

Spatio-temporal reconciliation of solar forecasts

Tommaso Di Fonzo^{*}, Daniele Girolimetto

Department of Statistical Sciences, University of Padua, Via C. Battisti 241, 35121 Padova, Italy

ARTICLE INFO

Keywords:

Forecasting
Photovoltaic power generation
Spatio-temporal forecast reconciliation
Sequential and iterative approaches
Non-negative forecasts

ABSTRACT

In recent works by Yang et al. (2017a, 2017b), and Yang et al. (2019), geographical, temporal, and sequential deterministic reconciliation of hierarchical photovoltaic (PV) power generation have been considered for a simulated PV dataset in California. In the first two cases, the reconciliations were carried out in spatial and temporal domains separately. To further improve forecasting accuracy, in the third case these two reconciliation approaches were applied sequentially. During the replication of the forecasting experiment, some issues emerged about non-negativity and coherency (in space and/or in time) of the sequentially reconciled forecasts. Furthermore, while the accuracy improvement of the considered approaches over the benchmark persistence forecasts is clearly visible at any data granularity, we argue that an even better performance may be obtained by a thorough exploitation of spatio-temporal hierarchies. To this end, in this paper the spatio-temporal point forecast reconciliation approach is applied to generate non-negative, fully coherent (both in space and time) forecasts. New spatio-temporal reconciliation approaches are adopted, exploiting for the first time some relationships between two-step, iterative and simultaneous spatio-temporal reconciliation procedures. Non-negativity issues of the final reconciled forecasts are discussed and correctly dealt with in a simple and effective way. The spatio-temporal reconciliation procedures are applied to the base forecasts with forecast horizon of 1 day, of PV generated power at different time granularities (1 h to 1 day), of a geographical hierarchy consisting of 324 series along 3 levels. The normalized Root Mean Square Error (nRMSE) and the normalized Mean Bias Error are used to measure forecasting accuracy, and a statistical multiple comparison procedure is performed to rank the approaches. In addition to assuring full coherence and non-negativity of the reconciled forecasts, the results show that for the considered dataset, spatio-temporal forecast reconciliation significantly improves on the sequential procedures proposed by Yang et al. (2019), at any level of the spatial hierarchy and for any temporal granularity. For example, the forecasted hourly PV generated power by the new spatio-temporal forecast reconciliation approaches improve on the NWP 3TIER forecasts in a range from 4.7% to 18.4% in terms of nRMSE.

1. Introduction

Traditional electricity relies heavily on fossil fuels such as coal and natural gas. Not only are they bad for the environment, but they are also limited resources. Net-zero emissions by 2050 are crucial to achieve the core Paris Agreement¹ goals of a global average temperature rise of 1.5 degrees Celsius (United Nations, 2015), and this in turn can only be achieved if global greenhouse gas emissions are halved by the end of this decade (European Commission, 2019, United Nations, 2022). Solar power is one of the crucial production methods in the move to clean energy, and as economies of scale drive prices down, its importance will undoubtedly increase. The deployment of solar-power generation is causing total installed capacity to increase at a very high

pace. Eurostat reports that the European Union added 18,224.8 MW of net capacity in 2020, compared to its 16,146.9 MW increase in 2019, registering a growth of 12.9%. At the end of 2020, the EU's photovoltaic base stood at 136,136.6 MW, which is a 15% year-on-year increase (EurObserv'ER, 2022, p. 14).

Using solar resource as a stable source of energy is not an easy task. Estimating the solar energy potential is a key target to ensure its management in a reliable and efficient way for its integration into an electrical power grid (Sengupta et al., 2017). The prediction of solar irradiation despite its variability is particularly important as it is a precondition for (i) the management of the solar photovoltaic production through storage systems to reduce the impact of the intermittent nature

^{*} Corresponding author.

E-mail address: difonzo@stat.unipd.it (T. Di Fonzo).

¹ The Paris Agreement is a legally binding international treaty on climate change. Its goal is to limit global warming to well below 2, preferably to 1.5 degrees Celsius, compared to pre-industrial levels. To achieve this long-term temperature goal, countries aim to reach global peaking of greenhouse gas emissions as soon as possible to achieve a climate neutral world (i.e., with net-zero greenhouse gas emissions) by mid-century.

<https://doi.org/10.1016/j.solener.2023.01.003>

Received 16 September 2022; Received in revised form 29 December 2022; Accepted 2 January 2023

Available online 12 January 2023

0038-092X/© 2023 International Solar Energy Society. Published by Elsevier Ltd. All rights reserved.

of the solar resource, and (ii) the integration of the solar resource into a power grid in order to meet the local energy needs and to cope with the load fluctuations (Antonanzas et al., 2016). Understanding of the need for short, mid or long term prediction (e.g., 1 h, 6 h or a day ahead forecasting) is growing as utilities and grid operators gain experience in dealing with solar-power sources. Increasing spatial and temporal resolution of the available forecasting models (Benavides Cesar et al., 2022, and references therein) would enable grid operators to better forecast how much solar energy will be added to the grid. These efforts will improve the management of solar power's variability and uncertainty, enabling its more reliable and cost-effective integration onto the grid.

Solar forecasting is a fast-growing sub-domain of energy forecasting (Yang et al., 2020, p. 20). We agree with the claim of Yang et al. (2022), p. 7, that a “common misconception is that the novelty in solar forecasting should be solely revolved around forecasting methodology. Indeed, forecasting methodology is an important aspect, but it is never the only one”. Nevertheless, we think that a clear assessment of the available forecasting procedures may help in moving forward the frontier of knowledge of this fundamental topic, generating beneficial effects on the activities of practitioners.

A major goal for solar forecasting is to provide information on future photovoltaic (PV) power generation at different locations, time scales, and horizons to power system operators (Yang et al., 2022). In recent works by Yang et al. (2017a,b), and Yagli et al. (2019), geographical, temporal, and sequential deterministic reconciliation of hierarchical PV power generation have been considered for a simulated PV dataset in California. In the first two cases, the reconciliations were carried out in spatial and temporal domains separately. To further improve prediction accuracy, in the third case these two reconciliation approaches were sequentially applied. During the replication of the forecasting experiment,² some issues emerged about non-negativity and coherency (in space and/or in time) of the sequentially reconciled forecasts. Furthermore, while the accuracy improvement of the considered approaches over the benchmark persistence forecasts is clearly visible at any data granularity, we think that an even better performance may be obtained by a thorough exploitation of both geographical (i.e., cross-sectional) and temporal hierarchies.

At a given time point, a cross-sectional hierarchy describes the accounting relationships linking the series of different levels of contemporaneous aggregation (in the simplest case, the total variable is equal to the sum of its most disaggregated component series). A temporal hierarchy describes the temporal aggregated time series of a single variable originally defined at a specific time frequency, obtained by non-overlapping sums of the high-frequency observations (e.g., a daily time series may be obtained as the sum of the twenty four values in each single day of an hourly time series). The idea of exploiting the aggregation relationships valid both in space (cross-sectional coherency), and for different time granularities (temporal coherency) to improve the forecast accuracy of base forecasts for a hierarchical time series, was discussed by Kourentzes and Athanasopoulos (2019), Yagli et al. (2019), Spiliotis et al. (2020), and Punia et al. (2020). Di Fonzo and Girolimetto (2023) have shown the potentiality and feasibility of a number of spatio-temporal³ reconciliation approaches. More recently, Di Fonzo and Girolimetto (2022a) have established some useful relationships between sequential, iterative and simultaneous spatio-temporal reconciliation procedures, and proposed

² The forecasting experiment grounds on the documentation files and data made available by Yang et al. (2017b).

³ In the cited references, the more general term ‘cross-temporal’ is used to denote this particular forecast reconciliation approach. Due to the geographical connotation of the hierarchical time series considered in the empirical application (Section 3), in this paper the expression ‘spatio-temporal forecast reconciliation’ is adopted.

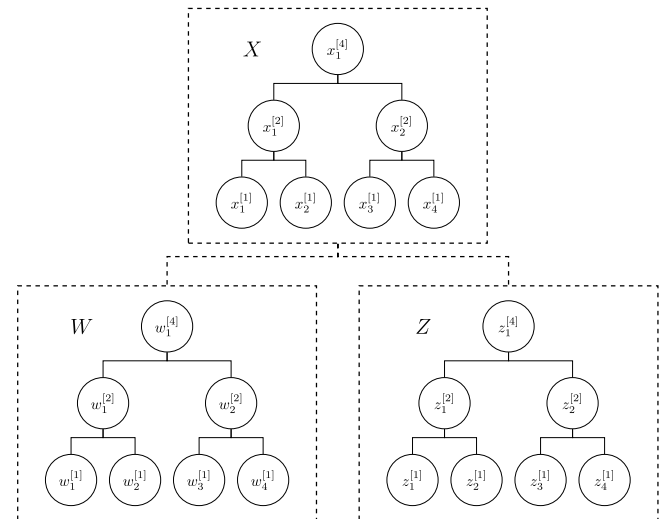


Fig. 1. Spatio-temporal hierarchy for a quarterly two-level spatial time series.

a simple and effective framework to ensure the non-negativity of the final reconciled forecasts, which turns out to be useful when dealing with naturally non-negative variables like PV power generation.

In this paper, spatio-temporal point forecast reconciliation is applied to generate non-negative, fully coherent (both in space and time) forecasts of PV generated power. In particular, the iterative and simultaneous approaches by Di Fonzo and Girolimetto (2023), and the heuristic procedure proposed by Kourentzes and Athanasopoulos (2019) are applied to a set of incoherent base forecasts with forecast horizon of 1 day, of PV generated power at different time granularities (1 h to 1 day), of a hierarchy comprising 324 series along 3 levels. The normalized Root Mean Square Error and the normalized Mean Bias Error are used to measure forecasting accuracy, and a statistical multiple comparison procedure is performed to rank the approaches.

The paper is organized as follows. The deterministic (point) spatio-temporal forecast reconciliation framework and some remarkable connections between apparently different approaches are described in Section 2. The forecasting experiment of Yagli et al. (2019) is replicated and discussed in Section 3, and the performance of the proposed forecasting approaches is presented in Section 4. Conclusions follow in Section 5. The on-line appendix contains supplementary tables and graphs related to the empirical application, and informations on the R scripts used for the forecasting experiments. All the forecast reconciliation procedures considered in this paper are available in the R package FoReco (Girolimetto and Di Fonzo, 2022).

2. Spatio-temporal point forecast reconciliation: a recap

To begin with, consider the very simple example of a two-level spatial hierarchy, where the top variable X is equal to the sum of two bottom series,⁴ W and Z . Further, assume that the highest time frequency the variables are observed at is quarterly, which means that semi-annual and annual time series may be obtained by simple non-overlapping temporal aggregation of quarterly time series. Fig. 1 gives a visual representation of such spatio-temporal hierarchy for a time cycle of 1 year.

The square boxes in the figure denote the nodes of a two-level spatial hierarchy, while the circles denote the nodes of the temporal

⁴ In this paper, we consider only genuine hierarchical/grouped time series, that share the same top- and bottom-level variables. The treatment of a general linearly constrained multiple time series is discussed in Di Fonzo and Girolimetto (2023).

hierarchies. Using a standard notation in the temporal forecast reconciliation literature (Athanasopoulos et al., 2017, Yang et al., 2017b), the superscript $[k]$ denotes the temporal aggregation order for each time granularity, i.e. annual ($k = 4$), semi-annual ($k = 2$), and quarterly ($k = 1$). The spatial hierarchy is described by the aggregation relationship $X = W + Z$, which is valid for any temporal aggregation order $k \in \mathcal{K} = \{4, 2, 1\}$ (i.e., $x_\tau^{[k]} = w_\tau^{[k]} + z_\tau^{[k]}$, $\tau = 1, \dots, 4/k$). Assuming v alternatively equal to x, w, z , the temporal hierarchies describing the relationships between different time granularities of a single time series may be expressed as

$$\begin{aligned}
 v_1^{[4]} &= v_1^{[2]} + v_2^{[2]} && \text{the annual value is the sum of the two semi-annual values,} \\
 v_1^{[2]} &= v_1^{[1]} + v_2^{[1]} && \text{the first half-year value is the sum of the first two quarters' values,} \\
 v_2^{[2]} &= v_3^{[1]} + v_4^{[1]} && \text{the second half-year value is the sum of the last two quarters' values,} \\
 \text{and thus} &&& \\
 v_1^{[4]} &= v_1^{[1]} + v_2^{[1]} + v_3^{[1]} + v_4^{[1]} && \text{the annual value is the sum of the four quarterly values.}
 \end{aligned}$$

All the relationships so far can be expressed in compact matrix form. Let $\mathbf{y}_\tau^{[k]} = [x_\tau^{[k]} \ w_\tau^{[k]} \ z_\tau^{[k]}]'$ be the (3×1) vector of the observations with temporal granularity $k \in \mathcal{K}$ of the variables in the spatial hierarchy at time $\tau = 1, \dots, 4/k$. The spatial aggregation relationships can be described as follows:

$$\begin{aligned}
 x_\tau^{[k]} &= \mathbf{C} \begin{bmatrix} w_\tau^{[k]} \\ z_\tau^{[k]} \end{bmatrix}, \quad \mathbf{y}_\tau^{[k]} = \mathbf{S} \begin{bmatrix} w_\tau^{[k]} \\ z_\tau^{[k]} \end{bmatrix}, \quad \mathbf{U}' \mathbf{y}_\tau^{[k]} = \mathbf{0}, \\
 \tau &= 1, \dots, \frac{4}{k}, \quad k \in \mathcal{K},
 \end{aligned} \tag{1}$$

where \mathbf{C} is the spatial aggregation matrix, \mathbf{S} is the spatial summing matrix, and \mathbf{U}' is the matrix expressing the spatial constraints in homogeneous form, respectively given by:

$$\mathbf{C} = \begin{bmatrix} 1 & 1 \\ \mathbf{I}_2 \end{bmatrix}, \quad \mathbf{S} = \begin{bmatrix} \mathbf{C} \\ \mathbf{I}_2 \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad \mathbf{U}' = [\mathbf{I}_1 \ -\mathbf{C}] = \begin{bmatrix} 1 & -1 & -1 \end{bmatrix}, \tag{2}$$

with \mathbf{I}_l denoting the identity matrix of order l . The complete temporal aggregation relationships linking the values of a single variable (say V) at different time granularities, may in turn be expressed through matrices

$$\begin{aligned}
 \mathbf{K} &= \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{bmatrix}, \quad \mathbf{R} = \begin{bmatrix} \mathbf{K} \\ \mathbf{I}_4 \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ & & & \mathbf{I}_4 \end{bmatrix}, \\
 \mathbf{Z}' &= [\mathbf{I}_3 \ -\mathbf{K}] = \begin{bmatrix} 1 & 0 & 0 & -1 & -1 & -1 & -1 \\ 0 & 1 & 0 & -1 & -1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & -1 & -1 \end{bmatrix},
 \end{aligned} \tag{3}$$

where \mathbf{K} is the temporal aggregation matrix, \mathbf{R} is the temporal summing matrix, and \mathbf{Z}' is the matrix expressing the temporal constraints in homogeneous form. It follows that:

$$\begin{bmatrix} v_1^{[4]} \\ v_1^{[2]} \\ v_2^{[2]} \end{bmatrix} = \mathbf{K} \mathbf{v}^{[1]}, \quad \mathbf{v} = \mathbf{R} \mathbf{v}^{[1]}, \quad \mathbf{Z}' \mathbf{v} = \mathbf{0}_{(3 \times 1)}, \tag{4}$$

where $\mathbf{v}^{[1]} = [v_1^{[1]} \ v_2^{[1]} \ v_3^{[1]} \ v_4^{[1]}]'$, and

$$\mathbf{v} = [v_1^{[4]} \ v_1^{[2]} \ v_2^{[2]} \ v_1^{[1]} \ v_2^{[1]} \ v_3^{[1]} \ v_4^{[1]}]'.$$

To simultaneously consider spatial and temporal aggregation relationships, all the nodes of the complete spatio-temporal hierarchy in Fig. 1 can be expressed in terms of the quarterly time series $w_\tau^{[1]}$ and

$z_\tau^{[1]}$, $\tau = 1, \dots, 4$, according to the structural representation:

$$\mathbf{y} = \mathbf{F} \mathbf{b}^{[1]}, \tag{5}$$

that is

$$\begin{bmatrix} x_1^{[4]} \\ x_1^{[2]} \\ x_2^{[2]} \\ x_1^{[1]} \\ x_2^{[1]} \\ x_3^{[1]} \\ x_4^{[1]} \\ w_1^{[4]} \\ w_1^{[2]} \\ w_2^{[2]} \\ w_1^{[1]} \\ w_2^{[1]} \\ w_3^{[1]} \\ w_4^{[1]} \\ z_1^{[4]} \\ z_1^{[2]} \\ z_2^{[2]} \\ z_1^{[1]} \\ z_2^{[1]} \\ z_3^{[1]} \\ z_4^{[1]} \end{bmatrix} = \underbrace{\begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}}_{\mathbf{F}} \underbrace{\begin{bmatrix} w_1^{[1]} \\ w_2^{[1]} \\ w_3^{[1]} \\ w_4^{[1]} \\ z_1^{[1]} \\ z_2^{[1]} \\ z_3^{[1]} \\ z_4^{[1]} \end{bmatrix}}_{\mathbf{b}^{[1]}}, \tag{6}$$

where $\mathbf{y} = [\mathbf{x}' \ \mathbf{w}' \ \mathbf{z}']'$ is the vector containing the data for all variables at any temporal granularity, $\mathbf{b}^{[1]} = [\mathbf{w}^{[1]'} \ \mathbf{z}^{[1]'}]'$ is the vector of the high-frequency bottom time series, and \mathbf{F} is the spatio-temporal summing matrix mapping $\mathbf{b}^{[1]}$ into \mathbf{y} . Expression (5) is the natural extension of the cross-sectional (spatial) structural representation firstly shown by Athanasopoulos et al. (2009). It relates the observations at the upper levels of both spatial and temporal hierarchies to the high-frequency bottom time series of the spatial hierarchy, which are the ‘very’ bottom time series in a spatio-temporal hierarchy (Di Fonzo and Girolimetto, 2023).

