Acta Univ. Sapientiae Informatica **14,** 2 (2022) 208–222

DOI: 10.2478/ausi-2022-0013

Hourly electricity price forecast for short- and long-term, using deep neural networks

Gergely DOMBI
Sonrisa Ltd.
Budapest, Hungary
email: dombi.gergely@sonrisa.hu

Tibor DULAI
University of Pannonia, Faculty of Information Technology
Veszprém, Hungary
email: dulai.tibor@mik.uni-pannon.hu

Abstract. Despite the practical importance of accurate long-term electricity price forecast with high resolution - and the significant need for that - only small percentage of the tremendous papers on energy price forecast attempted to target this topic. Its reason can be the high volatility of electricity prices and the hidden – and often unpredictable – relations with its influencing factors.

In our research, we performed different experiments to predicate hourly Hungarian electricity prices using deep neural networks, for short-term and long-term, too. During this work, investigations were made to compare the results of different network structures and to determine the effect of some environmental factors (meteorologic data and date/time - beside the historical electricity prices). Our results were promising, mostly for short-term forecasts - especially by using a deep neural network with one ConvLSTM encoder.

Computing Classification System 1998: I.2.1
Mathematics Subject Classification 2010: 68T07
Key words and phrases: hourly electricity price forecast, artificial intelligence, deep neural networks

1 Introduction

Energy price forecasting became a hot topic, especially in the past few years. The deregulation of the electricity market increased significantly the need for the accurate forecast of electricity prices, since the quality of the forecast has high importance and makes big impact on the risk management and the price-negotiation processes of the market participants [1]. On the other hand, with the increasing percentage of renewable energy sources in the system, trading strategy based on energy price forecast may influence the storage and control strategies, as well [6].

To obtain a good forecast, both a large amount and appropriately preprocessed/cleaned historical data, and a suitable mathematical model is required. They can be able to discover the relation between the energy price and its influencing factors, and ensure accurate and quick enough forecast. Prediction horizon has huge impact on the size of the data, that has to be chosen carefully: in case of small data size the performance of the trained model may be week, while large data size may make the training process much longer.

This paper presents our approaches to forecast electricity price with fine resolution from historical data, both for short-term and long-term, using deep neural networks. The impact of both some environmental factors and the network structure were investigated. Section 2 introduces the topic through some relevant papers. The applied methods are presented briefly in Section 3. Section 4 describes the experiment settings that led us to the results that are summarized in Section 5. Finally, Section 6 concludes our work.

2 Related works

The number of scientific papers on the topic of energy price forecast is increasing year-by-year. Lu et al. [5] summarized and compared 171 papers in their ,,decade review". Although, they dealt with not only electricity price forecast, but the price forecast for natural gas, crude oil and carbon, too, they concluded valuable results related electricity price forecast (from 54 publications). They found that - opposite to the other three energy sources - most of the papers on electricity price forecast had ultra-short prediction horizon: half-hourly or hourly. Because of the high volatility of electricity prices, and the fact that information from price drivers are usually annual estimates, hourly predictions are available mostly only for short-term. However, in a recent publication of Gabrielli et al. [3], an attempt was performed for producing finely resolved long-term electricity price forecast with a remarkable result (7.6%-19.3% Mean

Absolute Percentage Error for a single country, hourly price prediction). Their work applied Fourier analysis, and a regression model on the main frequencies of the Fourier series, that was based on the captured data both of the electricity market and price driver data. Back to the paper of Lu et al. [5], the reviewed publications showed that in the forecast, besides the time series of historical electricity prices, date (e.g., holidays, seasonal factors) and meteorological attributes (especially, the temperature) have huge role. They determined that the data cleaning – data decomposition – method (the most popular was Wavelet transformation, while Singular Spectral Analysis was on the second place) has much more importance than the chosen optimizer, in the improvement of the forecast accuracy. Considering the mathematical model, we can find papers that use statistical model, e.g., Uniejewski et al. [12] applied Ordinary Least Squares Regression for forecasting hourly electricity price; however, most of the papers on electricity price forecast apply machine learning models. For example, Romero et al. [8] worked with Random Forest method to forecast the spanish hourly electricity prices, Zhang et al. [15] applied a hybrid model that includes (next to Singular Spectral Analysis and Cuckoo Search) Support Vector Machine to forecast half-hourly electricity prices, while Windler et al. [13] used a Deep Feed-Forward Neural Network to forecast the hourly electricity prices for Germany and Austria. Based on the publications, neural networks seem to be the primary models for energy price forecast.

