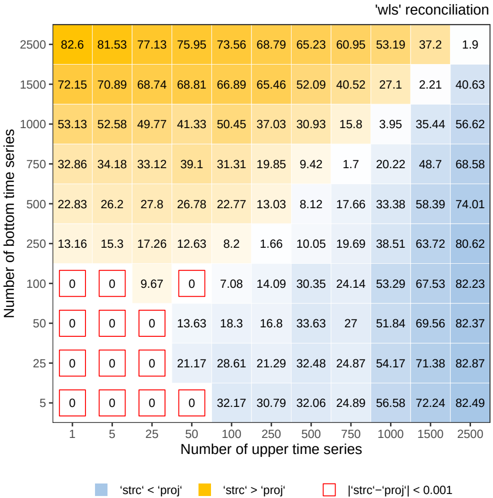

Two main factors:

- dimensions

  cs: n_a, n_b

  te: $\mathcal{K}$ (set of temporal aggregation orders)

  ct: n_a, n_b, $\mathcal{K}$
- computational cost of $\boldsymbol{W}^{-1}$

$$\mathbf{PI} = \left| \frac{time_{\mathrm{strc}} - time_{\mathrm{proj}}}{\max\left(time_{\mathrm{strc}}, time_{\mathrm{proj}}\right)} \right| 100$$

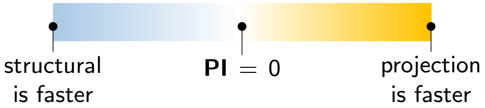

FoReco at work </>

Cross-sectional optimal forecast reconciliation

Input:

`base`	(12 × 525) monthly base forecasts
`comb`	A string specifying W_{cs} (e.g., `"ols"`, `"str"`, `"wls"`, `"shr"`, `"sam"`)
`res`	(228 × 525) in-sample residuals to compute the covariance matrix
`agg_mat`	(221 × 304) cross-sectional aggregation matrix A
`cons_mat`	(221 × 525) zero constraints cross-sectional matrix C

FoReco

```r
# Using the aggregation matrix A
rf_opt <- csrec(base = base, agg_mat = vnaggmat, approach = "proj", # (default)
                res = res, comb = "shr")
str(rf_opt, give.attr = FALSE)
#>  num [1:12, 1:525] 49160 21622 24815 29433 23260 ...

# Using the zero constraints matrix C: vnconsmat <- cbind(diag(221), -vnaggmat)
csrec(base = base, cons_mat = vnconsmat, res = res, comb = "shr")
```

Temporal framework

Athanasopoulos et al. (2017)

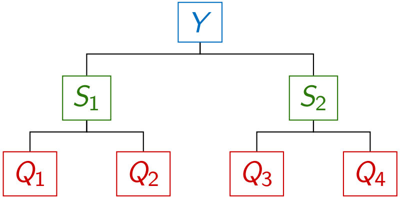

Quarterly hierarchy:
quarterly, semi-annual and annual series

Temporal hierarchy → non-overlapping aggregation of the observations of a time series (y) at regular intervals

$$x_j^{[k]} = \sum_{t=(j-1)k+1}^{jk} y_t$$

- Unlike cross-sectional hierarchies (n **variables at the same time index** are considered), in temporal hierarchies we have **one variable observed at different frequencies**
- Structural representation $\left(\boldsymbol{x}_\tau = \boldsymbol{S}_{te}\boldsymbol{x}_\tau^{[1]}\right)$ and zero-constrained representation $\left(\boldsymbol{C}_{te}\boldsymbol{x}_\tau = \boldsymbol{0}_{(k^*\times 1)}\right)$ still hold, and may be alternatively used for reconciliation

FoReco at work </>

Temporal optimal forecast reconciliation

Input:

`base`	(28×1) base forecasts ordered as $\left[x^{[12]}\ \boldsymbol{x}^{[6]}\ \dots\ \boldsymbol{x}^{[1]}\right]'$
`comb`	A string specifying $\boldsymbol{W}_{te}$ (e.g., `"ols"`, `"str"`, `"wlsv"`, ...)
`res`	(228×1) in-sample residuals to compute the covariance matrix
`agg_order`	max. order of temporal aggregation, $m = 12$ or a subset of temporal aggregation orders, e.g. `agg_order = c(12, 6, 1)`

FoReco

```r
rf_opt <- terec(base = base, agg_order = m, res = res, comb = "sar1")
str(rf_opt, give.attr = FALSE)
#>  Named num [1:28] 1.329 0.665 0.665 0.443 0.443 ...
```