Besides the number of variables forming the spatial hierarchy ($n = 3$ in the above example), two crucial aspects affecting the dimension of matrix \mathbf{F} are (i) the temporal frequency of the highest-frequency granularity ($k = 1$), and (ii) the amount of temporal granularities taken into account in the temporal hierarchy. For example, if one is interested in coherently forecasted hourly time series within a day-cycle, the complete spatio-temporal summing matrix \mathbf{R} defining all intra-day temporal granularities ($\mathcal{K} = \{24, 12, 8, 6, 4, 3, 2, 1\}$) has dimension (60×24) , and is equal to

$$\mathbf{R} = [\mathbf{I}_{24} \ \mathbf{I}_2 \otimes \mathbf{I}_{12} \ \mathbf{I}_3 \otimes \mathbf{I}_8 \ \mathbf{I}_4 \otimes \mathbf{I}_6 \ \mathbf{I}_6 \otimes \mathbf{I}_4 \ \mathbf{I}_8 \otimes \mathbf{I}_3 \ \mathbf{I}_{12} \otimes \mathbf{I}_2 \ \mathbf{I}_{24}]'. \tag{7}$$

Matrix \mathbf{F} is thus a large and sparse matrix⁵ of dimension $(60n \times 24n_b)$, where n is the total number of series, and n_b is the number of bottom time series in the spatial hierarchy, respectively (in the above example, $n = 3$ and $n_b = 2$). Just to give an idea, the total number of variables in the dataset of hourly time series analyzed in this paper (see Section 3) is $n = 324$, with $n_b = 318$ bottom time series, thus matrix \mathbf{F} has dimension $(19,440 \times 7,632)$.⁶

In this framework, by extending the seminal idea by Hyndman et al. (2011), a forecast reconciliation problem arises when, for the nodes of a spatio-temporal hierarchy, a set of base forecasts – however obtained, and usually not aggregate consistent either in space and/or in time – are wished to be revised to fulfill the coherency relationships in space and time valid for the target data. The purpose is to improve the accuracy of the initial forecasts by combining forecasts at different aggregation levels in space and time, and by incorporating in the final forecasts the information given by spatial and temporal constraints.

2.1. Notation

Suppose we want to forecast a n -variate high-frequency hierarchical time series $\{\mathbf{y}_t^{[1]}\}_{t=1}^T$, with forecast horizon equal to the seasonal cycle m , (e.g., month per year, $m = 12$, quarter per year, $m = 4$, hour per day, $m = 24$), or a multiple thereof. Given a factor k of m , we may consider a number of temporally aggregated versions of each component of $\mathbf{y}_t^{[1]}$, given by the non-overlapping sums of k successive values, each having seasonal period equal to $M_k = m/k$. To avoid ragged-edge data, we assume that the total number of observations involved in the non-overlapping aggregation is a multiple of m , and define N the number of the lowest-frequency series observations, i.e. $N = T/m$. Let \mathcal{K} be the set of p factors of m , in descending order, $\mathcal{K} = \{k_p, k_{p-1}, \dots, k_2, k_1\}$, where $k_p = m$ and $k_1 = 1$, and define $k^* = \sum_{j=2}^p k_j$.

Following Di Fonzo and Girolimetto (2023), denote $\mathbf{Y}_{N+h} \equiv \mathbf{Y}$ the $[n \times (k^* + m)]$ matrix of the target forecasts for any temporal granularity, with low-frequency temporal horizon h , given by:

$$\mathbf{Y} = \begin{bmatrix} \mathbf{Y}^{[m]} & \mathbf{Y}^{[k_{p-1}]} & \dots & \mathbf{Y}^{[k_2]} & \mathbf{Y}^{[1]} \end{bmatrix} = \begin{bmatrix} \mathbf{A} \\ \mathbf{B} \end{bmatrix} = \begin{bmatrix} \mathbf{A}^{[m]} & \mathbf{A}^{[k_{p-1}]} & \dots & \mathbf{A}^{[k_2]} & \mathbf{A}^{[1]} \\ \mathbf{B}^{[m]} & \mathbf{B}^{[k_{p-1}]} & \dots & \mathbf{B}^{[k_2]} & \mathbf{B}^{[1]} \end{bmatrix}, \tag{8}$$

where m is the highest available sampling frequency per seasonal cycle (i.e., max. order of temporal aggregation). Each matrix $\mathbf{Y}^{[k]} = \begin{bmatrix} \mathbf{A}^{[k]} \\ \mathbf{B}^{[k]} \end{bmatrix}$, $k \in \mathcal{K}$, contains the order- k temporal aggregates of the n_a upper time series ($\mathbf{A}^{[k]}$), and of the n_b bottom time series ($\mathbf{B}^{[k]}$) in the spatial hierarchy, respectively, with $n = n_a + n_b$. Accordingly, we define the matrix of base forecasts $\hat{\mathbf{Y}}$ as:

$$\hat{\mathbf{Y}} = \begin{bmatrix} \hat{\mathbf{Y}}^{[m]} & \hat{\mathbf{Y}}^{[k_{p-1}]} & \dots & \hat{\mathbf{Y}}^{[k_2]} & \hat{\mathbf{Y}}^{[1]} \end{bmatrix} = \begin{bmatrix} \hat{\mathbf{A}}^{[m]} & \hat{\mathbf{A}}^{[k_{p-1}]} & \dots & \hat{\mathbf{A}}^{[k_2]} & \hat{\mathbf{A}}^{[1]} \\ \hat{\mathbf{B}}^{[m]} & \hat{\mathbf{B}}^{[k_{p-1}]} & \dots & \hat{\mathbf{B}}^{[k_2]} & \hat{\mathbf{B}}^{[1]} \end{bmatrix}. \tag{9}$$

While the target forecasts are expected to be aggregate-consistent both in time and space, the base forecasts are in general spatially and/or

⁵ Sparse matrices require less memory than dense matrices, and allow some computations to be more efficient (Paige and Sanders, 1982, Davis, 2006, Bates et al., 2022).

⁶ If the interest in forecasting at certain time granularities is low, this dimensionality issue may be mitigated by considering only part of the temporal granularities between the highest and lowest temporal frequencies. For example, if one considers only hourly and daily forecasts, the reduced \mathbf{F} is a $(8,100 \times 7,632)$ matrix, with a decrease of about 58% in the amount of matrix entries wrt its complete counterpart. Possible losses in the forecasting accuracy of the reconciled forecasts according to reduced temporal hierarchies should however be evaluated. This issue is currently under study.

temporally incoherent, that is:

$$\mathbf{U}'\mathbf{Y} = \mathbf{0}_{[n_a \times (k^* + m)]} \quad \text{and} \quad \mathbf{Z}'\mathbf{Y}' = \mathbf{0}_{[k^* \times n]},$$

while $\mathbf{U}'\hat{\mathbf{Y}} \neq \mathbf{0}_{[n_a \times (k^* + m)]}$ and/or $\mathbf{Z}'\hat{\mathbf{Y}}' \neq \mathbf{0}_{[k^* \times n]}$,

where $\mathbf{U}' = [\mathbf{I}_{n_a} - \mathbf{C}]$ and $\mathbf{Z}' = [\mathbf{I}_{k^*} - \mathbf{K}]$ are zero-constraints matrices associated to the spatial and temporal constraints, respectively.

2.2. Spatio-temporal bottom-up reconciliation

Bottom-up is an old and classic approach in the forecast reconciliation literature (Dunn et al., 1976, Dangerfield and Morris, 1992). This approach simply consists in obtaining the upper-level series' forecasts by summing-up the base forecasts of the bottom level series in the hierarchy. The spatio-temporal bottom-up (st(bu)) reconciliation of $k^* + m$ hierarchical time series' base forecasts at different time granularities, may thus be represented as follows:

$$\text{vec}(\tilde{\mathbf{Y}}'_{\text{st}(bu)}) = \text{Fvec}(\hat{\mathbf{B}}^{[1]}) \iff \tilde{\mathbf{y}}_{\text{st}(bu)} = \mathbf{F}\hat{\mathbf{b}}^{[1]}, \tag{10}$$

where $\mathbf{F} = (\mathbf{S} \otimes \mathbf{R})$ is the spatio-temporal summing matrix, with \otimes denoting the Kronecker product, and $\hat{\mathbf{b}}^{[1]} = \text{vec}(\hat{\mathbf{B}}^{[1]})$ is the vector containing the base forecasts of the high-frequency bottom time series. The spatio-temporal bottom-up reconciliation can be thought of as a two-step sequential reconciliation approach, where either spatial reconciliation of the high-frequency bottom time series base forecasts is followed by temporal reconciliation, or vice-versa (Di Fonzo and Girolimetto, 2022a). This observation opens the way to 'partly bottom-up' spatio-temporal reconciliation approaches, where forecasts of the n time series for different time granularities, and aggregation coherent only along a single dimension, are subsequently spatio-temporally reconciled via simple bottom-up according to the other dimension. We call these spatio-temporal forecast reconciliation approaches either st(rec_{te}, bu_{sp}), or st(rec_{sp}, bu_{te}), where 'rec_{te}' and 'rec_{sp}' denote a generic forecast reconciliation approach in time and in space, respectively.

2.3. Optimal combination spatio-temporal forecast reconciliation

Let us consider the multivariate regression model

$$\hat{\mathbf{Y}} = \mathbf{Y} + \mathbf{E}, \tag{11}$$

where the involved matrices have each dimension $[n \times (k^* + m)]$ and contain, respectively, the base ($\hat{\mathbf{Y}}$) and the target forecasts (\mathbf{Y}), and the coherency errors (\mathbf{E}) for the n component variables of the hierarchical time series of interest. Consider now the vectorized version of the model, that is

$$\text{vec}(\hat{\mathbf{Y}}') = \text{vec}(\mathbf{Y}') + \text{vec}(\mathbf{E}') \iff \hat{\mathbf{y}} = \mathbf{y} + \boldsymbol{\eta}, \tag{12}$$

where $\boldsymbol{\eta} = \text{vec}(\mathbf{E}')$ is the spatio-temporal reconciliation error with zero mean and p.d. covariance matrix Ω_{st} . Assuming Ω_{st} known, the optimal combination reconciled forecasts $\tilde{\mathbf{y}}_{\text{ost}} = \text{vec}(\tilde{\mathbf{Y}}'_{\text{ost}})$ are found by linearly constrained minimization of the generalized least squares (GLS) objective function

$$(\mathbf{y} - \hat{\mathbf{y}})' \Omega_{\text{st}}^{-1} (\mathbf{y} - \hat{\mathbf{y}}) \quad \text{s.t.} \quad \mathbf{H}'\mathbf{y} = \mathbf{0}, \tag{13}$$

where

$$\mathbf{H}'\mathbf{y} = \begin{bmatrix} \mathbf{U}^* \\ \mathbf{I}_n \otimes \mathbf{Z} \end{bmatrix} \mathbf{y} = \mathbf{0}_{[n(k^* + m) \times 1]} \tag{14}$$

is a full row-rank spatio-temporal zero-constraint matrix, with $\mathbf{U}^* = \begin{bmatrix} \mathbf{0}_{(n_a \times m \times n_k^*)} & \mathbf{I}_m \otimes \mathbf{U}' \end{bmatrix}$, and \mathbf{P} is the $[n(k^* + m) \times n(k^* + m)]$ commutation matrix, such that $\mathbf{P}\text{vec}(\mathbf{Y}) = \mathbf{y}$ (Di Fonzo and Girolimetto, 2023).

The projection approach solution (Byron, 1978; see also van Erven and Cugliari, 2015, Wickramasuriya et al., 2019, Panagiotelis et al.,

Table 1

Approximations for the spatial (Hyndman et al., 2011, Hyndman et al., 2016, Wickramasuriya et al., 2019, Di Fonzo and Girolimetto, 2023) and temporal (Athanasopoulos et al., 2017, Di Fonzo and Girolimetto, 2023) covariance matrix to be used in a reconciliation approach*.

	Spatial framework	Temporal framework
identity	sp(ols): $\mathbf{W} = \mathbf{I}_n$	te(ols): $\Omega = \mathbf{I}_{k^*+m}$
structural	sp(struc): $\mathbf{W} = \text{diag}(\mathbf{S}\mathbf{I}_{mn_b})$	te(struc): $\Omega = \text{diag}(\mathbf{R}\mathbf{I}_m)$
series variance	sp(wlsv): $\mathbf{W} = \widehat{\mathbf{W}}_D = \mathbf{I}_n \odot \widehat{\mathbf{W}}$	te(wlsv): $\Omega = \widehat{\Omega}_{wlsv}$
MinT-shr	sp(shr): $\mathbf{W} = \hat{\lambda}\widehat{\mathbf{W}}_D + (1 - \hat{\lambda})\widehat{\mathbf{W}}$	te(shr): $\Omega = \hat{\lambda}\widehat{\Omega}_D + (1 - \hat{\lambda})\widehat{\Omega}$
MinT-sam	sp(sam): $\mathbf{W} = \widehat{\mathbf{W}}$	te(sam): $\Omega = \widehat{\Omega}$

* $\widehat{\mathbf{W}}$ ($\widehat{\Omega}$) is the covariance matrix of the spatial (temporal) one-step ahead in-sample forecast errors, $\widehat{\Omega}_{wlsv}$ is a diagonal matrix “which contains estimates of the in-sample one-step-ahead error variances across each level” (Athanasopoulos et al., 2017, p. 64), and $\widehat{\Omega}_D = \mathbf{I}_{k^*+m} \odot \widehat{\Omega}$.