Describing the reviewed training processes precisely, 5-35% of the data were used as test set, and only few publications used validation set. The Mean Absolute Percentage Error (MAPE) was in the range of 0.0003 and 0.192 throughout the 171 publications. As future research directions, the increase of the development and usage of deep learning models, the creation and application of hybrid models and the more precise identification and involvement of correlating external factors in the energy price forecast were identified.

Our research used deep neural networks for electricity price forecast with different length (1 hour – 4 weeks) of prediction horizon. Moreover, we investigated the role of meteorological data on the precision of the forecast, too.

3 Methods

We have investigated the performance of deep neural networks with different structure for the forecast of future electricity price. All the networks were built and teached in Python, using the Keras/Tensorflow framework. The simpler networks were realized in Google Colab, using an Nvidia Tesla K80 graphic accelerator with 11 GB memory. For the more complex neural networks a server computer was applied, equipped with an Intel Core i7-8700 processor, 16GB RAM, and an Nvidia GTX 1080 Ti graphic accelerator with 11GB memory.

The historical electricity prices – from 2011 to 2017 (and from July 2010 to Aug 2019, for the forecasts for longer than 1 hour prediction horizon) – were collected from the Hungarian Power Exchange webpage (hupx.hu). We dealt with only the Eur/MWh, hourly prices. To clean this data, we used the PyWavelets library.

The values of the meteorological factors - the external factors that were considered in our research - were collected from the site of National Oceanic and Atmospheric Administration (www.noaa.gov). This data included the wind speed, temperature, dewpoint and barometric data from six Hungarian meteorologic stations from 2011 to 2017. The downloaded data had quarter-hourly resolution, however, in several cases, with missing data. Thus, we have derived hourly data from them for our investigation. For the sake of data consistency, we had to unify the collected data (the historical electric price and the meteorologic data) both in time zone and related to the seasonal time change.

Finally, the cardinality of the training set was 61000 (8759 of them were used for validation), while the number of the test examples was 343. Thus, the size of the test set was 5.6% of the size of the whole data set.

In cases of forecast for the next one hour, for teaching the network, we used NAdam optimizer, and to measure the accuracy of the forecast, we calculated Mean Absolute Error (MAE), while for forecasts for longer periods, the optimizer was RAdam and the accuracy was characterized by Mean Squared Error (MSE).

4 Experiments

We have carried out two kinds of experiments related to the input data: with and without meteorologic data. In the prior case, the neural networks had five inputs: year, month, day, hour and the price of 1 MWh electricity (in EUR). In the latter case, the network had 24 more inputs: the 4 meteorologic factors for the 6 stations.

The data (except the time- and date-related data) was preprocessed: each value of the training set was decreased by its column's average value, and finally, divided by the column's maximum. For the test set, the same method was applied, with the average and the maximum value of the training set. It resulted values between -1 and 1, with 0 as expected value and 1 as standard deviation. It makes the training process of the network more efficient.

All the experiments we've done can be categorized into two groups: first, we've forecasted electricity prices only for the next one hour, second, we made forecast for longer periods (for the next 24 hours and for the next 3-4 weeks).

4.1 Forecasts for the next 1 hour

In our forecasts for the next 1 hour, we made experiments using historical data from different sized time windows. As time window size, we used 1 hour, 3 hours, 6 hours, 12 hours, 24 hours and 168 hours. Here, e.g., a 12-hour sized time window means that we have forecasted the electricity price of the next hour based on the data from the previous 12 hours. For each time window size, we applied two different approaches. In the first one, all the electricity prices of the time window were originated from the webpage as real prices. In the second one, only the first value of the electricity price of the time window was a real price; all the others were results of our forecast based on the previous data.

