

# Stealing accuracy: Predicting day-ahead electricity prices with Temporal Hierarchy Forecasting (THieF)

Arkadiusz Lipiecki<sup>1</sup>

Kaja Bilińska<sup>1</sup>

Nikolaos Kourentzes<sup>2</sup>

Rafał Weron<sup>3</sup>

<sup>1</sup> Department of Computational Social Science,  
Wrocław University of Science and Technology, Poland

<sup>2</sup> Skövde AI Lab, University of Skövde, 541-28 Skövde, Sweden

<sup>3</sup> Department of Operations Research and Business Intelligence,  
Wrocław University of Science and Technology, Poland

# Stealing Accuracy: Predicting Day-ahead Electricity Prices with Temporal Hierarchy Forecasting (THieF)

Arkadiusz Lipiecki, Kaja Bilińska, Nikolaos Kourentzes and Rafał Weron

**Abstract**—We introduce the concept of *temporal hierarchy forecasting* (THieF) in predicting day-ahead electricity prices and show that reconciling forecasts for hourly products, 2- to 12-hour blocks, and baseload contracts significantly (up to 13%) improves accuracy at all levels. These results remain consistent throughout a challenging 4-year test period (2021-2024) in the German power market and across model architectures, including linear regression, a shallow neural network, gradient boosting, and a state-of-the-art transformer. Given that (i) trading of block products is becoming more common and (ii) the computational cost of reconciliation is comparable to that of predicting hourly prices alone, we recommend using it in daily forecasting practice.

**Index Terms**—Electricity price, temporal hierarchy forecasting (THieF), forecast reconciliation, regression, machine learning

## I. INTRODUCTION

OPERATIONAL decisions often require tailored short-term forecasts that focus on different levels of detail and granularity [1]. For instance, models for hourly products in wholesale electricity markets can use different information sets than those for baseload prices [2]. These forecasts may not align, which can lead to suboptimal decisions. To cope with this, the forecasts from each temporal level of the hierarchy should be reconciled to be coherent.

The last decade has seen an unprecedented growth in interest in forecast reconciliation [3] and the introduction of *temporal hierarchy forecasting* (THieF) [4]. The latter can be applied to any time series by means of non-overlapping temporal aggregation – the predictions computed at all levels of the hierarchy are combined to yield temporally reconciled, accurate and robust forecasts. The concept is new in energy forecasting and there are only a handful of publications on predicting electric load [5], [6] or wind [7]–[9] and solar [10], [11] generation. More importantly, temporal hierarchies have not yet been applied to *electricity price forecasting* (EPF). The only related study predicts aggregated demand and supply curves by exploiting their intrinsic hierarchical structure [12], but is not concerned with THieF nor EPF.

To fill this gap and provide market participants with a universal tool to enhance predictions of hourly, block, and baseload prices, we conduct an extensive study involving

The study was partially supported by the National Science Center (NCN, Poland) through grant no. 2018/30/A/HS4/00444.

AL and KB are with the Department of Computational Social Science and RW is with the Department of Operations Research and Business Intelligence, Wrocław University of Science and Technology, 50-370 Wrocław, Poland. NK is with the Skövde AI Lab, University of Skövde, 541-28 Skövde, Sweden. E-mails: {arkadiusz.lipiecki, kaja.bilinska, rafal.weron}@pwr.edu.pl; nikolaos@kourentzes.com.

four classes of models and a 4-year test period from one of the largest power markets in Europe (Germany, 2021–2024). We find that temporal reconciliation leads to significant improvements in terms of the mean absolute error (MAE) and the root mean squared error (RMSE) across all considered model classes and hierarchy levels: from 1.2% to 5.5% for hourly and from 2.0 to 13.4% for baseload prices.

## II. TEMPORAL HIERARCHIES AND RECONCILIATION

THieF can be understood in three steps: (i) construct non-overlapping temporally aggregate time series (levels), (ii) generate base forecasts at these levels independently, (iii) reconcile the base forecasts to combine the diverse information [4]. Temporal aggregation is a moving average that filters and strengthens different aspects of the original signal, therefore the combination of the base forecasts helps recover information that would otherwise be difficult to estimate [13].