Cross-sectional + Temporal = Cross-temporal

A cross-temporal hierarchy of three quarterly time series ($T = A + B$)

cross-sectional $\longrightarrow$ temporal

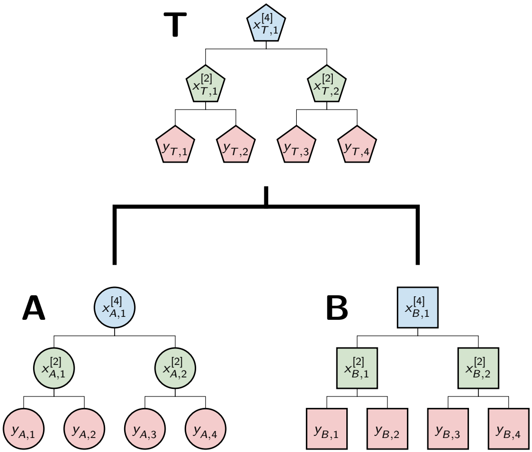

temporal $\longrightarrow$ cross-sectional

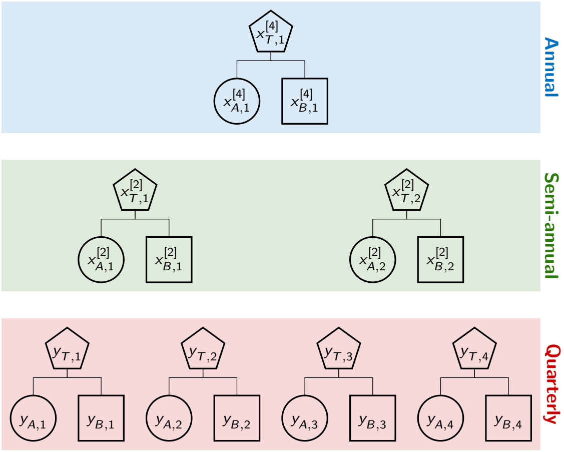

Annual: $j = 1$ **Semi-annual:** $j = 1, 2$ **Quarterly:** $j = 1, \ldots, 4$

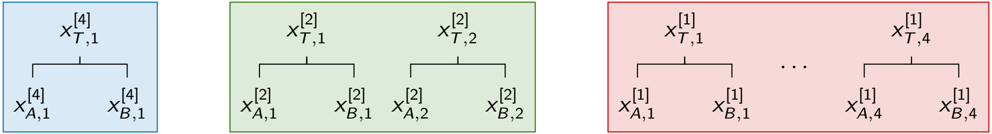

$$\left.\begin{array}{c} \boldsymbol{x}_{i,\tau} = \begin{bmatrix} x^{[4]}_{i,1} & x^{[2]}_{i,1} & x^{[2]}_{2} & x^{[1]}_{i,1} & \ldots & x^{[1]}_{i,4} \end{bmatrix}' \\ \text{for } i = T, A, B \end{array}\right\} \quad \Rightarrow \quad \boldsymbol{X}_\tau = \begin{bmatrix} \boldsymbol{x}'_{T,\tau} \\ \boldsymbol{x}'_{A,\tau} \\ \boldsymbol{x}'_{B,\tau} \end{bmatrix}$$

- Two dimensions to capture the complete nature of a multiple time series
- Any cross-temporal matrix may be constructed from the one-dimensional counterparts

  $\boldsymbol{C}_{ct} \rightarrow$ easy to compute as a function of $\boldsymbol{A}_{cs}/\boldsymbol{C}_{cs}$ and m

  $\boldsymbol{S}_{ct} \rightarrow$ fast to compute as $\boldsymbol{S}_{cs} \otimes \boldsymbol{S}_{te}$

FoReco at work </>

Cross-temporal optimal forecast reconciliation

Input:

`base`	(525×28) base forecasts matrix
`comb`	A string specifying $\boldsymbol{W}_{ct}$ (e.g., `"ols"`, `"wlsv"`, `"bdshr"`, ...)
`res`	(525×532) in-sample residuals to compute the covariance matrix
`agg_order`	max. order of temporal aggregation, $m = 12$
`agg_mat`	(221×304) cross-sectional aggregation matrix $\boldsymbol{A}$
`cons_mat`	(221×525) zero constraints cross-sectional matrix $\boldsymbol{C}$