In all of these experiments, we applied the same structure for the neural network. This structure is illustrated in Fig. 1. For the training and the forecast, we used the collected data from 2011 to 2017.

4.2 Forecasts for longer than 1 hour prediction horizon

Since the experiences showed that the forecast results are more accurate when meteorologic data are not considered, in the forecasts for periods longer than one hour only the date/hour and electricity prices were taken into account as input of the network. In these calculations, the input data set was extended: the electricity prices from the interval July 2010 - Aug. 2019 were used (altogether, electricity prices of 72000 hours were used for training).

First, forecasts were made for the next 24 hours while different network structures were investigated, then applying the network that gave the best result more investigations – even for longer prediction horizon – were made.

4.2.1 Forecasts for the next 24 hours

In all of these experiments, a time window with size of 48 hours was applied. In the investigations, four different types of network were used. All of these networks were similar to the network that was used in the forecasts for the

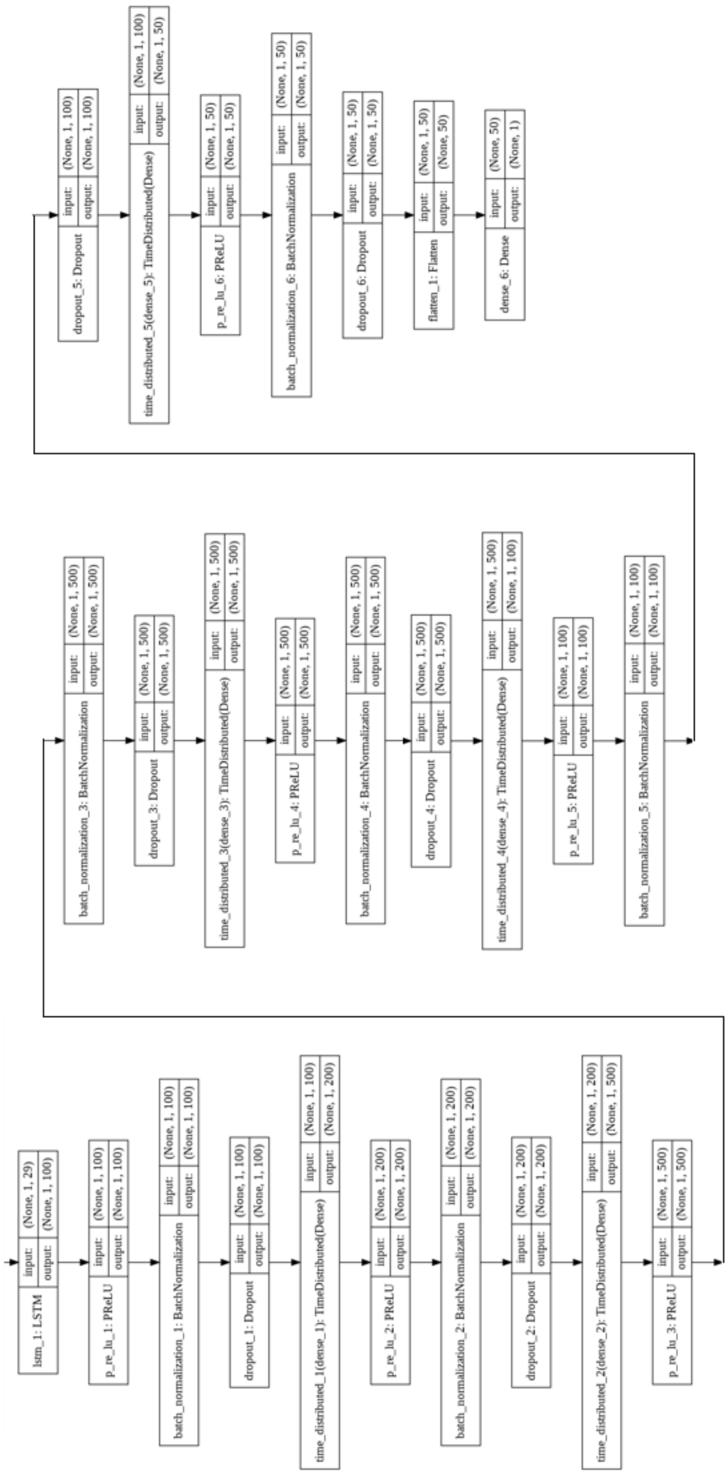

Figure 1: The neural network that was used for the 1 hour electricity price forecasts.

next one hour, however, they applied different methods to catch the useful patterns from the input data. These methods were:

1. including an LSTM encoder in the network
2. having a convolutional encoder in the network
3. containing a ConvLSTM encoder in the network
4. having two ConvLSTM layers in the network, where the first one returns also the inner variables from the loops of the LSTM neurons.

In all of these four variants, two different approaches were applied for the training: an univariate model, where only the the time series of the electricity price was used for training, and a multivariate model, where in addition to the electricity prices, the time- and date-related data (year, month, day, hour) were also used in the input.

4.2.2 Further experiments and forecasts for the next 3-4 weeks

In our further experiments, we kept the best constructions, so far. It means, that we used the deep neural network of the previous chapter that contains a ConvLSTM encoder, and we applied the univariate model – only the historical electricity prices were feeded as input. On this network, five more different approaches were applied:

