

Forecasting Electricity Prices: Autoregressive Hybrid Nearest Neighbors (ARHNN) method

Weronika Nitka¹[0000-0002-3360-4260], Tomasz Serafin¹[0000-0003-0518-9917], and
Dimitrios Sotiros^{1*}[0000-0002-9995-3906]

Department of Operations Research and Business Intelligence,
Faculty of Computer Science and Management,
Wroclaw University of Science and Technology, 50-370 Wroclaw, Poland
dimitrios.sotiros@pwr.edu.pl

Abstract. The ongoing reshape of electricity markets has significantly stimulated electricity trading. Limitations in storing electricity as well as on-the-fly changes in demand and supply dynamics, have led price forecasts to be a fundamental aspect of traders' economic stability and growth. In this perspective, there is a broad literature that focuses on developing methods and techniques to forecast electricity prices. In this paper, we develop a new hybrid method, called ARHNN, for electricity price forecasting (EPF) in day-ahead markets. A well performing autoregressive model, with exogenous variables, is the main forecasting instrument in our method. Contrarily to the traditional statistical approaches, in which the calibration sample consists of the most recent and successive observations, we employ the k -nearest neighbors (k -NN) instance-based learning algorithm and we select the calibration sample based on a similarity (distance) measure over a subset of the autoregressive model's variables. The optimal levels of the k -NN parameter are identified during the validation period in a way that the forecasting error is minimized. We apply our method in the EPEX SPOT market in Germany. The results reveal a significant improvement in accuracy compared to commonly used approaches.

Keywords: Electricity price forecasting · Day-ahead market · ARX · k -nearest neighbors.

1 Introduction

Electricity markets have witnessed significant changes over the last decades. Their deregulation, followed by the emergence of electrical power exchanges such as EPEX SPOT, OMIE and Nord Pool in Europe, or PJM in the USA, allowed for competitive electricity trading [23]. Electrical power exchanges usually consist of several markets. The market with the biggest volume of trade is the day-ahead (spot) market, which allows the traders to place bids and offers the day before the

* Corresponding author; the authors' names are ordered alphabetically

physical delivery of electricity. The day-ahead market is usually supplemented by intraday and balancing markets, which allow trading until a few minutes before delivery and target at providing more accurate offers. However, this is often associated with paying significant balancing fees.

Notably, electricity market clearing prices, defined by the supply and the demand curve, are characterized by high volatility. The cost of storing electricity at large scales as well as the transition of power generation from conventional to renewable sources, permeated with uncertainty in the production levels, lead to fluctuations in the supply. On the other hand, demand may vary on an hourly (peak and off-peak hours) and daily (weekends, weekdays and festivities) basis. These factors, along with the requirement of supply and demand to be precisely balanced in the power grid, lead to highly volatile prices in the electricity markets which can undergo extreme changes within a span of a single day.

Traders, ideally, aim to maximize their profit as well as to minimize the financial risk by selecting the most appropriate strategy in an imperfect market, where there is incomplete information. Given the high level of price volatility and the limitations in storing electricity, the selection of a wrong strategy, based on price misinformation, may lead to economic losses or even bankruptcy. On the contrary, utilizing accurate forecasts may increase profits or reduce the risk of economic losses [9, 14].

In this line of thought, there is a wide literature which focuses on providing accurate day-ahead electricity price forecasts. Extended literature reviews are provided in [1, 21, 23]. Two of the prominent broad classes of methods provided in the literature rely either on statistical approaches or on machine learning techniques. Statistical approaches utilize linear regression models or linear autoregressive models based on a set of variables related to observed prices and other exogenous variables (load, wind, solar, temperature) that may affect price levels. Differences in the implementation of the autoregressive models can be also identified in terms of the calibration window length, which can be predefined or estimated via more advanced econometric techniques [4, 11, 12]. However, in these cases, the calibration sample consists of the most recent and successive observations. Methods that rely on machine learning employ a variety of techniques such as artificial neural networks [24], support vector machines [27], clustering algorithms [22] or a combination of them [17]. It is worthy to mention that a hybrid approach that employs statistical and machine learning techniques has been also proposed in the literature. Specifically, in [18] three clustering algorithms and an autoregressive lag model were employed to predict consumers' energy consumption in a simulation suite. However, this approach was tested on simulated data.