2021, Di Fonzo and Girolimetto, 2023) is given by⁷

$$\tilde{\mathbf{y}}_{\text{ost}} = \left[\mathbf{I}_{n(k^*+m)} - \Omega_{\text{st}} \mathbf{H} (\mathbf{H}' \Omega_{\text{st}} \mathbf{H})^{-1} \mathbf{H}' \right] \hat{\mathbf{y}} = \mathbf{M}_{\text{st}} \hat{\mathbf{y}}. \quad (15)$$

Di Fonzo and Girolimetto (2023) considered the following approximations for the spatio-temporal covariance matrix (‘ost’ stands for ‘optimal spatio-temporal’):

- ost(ols) - identity: $\Omega_{\text{st}} = \mathbf{I}_{n(k^*+m)}$
- ost(struc) - structural: $\Omega_{\text{st}} = \text{diag}(\mathbf{F}\mathbf{I}_{mn_b})$
- ost(wlsv) - series variance scaling: $\Omega_{\text{st}} = \widehat{\Omega}_{\text{st},wlsv}$, that is a straightforward extension of the series variance scaling matrix presented by Athanasopoulos et al. (2017) in the temporal framework
- ost(bdshr) - block-diagonal shrunk cross-covariance scaling: $\Omega_{\text{st}} = \widehat{\mathbf{P}} \widehat{\mathbf{W}}_{\text{st},shr}^{BD} \widehat{\mathbf{P}}'$
- ost(bdsam) - block-diagonal cross-covariance scaling: $\Omega_{\text{st}} = \widehat{\mathbf{P}} \widehat{\mathbf{W}}_{\text{st},sam}^{BD} \widehat{\mathbf{P}}'$
- ost(shr) - MinT-shr: $\Omega_{\text{st}} = \hat{\lambda}\widehat{\Omega}_{\text{st},D} + (1 - \hat{\lambda})\widehat{\Omega}_{\text{st}}$
- ost(sam) - MinT-sam: $\Omega_{\text{st}} = \widehat{\Omega}_{\text{st}}$

where the symbol \odot denotes the Hadamard product, $\hat{\lambda}$ is an estimated shrinkage coefficient (Ledoit and Wolf, 2004), $\widehat{\Omega}_{\text{st},D} = \mathbf{I}_{n(k^*+m)} \odot \widehat{\Omega}_{\text{st}}$, and $\widehat{\Omega}_{\text{st}}$ is the covariance matrix of the spatio-temporal one-step ahead in-sample forecast errors. The spatial point forecast reconciliation formula is obtained by assuming $m = 1$ (then $k^* = 0$, and $\Omega_{\text{st}} = \mathbf{W}$ is an $(n \times n)$ p.d. matrix):

$$\tilde{\mathbf{y}} = \left[\mathbf{I}_n - \mathbf{W}\mathbf{U} (\mathbf{U}'\mathbf{W}\mathbf{U})^{-1} \mathbf{U}' \right] \hat{\mathbf{y}} = \mathbf{M}_{\text{sp}} \hat{\mathbf{y}}. \quad (16)$$

The reconciled forecasts through temporal hierarchies for a single time series (Athanasopoulos et al., 2017) are in turn obtained by setting $n = 1$ (i.e., $n_a = 0$ and $n_b = 1$; $\Omega_{\text{st}} = \Omega$ is a $(k^* + m \times k^* + m)$ p.d. matrix):

$$\tilde{\mathbf{y}} = \left[\mathbf{I}_{(k^*+m)} - \Omega \mathbf{Z} (\mathbf{Z}'\Omega \mathbf{Z})^{-1} \mathbf{Z}' \right] \hat{\mathbf{y}} = \mathbf{M}_{\text{te}} \hat{\mathbf{y}}. \quad (17)$$

Table 1 presents some approximations for the spatial and the temporal covariance matrices. Other alternatives for temporal reconciliation, exploiting possible information in the residuals’ autocorrelation, can be found in Nystrup et al. (2020) and Di Fonzo and Girolimetto (2023).

2.4. Heuristic and iterative spatio-temporal reconciliation

Kourentzes and Athanasopoulos (2019) proposed an ensemble forecasting procedure (denoted KA), that exploits the simple averaging of

different forecasts. It consists in the following steps (for further details, see Di Fonzo and Girolimetto, 2023):

KA-Step 1 compute the temporally reconciled forecasts for each variable $i \in \{1, \dots, n\}$, and arrange them in the $[n \times (k + m)]$ matrix $\tilde{\mathbf{Y}}_{\text{te}}$;

KA-Step 2 starting from $\tilde{\mathbf{Y}}_{\text{te}}$, compute the time-by-time spatially reconciled forecasts for all the temporal aggregation levels ($\tilde{\mathbf{Y}}_{\text{sp}}$), and collect all the $(n \times n)$ projection matrices used to reconcile forecasts of k -level temporally aggregated time series, $\mathbf{M}_{\text{sp}}^{[k]}$, $k \in \mathcal{K}$;

KA-Step 3 transform the step 1 forecasts once more, by computing time-by-time spatially reconciled forecasts for all temporal aggregation levels using the $(n \times n)$ matrix $\bar{\mathbf{M}}$, given by the average of the matrices $\mathbf{M}_{\text{sp}}^{[k]}$:

$$\tilde{\mathbf{Y}}_{\text{KA}} = \left(\frac{1}{p} \sum_{k \in \mathcal{K}} \mathbf{M}_{\text{sp}}^{[k]} \right) \tilde{\mathbf{Y}}_{\text{te}} = \bar{\mathbf{M}} \tilde{\mathbf{Y}}_{\text{te}}. \quad (18)$$

Di Fonzo and Girolimetto (2023) presented an iterative approach, that produces spatio-temporal reconciled forecasts by alternating forecast reconciliation along one dimension (spatial or temporal), based on the first two steps of the KA approach. The iteration $j \geq 1$ can be described as follows:

Step 1 compute the temporally reconciled forecasts ($\tilde{\mathbf{Y}}_{\text{te}}^{(j)}$) for each variable $i \in \{1, \dots, n\}$ of $\tilde{\mathbf{Y}}_{\text{sp}}^{(j-1)}$;

Step 2 compute the time-by-time spatially reconciled forecasts ($\tilde{\mathbf{Y}}_{\text{sp}}^{(j)}$) for all the temporal aggregation levels of $\tilde{\mathbf{Y}}_{\text{te}}^{(j)}$.

At $j = 0$, the starting values are given by $\tilde{\mathbf{Y}}_{\text{sp}}^{(0)} = \hat{\mathbf{Y}}$, and the iterates end when the entries of matrix $\mathbf{D}_{\text{te}} = \mathbf{Z}' \tilde{\mathbf{Y}}_{\text{sp}}^{(j)'} - \tilde{\mathbf{Y}}_{\text{te}}^{(j)}$, containing all the temporal discrepancies, are small enough according to a suitable convergence criterion (Di Fonzo and Girolimetto, 2023). In the description above, *temporal-then-spatial* reconciliation is iteratively performed (ite(rec_{te}, rec_{sp})), otherwise the order may be reversed, thus generating *spatial-then-temporal* reconciliation (ite(rec_{sp}, rec_{te})).

3. Replication and assessment of the forecasting experiment of Yagli et al. (2019)

The dataset used in this study, called PV324, is the same used by Yang et al. (2017a,b), and Yagli et al. (2019). It refers to 318 simulated PV plants in California, whose hourly irradiation data are organized in three levels (Fig. 2):

- \mathcal{L}_0 : 1 time series for the Independent System Operator (ISO), given by the sum of the 318 plant series;
- \mathcal{L}_1 : 5 time series for the Transmission Zones (TZ), each given by the sum of 27, 73, 101, 86, and 31 plants, respectively;
- \mathcal{L}_2 : 318 bottom time series at plant level (P).

Following Yang et al. (2017b) and Yagli et al. (2019), we perform a forecasting experiment with fixed length window of 14 days (i.e., 336 h), forecast horizon of two days, and forecasting evaluation taking into account only the day-2 forecasts. These settings are coherent with the forecast operational submission requirements of CAISO, the public corporation managing power grid operations in California (Makarov et al., 2011, Kleissl, 2013). For the 318 hourly time series at plant level, numerical weather prediction (NWP) forecasts generated by 3TIER (3TIER, 2010) are used as base forecasts. All the remaining base forecasts, for the six \mathcal{L}_0 and \mathcal{L}_1 time series at any time granularity $k \in \mathcal{K}$, and for the \mathcal{L}_2 2-3-4-6-8-12-24 h time series, are computed using the automatic ETS forecasting procedure of the R-package forecast (Hyndman et al., 2021), not controlling for possible negative forecasts. Furthermore, following Yagli et al. (2019), day-ahead persistence is

⁷ Alternatively, the spatio-temporally reconciled forecasts may be obtained according to the spatio-temporal extension of the structural approach developed by Hyndman et al. (2011) for the cross-sectional framework (Di Fonzo and Girolimetto, 2022a).

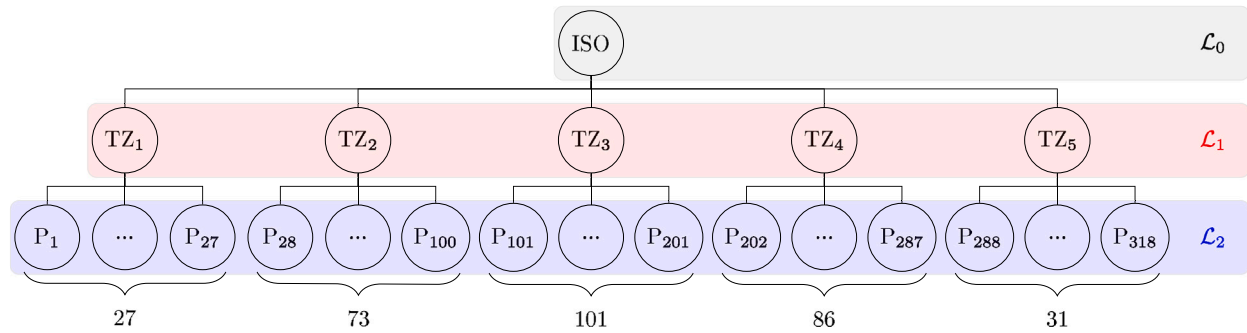


Fig. 2. PV324 hierarchy: Independent System Operator level (1 time series, \mathcal{L}_0), Transmission Zones level (5 time series, \mathcal{L}_1), and Plant level (318 time series, \mathcal{L}_2).




Fig. 3. Base forecasts and sequential $TSR_{\mathcal{L}_2}$ (Yagli et al., 2019) reconciled forecasts. Boxplots of the distribution in the 350 replications of the forecasting experiment of the spatial (in red) and temporal (in blue) gross discrepancies, as defined in (22). For spatio-temporal reconciled forecasts, both discrepancies are expected to be zero. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 2

Summary informations on the negative forecasts produced by the procedures considered by Yagli et al. (2019) in the forecasting experiment. Replications with at least a negative forecast (# rep), number of series out of 324 (# series) with at least a negative forecast in a single replication (min and max), and min and max negative values found in all replications (values). Hourly and daily forecasts, forecast horizon: operating day.

Approach	# rep (350)	# series		Values		# rep (350)	# series		values	
		min	max	min	max		min	max	min	max
Hourly forecasts										
base ^a	350	4	6	-15.617	-0.000	11	0	35	-51.205	-0.006
ost(ols)	350	109	324	-69.165	-0.000	11	0	60	-29.615	-0.001
ost(ols _{sp} , struc _{te})	350	91	324	-17.961	-0.000	6	0	28	-8.915	-0.008
ost(struc _{sp} , ols _{te})	350	137	324	-20.739	-0.000	11	0	33	-10.363	-0.000
ost(struc)	350	89	324	-14.292	-0.000	4	0	11	-1.020	-0.026
te(ols ₂) + sp(ols) ^b	350	42	324	-12.993	-0.000	10	0	72	-15.037	-0.018
te(struc ₂) + sp(ols) ^b	350	36	324	-12.994	-0.000	10	0	71	-14.791	-0.005
te(ols ₂) + sp(struc) ^b	350	144	324	-7.526	-0.000	10	0	43	-12.969	-0.005
te(struc ₂) + sp(struc) ^b	350	124	324	-7.642	-0.000	5	0	31	-7.042	-0.002
Daily forecasts										

^aThe approach produces spatial and temporal incoherent forecasts.

^bThe approach produces temporal incoherent forecasts.

used as the reference model (PERS)⁸:

$$\hat{y}_{T+h|T, PERS}^{[1]} = y_{T+h-48}^{[1]} \quad (19)$$

Given the lead time of 48 h, the day-ahead persistence takes the measurements made at day -2 as the forecasts for the operating day (Yagli et al., 2019, p. 394). Benchmark forecasts at any level of the spatial hierarchy and for any temporal granularity are obtained through spatio-temporal bottom-up of the 318 hourly bottom time series, that

is:

$$\tilde{y}_{PERS_{bu}} = \mathbf{F}\hat{\mathbf{b}}_{PERS}^{[1]} = \text{vec}(\tilde{\mathbf{Y}}'_{PERS_{bu}}). \quad (20)$$

It is worth noting that the benchmark forecasts are always non-negative, and both spatially and temporally coherent. These important properties are valid also for the forecasts obtained by spatio-temporal bottom-up reconciliation of the 318 hourly bottom time series' NWP forecasts 3TIER:

$$\tilde{y}_{3TIER_{bu}} = \mathbf{F}\hat{\mathbf{b}}_{3TIER}^{[1]} = \text{vec}(\tilde{\mathbf{Y}}'_{3TIER_{bu}}). \quad (21)$$

In light of the theorem 1 in Di Fonzo and Girolimetto (2022a), the STR and TSR sequential reconciliations proposed by Yagli et al.

⁸ On the choice of the standard reference forecasting method, see Yang (2019).




Fig. 4. MCB Nemenyi test results: average ranks and 95% confidence intervals. The unconstrained reconciliation approaches considered by Yagli et al. (2019) are sorted vertically according to the nRMSE(%) mean rank. Hourly (top panel) and Daily (bottom panel) forecasts for $\mathcal{L}_0, \mathcal{L}_1, \mathcal{L}_2$ levels (324 series). Forecast horizon: operating day. The mean rank of each approach is displayed to the right of their names. If the intervals of two forecast reconciliation approaches do not overlap, this indicates a statistically different performance. Thus, approaches that do not overlap with the green interval are considered significantly worse than the best, and vice-versa. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

(2019) reduce to the following approaches: $ost(ols)$, $ost(ols_{sp}, struc_{te})$, $ost(struc_{sp}, ols_{te})$, $ost(struc)$. In addition, Yagli et al. (2019) consider other four Temporal-then-Spatial-Reconciliation approaches, called $TSR_{\mathcal{L}_2}$, where the temporal reconciliation is applied only to the 318 plant level series’ base forecasts. In this case, although constant matrices are used in either reconciliation steps, the theorem so far no longer holds, so the obtained forecasts are temporally incoherent, as we show in the following. In order to distinguish these approaches from the conventional sequential techniques, we call them $te(ols_2)+sp(ols)$, $te(struc_2)+sp(ols)$, $te(ols_2)+sp(struc)$, $te(struc_2)+sp(struc)$, respectively.

3.1. Non-negativity and aggregation consistency issues

Standard forecast reconciliation, both in space and/or in time, may produce negative revised forecasts (Yang et al., 2017a,b, Yagli et al., 2019), unless specific non-negative reconciliation approaches are applied (Wickramasuriya et al., 2020, Di Fonzo and Girolimetto, 2022a, Girolimetto and Di Fonzo, 2022). Thus it is not surprising that, with the exception of $PERS_{bu}$ and $3TIER_{bu}$, the other approaches considered by Yagli et al. (2019) produce some negative reconciled forecasts. Details on this issue are shown in Table 2.

Negative hourly forecasts are obtained in all 350 replications of the forecasting experiment, and there are some cases (in a range within 4 and 11 replications out of 350) where negative daily forecasts are produced as well.⁹ Furthermore, the base forecasts are incoherent both in space and time, and the sequential $TSR_{\mathcal{L}_2}$ approaches proposed

⁹ The relatively low number of negative hourly base forecasts is explained by the fact that the hourly base forecasts of the 318 bottom time series are NWP 3TIER forecasts, that are always non-negative. Negative values are

by Yagli et al. (2019) are temporally incoherent. This can be visually appreciated from Fig. 3, showing the boxplots from the distribution of the spatial and temporal gross discrepancies registered in the 350 replications of the forecasting experiment, computed as (Di Fonzo and Girolimetto, 2023):

$$\begin{aligned} \text{Spatial gross discrepancy: } d_{sp} &= \|\mathbf{U}'\hat{\mathbf{Y}}\|_1 \\ \text{Temporal gross discrepancy: } d_{te} &= \|\mathbf{Z}'\hat{\mathbf{Y}}\|_1 \end{aligned} \quad (22)$$

where $\|\mathbf{X}\|_1 = \sum_{i,j} |x_{i,j}|$. For truly spatio-temporal reconciled forecasts, neither spatial nor temporal discrepancies are present, that is $d_{sp} = d_{te} = 0$.