1. First, noise filtering was applied on the training electricity price data. This method was tried because in some cases the sequence of the electricity prices contained significantly strange offset values. The applied noise filter was as follows: the discrete Haar wavelet of the electricity price was decreased by the standard deviation, then it was retransformed into its normal interval. As previously, forecast was made for the next 24 hours, based on the data of the previous 48 hours.

2. Second, stateful LSTM and ConvLSTM neurons were used. Here, statefulness means that the network keeps also the values got in the previous batch ,,in mind" (beside the change), this way it could catch more information about the periodicity of the electricity price. In our experiments, statefulness was first used only in the encoder layer, then only in the decored layer, and finally, in both of them. Here, the forecast was also made for the next 24 hours and the time window size was 48 hours.

	With meteorologic data		Without meteorologic data	
Window size	**min. MAE**	**epoch of min. MAE**	**min. MAE**	**epoch of min. MAE**
1 hour	0.02308	351	0.02468	90
3 hours	0.02103	34	0.01780	34
6 hours	0.02314	102	0.01735	26
12 hours	0.02260	56	0.01999	95
24 hours	0.02125	57	**0.01493**	**28**
168 hours	0.02668	78	**0.01404**	**43**

Table 1: Mean Absolute Error values of our experiments forecasting electricity price for the next 1 hour

3. Third, the prediction horizon was extended to 3 weeks, and a 4 weeks length time window was used. Because of the huge amount of data, the applied neural network – that uses a ConvLSTM encoder – had more filters than in the previous subchapter.

4. Fourth, the prediction horizon was extended to 4 weeks. The window size remained 4 weeks, however, the neural network was modified a bit: beside the increased number of filters, the size of the kernel of the convolution was changed from 6 to 8.

5. Fifth, we attempted to eliminate autocorrelation. For that purpose, we modified the training set as follows: iterating backwards on the sequence of the historical electricity prices, for each hour, its logarithmic yield was calculated by obtaining the natural logarithm of the current item divided by its prior item. The resulted data was used to train the network that had both a ConvLSTM encoder and an LSTM decoder.

5 Results

5.1 Results of the forecasts for the next 1 hour

To evaluate the success of the forecasts, we have calculated Mean Absolute Error (MAE). We have summarized the accuracy of all of our forecasts for one hour in Table 1.