In this paper we build on the bridge between the two aforementioned classes of methods and we propose a new hybrid method, called *autoregressive hybrid nearest neighbors* (ARHNN), for forecasting spot electricity prices. We generate one-day-ahead forecasts using a linear ARX (autoregressive with exogenous variables) model with parameters calibrated on samples selected with the k -nearest neighbors (k -NN) algorithm. ARX models are well-established in electricity price

forecasting (EPF), as noted in [10, 23]. The k -nearest neighbors algorithm has been found to be successful in the field of electricity market forecasts, mainly in forecasting electricity price and load [2, 3, 8, 13, 19, 26] and renewable energy sources (RES) generation [25]. The proposed method is applied to the EPEX SPOT market in Germany. The results show a significant improvement in accuracy compared to commonly used benchmark approaches, while low increase in the computational load is ensured.

The rest of this paper unfolds as follows. Section 2 describes the most important features of the data used in this analysis. Section 3 provides an in-depth explanation of the proposed method. Section 4 illustrates the results of the proposed method applied to the EPEX SPOT data and provides comparison with commonly used benchmark models. Finally, conclusions are drawn in Section 5.

2 Data

To illustrate our method, we use data describing the day-ahead electricity prices in the EPEX SPOT market in Germany. As described in the Introduction, the day-ahead market is the most important market in terms of traded volume. The dataset, published by the transmission system operator (TSO), comprises four variables: the electricity price in EUR/MWh and the corresponding official TSO forecasts of total electrical load, wind energy generation and photovoltaic energy generation, expressed in GWh.

The dataset spans six full years, from January 2015 until December 2020, with hourly data (see Figure 1). To evaluate the performance of the proposed algorithm, the data is divided into three periods with lengths of approximately two years each. The first 728-day period is reserved for the initial calibration window. Then, the middle period, of the same length, is utilized for validation and tuning the hyperparameters of the model as described in Subsection 3.2. Finally, the procedure is tested on the last period with length equal to 736 days.

The time series of the price and the load forecasts, as well as the division into calibration, validation and testing periods, are depicted in Figure 1. It can be seen that the spot prices are indeed highly volatile, with frequent upward and downward spikes multiple times greater in magnitude than the average price range. However, load is relatively predictable, exhibiting both weekly and yearly seasonality, which needs to be addressed by the predictive model.

3 Methods and algorithms

As shown by numerous studies in the EPF [7, 14], the selection of the calibration sample impacts the overall forecasting accuracy of the autoregressive model. While the majority of authors consider the longest possible portion of data for the model calibration, averaging predictions obtained from calibration samples of different lengths [16] or utilizing more sophisticated statistical methods [15] allows for the significant reduction of forecasting errors. In this paper, we propose a new method for the selection of the calibration sample, based on the k -nearest




Fig. 1: Time series plot of the electricity spot prices (upper panel) and TSO load forecast (lower panel) from the EPEX SPOT market. Dashed lines indicate the split into calibration, validation and testing periods.

neighbors algorithm. The aforementioned methods rely on the time dimension to select the calibration sample, i.e. the most recent successive observations compose the calibration sample. On the contrary, in our method we define the calibration sample on the basis of a similarity measure over a set of features.

3.1 Predictive model

To predict the spot prices in hour h of day $d + 1$ we use an expert ARX model with a specification well-established in the electricity price forecasting literature [7, 20]. Due to the idiosyncratic nature of the electricity market, every hour of the day is treated as a distinct market product and separate forecasts are implemented for each hour, i.e. predicting the prices for the entire day $d + 1$ requires estimating 24 independent parameter sets. The models for every hour have an identical specification, incorporating an autoregressive structure with lags corresponding to two preceding days and a week, notated as $P_{d+1-p,h}$ where $p \in \{1, 2, 7\}$. The price dynamics are further captured by including the minimal and the maximal price from the previous day (respectively $P_{d,min}$ and $P_{d,max}$) as well as that day's price in hour 24 ($P_{d,24}$) – the previous day's last known price.

Finally, the model incorporates the publicly available forecasts of three exogenous variables relevant to the price levels: total electrical load (\hat{L}), wind energy generation (\hat{W}) and photovoltaic energy generation (\hat{S}). The complete model takes the form

$$\begin{aligned}
 P_{d+1,h} = & \alpha_h D_{d+1} + \underbrace{\sum_{p \in \{1,2,7\}} \beta_{h,p} P_{d+1-p,h}}_{\text{AR component}} + \underbrace{\theta_{h,1} P_{d,\min} + \theta_{h,2} P_{d,\max}}_{\text{Daily statistics}} \\
 & + \underbrace{\theta_{h,3} P_{d,24}}_{\text{Last known price}} + \underbrace{\theta_{h,4} \hat{L}_{d+1,h} + \theta_{h,5} \hat{W}_{d+1,h} + \theta_{h,6} \hat{S}_{d+1,h}}_{\text{Exogenous variables}} + \varepsilon_{d+1,h},
 \end{aligned} \tag{1}$$

where D_{d+1} is the 1×7 vector of dummy variables representing days of the week and $\varepsilon_{d+1,h}$ is the Gaussian white noise. Henceforth, by $\hat{P}_{d+1,h}(\tau)$ we denote the prediction obtained from model (1) calibrated on a sample containing the τ most recent observations.