FoReco

```r
rf_opt <- ctrec(base = base, agg_mat = vnaggmat, # or cons_mat = vnconsmat,
                agg_order = m, res = res, comb = "wlsv", approach = "strc")
str(rf_opt, give.attr=FALSE)
#>  num [1:525, 1:28] 314297 97383 62631 77793 19695 ...
```

FoReco at work </>

Two-step cross-temporal reconciliation, Kourentzes and Athanasopoulos (2019)

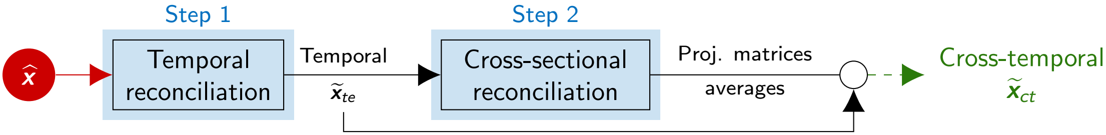

- The final temporally **and** cross-sectionally coherent reconciled forecasts are a transformation of the step 1 forecasts through the step 2 reconciliation matrices
- `tcsrec( )` → *first-temporal-then-cross-sectional reconciliation*
  `cstrec( )` → *first-cross-sectional-then-temporal reconciliation*

FoReco

```r
tcsrec(base = base, res = res,
       cslist = list(agg_mat = vnaggmat, comb = "shr"),
       telist = list(agg_order = m, comb = "wlsv"))
```

FoReco at work </>

Iterative cross-temporal reconciliation, Di Fonzo and Girolimetto (2023a)

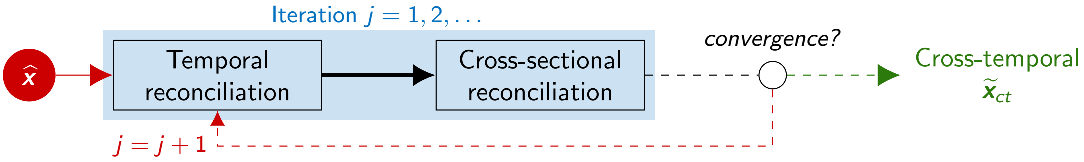

- Each iteration consists in the first two steps of the heuristic KA procedure, until a convergence criterion is met

FoReco

```r
iterec(base = base, res = res,
       type = "tcs", # → first-temporal-then-cross-sectional reconciliation
       # or type = "cst", → first-cross-sectional-then-temporal reconciliation
       cslist = list(agg_mat = vnaggmat, comb = "shr"),
       telist = list(agg_order = m, comb = "wlsv"))
```

Additional reconciliation approaches

Classical and LCC approaches for cross-sectional, temporal and cross-temporal frameworks

- **Classical reconciliation** (Dunn et al., 1976; Gross and Sohl, 1990; Athanasopoulos et al., 2009)

	Cross-sectional	*Temporal*	*Cross-Temporal*
Top-down, *`td( )`	`cstd( )`	`tetd( )`	`cttd( )`
Bottom-up, *`bu( )`	`csbu( )`	`tebu( )`	`ctbu( )`
Middle-out, *`mo( )`	`csmo( )`	`temo( )`	`ctmo( )`

- **Level conditional coherent reconciliation** (Hollyman et al., 2021; Di Fonzo and Girolimetto, 2024)

	Cross-sectional	*Temporal*	*Cross-Temporal*
LCC, *`lcc( )`	`cslcc( )`	`telcc( )`	`ctlcc( )`

Inside the engine room

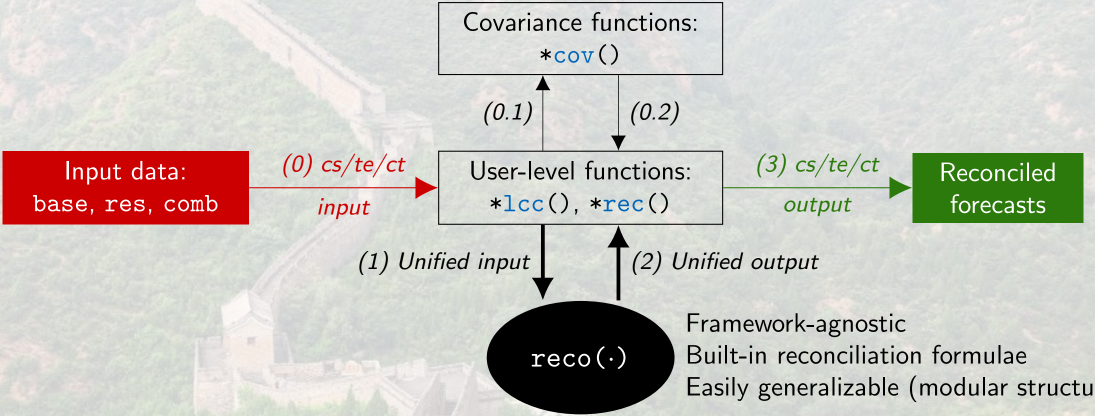

Framework-agnostic
Built-in reconciliation formulae
Easily generalizable (modular structure)