The summing matrix for a hierarchy of hourly prices and block prices is given by  $\mathbf{S} = [\mathbf{S}_b' \ \mathbf{I}_{24}']'$ , where:

$$\mathbf{S}_b = \begin{bmatrix} \frac{1}{24} & \dots & & & \frac{1}{24} \\ 0 & \dots & 0 & \frac{1}{12} & \dots & \frac{1}{12} \\ \frac{1}{12} & \dots & \frac{1}{12} & 0 & \dots & 0 \\ & & & & \vdots & \\ 0 & 0 & \frac{1}{2} & \frac{1}{2} & 0 & \dots & 0 \\ \frac{1}{2} & \frac{1}{2} & 0 & \dots & & & 0 \end{bmatrix} \quad (36,24)$$

and  $\mathbf{A}'$  is the transpose of  $\mathbf{A}$ . Let  $\mathbf{p}_d$  be a column vector that contains the hourly observations for day  $d$ , then  $\mathbf{P}_d = \mathbf{S}\mathbf{p}_d$  implements step (i). Note that our formulation differs from [4] by using the mean for the aggregation, as in [13].

Further, let  $\tilde{\mathbf{P}}_d$  be a vector of unreconciled base forecasts → step (ii). Following [4], THieF reconciliation is performed by computing  $\tilde{\mathbf{P}}_d = \mathbf{S}(\mathbf{S}^T \mathbf{W}^{-1} \mathbf{S})^{-1} \mathbf{S}^T \mathbf{W}^{-1} \tilde{\mathbf{P}}_d$ , where  $\mathbf{W}$  is the covariance matrix of the base forecast errors, estimated by shrinking the training sample errors [14] → step (iii). The above reformulation of  $\mathbf{S}_b$  only normalizes  $\mathbf{W}$  and does not impact  $\tilde{\mathbf{P}}_d$ . The reconciliation is updated daily.

Figure 1 illustrates the benefits of THieF. It plots the out-of-sample *extreme Gradient Boosting* (XGB) base (→ dashed orange) and reconciled (→ solid blue) forecasts at four aggregation levels for 07.02.2021. THieF blends the information of the base forecasts, and the relatively low bias of the 24-hour base forecast helps to reduce the bias at all other levels.

## III. COMPUTING BASE FORECASTS

To demonstrate the versatility of our approach, we generate base forecasts using four models with four distinct architectures: linear regression, a shallow neural network, gradient

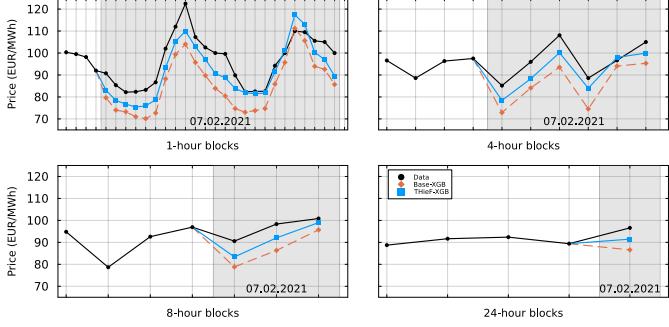


Fig. 1. Stylized example of the impact of THieF for the XBG model on one test day (07.02.2021) and four levels of the hierarchy (1H, 4H, 8H, 24H).

boosting, and a transformer. The first two are often used as benchmarks in EPF [15]–[19]. The third has been reported to perform well in numerous forecasting competitions, including M5 [20]. The fourth uses state-of-the-art transformer architecture, which is still rare in EPF [21]. It is AutoGluon’s new tabular foundation model called *Mitra*, which excels on datasets with less than 5,000 samples and 100 features [22].

All four models compute the electricity price forecast  $\hat{p}_{d,h}$  for day  $d$  and block  $h$  as a function of 20 features, as in [19]:

$$\hat{p}_{d,h} = f(p_{d-1,h}, \dots, p_{d-7,h}, p_{d-1}^{\min}, p_{d-1}^{\max}, \hat{L}_{d,h}, \hat{W}_{d,h}, \text{API}_{d-2}, \text{TTF}_{d-2}, D_d^{(1)}, \dots, D_d^{(7)}), \quad (2)$$

where  $p_{d-i,h}$  are the lagged prices of the same block in the last seven days  $i = 1, \dots, 7$ ,  $p_{d-1}^{\min}$  and  $p_{d-1}^{\max}$  are the minimum and maximum hourly prices of the previous day,  $\hat{L}_{d,h}$  and  $\hat{W}_{d,h}$  are the day-ahead load and wind generation forecasts for the target block (source: ENTSO-E Transparency),  $\text{API}_{d-2}$  and  $\text{TTF}_{d-2}$  are the closing prices of the nearest to delivery monthly coal (API2) and yearly natural gas futures (TTF) from day  $d-2$  (source: Investing.com), and  $D_d^{(i)}$  are the weekday dummies. All models are trained independently for each block  $h$  using a 3-year window of past values; each day the window is rolled forward by one day. Overall, we consider 60 blocks:  $24 \times 1\text{H}$  (i.e., 24 one-hour blocks),  $12 \times 2\text{H}$ ,  $8 \times 3\text{H}$ ,  $6 \times 4\text{H}$ ,  $4 \times 6\text{H}$ ,  $3 \times 8\text{H}$ ,  $2 \times 12\text{H}$ , and  $1 \times 24\text{H}$  (i.e., baseload).