3.2 ARHNN calibration sample selection

The k -Nearest Neighbors is an instance-based learning algorithm that can be used either for classification or regression. In the former case, an observation is assigned to the most common class label shared by its k -nearest neighbors. In the second case, the property value for an observation derives from the average of the k -nearest neighbors' values. In both cases, a neighbor weighting function can be employed [6].

To explain the applicability of the k -NN algorithm in our case study and the differentiation of our method, suppose that at day d we want to forecast the electricity price for the day ahead (day $d+1$). We denote by x_d the vector of the explanatory variables from model (1) for a given day d , after omitting dummy variables, 2-day and 7-day lagged prices and random noise. Within the matrix $X_{d+1} = (x_{d-726}; \dots; x_{d+1})$, it is evident that the most recent information we possess, x_{d+1} , provides the most accurate outlook at the current market state, i.e. prices from previous days as well as forecasts for day $d+1$. Notably, the proposed statistical methods in the literature, rely on this assumption and they further extend it. Specifically, they assume that the most recent observations will provide the most accurate forecast and thus, they should compose the calibration sample. However, in case structural breaks exist among the last observations, the selected calibration sample will lead to a decreased forecasting accuracy. In addition, this approach relies exclusively on the last observations (in terms of time) and does not exploit information from other past data.

The main idea of our method is to identify past observations that resemble x_{d+1} as closely as possible and use them to estimate the parameters of the forecasting model. To this end, we employ the k -NN algorithm to select a calibration sample for the ARX model (1) consisting of the k -nearest neighbors of the point x_{d+1} (see Figure 2), in terms of the Euclidean distance. In a sense, we invert the rationale of the k -NN method - instead of classifying the latest observation

based on its neighboring points, we assume that the closest neighbors (in terms of the distance, not time) of x_{d+1} belong to the same market “regime”.

Analogously to the notation in Section 3.1, we denote the price prediction for day $d + 1$ and hour h , obtained by calibrating the forecasting model (1) on the sample consisting of k closest observations, by $\hat{P}_{d+1,h}^*(k)$. Note that for the clarity of notation, forecasts corresponding to the ARHNN method are marked with an asterisk.




Fig. 2: The optimal (i.e. producing the lowest absolute prediction error) selection of the calibration sample ($\bar{k}_i = 181$) based on the matrix X_{d+1} for a specific day ($d + 1$). The upper panel illustrates the sample selection, presented on three key variables, i.e. preceding day’s price as well as forecasts of load and wind generation; while the lower panel depicts the corresponding selection in the time dimension. The most recent observation is marked with a red dot, while the observations selected for the model calibration are depicted with blue points.

Obviously, the choice of the parameter k has a direct impact on the forecasting accuracy of the model. Disentangling its effects, is one of the main challenges that we address in the paper. As discussed in Section 2, in the validation period, we use the 728-day rolling window to identify the optimal values of the k parameter, which is responsible for the number of observations in the calibration sample. For each of the 728 days in the validation period, the procedure identifies (ex-post) the optimal value (i.e. the one that produced the lowest absolute prediction error for a certain day; see Figures 2, 3) of the parameter, \bar{k}_i , $i = 1, \dots, 728$. Next, in the evaluation (testing) procedure, instead of selecting only one value of k for each day, we consider 728 calibration samples, based on the set of past optimal values $(\bar{k}_1, \dots, \bar{k}_{728})$. In such way, we obtain 728 price predictions for day $d+1$, i.e. $(\hat{P}_{d+1,h}^*(\bar{k}_1), \dots, \hat{P}_{d+1,h}^*(\bar{k}_{728}))$. Eventually, inspired by [16], we obtain the final price prediction for day $d+1$ and hour h from the average of these forecasts:

$$\hat{P}_{d+1,h} = \frac{1}{728} \sum_{i=1}^{728} \hat{P}_{d+1,h}^*(\bar{k}_i). \quad (2)$$

Notably, there may be cases where the values of \bar{k}_i , $i = 1, \dots, 728$ coincide, i.e. $\bar{k}_i = \bar{k}_j$ for $i \neq j$. Therefore, the above expression is translated to the weighted average of forecasts calibrated to different samples, where the weight corresponding to a certain prediction $\hat{P}_{d+1,h}^*(\bar{k}_i)$ depicts the relative frequency of \bar{k}_i in $(\bar{k}_1, \dots, \bar{k}_{728})$. This can be interpreted as a weighting function which reflects the “relative significance” of the \bar{k}_i values.