Practical challenges and other features

Cross-sectional, temporal and cross-temporal forecast reconciliation

- **Non-negative forecast reconciliation**

  `nn = "osqp"`: quadratic programming optimization (Stellato et al., 2020)

  `nn = "bpv"` : quadratic programming optimization (Wickramasuriya et al., 2020) NEW

  `nn = "sntz"`: heuristic "set-negative-to-zero" (Di Fonzo and Girolimetto, 2023b)
  Simple, quick and effective!

- **Probabilistic forecast reconciliation** (Panagiotelis et al., 2023; Girolimetto et al., 2024)
  - Non-parametric approach: `*boot( )` + `*rec( )`
  - Parametric approach (samples): `MASS::mvrnorm( )` + `*rec( )`

- Reconciliation with immutable forecasts (Zhang et al., 2023) using the `immutable` argument

Beyond FoReco

- github.com/daniGiro/FoReco
- danigiro.github.io/FoReco
- cran.r-project.org/package=FoReco

FoReco is a valuable tool for researchers and practitioners in the field of forecasting

Futures steps: a Python version is on the way! COMING SOON

Vision and next goals

- Organize (with simple input and output structure) a range of post-forecasting processes related to forecast combination and reconciliation
- Offer a flexible and comprehensive design that makes it a versatile solution for a wide range of forecasting applications (Sales, Energy, Healthcare, Supply chain, ...)
- Provide a foundational structure that can support other front-end-oriented packages

FoCo2*: Coherent Forecast Combination NEW

Girolimetto and Di Fonzo (2024a)

- **The problem**

Bates and Granger (1969): linear forecast combination (one variable and p experts) **+** Stone et al. (1942): constrained multivariate least-squares adjustment (n variables and one expert)

- **Our proposed solutions** (n variables and p experts)
  1. multi-task forecast combination → `csmtc( )`
  2. coherent forecast combination → `csscr( )`, `csscr( )`, and `csocc( )`

- **Links:**
  - `github.com/danigiro/FoCo2`
  - `danigiro.github.io/FoCo2`
  - `cran.r-project.org/package=FoCo2`

* *"S'i' fosse foco, arderei 'l mondo"* Sonetti (86), Cecco Angiolieri, Italian poet

Under development: Preview of future releases

nlcReco – Forecast reconciliation with non-linear constraints COMING SOON

nlcReco

$g(y)=0$

- Optimal combination forecast reconciliation procedures for multiple time series with linear and non-linear constraints in cross-sectional framework
- ISF2025 keynote "Forecast reconciliation: Geometry, optimisation, and insights beyond hierarchical time series" by Prof. Anastasios Panagiotelis
- Work in progress with Panagiotelis A., Li H., Di Fonzo T.

Under development: Preview of future releases

FoRecoML (?) – Forecast Reconciliation with Machine Learning COMING SOON

- Non-linear hierarchical forecast reconciliation approaches that produces cross-sectional/ temporal/cross-temporal reconciled forecasts in a direct and automated way through the use of popular machine learning methods
- Work in progress with Rombouts J., Wilms I., Yang Y.F.
- **References**:

  - Spiliotis, E., Abolghasemi, M., Hyndman, R. J., Petropoulos, F., & Assimakopoulos, V. (2021). Hierarchical forecast reconciliation with machine learning. Applied Soft Computing, 112, 107756
  - Cross-temporal framework: Rombouts, J., Ternes, M., & Wilms, I. (2024). Cross-temporal forecast reconciliation at digital platforms with machine learning. International Journal of Forecasting, 41(1), 321-344

Thank you!

Daniele Girolimetto

Department of Statistical Sciences, University of Padova

danigiro.github.io

danigiro

daniele.girolimetto@unipd.it

References I

Athanasopoulos, G., Ahmed, R. A., & Hyndman, R. J. (2009). Hierarchical forecasts for Australian domestic tourism. *International Journal of Forecasting*, *25*(1), 146–166.