For more robust parameter estimation, following [15], before training the models, we preprocess the inputs using the area hyperbolic sine transformation:  $\text{asinh}((y - \hat{\mu}_y)/\hat{\sigma}_y)$ , where  $\hat{\mu}_y$  and  $\hat{\sigma}_y$  are the sample mean and sample standard deviation of  $y$  estimated on the training set. More precisely,  $p_{d-i,h}$  for all  $i = 1, \dots, 7$  and each selected block  $h$  are transformed with a common  $\hat{\mu}_y$  and  $\hat{\sigma}_y$  calculated for the vector  $[p_{d-1092,h}, \dots, p_{d-1,h}]$ . The extremes  $p_{d-1}^{\min}$  and  $p_{d-1}^{\max}$  as well as the exogenous variables  $\hat{L}_{d,h}$ ,  $\hat{W}_{d,h}$ ,  $\text{API}_{d-2}$  and  $\text{TTF}_{d-2}$  are transformed independently, also using 3-year vectors of past values; weekday dummies are not transformed.

1) *AutoRegression with eXogenous inputs (ARX)*: This expert – in the sense of [23] – model is estimated via ordinary least squares (OLS) and uses the same inputs as [19] to allow for direct comparisons; the dataset is identical.

2) *Nonlinear ARX (NARX)*: The nonlinear counterpart of ARX approximates  $f(\cdot)$  in Eq. (2) with a shallow feedforward

neural network using the series-parallel architecture [24]. As in [16], the hidden layer consists of 5 neurons and uses hyperbolic tangent activation, with a linear function in the output layer. The weights are calculated in Matlab R2025a using the Levenberg-Marquadt algorithm with early stopping based on a 10% validation set. To mitigate the uncertainty of parameter estimation, the output is clipped to  $[-3, 3]$  before applying the inverse transformation, i.e., the hyperbolic sine, and the final  $\hat{p}_{d,h}$  is obtained by training the network 10 times for each  $d$  and  $h$ , and averaging the 10 price forecasts; no hyperparameter optimization is performed.

3) *Extreme Gradient Boosting (XGB)*: The third model is an ensemble of gradient-boosted decision trees (GBDT) [25]. We use eXtreme GBDT implemented in the XGBoost (v1.7) Python package [26], with MSE as the loss function for training and hyperparameter optimization. The latter is carried out 10 times at the beginning of each year (independently for each block  $h$ ) using 3 years of past data and Bayesian sequential optimization [27] with early stopping for the number of trees ( $\leq 1000$ ).  $\hat{p}_{d,h}$  is obtained by averaging the forecasts generated using these 10 sets of hyperparameters.

4) *Mitra*: The final forecaster is Amazon’s state-of-the-art tabular foundation model released with AutoGluon v1.4 in July 2025, based on a 12-layer 72 million-parameter transformer architecture [22]. Since it is pre-trained on purely synthetic data, evaluating its performance on historic time series does not pose data contamination issues. Although *Mitra* is not specifically tailored to time series, our forecasting task requires only one-step-ahead predictions, with Eq. (2) being equivalent to a tabular regression problem. We generate forecasts in a *zero-shot* mode, relying on a model pre-trained for regression tasks. At each timestep, the training set is fed as support examples for in-context learning. *Mitra* then generates a single prediction based on the input variables acting as a query.

#### IV. EMPIRICAL RESULTS

To ensure a sound assessment of the THieF approach, we consider one of the largest power markets in Europe (EPEX-DE, Germany) and a dataset that spans 7 years (05.01.2018–31.12.2024); the same as in [19]. The first 1092 days (until 31.12.2020) are the initial training window; each day it is rolled forward by 24 hours. The remaining 4-year period (starting 01.01.2021) is a challenging test set that includes the COVID-19 pandemic (the first day of hard lockdown in Germany was 15.12.2020), the Russian invasion of Ukraine (24.02.2022) and the soaring natural gas prices (Q4 2021–Q4 2022), as well as the appearance of negative price spikes due to low demand and high renewable generation (e.g., on Sunday 02.07.2023 at 3 p.m. the price dropped to  $-500$  EUR/MWh).