3.3 Benchmark approaches

We evaluate the effectiveness of selecting the calibration period with the proposed ARHNN procedure by comparing it to a number of literature benchmarks. While all of them use Model (1) for computing the forecasts themselves, they differ in the selection of the calibration sample and in the forecasts post-processing. The first group of benchmark approaches provides forecasts obtained by using a single calibration window length throughout the entire test period. The calibration windows include from 56 to 728 days of the most recent data up to the moment of forecasting. The second group utilizes two additional approaches following [7]: the arithmetic mean of all the forecasts within the first group (673 predictions obtained from calibration windows of different lengths), and the average of forecasts from six hand-picked calibration windows: three short ones (56, 84, 112 days) and three long ones (714, 721, 728 days).

We assume the following convention to notate the aforementioned benchmark methods: the single-length windows with length τ are denoted as $\text{Win}(\tau)$. The forecast averages are named using the MATLAB sequence convention, respectively becoming $\text{Av}(56:728)$ and $\text{Av}(56:28:112, 714:7:728)$.

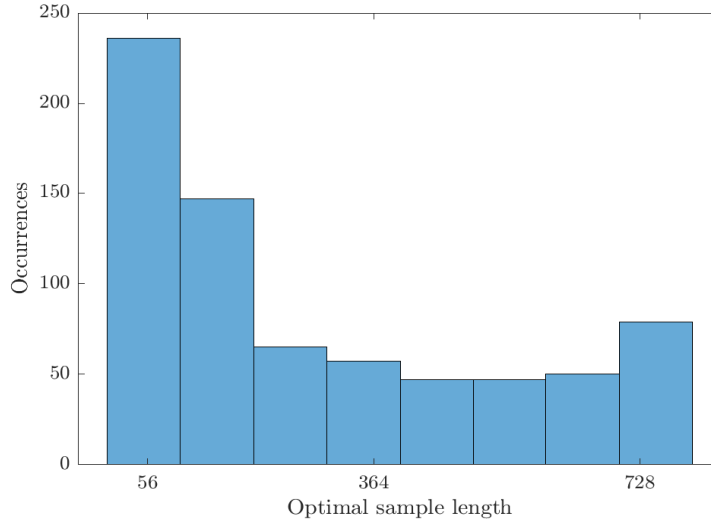


Fig. 3: Histogram of the optimal calibration sample lengths within the validation period (728 calibration sample lengths in total) for hour 18.

4 Results

We evaluate the accuracy of the forecasts obtained from different approaches with the use of the *root mean squared error* (RMSE). The reported error is calculated across all hours and days of the 736-day out-of-sample period. The results are presented in Figure 4 and Table 1. The performance of single calibration window benchmarks (i.e. models trained on samples comprising a fixed amount of most recent observations) is presented with gray dots. In this approach, although the average error generally diminishes with the increase of the calibration window length τ and the longest window turns out to be the best choice, the decrease is not monotonic as we may expect. As shown by [16], for certain datasets, the error may even increase alongside with the calibration window length.

Table 1: The RMSE values of the selected benchmarks and the ARHNN method.

Method	RMSE
Win(364)	8.4443
Win(728)	8.2860
Av(56:728)	8.0584
Av(56:28:112, 714:7:728)	8.0286
ARHNN	7.8604

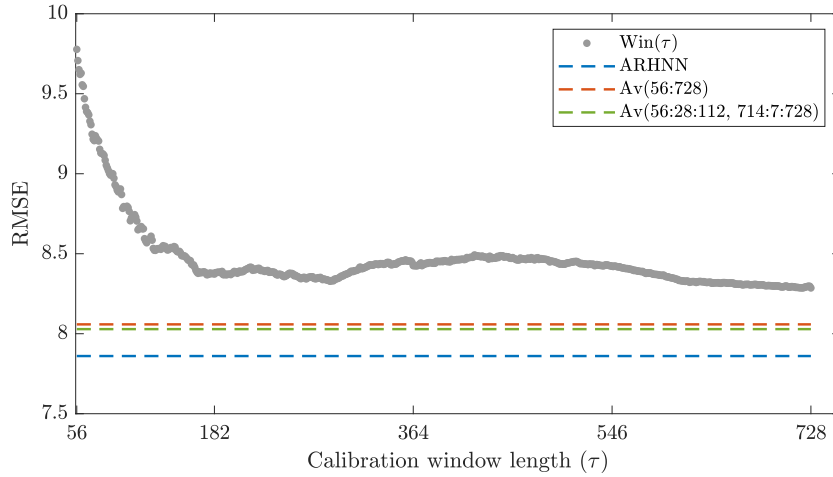


Fig. 4: The RMSE values as a function of calibration window length for the benchmark approaches and the ARHNN method.