We have experienced, that the best results were obtained when meteorologic data were not taken into account in the forecast. Its reason could be that the ratio of the electricity originated from renewable energy sources is low in Hungary, yet. It can be also seen, that the more the input data (the highest the time window) is, the more precise the forecast we get. Another conclusion that we can obtain is that when the input data is also a forecasted value

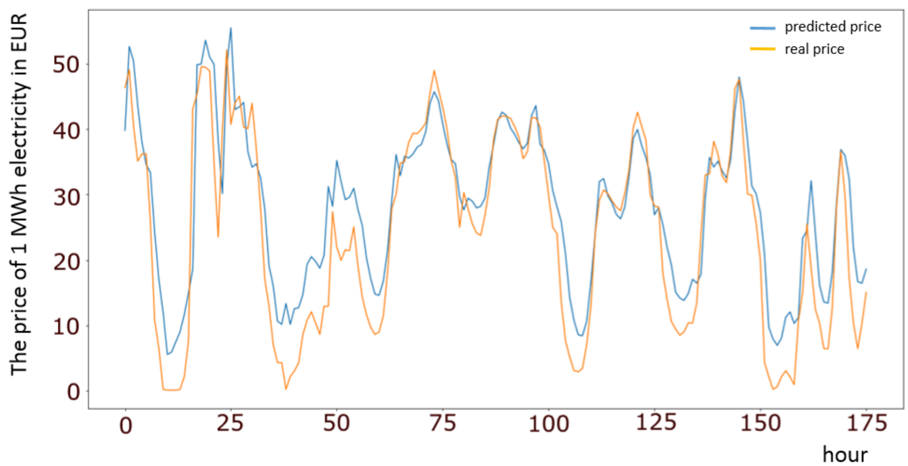

Figure 2: Comparison of the real and the forecasted electricity prices in our best case.

(only the first electricity price of the time window is real and all the others are forecasted), the curve of the forecast looses quickly its characteristics. Its reason can be that the forecast may introduce differences from the real values.

The difference between the curve of the real electricity prices and the curve of the forecasted electricity prices in the best case (168 hours time window, forecast for the next 1 hour, without meteorologic data) can be seen in Figure 2.

5.2 Results of the forecasts for longer than 1 hour prediction horizon

In these experiments, to evaluate the quality of the forecasts, Main Squared Error (MSE) was calculated. In Table 2, the minimum value of MSE on the valudation set, and the average of the square root of MSE on the test set are illustrated for all the four network structures, both on the univariate (the input is only the electricity price) and the multivariate (the input includes date and hour, too) model.

We can conclude, that the best results were obtained with one ConvLSTM encoder in the network, when input values were only the electricity price values. Still, – as Figure 3 shows – the curve of the forecasted electricity price

Network	Model	Min. MSE on the validation set	epoch of min. MSE	Avg. of square root of MSE on the test set
with LSTM encoder	**univariate**	0.00254	30	15.868
	multivariate	0.00274	21	16.821
with conv. encoder	**univariate**	0.00312	44	18.527
	multivariate	0.00313	34	18.753
with ConvLSTM encoder	**univariate**	0.00231	33	15.023
	multivariate	0.00302	64	17.856
with two ConvLSTM layers	**univariate**	0.00357	31	20.127
	multivariate	0.00369	46	20.140

Table 2: Mean Squared Error values of our experiments forecasting electricity price for the next 24 hours

Experiment	Subtype	Min. MSE on the validation set	epoch of min. MSE	Avg. of square root of MSE on the test set
Noise filtering		0.00229	23	18.046
Statefulness	**only in the decoder layer**	0.00237	36	15.959
	only in the encoder layer	0.00265	33	16.422
	in both of these layers	0.00223	41	17.762
3 weeks prediction horizon		0.00273	448	17.949
4 weeks prediction horizon		0.00244	187	17.935
autocorrelation elimination		0.00069	49	44.154

Table 3: Evaluation of the further experiments

was more precise in the multivariate case, when date and time were also parts of the input of the network.