In Table I we report the mean absolute errors (MAE) and the root mean squared errors (RMSE) of the four models over the entire test period, as well as the respective gains from forecast reconciliation (in %). The latter are all significant at the 5%

<sup>1</sup>The search ranges are:  $\text{max\_depth} \in [2, 10]$ ,  $\text{learning\_rate} \in [0.0001, 1.0]$  (log scale),  $\text{subsample} \in [0.5, 1.0]$ ,  $\text{min\_child\_weight} \in [0, 10]$ ,  $\gamma \in [0.0, 0.5]$ ,  $\lambda \in [0.001, 10.0]$  (log scale), and  $\alpha \in [0.001, 10.0]$  (log scale). The remaining parameters use default values.

TABLE I

THE MAE AND RMSE ERRORS OF THE BASE FORECASTS FOR SELECTED HIERARCHY LEVELS (1H, 4H, 8H, 24H) OVER THE 4-YEAR TEST PERIOD AND THE RESPECTIVE GAINS FROM RECONCILIATION (IN %).

Level	Model	MAE	%gain	RMSE	%gain
1H	ARX	26.93	2.8%	39.83	2.7%
	NARX	23.06	3.3%	35.62	2.8%
	XGB	23.34	4.4%	37.05	5.5%
	Mitra	21.37	1.2%	33.31	1.2%
4H	ARX	25.65	3.1%	37.50	3.1%
	NARX	21.75	5.0%	33.27	4.6%
	XGB	22.01	5.8%	34.25	6.0%
	Mitra	20.13	3.7%	30.87	3.5%
8H	ARX	24.37	2.9%	35.59	3.5%
	NARX	20.54	5.5%	31.20	4.8%
	XGB	21.07	6.8%	32.94	8.2%
	Mitra	19.45	6.3%	29.34	5.5%
24H	ARX	20.92	2.0%	30.29	3.4%
	NARX	17.99	8.2%	27.24	7.9%
	XGB	18.80	11.7%	29.50	13.4%
	Mitra	16.56	8.1%	24.81	8.5%

level, as measured by the multivariate variant of the Diebold-Mariano test [23]; this is true not only for the four hierarchy levels (1H, 4H, 8H, 24H) in Table I, but for all considered levels, see Sec. III. The fact that the results are consistent across architectures ranging from parsimonious linear regression with 20 parameters to a state-of-the-art transformer with up to 72 million parameters is a strong argument in favor of using the THieF approach in day-ahead EPF.

## V. CONCLUSIONS

We demonstrate the benefit of *temporal hierarchy forecasting* (THieF) for predicting electricity prices. We evidence this in the German power market, using a variety of base forecasting methods with increasing complexity and diverse modeling assumptions. We found a consistent improvement in all setups (up to 13%; on average above 5% for both MAE and RMSE), with increasing benefits for longer forecast horizons.

There are various extensions that stem from this work. We applied the same base forecast to all aggregation levels, however, more diverse model selection can be beneficial, better leveraging the properties of each level, and thus minimizing computational overheads. Likewise, investigating the impact of THieF on probabilistic EPF is a useful extension. Beyond any accuracy gains, THieF provides reconciled forecasts across all levels, which can support the alignment of different decisions and potentially new strategies in energy markets. These are fruitful avenues for future research.