As can be seen from Figure 4, the ARHNN method as well as the averaging schemes outperform every approach based on a single, fixed calibration window length in terms of RMSE. Methods based on forecasts averaging, Av(56:728) and Av(56:28:112, 714:7:728), managed to outperform the predictive accuracy of the longest, 728-day calibration window, approximately by 3%. The forecasts obtained from the introduced ARHNN method exhibit over 5% lower error comparing to the best performing single calibration window length. The method also gains over 2% in terms of the forecasting accuracy compared to the well-performing literature benchmarks Av(56:728) and Av(56:28:112, 714:7:728). Since these results are not sufficient for determining the statistical significance of the difference between forecasts obtained from different approaches, we decided to use the Diebold and Mariano (DM) [5] test. First, for each pair of methods X and Y , we create a vector of errors for each day of the out-of-sample period. Here we consider two different perspectives - univariate and multivariate as classified by [28]. In the first one (multivariate), we consider 24-dimensional error vectors for each day:

$$\Delta_{X,Y,d} = \|\bar{\varepsilon}_{X,d}\| - \|\bar{\varepsilon}_{Y,d}\|, \quad (3)$$

where $\bar{\varepsilon}_{X,d} = \sqrt{\frac{1}{24} \sum_{h=1}^{24} \varepsilon_{X,d,h}^2}$ and $\varepsilon_{X,d,h}$ is the error of forecasts obtained with method X for day d and hour h . In the second approach (univariate), instead of considering 24 hours jointly, we are looking at each of them separately. More precisely:

$$\Delta_{X,Y,d,h} = |\varepsilon_{X,d,h}| - |\varepsilon_{Y,d,h}|. \quad (4)$$

For each pair of approaches, we compute the p -value of the DM test with null hypothesis $H_0: \mathbb{E}(\Delta_{X,Y,d}) \leq 0$ (or $H_0: \mathbb{E}(\Delta_{X,Y,d,h}) \leq 0$ in case of the univariate

approach) and additionally perform a complementary test with the reverse null hypothesis, $H_0^R: \mathbb{E}(\Delta_{X,Y,d}) \geq 0$ (or $H_0^R: \mathbb{E}(\Delta_{X,Y,d,h}) \geq 0$).

In Figure 5 and Figure 6, we present the p -values of the test. We use a heatmap to indicate the span of p -values. The closer they are to zero (dark green), the more significant is the difference between forecasts obtained with the approach from X-axis (superior) and predictions from the method in the Y-axis (inferior) [7, 16, 15]. The “chessboard” in Figure 5 corresponds to the results of the multivariate approach, considering 24-dimensional error vectors (see Equation 3). It turns out, that forecasts from the ARHNN method were able to significantly outperform predictions from nearly all benchmarks. The well-performing averaging scheme Av(56:28:112, 714:7:728) was neither significantly worse nor better than the proposed approach. Two “chessboards” in Figure 6, correspond to the results of the univariate DM test for two exemplary hours. The selected Hour 9 and Hour 15 correspond to the worst and the best performance of the ARHNN method across all hours, respectively. For Hour 9, the forecasts based on the ARHNN approach were not able to statistically outperform predictions from any other method. Additionally, they are outperformed by the forecasts based on the Av(56:728) averaging scheme. When it comes to the results for Hour 15, the predictions from the proposed ARHNN method significantly outperform forecasts from all benchmarks, with p -values of the DM test close to zero. In general, the performance of the ARHNN approach across 24 hours of the day is shown in Table 2. The columns, corresponding to 24 hours are marked with five different colors, each of them representing a certain result of the DM test:

- **Sharp green** (Hours 2, 6, 13, 14, 15, 16, 17) - forecasts from the ARHNN method significantly outperform predictions from all benchmarks and are not outperformed by any of them,
- **Green** (Hours 1, 3, 4, 5) - forecasts from the ARHNN method significantly outperform predictions from three out of four benchmarks and are not outperformed by any of them,
- **Yellow** (Hours 7, 8, 10, 11, 12, 18) - forecasts from the ARHNN method significantly outperform predictions from two out of four benchmarks and are not outperformed by any of them,
- **Orange** (Hours 22, 23, 24) - forecasts from the ARHNN method significantly outperform predictions from two out of four benchmarks and are outperformed by one of them,
- **Sharp orange** (Hour 19) - forecasts from the ARHNN method do not significantly outperform predictions from any benchmark and are not outperformed by any of them,
- **Red** (Hours 9, 20, 21) - forecasts from the ARHNN method do not significantly outperform predictions from any benchmark and are outperformed by one of them.