In this case, forecast reconciliation may thus generate physically unreasonable values. Furthermore, if coherency is wished, the apparently innocuous practice of setting possible negative forecasts to zero is not advisable, since incoherence in spatial and/or time dimensions would be produced. To overcome the above limitations, in Section 4 we apply the simple operational strategy proposed by Di Fonzo and Girolimetto (2022a), able to generate fully reconciled non-negative forecasts.

3.2. Forecast evaluation

Following Yagli et al., 2019, the accuracy of the considered approaches is measured in terms of normalized Root Mean Square Error (nRMSE), normalized Mean Bias Error (nMBE), and Forecast Skill score

instead present in the ETS base forecasts for the aggregated series. Details can be found in the on-line appendix.

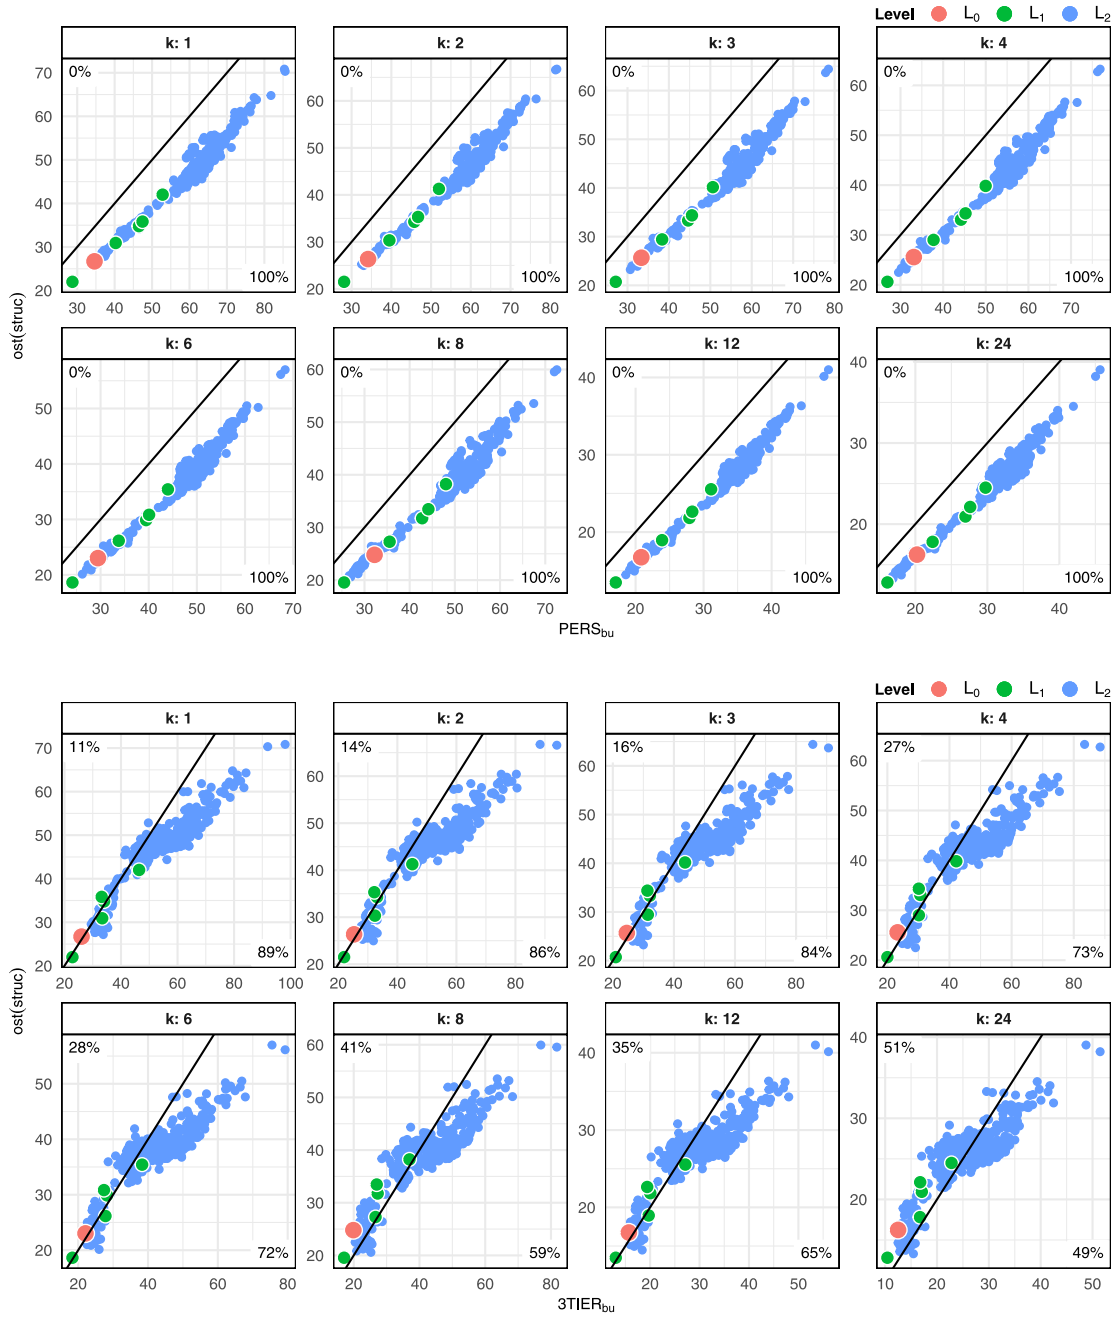


Fig. 5. Comparison of nRMSE(%) between $PERS_{bu}$ and $ost(struc)$ (top panel), and between $3TIER_{bu}$ and $ost(struc)$ (bottom panel). The black line represents the bisector, where the nRMSE's for both approaches are equal. On the top-left (bottom-right) corner of each graph, the percentage of points above (below) the bisector is reported. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article).

(FS):

$$nRMSE_{i,j}^{[k]} = \sqrt{\frac{\frac{1}{L} \sum_{l=1}^L (\hat{y}_{i,j,l}^{[k]} - y_{i,l}^{[k]})^2}{\frac{1}{L} \sum_{l=1}^L y_{i,l}^{[k]}}}, \quad (23)$$

$$nMBE_{i,j}^{[k]} = \frac{\frac{1}{L} \sum_{l=1}^L (\hat{y}_{i,j,l}^{[k]} - y_{i,l}^{[k]})}{\frac{1}{L} \sum_{l=1}^L y_{i,l}^{[k]}}, \quad FS_{i,j}^{[k]} = 1 - \frac{nRMSE_{i,j}^{[k]}}{nRMSE_{i,0}^{[k]}}$$

where $i = 1, \dots, n$, denotes the series, $k \in \mathcal{K}$, $j = 0, \dots, J$, denotes the forecasting approach ($j = 0$ for the reference model $PERS$), and $L = nrep \cdot \frac{m}{k}$, where $nrep = 350$ is the number of the forecasting experiment replications. Whereas nRMSE penalizes large errors, nMBE reveals over- and under-prediction through the sign of the metric (Yagli

et al., 2019, p. 394). Forecast Skill can be either negative (approach is worse than the reference model) or positive (approach is better than the reference model).

It should be noted that for spatio-temporal coherent forecasts the index $nMBE_{i,j}^{[k]}$ does not depend on the temporal aggregation order $k \in \mathcal{K}$ (see Appendix A). This means that, if j denotes a fully coherent spatio-temporal forecast reconciliation approach, the following identities hold:

$$nMBE_{i,j}^{[m]} = nMBE_{i,j}^{[k_{p-1}]} = \dots = nMBE_{i,j}^{[k_2]} = nMBE_{i,j}^{[1]}, \quad i = 1, \dots, n. \quad (24)$$

From Table 3 it appears that on average all the considered forecasting approaches improve on the benchmark $PERS_{bu}$, in a range between

Table 3

Forecast accuracies in terms of nRMSE(%), nMBE(%), and forecast skill over the PERS_{bu} benchmark of base forecasts and sequential reconciliation approaches. Unconstrained reconciliation procedures considered by Yagli et al. (2019), Tables 2, 3, p. 395. Hourly (H) and Daily (D) forecasts, forecast horizon: operating day. Bold entries and italic entries identify the best and the second best performing approaches, respectively. For hourly and daily forecasts produced by fully coherent spatio-temporal reconciliation approaches, the nMBE(%) indices are equal (see Eq. (24)).

Approach	\mathcal{L}_0		\mathcal{L}_1		\mathcal{L}_2	
	H	D	H	D	H	D
nRMSE(%)						
PERS _{bu}	34.62	20.23	43.15	24.57	59.75	30.65
3TIER _{bu}	26.03	12.48	33.95	16.75	53.46	25.19
base ^a	27.85	18.17	34.24	20.94	53.46	25.82
ost(<i>ols</i>)	30.69	17.91	39.17	21.74	51.54	26.73
ost(<i>ols</i> _{sp} , <i>struc</i> _{te})	28.26	16.96	35.19	20.61	48.20	25.46
ost(<i>struc</i> _{sp} , <i>ols</i> _{te})	28.74	17.33	35.48	20.82	49.32	26.14
ost(<i>struc</i>)	26.71	16.24	<i>33.11</i>	<i>19.64</i>	<i>46.74</i>	24.73
te(<i>ols</i> ₂)+sp(<i>ols</i>) ^b	27.80	17.96	34.36	21.80	48.52	26.71
te(<i>struc</i> ₂)+sp(<i>ols</i>) ^b	27.80	17.95	34.34	21.79	47.77	26.36
te(<i>ols</i> ₂)+sp(<i>struc</i>) ^b	27.02	17.24	33.40	20.62	47.76	25.98
te(<i>struc</i> ₂)+sp(<i>struc</i>) ^b	26.17	16.68	32.44	19.96	46.25	25.03
nMBE(%)						
PERS _{bu}	0.075	0.075	<i>0.058</i>	0.058	0.075	0.075
3TIER _{bu}	-4.213	-4.213	-3.745	-3.745	-4.362	-4.362
base ^a	-0.394	0.725	-0.131	1.000	-4.362	0.902
ost(<i>ols</i>)	0.593	0.593	0.557	0.557	0.590	0.590
ost(<i>ols</i> _{sp} , <i>struc</i> _{te})	0.359	<i>0.359</i>	0.339	<i>0.339</i>	0.354	<i>0.354</i>
ost(<i>struc</i> _{sp} , <i>ols</i> _{te})	0.735	0.735	0.744	0.744	0.733	0.733
ost(<i>struc</i>)	0.385	0.385	0.426	0.426	0.374	0.374
te(<i>ols</i> ₂)+sp(<i>ols</i>) ^b	-0.353	0.769	-0.369	0.718	-0.352	0.754
te(<i>struc</i> ₂)+sp(<i>ols</i>) ^b	-0.355	0.768	-0.369	0.718	-0.360	0.747
te(<i>ols</i> ₂)+sp(<i>struc</i>) ^b	0.100	0.858	0.119	0.863	0.099	0.850
te(<i>struc</i> ₂)+sp(<i>struc</i>) ^b	-0.093	0.666	-0.048	0.695	-0.104	0.647
forecast skill						
3TIER _{bu}	0.248	0.383	0.213	0.318	0.105	0.178
base ^a	0.196	0.102	0.206	0.148	0.105	0.158
ost(<i>ols</i>)	0.113	0.115	0.092	0.115	0.137	0.128
ost(<i>ols</i> _{sp} , <i>struc</i> _{te})	0.184	0.162	0.184	0.161	0.193	0.169
ost(<i>struc</i> _{sp} , <i>ols</i> _{te})	0.170	0.143	0.178	0.152	0.175	0.147
ost(<i>struc</i>)	0.228	<i>0.197</i>	<i>0.233</i>	<i>0.201</i>	<i>0.218</i>	0.193
te(<i>ols</i> ₂)+sp(<i>ols</i>) ^b	0.197	0.112	0.204	0.113	0.188	0.129
te(<i>struc</i> ₂)+sp(<i>ols</i>) ^b	0.197	0.113	0.204	0.113	0.201	0.140
te(<i>ols</i> ₂)+sp(<i>struc</i>) ^b	0.219	0.148	0.226	0.161	0.201	0.152
te(<i>struc</i> ₂)+sp(<i>struc</i>) ^b	<i>0.244</i>	0.176	0.248	0.188	0.226	<i>0.183</i>

^aThe approach produces spatial and temporal incoherent forecasts.

^bThe approach produces temporal incoherent forecasts.

9.2% (ost(*ols*) for the 5 \mathcal{L}_1 series' hourly forecasts) and 38.3% (3TIER_{bu} for the \mathcal{L}_0 total series' daily forecasts).¹⁰

The forecasting accuracy indices for each Transmission Zone forecasts are reported in Table 4.

We observe that:

- 3TIER_{bu} performs the best for the total series (\mathcal{L}_0) at any temporal granularity, the second best being te(*struc*₂)+sp(*struc*).
- The approach te(*struc*₂)+sp(*struc*) ranks first for 4 out of the 5 hourly upper series at level \mathcal{L}_1 , whereas at daily level 3TIER_{bu} 'wins' again. The performance of ost(*struc*) appears very close to that of te(*struc*₂)+sp(*struc*).
- te(*struc*₂)+sp(*struc*) and ost(*struc*) show the best performance for the 318 bottom time series at any temporal granularity, with a slight prevalence of ost(*struc*) (for $k \geq 3$). It should be noted that, unlike te(*struc*₂)+sp(*struc*), ost(*struc*) forecasts are spatio-temporally coherent.
- Unlike te(*struc*₂)+sp(*struc*), 3TIER_{bu} forecasts are always non-negative, and coherent both in space and time at any granularity.

¹⁰ The results for all temporal aggregation orders, $k \in \{24, 12, 8, 6, 4, 3, 2, 1\}$, are available in the on-line appendix.

- For these reasons, 3TIER_{bu} should be considered as a challenging competitor in the evaluation of the new proposed procedures (Section 4).

To give a complete picture of the evaluation results for hourly and daily forecasts, in Fig. 4 the Multiple Comparison with the Best (MCB) Nemeyi tests are shown (Koning et al., 2005, Kourentzes and Athanasopoulos, 2019, Makridakis et al., 2022). This allows to establish if the forecasting performances of the considered techniques are significantly different.

At daily level, ost(*struc*) ranks first, and is significantly better than the other forecasting approaches, with te(*struc*₂)+sp(*struc*) at the second place. This result is reversed at hourly level: te(*struc*₂)+sp(*struc*) ranks first and is significantly better than all the other approaches, with ost(*struc*) at the second place. However, it should be recalled that the forecasts produced by te(*struc*₂)+sp(*struc*) are not temporally coherent, which means that the sum of the hourly forecasts does not match with the daily forecast.

Limiting ourselves to consider fully coherent reconciled forecasts, the scatter plots of the 324 couples of nRMSE(%) for ost(*struc*) vs., respectively, PERS_{bu} and 3TIER_{bu} (Fig. 5), show that the most performing regression-based spatio-temporal reconciliation approach improves uniformly on the benchmark, and in the majority of cases on 3TIER_{bu}, particularly at hourly level ($k = 1$), where 89% of the variables are observed to show an improved nRMSE. However, it is worth noting that at daily level, this value decreases to 49%, which means that the NWP forecasts may still play a role at lower time granularity.

4. Extended analysis: non-negative cross-temporal reconciliation

In this section, we explore the performance of forecast reconciliation approaches able to produce non-negative PV forecasts, both temporally and spatially coherent. For this reason, among the approaches proposed by Yagli et al. (2019), we consider the reference benchmark PERS_{bu}, the NWP base forecasts 3TIER_{bu}, and ost(*struc*). These approaches are then compared with the following 7 spatio-temporal forecasting procedures:

- KA(*wlsv*_{te}, *wlsv*_{sp}): the heuristic approach by Kourentzes and Athanasopoulos (2019), using te(*wlsv*) in the first step, and sp(*wlsv*) in the second, respectively;
- st(*wlsv*_{te}, *bu*_{sp}), st(*struc*_{te}, *bu*_{sp}), st(*wlsv*_{sp}, *bu*_{te}), and st(*struc*_{sp}, *bu*_{te}): partly bottom-up spatio temporal reconciliation (see Section 2.2) according to, respectively, te(*wlsv*), te(*struc*), sp(*wlsv*), and sp(*struc*);
- ost(*wlsv*) and ost(*bdshr*): optimal (in least squares sense) spatio-temporal forecast reconciliation approaches using the in-sample forecast errors (Di Fonzo and Girolimetto, 2023).