## REFERENCES

- [1] F. Petropoulos, D. Apiletti, V. Assimakopoulos, M. Babai, D. Barrow *et al.*, “Forecasting: Theory and practice,” *International Journal of Forecasting*, vol. 38, no. 3, pp. 705–781, 2022.
- [2] K. Maciejowska and R. Weron, “Short- and mid-term forecasting of baseload electricity prices in the U.K.: The impact of intra-day price relationships and market fundamentals,” *IEEE Transactions on Power Systems*, vol. 31, no. 2, pp. 994–1005, 2016.
- [3] G. Athanasopoulos, R. J. Hyndman, N. Kourentzes, and A. Panagiotelis, “Forecast reconciliation: A review,” *International Journal of Forecasting*, vol. 40, no. 2, pp. 430–456, 2024.
- [4] G. Athanasopoulos, R. J. Hyndman, N. Kourentzes, and F. Petropoulos, “Forecasting with temporal hierarchies,” *European Journal of Operational Research*, vol. 262, no. 1, pp. 60–74, 2017.
- [5] P. Nystrup, E. Lindström, P. Pinson, and H. Madsen, “Temporal hierarchies with autocorrelation for load forecasting,” *European Journal of Operational Research*, vol. 280, no. 3, pp. 876–888, 2020.
- [6] J. K. Møller, P. Nystrup, and H. Madsen, “Likelihood-based inference in temporal hierarchies,” *International Journal of Forecasting*, vol. 40, no. 2, pp. 515–531, 2024.
- [7] J. Jeon, A. Panagiotelis, and F. Petropoulos, “Probabilistic forecast reconciliation with applications to wind power and electric load,” *European Journal of Operational Research*, vol. 279, no. 2, pp. 364–379, 2019.
- [8] L. English and M. Abolghasemi, “Improving the forecast accuracy of wind power by leveraging multiple hierarchical structure,” *Sustainable Energy, Grids and Networks*, vol. 40, p. 101517, 2024.
- [9] N. Sharma, R. Bhakar, and P. Jain, “Optimal reconciliation of hierarchical wind energy forecasts utilizing temporal correlation,” *Energy Conversion and Management*, vol. 301, p. 118053, 2024.
- [10] D. Yang, H. Quan, V. R. Disfani, and C. D. Rodríguez-Gallegos, “Reconciling solar forecasts: Temporal hierarchy,” *Solar Energy*, vol. 158, pp. 332–346, 2017.
- [11] T. Di Fonzo and D. Girolimetto, “Spatio-temporal reconciliation of solar forecasts,” *Solar Energy*, vol. 251, pp. 13–29, 2023.
- [12] P. Ghelasi and F. Ziel, “Hierarchical forecasting for aggregated curves with an application to day-ahead electricity price auctions,” *International Journal of Forecasting*, vol. 40, no. 2, pp. 581–596, 2024.
- [13] N. Kourentzes, F. Petropoulos, and J. R. Trapero, “Improving forecasting by estimating time series structural components across multiple frequencies,” *International Journal of Forecasting*, vol. 30, no. 2, pp. 291–302, 2014.
- [14] O. Ledoit and M. Wolf, “Honey, I shrunk the sample covariance matrix,” *Journal of Portfolio Management*, vol. 30, no. 4, p. 110–119, 2004.
- [15] B. Uniejewski, R. Weron, and F. Ziel, “Variance stabilizing transformations for electricity spot price forecasting,” *IEEE Transactions on Power Systems*, vol. 33, no. 2, pp. 2219–2229, 2018.
- [16] K. Hubicka, G. Marcjasz, and R. Weron, “A note on averaging day-ahead electricity price forecasts across calibration windows,” *IEEE Transactions on Sustainable Energy*, vol. 10, no. 1, pp. 321–323, 2019.
- [17] A. Billé, A. Gianfreda, F. Del Grossi, and F. Ravazzolo, “Forecasting electricity prices with expert, linear, and nonlinear models,” *International Journal of Forecasting*, vol. 39, no. 2, pp. 570–586, 2023.
- [18] C. Zhang and Y. Fu, “Probabilistic electricity price forecast with optimal prediction interval,” *IEEE Transactions on Power Systems*, vol. 39, no. 1, p. 442–452, 2024.
- [19] T. Serafin and R. Weron, “Loss functions in regression models: Impact on profits and risk in day-ahead electricity trading,” *Energy Economics*, vol. 148, p. 108596, 2025.
- [20] T. Januschowski, Y. Wang, K. Torkkola, T. Erkkilä, H. Hasson, and J. Gasthaus, “Forecasting with trees,” *International Journal of Forecasting*, vol. 38, no. 4, pp. 1473–1481, 2022.
- [21] J. Bottieau, Y. Wang, Z. De Grève, F. Vallée, and J.-F. Toubeau, “Interpretable transformer model for capturing regime switching effects of real-time electricity prices,” *IEEE Transactions on Power Systems*, vol. 38, no. 3, pp. 2162–2176, 2023.
- [22] X. Zhang and D. M. Robinson, “Mitra: Mixed synthetic priors for enhancing tabular foundation models,” Amazon Science, Tech. Rep., 2025. [Online]. Available: <https://www.amazon.science/blog/mitra-mixed-synthetic-priors-for-enhancing-tabular-foundation-models>
- [23] F. Ziel and R. Weron, “Day-ahead electricity price forecasting with high-dimensional structures: Univariate vs. multivariate modeling frameworks,” *Energy Economics*, vol. 70, pp. 396–420, 2018.
- [24] H. Xie, H. Tang, and Y.-H. Liao, “Time series prediction based on NARX neural networks: An advanced approach,” in *2009 International Conference on Machine Learning and Cybernetics*, vol. 3, 2009, pp. 1275–1279.
- [25] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer (2nd ed.), 2009.
- [26] T. Chen and C. Guestrin, “XGBoost: A scalable tree boosting system,” in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016, pp. 785–794.
- [27] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama, “Optuna: A next-generation hyperparameter optimization framework,” in *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2019, pp. 2623–2631.