Looking at the results of the Diebold-Mariano test it can be observed that forecasts from the ARHNN approach exhibit very satisfactory predictive accuracy compared to forecasts from the selected benchmarks. For eleven hours,

Table 2: Results of the statistical significance test between forecasts from the ARHNN approach and the selected benchmarks for all 24 hours. Each color represents a certain result of the DM test.

Hour	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
------	---	---	---	---	---	---	---	---	---	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----

ARHNN forecasts were able to significantly outperform predictions from at least three out of four benchmarks and, for twenty hours, at least two out of four. Although the forecasts exhibit the worst performance for hours 9, 19, 20 and 21, they were significantly outperformed by at most one benchmark approach and, in the remaining twenty one hours of the day, by none of them.

5 Conclusions and discussion

In this paper we introduced a hybrid method for electricity price forecasting in day ahead markets. We employed a linear autoregressive model, with exogenous

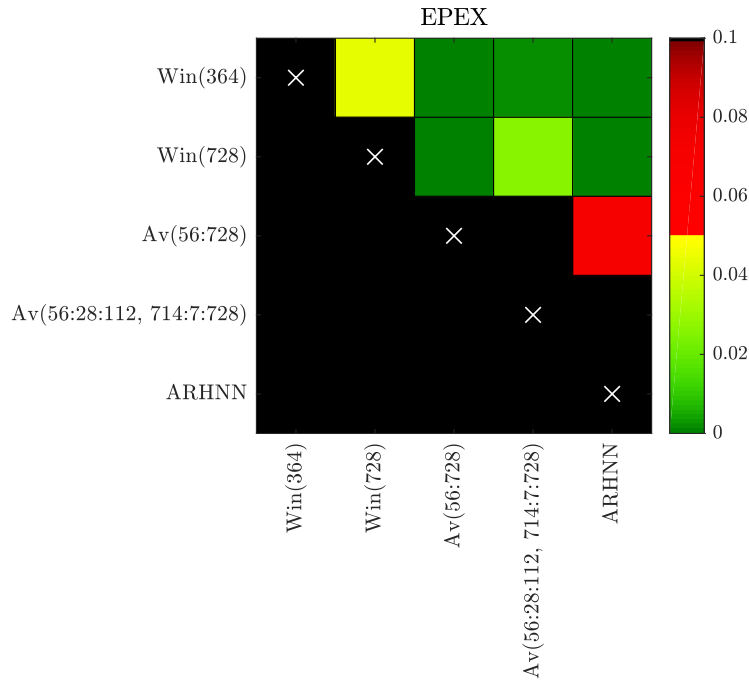


Fig. 5: Results of the multivariate approach to the pairwise Diebold-Mariano test between ARHNN method and the selected benchmarks. We illustrate the range of p -values using a heatmap: green squares indicate a statistically significant superiority of the forecasts from the method on the X-axis over the ones from the method on the Y-axis.