4.1. Non negative forecast reconciliation: *sntz*

Each approach, when used in its 'free' version, i.e., without considering non-negative constraints in the linearly constrained quadratic program (13), is not guaranteed to always produce non-negative reconciled forecasts. This fact may be an issue for the analyst, since in many practical situations negative forecasts could have no meaning, thus undermining the quality of the results found and the conclusions thereof. In what follows, we consider a simple heuristic strategy to avoid negative reconciled forecasts, without using any sophisticated, and time consuming, numerical optimization procedure.¹¹ More precisely, possible negative values of the unconstrained reconciled high-frequency bottom time series forecasts are set to zero. Denote $\tilde{\mathbf{B}}_0^{[1]}$ the matrix containing the non-negative reconciled forecasts produced by a 'free' approach, and the 'zeroed' ones. The complete vector of non-negative

¹¹ Recent contributions on this topic in the hierarchical forecasting field are Wickramasuriya et al. (2020), Girolimetto and Di Fonzo (2022) and Di Fonzo and Girolimetto (2022b).

Table 4

Forecast accuracies in terms of nRMSE(%), nMBE(%), and forecast skill over the PERS_{bu} benchmark of base forecasts and sequential reconciliation approaches for the series at L₁ level (Transmission Zones). Unconstrained reconciliation procedures considered by Yagli et al. (2019), Tables 2, 3, p. 395. Hourly (H) and Daily (D) forecasts, forecast horizon: operating day. Bold entries and italic entries identify the best and the second best performing approaches, respectively. For hourly and daily forecasts produced by fully coherent spatio-temporal reconciliation approaches, the nMBE(%) indices are equal (see Eq. (24)).

Approach	TZ ₁		TZ ₂		TZ ₃		TZ ₄		TZ ₅	
	H	D	H	D	H	D	H	D	H	D
nRMSE(%)										
PERS _{bu}	28.72	16.12	40.27	22.40	46.48	26.96	52.82	29.75	47.44	27.62
3TIER _{bu}	22.81	10.43	33.34	16.72	<i>34.01</i>	17.06	46.40	22.80	33.16	16.75
base ^a	22.41	13.55	32.05	18.70	<i>35.14</i>	21.95	44.94	26.89	36.67	23.58
ost(<i>ols</i>)	29.27	14.89	34.58	19.25	39.55	22.91	47.30	26.50	45.16	25.15
ost(<i>ols</i> _{sp} , <i>struc</i> _{te})	24.26	13.93	32.32	18.47	36.69	21.78	43.89	25.30	38.81	23.59
ost(<i>struc</i> _{sp} , <i>ols</i> _{te})	23.44	13.50	32.89	18.83	37.42	22.33	44.76	25.97	38.90	23.49
ost(<i>struc</i>)	22.00	<i>12.81</i>	<i>30.92</i>	<i>17.82</i>	34.79	<i>20.94</i>	<i>42.02</i>	<i>24.50</i>	35.83	<i>22.12</i>
te(<i>ols</i> ₂) + sp(<i>ols</i>) ^b	22.95	15.46	32.06	19.23	35.12	22.38	44.93	26.91	36.73	25.02
te(<i>struc</i> ₂) + sp(<i>ols</i>) ^b	22.93	15.46	32.05	19.22	35.11	22.37	44.92	26.90	36.70	25.00
te(<i>ols</i> ₂) + sp(<i>struc</i>) ^b	22.11	13.36	31.40	18.70	34.70	21.93	42.97	26.06	35.82	23.06
te(<i>struc</i> ₂) + sp(<i>struc</i>) ^b	21.58	13.00	30.48	18.09	33.58	21.18	41.86	25.21	<i>34.69</i>	22.32
nMBE(%)										
PERS _{bu}	-0.013	-0.013	<i>0.099</i>	<i>0.099</i>	<i>0.127</i>	0.127	0.030	0.030	0.050	0.050
3TIER _{bu}	-1.967	-1.967	-3.215	-3.215	-3.636	-3.636	-7.364	-7.364	-2.542	-2.542
base ^a	-0.101	0.889	-0.419	0.842	-0.134	1.613	-0.119	0.159	<i>0.118</i>	1.499
ost(<i>ols</i>)	0.250	0.250	0.307	0.307	1.006	1.006	0.345	0.345	0.876	0.876
ost(<i>ols</i> _{sp} , <i>struc</i> _{te})	0.078	<i>0.078</i>	0.091	0.091	0.577	0.577	0.326	0.326	0.625	<i>0.625</i>
ost(<i>struc</i> _{sp} , <i>ols</i> _{te})	0.660	0.660	0.478	0.478	0.968	0.968	0.591	0.591	1.026	1.026
ost(<i>struc</i>)	0.442	0.442	0.203	0.203	0.541	<i>0.541</i>	0.185	0.185	0.759	0.759
te(<i>ols</i> ₂) + sp(<i>ols</i>) ^b	-0.492	0.428	-0.577	0.653	-0.248	1.471	-0.268	-0.007	-0.258	1.047
te(<i>struc</i> ₂) + sp(<i>ols</i>) ^b	-0.490	0.429	-0.578	0.653	-0.249	1.470	-0.273	<i>-0.012</i>	-0.256	1.049
te(<i>ols</i> ₂) + sp(<i>struc</i>) ^b	0.135	0.794	-0.123	0.685	0.182	1.233	<i>0.114</i>	0.446	0.289	1.157
te(<i>struc</i> ₂) + sp(<i>struc</i>) ^b	<i>0.047</i>	0.706	-0.249	0.559	0.009	1.060	-0.242	0.090	0.194	1.063
forecast skill										
3TIER _{bu}	0.206	0.353	0.172	0.254	<i>0.268</i>	0.367	0.121	0.234	0.301	0.394
base ^a	0.220	0.159	0.204	0.165	0.244	0.186	0.149	0.096	0.227	0.146
ost(<i>ols</i>)	-0.019	0.076	0.141	0.141	0.149	0.150	0.104	0.109	0.048	0.089
ost(<i>ols</i> _{sp} , <i>struc</i> _{te})	0.155	0.136	0.197	0.176	0.211	0.192	0.169	0.150	0.182	0.146
ost(<i>struc</i> _{sp} , <i>ols</i> _{te})	0.184	0.163	0.183	0.159	0.195	0.172	0.153	0.127	0.180	0.149
ost(<i>struc</i>)	<i>0.234</i>	<i>0.205</i>	<i>0.232</i>	<i>0.205</i>	0.251	<i>0.223</i>	<i>0.204</i>	<i>0.177</i>	0.245	<i>0.199</i>
te(<i>ols</i> ₂) + sp(<i>ols</i>) ^b	0.201	0.041	0.204	0.141	0.244	0.170	0.149	0.095	0.226	0.094
te(<i>struc</i> ₂) + sp(<i>ols</i>) ^b	0.202	0.041	0.204	0.142	0.245	0.170	0.150	0.096	0.227	0.095
te(<i>ols</i> ₂) + sp(<i>struc</i>) ^b	0.230	0.171	0.220	0.165	0.253	0.187	0.187	0.124	0.245	0.165
te(<i>struc</i> ₂) + sp(<i>struc</i>) ^b	0.249	0.194	0.243	0.192	0.277	0.215	0.208	0.153	<i>0.269</i>	0.192

^aThe approach produces spatial and temporal incoherent forecasts.

^bThe approach produces temporal incoherent forecasts.

spatio-temporal reconciled forecasts is computed as the spatio-temporal bottom-up aggregation of $\tilde{\mathbf{b}}_0^{[1]} = \text{vec} \left[\left(\tilde{\mathbf{B}}_0^{[1]} \right)' \right]$, that is, according to expression (10):

$$\tilde{\mathbf{y}}_0 = \tilde{\mathbf{F}} \tilde{\mathbf{b}}_0^{[1]} \tag{25}$$

We call set-negative-to-zero (sntz) this simple, and quick device to obtain non negative reconciled forecasts. While it certainly increases the forecasting accuracy of the high-frequency bottom time series forecasts wrt the ‘free’ counterparts, this does not hold true in general for the upper level series forecasts. In addition, even if the originally reconciled forecasts are obtained according to an unbiased approach, the sntz-reconciled forecasts are no more unbiased, like the non-negative forecasts obtained through numerical optimization procedures (Wickramasuriya et al., 2020). However, in practical situations the differences between the results produced by the sntz heuristic, and those obtained through a state-of-the-art numerical optimization procedure like osqp (Stellato et al., 2020) implemented in FoReco (Girolimetto and Di Fonzo, 2022), could be negligible. For example, Fig. 6 shows the graphs of one day of hourly forecasts computed by unconstrained ost(*struc*), ost(*struc*)_{osqp} and ost(*struc*)_{sntz}, respectively, for the aggregated series (ISO and five Transmission Zones), and for two bottom variables at plant level (P₂₉₅ and P₃₁₅) with a large number of negative reconciled forecasts. For each variable we consider the day

with the highest number of negative forecasts. It appears that the physically feasible forecasts produced by the non-negative reconciliation approaches are about the same at the aggregate levels, whereas some difference is visible at Plant level.

Table 5 shows the indices nRMSE(%) of the reconciled forecasts produced by the ost(*struc*) approach according to unconstrained and non-negative (both sntz and osqp) variants. It is worth noting that the sntz heuristic always gives the lowest nRMSE(%), independently of the temporal granularity of the forecasts.¹² This result is visually confirmed by the graphs in Fig. 7, showing the scatter plots of the 324 couples of nRMSE(%) for ost(*struc*)_{osqp} vs. ost(*struc*)_{sntz}. For this dataset, ost(*struc*)_{sntz} beats ost(*struc*)_{osqp} in no less than 94% of the 324 series for any temporal granularity. However, looking at indices nMBE(%) in Table 5, it emerges also that the bias of the ost(*struc*)_{sntz} forecasts is more pronounced than ost(*struc*)_{osqp}. This seems to be a price to pay for using such a simple heuristic.

4.2. Forecast accuracy of the selected approaches

Table 6 shows the indices nRMSE(%) and nMBE(%) of the considered non-negative spatio-temporal forecast reconciliation approaches,

¹² Similar results were found for all the considered reconciliation approaches. The details are available at request from the authors.

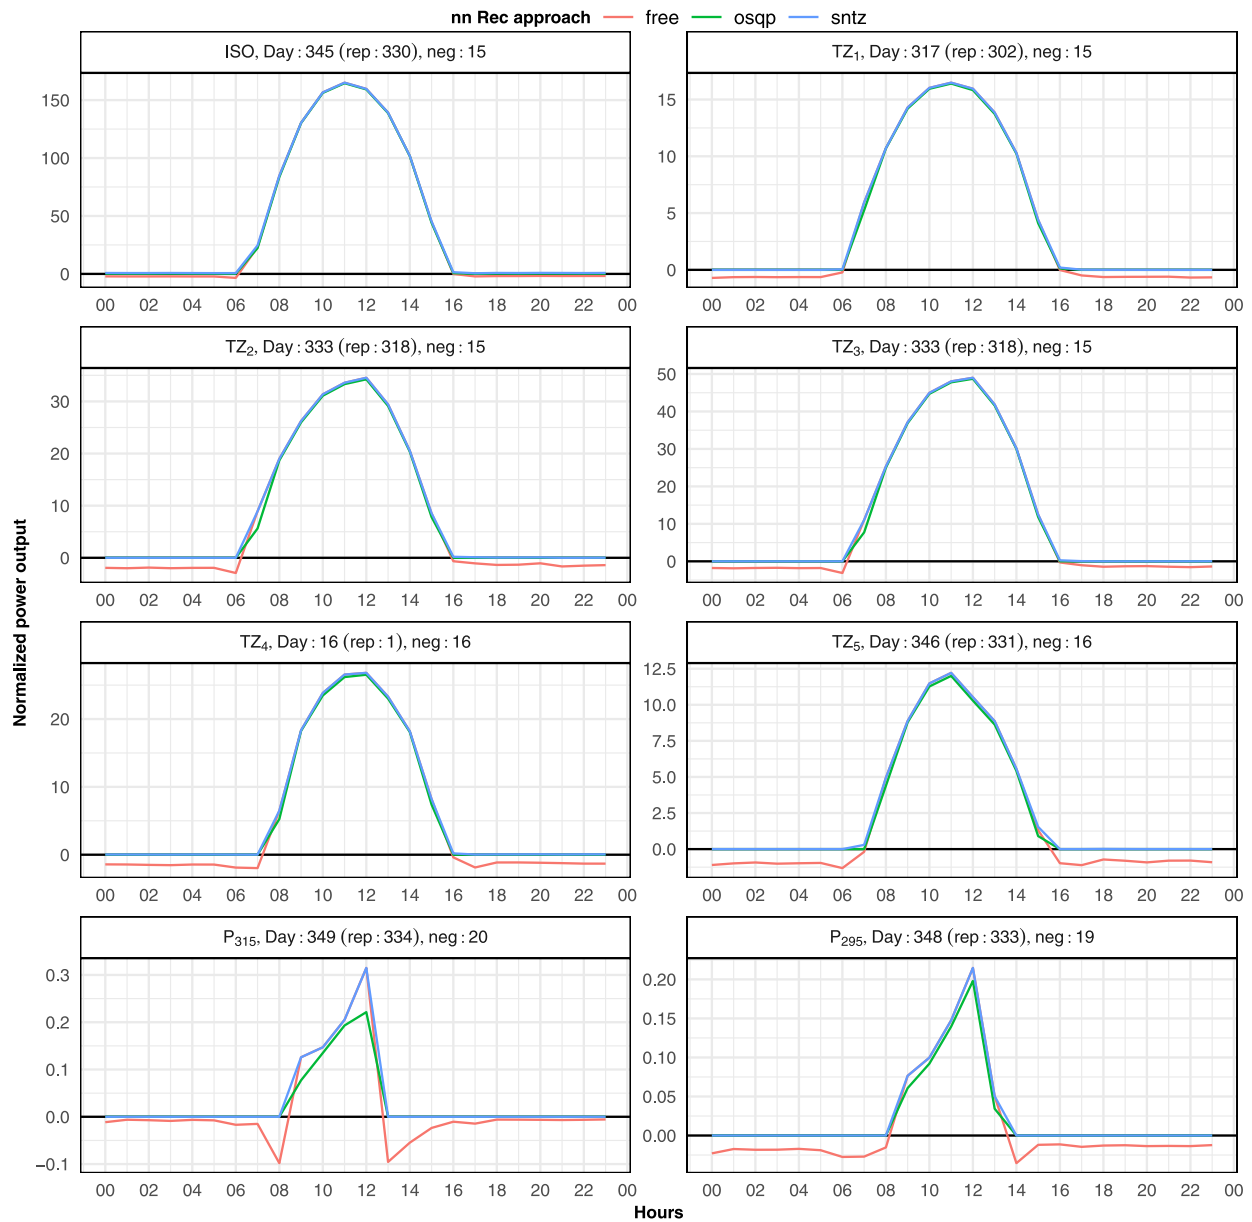


Fig. 6. One day of hourly reconciled forecasts for two of the 318 bottom variables (Plants P_{295} and P_{315} , component of TZ_5), and for the six upper time series (5 Transmission Zones and the Total ISO). For each series, it is shown the day with the highest number of negative forecasts produced by the reconciliation approach $ost(struc)$ (in red). The non-negative forecasts are obtained by $ost(struc)_{osqp}$ (in green) and $ost(struc)_{sntz}$ (in blue). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 5

Forecast accuracy in terms of nRMSE(%) and nMBE(%) of unconstrained and non-negative reconciled forecasts using the $ost(struc)$ approach. All temporal aggregation orders are considered, from hourly ($k = 1$) to daily ($k = 24$). Forecast horizon: operating day. Bold entries identify the best approach. As the forecasts are produced by fully coherent spatio-temporal reconciliation approaches, the nMBE(%) indices are equal for any temporal aggregation order k (see Eq. (24)).