Athanasopoulos, G., Hyndman, R. J., Kourentzes, N., & Petropoulos, F. (2017). Forecasting with temporal hierarchies. *European Journal of Operational Research*, *262*(1), 60–74.

Bates, J. & Granger, C. W. (1969). The combination of forecasts. *Operations Research Quarterly*, *20*(4), 451–468.

Di Fonzo, T. & Girolimetto, D. (2023a). Cross-temporal forecast reconciliation: Optimal combination method and heuristic alternatives. *International Journal of Forecasting*, *39*(1), 39–57.

Di Fonzo, T. & Girolimetto, D. (2023b). Spatio-temporal reconciliation of solar forecasts. *Solar Energy*, *251*, 13–29.

Di Fonzo, T. & Girolimetto, D. (2024). Forecast combination-based forecast reconciliation: Insights and extensions. *International Journal of Forecasting*, *40*(2), 490–514.

Dunn, D. M., Williams, W. H., & Dechaine, T. L. (1976). Aggregate versus Subaggregate Models in Local Area Forecasting. *Journal of the American Statistical Association*, *71*(353), 68–71.

References II

Girolimetto, D., Athanasopoulos, G., Di Fonzo, T., & Hyndman, R. J. (2024). Cross-temporal probabilistic forecast reconciliation: Methodological and practical issues. *International Journal of Forecasting*, *40*(3), 1134–1151.

Girolimetto, D. & Di Fonzo, T. (2024a). Coherent forecast combination for linearly constrained multiple time series. *arXiv*, (pp. 1–61).

Girolimetto, D. & Di Fonzo, T. (2024b). Point and probabilistic forecast reconciliation for general linearly constrained multiple time series. *Statistical Methods & Applications*, *33*, 581–607.

Gross, C. W. & Sohl, J. E. (1990). Disaggregation methods to expedite product line forecasting. *Journal of Forecasting*, *9*(3), 233–254.

Hollyman, R., Petropoulos, F., & Tipping, M. E. (2021). Understanding forecast reconciliation. *European Journal of Operational Research*, *294*(1), 149–160.

Hyndman, R. J., Ahmed, R. A., Athanasopoulos, G., & Shang, H. L. (2011). Optimal combination forecasts for hierarchical time series. *Computational Statistics & Data Analysis*, *55*(9), 2579–2589.

Kourentzes, N. & Athanasopoulos, G. (2019). Cross-temporal coherent forecasts for Australian tourism. *Annals of Tourism Research*, *75*, 393–409.

References III

Panagiotelis, A., Athanasopoulos, G., Gamakumara, P., & Hyndman, R. J. (2021). Forecast reconciliation: A geometric view with new insights on bias correction. *International Journal of Forecasting*, *37*(1), 343–359.

Panagiotelis, A., Gamakumara, P., Athanasopoulos, G., & Hyndman, R. J. (2023). Probabilistic forecast reconciliation: Properties, evaluation and score optimisation. *European Journal of Operational Research*, *306*(2), 693–706.

R Core Team (2024). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria `https://www.R-project.org/`.

Rombouts, J., Ternes, M., & Wilms, I. (2025). Cross-temporal forecast reconciliation at digital platforms with machine learning. *International journal of forecasting*, *41*(1).

Spiliotis, E., Abolghasemi, M., Hyndman, R. J., Petropoulos, F., & Assimakopoulos, V. (2021). Hierarchical forecast reconciliation with machine learning. *Applied soft computing*, *112*, 107756.

Stellato, B., Banjac, G., Goulart, P., Bemporad, A., & Boyd, S. (2020). OSQP: an operator splitting solver for quadratic programs. *Mathematical Programming Computation*, *12*(4), 637–672.

References IV

Stone, R., Champernowne, D. G., & Meade, J. E. (1942). The precision of national income estimates. *The Review of Economic Studies*, *9*(2), 111–125.

Wickramasuriya, S. L., Athanasopoulos, G., & Hyndman, R. J. (2019). Optimal Forecast Reconciliation for Hierarchical and Grouped Time Series Through Trace Minimization. *Journal of the American Statistical Association*, *114*(526), 804–819.

Wickramasuriya, S. L., Turlach, B. A., & Hyndman, R. J. (2020). Optimal non-negative forecast reconciliation. *Statistics and Computing*, *30*(5), 1167–1182.

Zhang, B., Kang, Y., Panagiotelis, A., & Li, F. (2023). Optimal reconciliation with immutable forecasts. *European Journal of Operational Research*, *308*(2), 650–660.