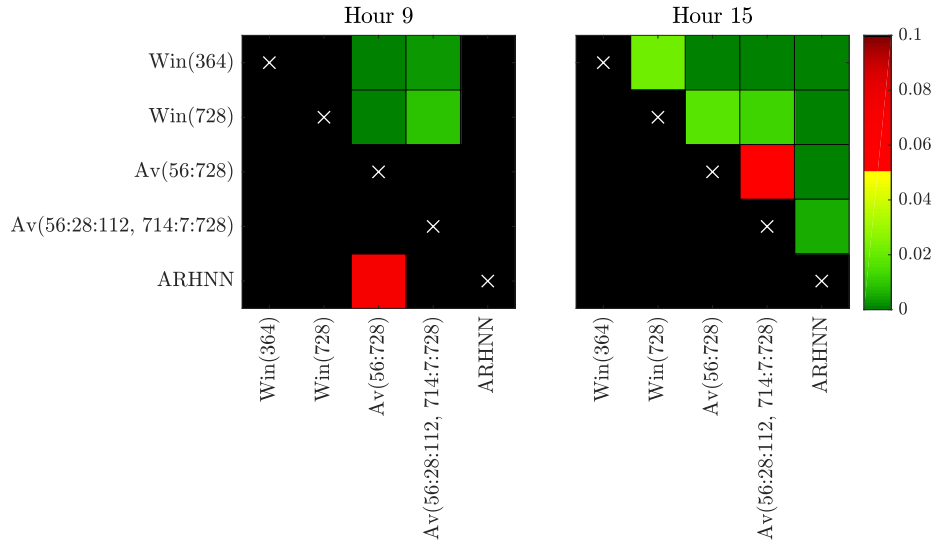


Fig. 6: Sample results of the univariate approach to the pairwise Diebold-Mariano test between ARHNN model and the selected benchmarks. We illustrate the range of p -values using a heatmap: green squares indicate a statistically significant superiority of the forecasts from the method on the X-axis over the ones from the method on the Y-axis.

variables (total electrical load, wind energy generation and photovoltaic energy generation), as the underlying instrument for forecasts. Our novelty lies in the selection of the calibration sample which is achieved via a machine learning algorithm. Specifically, we utilized the k -NN instance-based learning algorithm to select the calibration sample based on a similarity (distance) measure between the most recent information and past observations, over a subset of the autoregressive model’s variables. Our aim was to identify past observations that belong to the same “regime” with the latest available information.

The advantage of our method is therefore twofold. The selection of the calibration sample relies on a similarity measure over a set of variables rather than on the time dimension (i.e. to include only the most recent observations). With this type of selection, homogeneity within the calibration sample is secured and structural breaks are avoided. In addition, information from past observations is exploited and consequently, the selected calibration sample is expected to provide more accurate forecasts.

We applied our method on the EPEX SPOT market and we provided comparison with commonly used literature benchmarks. The results show that our proposed method achieves a statistically significant reduction in the forecasting error compared to the rest of the approaches, while remaining highly interpretable and meaningful. The accuracy of the proposed method in other mar-

kets, the adoption of other machine learning techniques as well as comparison with other methods relying exclusively on them, are subjects for future research. Nevertheless, our findings signify the importance and benefits of interdisciplinary research in this field.

Acknowledgments

This work is partially supported by the National Science Center (NCN, Poland) through MAESTRO grant no. 2018/30/A/HS4/00444 (D.S.). Also, it is partially supported by the Ministry of Science and Higher Education (MNiSW, Poland) through Diamond Grant no. 0027/DIA/2020/49 (W.N.) and Diamond Grant no. 0009/DIA/2020/49 (T.S.).

References

1. Aggarwal, S.K., Saini, L.M., Kumar, A.: Electricity price forecasting in deregulated markets: A review and evaluation. *International Journal of Electrical Power & Energy Systems* **31**(1), 13–22 (2009)
2. Ashfaq, T., Javaid, N.: Short-Term Electricity Load and Price Forecasting using Enhanced KNN. In: 2019 International Conference on Frontiers of Information Technology (FIT). pp. 266–2665 (2019)
3. Chaudhury, P., Tyagi, A., Shanmugam, P.K.: Comparison of Various Machine Learning Algorithms for Predicting Energy Price in Open Electricity Market. In: 2020 International Conference and Utility Exhibition on Energy, Environment and Climate Change (ICUE). pp. 1–7 (2020)
4. Chow, G.C.: Tests of Equality Between Sets of Coefficients in Two Linear Regressions. *Econometrica* **28**(3), 591–605 (1960)
5. Diebold, F.X., Mariano, R.S.: Comparing predictive accuracy. *Journal of Business & economic statistics* **20**(1), 134–144 (2002)
6. Dudani, S.A.: The Distance-Weighted k-Nearest-Neighbor Rule. *IEEE Transactions on Systems, Man, and Cybernetics* **SMC-6**(4), 325–327 (1976)
7. Hubicka, K., Marcjasz, G., Weron, R.: A Note on Averaging Day-Ahead Electricity Price Forecasts Across Calibration Windows. *IEEE Transactions on Sustainable Energy* **10**(1), 321–323 (2019)
8. Jawad, M., Nadeem, M.S.A., Shim, S.O., Khan, I.R., Shaheen, A., Habib, N., Hussain, L., Aziz, W.: Machine Learning Based Cost Effective Electricity Load Forecasting Model Using Correlated Meteorological Parameters. *IEEE Access* **8**, 146847–146864 (2020)
9. Kath, C., Nitka, W., Serafin, T., Weron, T., Zaleski, P., Weron, R.: Balancing Generation from Renewable Energy Sources: Profitability of an Energy Trader. *Energies* **13**(1), 205 (2020)
10. Kiesel, R., Paraschiv, F.: Econometric analysis of 15-minute intraday electricity prices. *Energy Economics* **64**, 77–90 (2017)
11. Killick, R., Fearnhead, P., Eckley, I.A.: Optimal Detection of Changepoints With a Linear Computational Cost. *Journal of the American Statistical Association* **107**(500), 1590–1598 (2012)
12. Lavielle, M.: Using penalized contrasts for the change-point problem. *Signal Processing* **85**(8), 1501–1510 (2005)