Level (series)	Non-negative reconciliation	nRMSE(%)								nMBE(%)
		1	2	3	4	6	8	12	24	1–24
\mathcal{L}_0 (1)	free	26.71	26.35	25.67	25.57	23.03	24.80	16.74	16.24	0.385
	sntz	26.64	26.27	25.56	25.48	22.86	24.71	16.25	15.73	-1.079
	osqp	26.74	26.36	25.67	25.55	23.01	24.77	16.36	15.85	-0.398
\mathcal{L}_1 (5)	free	33.11	32.54	31.60	31.37	28.19	30.06	20.49	19.64	0.426
	sntz	32.99	32.41	31.45	31.23	27.96	29.91	19.88	19.00	-0.973
	osqp	33.11	32.52	31.58	31.32	28.15	29.99	20.01	19.14	-0.311
\mathcal{L}_2 (318)	free	46.74	44.02	42.05	41.07	36.67	38.16	26.56	24.73	0.374
	sntz	46.51	43.80	41.80	40.83	36.34	37.90	25.85	23.97	-1.109
	osqp	46.63	43.93	41.94	40.93	36.53	37.99	25.97	24.10	-0.422

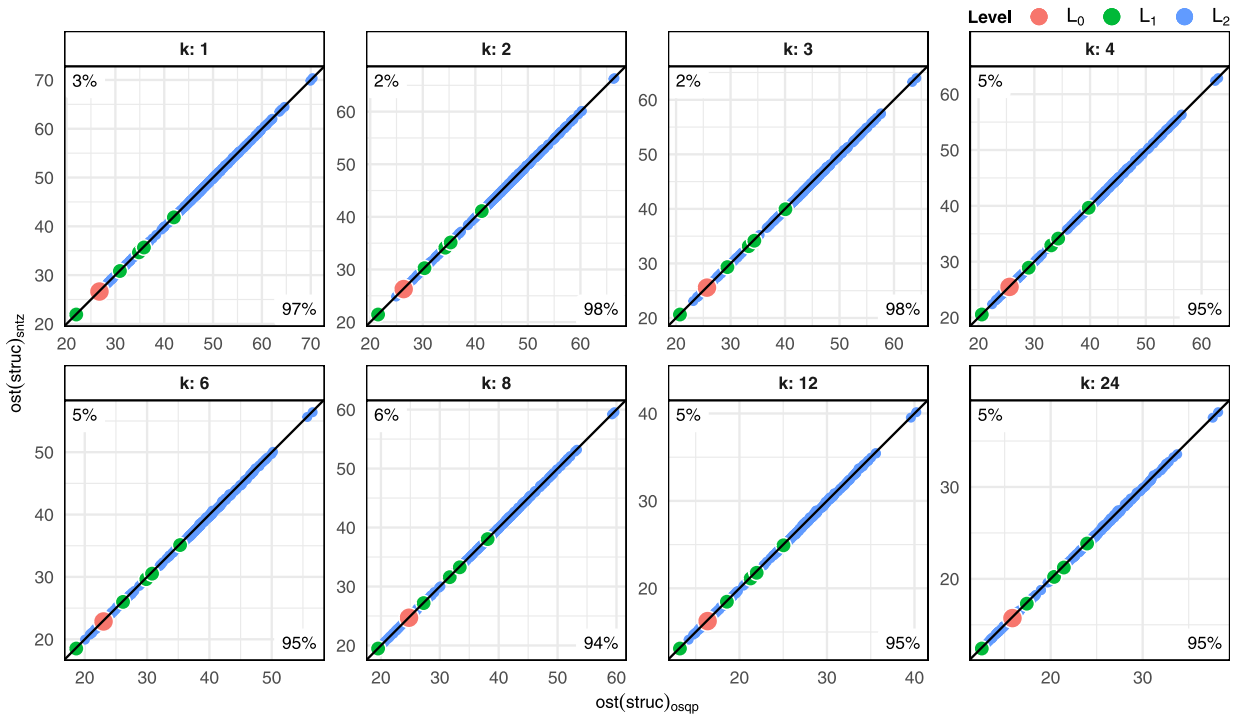


Fig. 7. Comparison of nRMSE(%) between sntz and osqp non-negative forecast reconciliation using the $ost(struc)$ approach. The black line represents the bisector, where the nRMSE's for $ost(struc)_{osqp}$ and $ost(struc)_{sntz}$ are equal. On the top-left (bottom-right) corner of each graph, the percentage of points above (below) the bisector is reported. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article).

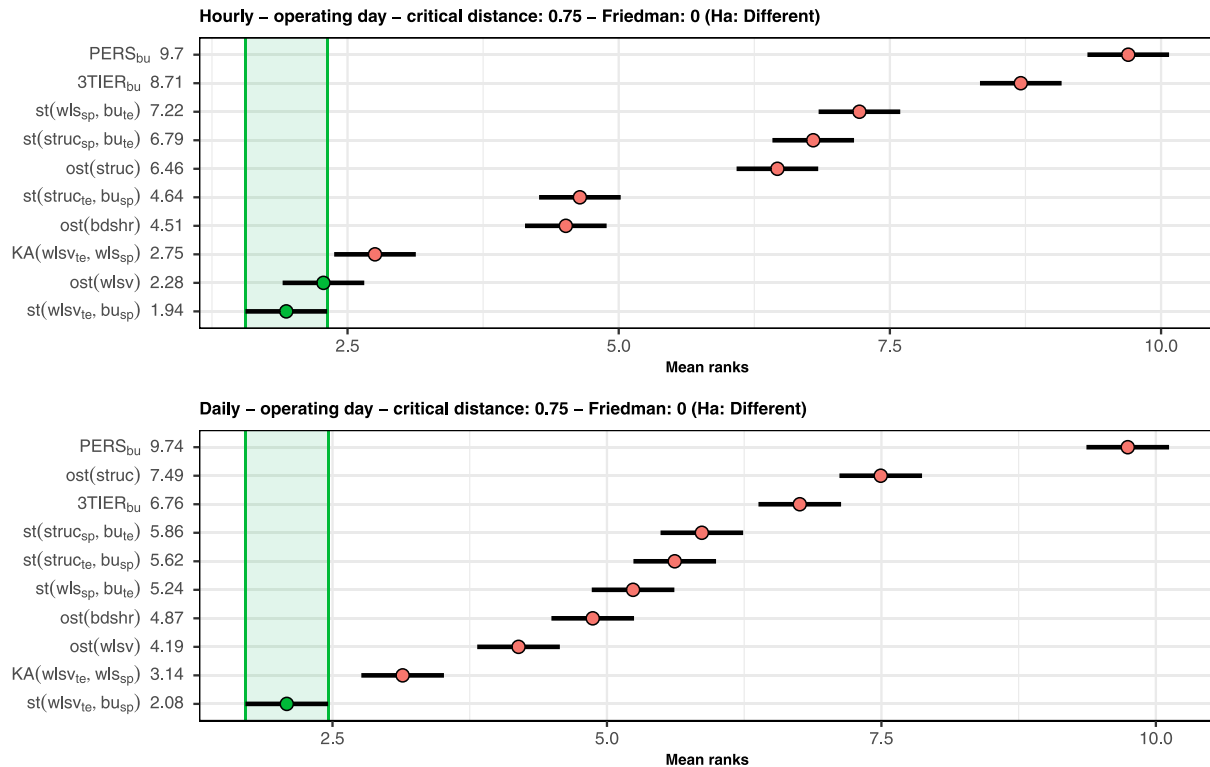


Fig. 8. MCB-Nemenyi test on selected non-negative spatio-temporal reconciliation approaches with operating day forecast horizon. L_0, L_1, L_2 levels (324 series). Top panel: hourly forecasts; Bottom panel: daily forecasts. The mean rank of each approach is displayed to the right of their names. If the intervals of two forecast reconciliation approaches do not overlap, this indicates a statistically different performance. Thus, approaches that do not overlap with the green interval are considered significantly worse than the best, and vice-versa. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

using the sntz heuristic, and the corresponding forecast skills over the benchmark forecasts $PERS_{bu}$. Overall, the accuracy improvements of

the new proposed approaches over the persistence model are in the range 16.5%–46.5%, whereas $3TIER_{bu}$ improvements are in the range

Table 6

Forecast accuracy of selected non-negative spatio-temporal reconciliation approaches and base forecasts in terms of nRMSE(%), nMBE(%), and forecast skill over the PERS_{bu} benchmark. Hourly and daily forecasts, forecast horizon: operating day. Bold entries and italic entries identify the best and the second best performing approaches, respectively. For hourly and daily forecasts produced by fully coherent spatio-temporal reconciliation approaches, the nMBE(%) indices are equal (see Eq. (24)).

Approach	\mathcal{L}_0		\mathcal{L}_1		\mathcal{L}_2	
	H	D	H	D	H	D
	nRMSE(%)					
PERS _{bu}	34.62	20.23	43.15	24.57	59.75	30.65
3TIER _{bu}	26.03	12.48	33.95	16.75	53.46	25.19
ost(<i>struc</i>)	26.64	15.73	32.99	19.00	46.51	23.97
KA(<i>wlsv</i> _{te} , <i>wlsv</i> _{sp})	23.59	13.55	30.23	16.91	44.48	22.21
st(<i>struc</i> _{sp} , <i>bu</i> _{te})	21.99	12.38	28.80	15.82	49.76	24.33
st(<i>wlsv</i> _{sp} , <i>bu</i> _{te})	21.23	10.82	28.97	15.01	49.88	23.66
st(<i>struc</i> _{te} , <i>bu</i> _{sp})	24.81	14.43	31.41	17.80	45.30	22.93
st(<i>wlsv</i> _{te} , <i>bu</i> _{sp})	23.50	13.42	30.13	16.78	44.42	22.12
ost(<i>wlsv</i>)	23.63	13.64	30.21	17.01	44.44	22.27
ost(<i>bdshr</i>)	24.13	13.92	30.45	17.12	45.24	22.54
	nMBE(%)					
PERS _{bu}	0.075		0.058		0.075	
3TIER _{bu}	-4.213		-3.745		-4.362	
ost(<i>struc</i>)	-1.079		-0.973		-1.109	
KA(<i>wlsv</i> _{te} , <i>wlsv</i> _{sp})	-0.691		-0.823		-1.399	
st(<i>struc</i> _{sp} , <i>bu</i> _{te})	-1.834		-1.593		-1.935	
st(<i>wlsv</i> _{sp} , <i>bu</i> _{te})	-2.810		-2.418		-2.910	
st(<i>struc</i> _{te} , <i>bu</i> _{sp})	-1.435		-1.308		-1.473	
st(<i>wlsv</i> _{te} , <i>bu</i> _{sp})	-1.429		-1.320		-1.454	
ost(<i>wlsv</i>)	-1.312		-1.195		-1.335	
ost(<i>bdshr</i>)	-0.734		-0.734		-0.728	
	forecast skill					
3TIER _{bu}	0.248	0.383	0.213	0.318	0.105	0.178
ost(<i>struc</i>)	0.231	0.222	0.235	0.227	0.222	0.218
KA(<i>wlsv</i> _{te} , <i>wlsv</i> _{sp})	0.319	0.330	0.299	0.312	0.256	0.275
st(<i>struc</i> _{sp} , <i>bu</i> _{te})	0.365	0.388	0.333	0.356	0.167	0.206
st(<i>wlsv</i> _{sp} , <i>bu</i> _{te})	0.387	0.465	0.329	0.389	0.165	0.228
st(<i>struc</i> _{te} , <i>bu</i> _{sp})	0.283	0.287	0.272	0.276	0.242	0.252
st(<i>wlsv</i> _{te} , <i>bu</i> _{sp})	0.321	0.337	0.302	0.317	0.257	0.278
ost(<i>wlsv</i>)	0.317	0.326	0.300	0.308	0.256	0.273
ost(<i>bdshr</i>)	0.303	0.312	0.294	0.303	0.243	0.265

10.5%–38.3%. Partly bottom-up approaches, with spatial reconciliation at the first step (i.e., st(*wlsv*_{sp}, *bu*_{te}) and st(*struc*_{sp}, *bu*_{te})) show a good performance for the six series at levels \mathcal{L}_0 and \mathcal{L}_1 . For the 318 disaggregated series at the Plant level, st(*wlsv*_{te}, *bu*_{sp}) and ost(*wlsv*) rank first and second, respectively, and their accuracy indices are very close each other, whereas the nMBE(%) of ost(*wlsv*) is lower than st(*wlsv*_{te}, *bu*_{sp}). In addition, all the considered non-negative forecast reconciliation approaches, even using the *sntz* heuristic (see Section 4.1), have smaller bias than the NWP forecasts 3TIER.

From Table 7 it emerges that even for distinct \mathcal{L}_1 series (Transmission Zones), the new approaches perform better than all the approaches considered by Yagli et al. (2019), reported in Table 4.

Furthermore, almost all the new approaches significantly outperform PERS_{bu}, 3TIER_{bu}, and ost(*struc*) for both hourly and daily forecasts (Fig. 8).

For a more compelling comparison of different forecasting methods, in Table 8 are shown the forecast skills in terms of nRMSE of the new spatio-temporal reconciliation approaches over the more challenging NWP 3TIER_{bu} forecasts. All the forecasts skills are positive, and in a range between 0.047 and 0.184, for the hourly PV generated power forecasts at any spatial level. In this case it turns out that the new forecast reconciliation approaches, in addition to assuring full coherence and non-negativity of the revised forecasts, have a better forecasting accuracy. For daily forecasts, however, the picture is less clear-cut. The new forecasting approaches always improve for the 318 disaggregate series at the Plant level (\mathcal{L}_2), whereas for ISO and Transmission Zones this happens only for two partly bottom-up approaches.

Finally, focusing on the two best performing approaches of the forecasting experiment (st(*wlsv*_{te}, *bu*_{sp}) and ost(*wlsv*)), and looking at the nRMSEs of the individual series, it is worth noting that:

- st(*wlsv*_{te}, *bu*_{sp}) always produces more accurate forecasts than ost(*struc*), for any series and granularity (Fig. 9);
- the accuracy increases of st(*wlsv*_{te}, *bu*_{sp}) over the NWP approach 3TIER_{bu} are still clear (Fig. 10), even though for daily forecasts 3TIER_{bu} performs better in about 1 case out 4;
- the accuracy of st(*wlsv*_{te}, *bu*_{sp}) is practically indistinguishable from that of ost(*wlsv*) (Fig. 11).

We may thus conclude that, for the PV324 dataset considered in this work, a thorough exploitation of spatio-temporal hierarchies significantly improves the forecasting accuracy over the approaches considered in Yagli et al. (2019). In particular, the use of the in-sample base forecast errors, even in the simple diagonal versions of te(*wlsv*) (first step of st(*wlsv*_{te}, *bu*_{sp})), and ost(*wlsv*), increases the forecasting accuracy at different time granularities.

5. Conclusion

Renewable energy is providing increasingly more energy to the grid all over the world. But grid operators must carefully manage the balance between the generation and consumption of energy to make the best use of abundant renewable energy. For an ISO, this provides greater grid stability, higher revenue and better use of what sun is available at any one time. Better short-term solar energy forecasts mean lower-emissions, cheaper energy and a more stable electricity grid. Solar forecasting is thus a key tool to achieve these results. In this paper, spatio-temporal point forecast reconciliation has been applied to generate non-negative, fully coherent (both in space and time) forecasts of PV generated power. Both methodological and practical issues have been tackled, in order to develop effective and easy-to-handle spatio-temporal forecasting approaches. In addition to assuring both spatial and temporal coherence, and non-negativity of the reconciled forecasts, the results show that for the considered dataset, spatio-temporal forecast reconciliation significantly improves on the sequential procedures proposed by Yagli et al. (2019), at any geographical level of the hierarchy and for any temporal aggregation order. It is worth noting that for the hourly PV generated power at any spatial level, all forecasts skills are positive and range between 4.7% and 18.4% when the NWP 3TIER is used as reference model.