13. Li, W., Kong, D., Wu, J.: A Novel Hybrid Model Based on Extreme Learning Machine, k-Nearest Neighbor Regression and Wavelet Denoising Applied to Short-Term Electric Load Forecasting. *Energies* **10**(5), 694 (2017)
14. Maciejowska, K., Nitka, W., Weron, T.: Day-Ahead vs. Intraday—Forecasting the Price Spread to Maximize Economic Benefits. *Energies* **12**(4), 631 (2019)
15. Maciejowska, K., Uniejewski, B., Serafin, T.: PCA Forecast Averaging—Predicting Day-Ahead and Intraday Electricity Prices. *Energies* **13**(14), 3530 (2020)
16. Marcjasz, G., Serafin, T., Weron, R.: Selection of Calibration Windows for Day-Ahead Electricity Price Forecasting. *Energies* **11**(9), 2364 (2018)
17. de Marcos, R.A., Bunn, D.W., Bello, A., Reneses, J.: Short-Term Electricity Price Forecasting with Recurrent Regimes and Structural Breaks. *Energies* **13**(20), 5452 (2020)
18. Natividad, F., Folk, R.Y., Yeoh, W., Cao, H.: On the Use of Off-the-Shelf Machine Learning Techniques to Predict Energy Demands of Power TAC Consumers. In: Ceppi, S., David, E., Hajaj, C., Robu, V., Vetsikas, I.A. (eds.) *Agent-Mediated Electronic Commerce. Designing Trading Strategies and Mechanisms for Electronic Markets*. pp. 112–126. *Lecture Notes in Business Information Processing*, Springer International Publishing, Cham (2017)
19. Nawaz, M., Javaid, N., Mangla, F.U., Munir, M., Ihsan, F., Javaid, A., Asif, M.: An Approximate Forecasting of Electricity Load and Price of a Smart Home Using Nearest Neighbor. In: Barolli, L., Hussain, F.K., Ikeda, M. (eds.) *Complex, Intelligent, and Software Intensive Systems*. pp. 521–533. *Advances in Intelligent Systems and Computing*, Springer International Publishing, Cham (2020)
20. Nowotarski, J., Raviv, E., Trück, S., Weron, R.: An empirical comparison of alternative schemes for combining electricity spot price forecasts. *Energy Economics* **46**, 395–412 (2014)
21. Nowotarski, J., Weron, R.: Recent advances in electricity price forecasting: A review of probabilistic forecasting. *Renewable and Sustainable Energy Reviews* **81**, 1548–1568 (2018)
22. Rocha, H.R.O., Honorato, I.H., Fiorotti, R., Celeste, W.C., Silvestre, L.J., Silva, J.A.L.: An Artificial Intelligence based scheduling algorithm for demand-side energy management in Smart Homes. *Applied Energy* **282**, 116145 (2021)
23. Weron, R.: Electricity price forecasting: A review of the state-of-the-art with a look into the future. *International Journal of Forecasting* **30**(4), 1030–1081 (2014)
24. Yamin, H.Y., Shahidehpour, S.M., Li, Z.: Adaptive short-term electricity price forecasting using artificial neural networks in the restructured power markets. *International Journal of Electrical Power & Energy Systems* **26**(8), 571–581 (2004)
25. Yesilbudak, M., Sagioglu, S., Colak, I.: A novel implementation of kNN classifier based on multi-tupled meteorological input data for wind power prediction. *Energy Conversion and Management* **135**, 434–444 (2017)
26. Zhang, R., Xu, Y., Dong, Z.Y., Kong, W., Wong, K.P.: A composite k-nearest neighbor model for day-ahead load forecasting with limited temperature forecasts. In: 2016 IEEE Power and Energy Society General Meeting (PESGM). pp. 1–5 (2016)
27. Zhao, J.H., Dong, Z.Y., Xu, Z., Wong, K.P.: A Statistical Approach for Interval Forecasting of the Electricity Price. *IEEE Transactions on Power Systems* **23**(2), 267–276 (2008)
28. Ziel, F., Weron, R.: Day-ahead electricity price forecasting with high-dimensional structures: Univariate vs. multivariate modeling frameworks. *Energy Economics* **70**, 396–420 (2018)