However, these findings should be considered neither conclusive, nor valid in general. Other forecasting experiments, through simulations and using other datasets, as well as in solar forecasting and in other application fields, should be performed to empirically robustify the results shown so far. In addition, other research queries could raise in the field of spatio-temporal PV forecast reconciliation, like using reduced temporal hierarchies, as well as forecasting through appropriate Machine Learning end-to-end approaches (Stratigakos et al., 2022), and probabilistic instead of deterministic (point) forecasting (Panamtash and Zhou, 2018, Jeon et al., 2019, Yang, 2020, Yagli et al., 2020, Ben Taieb et al., 2021, Panagiotelis et al., 2022). All these topics are left for future research.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgment

Parts of this research were carried in the frame of the PRIN2017 project “HiDEA: Advanced Econometrics for High-frequency Data”, 2017RSMPZZ.

Table 7

Forecast accuracy of selected non-negative spatio-temporal reconciliation approaches and base forecasts in terms of nRMSE(%), nMBE(%), and forecast skill over the PERS_{bu} benchmark, for the series at L₁ level (Transmission Zones). Hourly and daily forecasts, forecast horizon: Operating day. Bold entries and italic entries identify the best and the second best performing approaches, respectively. For hourly and daily forecasts produced by fully coherent spatio-temporal reconciliation approaches, the nMBE(%) indices are equal (see Eq. (24)).

Approach	TZ ₁		TZ ₂		TZ ₃		TZ ₄		TZ ₅	
	H	D	H	D	H	D	H	D	H	D
nRMSE(%)										
PERS _{bu}	28.72	16.12	40.27	22.40	46.48	26.96	52.82	29.75	47.44	27.62
3TIER _{bu}	22.81	<i>10.43</i>	33.34	16.72	34.01	17.06	46.40	22.80	33.16	16.75
ost(<i>struc</i>)	21.92	12.41	30.84	17.30	34.67	20.21	41.84	23.84	35.66	21.23
KA(<i>wlsv_{te}</i> , <i>wlsv_{sp}</i>)	20.23	11.10	28.57	15.43	30.98	17.45	39.72	22.09	31.64	18.51
st(<i>struc_{sp}</i> , <i>bu_{te}</i>)	<i>20.15</i>	10.92	27.01	<i>14.37</i>	28.74	<i>15.74</i>	38.12	<i>20.89</i>	29.99	17.17
st(<i>wlsv_{sp}</i> , <i>bu_{te}</i>)	19.33	9.72	28.11	14.18	29.55	15.10	39.20	20.17	28.68	15.88
st(<i>struc_{te}</i> , <i>bu_{sp}</i>)	21.04	11.63	29.55	16.24	32.57	18.73	40.51	22.67	33.37	19.75
st(<i>wlsv_{te}</i> , <i>bu_{sp}</i>)	20.20	10.98	28.53	15.35	30.89	17.37	39.69	22.01	31.32	18.19
ost(<i>wlsv</i>)	20.18	11.18	28.53	15.50	30.97	17.53	39.56	22.16	31.84	18.69
ost(<i>bdshr</i>)	21.00	11.74	29.25	16.02	31.63	17.85	39.73	22.27	30.65	17.74
nMBE(%)										
PERS _{bu}		-0.013		-0.099		0.127		0.030		0.050
3TIER _{bu}		-1.967		-3.215		-3.636		-7.364		-2.542
ost(<i>struc</i>)		-0.422		-1.082		-0.864		-1.707		-0.790
KA(<i>wlsv_{te}</i> , <i>wlsv_{sp}</i>)		-0.575		-0.959		-0.814		-1.348		-0.417
st(<i>struc_{sp}</i> , <i>bu_{te}</i>)		-0.702		-1.434		-1.514		-3.390		-0.926
st(<i>wlsv_{sp}</i> , <i>bu_{te}</i>)		-1.437		-2.237		-2.530		-4.937		-0.949
st(<i>struc_{te}</i> , <i>bu_{sp}</i>)		-0.620		-1.300		-1.224		-2.243		-1.153
st(<i>wlsv_{te}</i> , <i>bu_{sp}</i>)		-0.890		-1.417		-1.246		-2.039		-1.011
ost(<i>wlsv</i>)		-0.768		-1.330		-1.169		-1.865		-0.843
ost(<i>bdshr</i>)		-0.652		-0.727		-0.676		-0.828		-0.778
forecast skill										
3TIER _{bu}	0.206	0.353	0.172	0.254	0.268	0.367	0.121	0.234	0.301	0.394
ost(<i>struc</i>)	0.237	0.230	0.234	0.228	0.254	0.250	0.208	0.199	0.248	0.231
KA(<i>wlsv_{te}</i> , <i>wlsv_{sp}</i>)	0.296	0.312	0.291	0.311	0.334	0.353	0.248	0.258	0.333	0.330
st(<i>struc_{sp}</i> , <i>bu_{te}</i>)	0.299	0.323	0.329	0.359	0.382	<i>0.416</i>	0.278	<i>0.298</i>	0.368	0.378
st(<i>wlsv_{sp}</i> , <i>bu_{te}</i>)	0.327	0.397	0.302	0.367	<i>0.364</i>	0.440	<i>0.258</i>	0.322	0.396	0.425
st(<i>struc_{te}</i> , <i>bu_{sp}</i>)	0.268	0.279	0.266	0.275	0.299	0.305	0.233	0.238	0.297	0.285
st(<i>wlsv_{te}</i> , <i>bu_{sp}</i>)	0.297	0.319	0.292	0.315	0.335	0.356	0.249	0.260	0.340	0.341
ost(<i>wlsv</i>)	0.297	0.307	0.292	0.308	0.334	0.350	0.251	0.255	0.329	0.323
ost(<i>bdshr</i>)	0.269	0.272	0.274	0.285	0.320	0.338	0.248	0.251	0.354	0.358

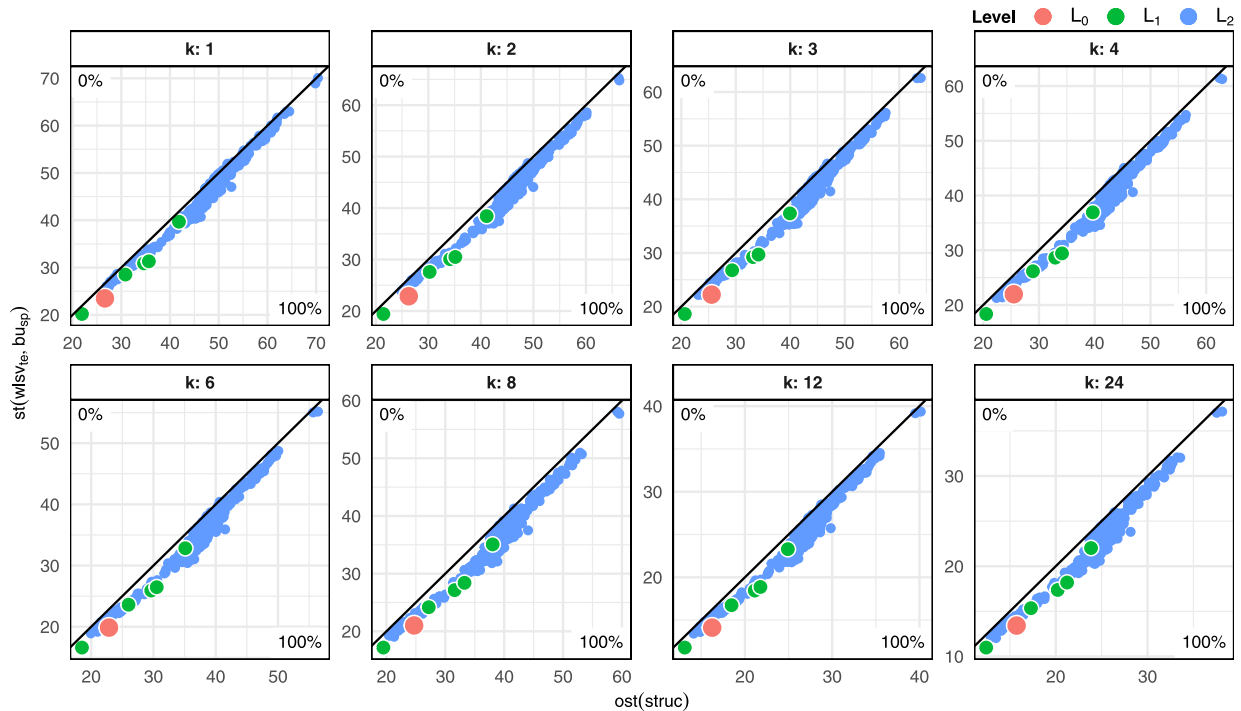


Fig. 9. Comparison of nRMSE(%) between non-negative reconciliation approaches: ost(*struc*) and st(*wlsv_{te}*, *bu_{sp}*). Forecast horizon: operating day. The black line represents the bisector, where the nRMSE's for both approaches are equal. On the top-left (bottom-right) corner of each graph, the percentage of points above (below) the bisector is reported. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article).

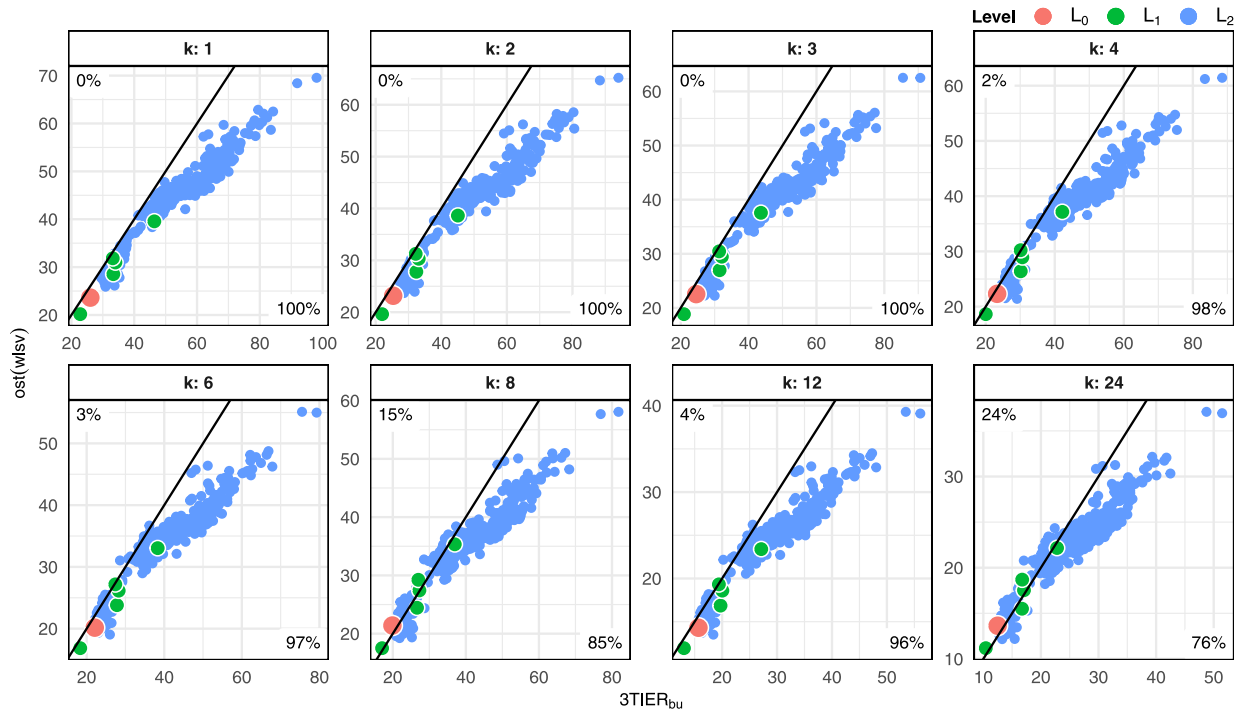


Fig. 10. Comparison of nRMSE(%) between non-negative reconciliation approaches: $3TIER_{bu}$ and $ost(wlsv_{te}, bu_{sp})$. Forecast horizon: operating day. The black line represents the bisector, where the nRMSE's for both approaches are equal. On the top-left (bottom-right) corner of each graph, the percentage of points above (below) the bisector is reported. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article).

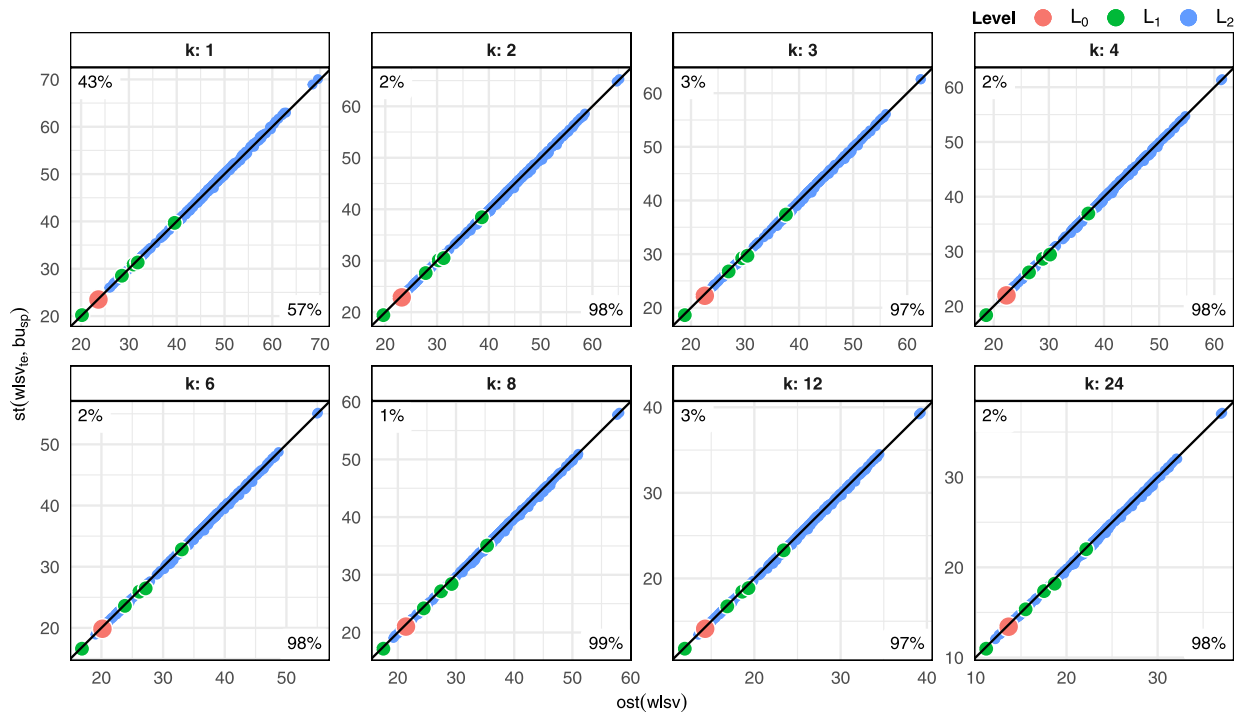


Fig. 11. Comparison of nRMSE(%) between non-negative reconciliation approaches: $ost(wlsv)$ and $st(wlsv_{te}, bu_{sp})$. Forecast horizon: operating day. The black line represents the bisector, where the nRMSE's for both approaches are equal. On the top-left (bottom-right) corner of each graph, the percentage of points above (below) the bisector is reported. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article).

Table 8

Forecast skills over the NWP 3TIER_{bu} forecasts of ost(*struc*) and of the new spatio-temporal forecast reconciliation approaches. Red entries identify negative skills. (For interpretation of the references to color in the entries of this table, the reader is referred to the web version of this article).

Approach	Hourly forecasts			Daily forecasts		
	\mathcal{L}_0	\mathcal{L}_1	\mathcal{L}_2	\mathcal{L}_0	\mathcal{L}_1	\mathcal{L}_2
ost(<i>struc</i>)	-0.023	0.028	0.130	-0.260	-0.134	0.048
KA(<i>wlsv_{te}</i> , <i>wlsv_{sp}</i>)	0.094	0.110	0.168	-0.085	-0.010	0.118
st(<i>struc_{sp}</i> , <i>bu_{te}</i>)	0.155	0.152	0.069	0.008	0.056	0.034
st(<i>wlsv_{sp}</i> , <i>bu_{te}</i>)	0.184	0.147	0.067	0.133	0.104	0.061
st(<i>struc_{te}</i> , <i>bu_{sp}</i>)	0.047	0.075	0.153	-0.156	-0.063	0.089
st(<i>wlsv_{te}</i> , <i>bu_{sp}</i>)	0.097	0.113	0.169	-0.075	-0.002	0.122
ost(<i>wlsv</i>)	0.092	0.110	0.169	-0.093	-0.016	0.116
ost(<i>bdshr</i>)	0.073	0.103	0.154	-0.115	-0.022	0.105

Appendix A. The index nMBE for different temporal aggregation orders with fully coherent spatio-temporal forecast reconciliation approaches

Let $\tilde{y}_t^{[k]}$, $t = 1, \dots, N_k$, $k \in \mathcal{K}$, be the spatio-temporal reconciled forecasts for a single time series in a complete time cycle (see Section 2). The temporal coherency implies that

$$\tilde{y}_t^{[k]} = \sum_{j=k(t-1)+1}^{kt} \tilde{y}_j^{[1]}, \quad \forall k \in \mathcal{K}, \tag{A.1}$$

where $\tilde{y}_j^{[1]}$, $j = 1, \dots, N_1$, denotes the forecasts for the series observed at the highest time frequency.

As the realizations of the time series are by definition temporally coherent (i.e., $y_t^{[k]} = \sum_{j=k(t-1)+1}^{kt} y_j^{[1]}$), the normalized Mean Bias Error defined in Eq. (23) can be written as

$$\begin{aligned} \text{nMBE}^{[k]} &= \frac{\frac{1}{N_k} \sum_{t=1}^{N_k} (\tilde{y}_t^{[k]} - y_t^{[k]})}{\frac{1}{N_k} \sum_{t=1}^{N_k} y_t^{[k]}} \\ &= \frac{\sum_{t=1}^{N_k} (\sum_{j=k(t-1)+1}^{kt} \tilde{y}_j^{[1]} - \sum_{j=k(t-1)+1}^{kt} y_j^{[1]})}{\sum_{t=1}^{N_k} \sum_{j=k(t-1)+1}^{kt} y_j^{[1]}} \\ &= \frac{\sum_{j=1}^{N_1} (\tilde{y}_j^{[1]} - y_j^{[1]})}{\sum_{j=1}^{N_1} y_j^{[1]}} = \text{nMBE}^{[1]}, \quad k \in \{m, k_{p-1}, \dots, k_2\}. \end{aligned}$$

Appendix B. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.solener.2023.01.003>.

References

3TIER, 2010. Development of regional wind resource and wind plant output datasets: final subcontract report. In: Electricity generation, Energy storage, Forecasting wind, Solar energy, Wind energy, Wind farms, Wind modeling. URL: https://digitalscholarship.unlv.edu/renew_pubs/4. (visited on 8 September 2022).

Antonanzas, J., Osorio, N., Escobar, R., Urraca, R., Martinez-de Pison, F., Antonanzas-Torres, F., 2016. Review of photovoltaic power forecasting. *Sol. Energy* 136, 78–111. <http://dx.doi.org/10.1016/j.solener.2016.06.069>.

Athanasopoulos, G., Ahmed, R.A., Hyndman, R.J., 2009. Hierarchical forecasts for Australian domestic tourism. *Int. J. Forecast.* 25 (1), 146–166. <http://dx.doi.org/10.1016/j.ijforecast.2008.07.004>.

Athanasopoulos, G., Hyndman, R.J., Kourentzes, N., Petropoulos, F., 2017. Forecasting with temporal hierarchies. *European J. Oper. Res.* 262 (1), 60–74. <http://dx.doi.org/10.1016/j.ejor.2017.02.046>.

Bates, D., Maechler, M., Jagan, M., Davis, T.A., Oehlschlagel, J., Riedy, J., 2022. Package matrix: sparse and dense matrix classes and methods. Version 1.4-1. URL: <https://cran.r-project.org/web/packages/Matrix/index.html>.

Ben Taieb, S., Taylor, J.W., Hyndman, R.J., 2021. Hierarchical probabilistic forecasting of electricity demand with smart meter data. *J. Amer. Statist. Assoc.* 116 (533), 27–43. <http://dx.doi.org/10.1080/01621459.2020.1736081>.

Benavides Cesar, L., Amaro e Silva, R., Manso Callejo, M.A., Cira, C.I., 2022. Review on spatio-temporal solar forecasting methods driven by in situ measurements or their combination with satellite and numerical weather prediction (NWP) estimates. *Engines* 15 (12), 4341. <http://dx.doi.org/10.3390/en15124341>.

Byron, R.P., 1978. The estimation of large social account matrices. *J. R. Statist. Soc. Ser. A* 141 (3), 359–367. <http://dx.doi.org/10.2307/2344807>.

Dangerfield, B., Morris, J., 1992. Top-down or bottom-up: Aggregate versus disaggregate extrapolations. *Int. J. Forecast.* 6 (2), 233–241. [http://dx.doi.org/10.1016/0169-2070\(92\)90121-O](http://dx.doi.org/10.1016/0169-2070(92)90121-O).

Davis, T.A., 2006. *Direct Methods for Sparse Linear Systems*. Society for Industrial and Applied Mathematics.

Di Fonzo, T., Girolimetto, D., 2022a. Enhancements in cross-temporal forecast reconciliation, with an application to solar irradiance forecasting. <http://dx.doi.org/10.48550/arXiv.2209.07146>.

Di Fonzo, T., Girolimetto, D., 2022b. Forecast combination-based forecast reconciliation: Insights and extensions. *Int. J. Forecast.* <http://dx.doi.org/10.1016/j.ijforecast.2022.07.001>.

Di Fonzo, T., Girolimetto, D., 2023. Cross-temporal forecast reconciliation: Optimal combination method and heuristic alternatives. *Int. J. Forecast.* 39 (1), 39–57. <http://dx.doi.org/10.1016/j.ijforecast.2021.08.004>.

Dunn, D.M., Williams, W.H., Dechaine, T.L., 1976. Aggregate versus subaggregate models in local area forecasting. *J. Amer. Statist. Assoc.* 71 (353), 68–71. <http://dx.doi.org/10.1080/01621459.1976.10481478>.

van Erven, T., Cugliari, J., 2015. Game-theoretically optimal reconciliation of contemporaneous hierarchical time series forecasts. In: Antoniadis, A., Poggi, J.-M., Brossat, X. (Eds.), *Modeling and Stochastic Learning for Forecasting in High Dimensions*. Vol. 217, Springer Science and Business Media, LLC, pp. 297–317. http://dx.doi.org/10.1007/978-3-319-18732-7_15.

EurObserv’ER, 2022. The State of Renewable Energies in Europe. 20th EurObserv’ER Report. Edition 2021, URL: <https://www.eurobserv-er.org/pdf/20th-annual-overview-barometer/>. (visited 8 September 2022).

European Commission, 2019. Going climate-neutral by 2050: a strategic long-term vision for a prosperous, modern, competitive and climate-neutral EU economy. Dir.-Gen. Clim. Action <http://dx.doi.org/10.2834/02074>.

Girolimetto, D., Di Fonzo, T., 2022. Package FoReco: point forecast reconciliation. Version 0.2.5. URL: <https://cran.r-project.org/package=FoReco>.

Hyndman, R.J., Ahmed, R.A., Athanasopoulos, G., Shang, H.L., 2011. Optimal combination forecasts for hierarchical time series. *Comput. Statist. Data Anal.* 55 (9), 2579–2589. <http://dx.doi.org/10.1016/j.csda.2011.03.006>.

Hyndman, R.J., Athanasopoulos, G., Bergmeir, C., Caceres, G., Chhay, L., O’Hara-Wild, M., Petropoulos, F., Razbash, S., Wang, E., Yatsmeen, F., 2021. Package forecast: Forecasting functions for time series and linear models. Version 8.15. URL: <https://pkg.robjhyndman.com/forecast/>.

Hyndman, R.J., Lee, A.J., Wang, E., 2016. Fast computation of reconciled forecasts for hierarchical and grouped time series. *Comput. Statist. Data Anal.* 97, 16–32. <http://dx.doi.org/10.1016/j.csda.2015.11.007>.

Jeon, J., Panagiotelis, A., Petropoulos, F., 2019. Probabilistic forecast reconciliation with applications to wind power and electric load. *European J. Oper. Res.* 279 (2), 364–379. <http://dx.doi.org/10.1016/j.ejor.2019.05.020>.

Kleissl, J., 2013. *Solar Energy Forecasting and Resource Assessment*. Academic Press.

Koning, A.J., Franses, P.H., Hibon, M., Stekler, H.O., 2005. The M3 competition: Statistical tests of the results. *Int. J. Forecast.* 21 (3), 397–409. <http://dx.doi.org/10.1016/j.ijforecast.2004.10.003>.

Kourentzes, N., Athanasopoulos, G., 2019. Cross-temporal coherent forecasts for Australian tourism. *Ann. Tourism Res.* 75, 393–409. <http://dx.doi.org/10.1016/j.annals.2019.02.001>.

Ledoit, O., Wolf, M., 2004. A well-conditioned estimator for large-dimensional covariance matrices. *J. Multivariate Anal.* 88 (2), 365–411. [http://dx.doi.org/10.1016/S0047-259X\(03\)00096-4](http://dx.doi.org/10.1016/S0047-259X(03)00096-4).

Makarov, Y.V., Etingov, P.V., Ma, J., Huang, Z., Subbarao, K., 2011. Incorporating uncertainty of wind power generation forecast into power system operation, dispatch, and unit commitment procedures. *IEEE Trans. Sustain. Energy* 4 (1), 433–442. <http://dx.doi.org/10.1109/TSTE.2011.2159254>.

Makridakis, S., Spiliotis, E., Assimakopoulos, V., 2022. The M5 accuracy competition: Results, findings and conclusions. *Int. J. Forecast.* 38 (4), 1346–1364. <http://dx.doi.org/10.1016/j.ijforecast.2021.11.013>.

Nystrup, P., Lindstrom, E., Pinson, P., Madsen, H., 2020. Temporal hierarchies with autocorrelation for load forecasting. *European J. Oper. Res.* 280 (3), 876–888. <http://dx.doi.org/10.1016/j.ejor.2019.07.061>.

Paige, C.C., Sanders, M.A., 1982. LSQR: an algorithm for sparse linear equations and sparse least squares. *ACM Trans. Math. Software* 8 (1), 43–71. <http://dx.doi.org/10.1145/355984.355989>.

Panagiotelis, A., Athanasopoulos, G., Gamakumara, P., Hyndman, R.J., 2021. Forecast reconciliation: A geometric view with new insights on bias correction. *Int. J. Forecast.* 37 (1), 343–359. <http://dx.doi.org/10.1016/j.ijforecast.2020.06.004>.

Panagiotelis, A., Gamakumara, P., Athanasopoulos, G., Hyndman, R.J., 2022. Probabilistic forecast reconciliation: properties, evaluation and score optimisation. *European J. Oper. Res.* <http://dx.doi.org/10.1016/j.ejor.2022.07.040>.

Panamtash, H., Zhou, Q., 2018. Coherent probabilistic solar power forecasting. In: 2018 IEEE International Conference on Probabilistic Methods Applied to Power Systems. PMAPS, <http://dx.doi.org/10.1109/PMAPS.2018.8440483>.

Punia, S., Singh, S.P., Madaan, J.K., 2020. A cross-temporal hierarchical framework and deep learning for supply chain forecasting. *Comput. Ind. Eng.* 149, 106796. <http://dx.doi.org/10.1016/j.cie.2020.106796>.

- Sengupta, M., Habte, A., Gueymard, C., Wilbert, S., Renné, D., Stoffel, T., 2017. Best Practices Handbook for the Collection and Use of Solar Resource Data for Solar Energy Applications: Second Edition. Technical Report NREL/TP-5D00-68886, NREL, <http://dx.doi.org/10.2172/1411856>.
- Spiliotis, E., Petropoulos, F., Kourentzes, N., Assimakopoulos, V., 2020. Cross-temporal aggregation: Improving the forecast accuracy of hierarchical electricity consumption. *Appl. Energy* 261, 114339. <http://dx.doi.org/10.1016/j.apenergy.2019.114339>.
- Stellato, B., Banjac, Goran, P., Bemporad, A., Boyd, S., 2020. OSQP: an operator splitting solver for quadratic programs. *Math. Program. Comput.* 12 (4), 637–672. <http://dx.doi.org/10.1007/s12532-020-00179-2>.
- Stratigakos, A., van der Meer, D., Camal, S., Kariniotakis, G., 2022. End-to-end learning for hierarchical forecasting of renewable energy production with missing values. URL: <https://hal.archives-ouvertes.fr/hal-03527644v2>.
- United Nations, 2015. Paris agreement. URL: <https://unfccc.int/process-and-meetings/the-paris-agreement/the-paris-agreement>. (visited 8 September 2022).
- United Nations, 2022. The race to zero emissions accelerates in Asia. UN Climate Change News, 27 April 2022, URL: <https://unfccc.int/news/the-race-to-zero-emissions-accelerates-in-asia>. (visited 8 September 2022).
- Wickramasuriya, S.L., Athanasopoulos, G., Hyndman, R.J., 2019. Optimal forecast reconciliation for hierarchical and grouped time series through trace minimization. *J. Amer. Statist. Assoc.* 114 (526), 804–819. <http://dx.doi.org/10.1080/01621459.2018.1448825>.
- Wickramasuriya, S.L., Turlach, B.A., Hyndman, R.J., 2020. Optimal non-negative forecast reconciliation. *Stat. Comput.* 30 (5), 1167–1182. <http://dx.doi.org/10.1007/s11222-020-09930-0>.
- Yagli, G.M., Yang, D., Srinivasan, D., 2019. Reconciling solar forecasts: Sequential reconciliation. *Sol. Energy* 179 (December 2018), 391–397. <http://dx.doi.org/10.1016/j.solener.2018.12.075>.
- Yagli, G.M., Yang, D., Srinivasan, D., 2020. Reconciling solar forecasts: Probabilistic forecasting with homoscedastic Gaussian errors on a geographical hierarchy. *Sol. Energy* 210, 59–67. <http://dx.doi.org/10.1016/j.solener.2020.06.005>.
- Yang, D., 2019. Standard of reference in operational day-ahead deterministic solar forecasting. *J. Renew. Sustain. Energy* 11, 053702. <http://dx.doi.org/10.1063/1.5114985>.
- Yang, D., 2020. Reconciling solar forecasts: Probabilistic forecast reconciliation in a nonparametric framework. *Sol. Energy* 210, 49–58. <http://dx.doi.org/10.1016/j.solener.2020.03.095>.
- Yang, D., Alessandrini, S., Antonanzas, J., Antonanzas-Torres, F., Badescu, V., Beyer, H., Blaga, R., Boland, J., Bright, J., Coimbra, C., David, M., Frimane, A., Gueymard, C., Hong, T., Kay, M., Killinger, S., Kleissl, J., Lauret, P., Lorenz, E., van der Meer, D., Pulescu, M., Perez, R., Perpiñán-Lamigueiro, O., Peters, I., Reikard, G., Renné, D., Saint-Drenan, Y., Shuai, Y., Urraca, R., Verbois, H., Vignola, F., Voyant, C., Zhang, J., 2020. Verification of deterministic solar forecasts. *Sol. Energy* 210, 20–37. <http://dx.doi.org/10.1016/j.solener.2020.04.019>.
- Yang, D., Quan, H., Disfani, V.R., Liu, L., 2017a. Reconciling solar forecasts: Geographical hierarchy. *Sol. Energy* 146 (April), 276–286. <http://dx.doi.org/10.1016/j.solener.2017.02.010>.
- Yang, D., Quan, H., Disfani, V.R., Rodríguez-Gallegos, C.D., 2017b. Reconciling solar forecasts: Temporal hierarchy. *Sol. Energy* 158 (June), 332–346. <http://dx.doi.org/10.1016/j.solener.2017.09.055>.
- Yang, D., Wang, W., Xia, X., 2022. A concise overview on solar resource assessment and forecasting. *Adv. Atmospheric Sci.* <http://dx.doi.org/10.1007/s00376-021-1372-8>.