5.3 Results for the further experiments

To evaluate the results of the five further experiments (noise filtering, statefulness, 3 and 4 weeks of prediction horizon and elimination of autocorrelation), Mean Squared Error (MSE) was calculated. Table 3 shows the results.

In case of noise filtering, the error in the validation phase was quite small, however, the error was higher for the test data than for the best result so far.

All the three experiments with statefulness gave worse results than the best experiment without it. Thus, it can be concluded that stateful networks are not suitable to solve this type of problems.

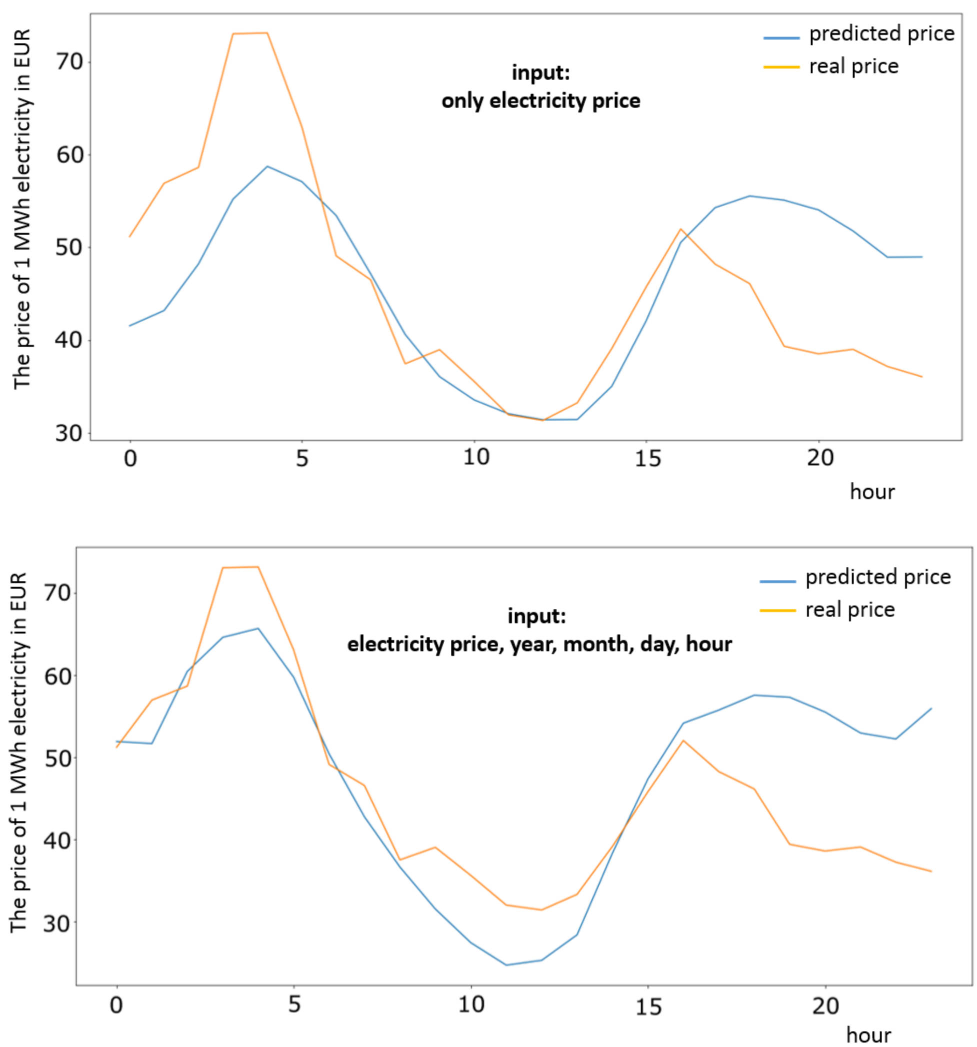

Figure 3: Comparison of the real and the forecasted electricity prices in case of forecast for the next 24 hours (for 25^{th} March 2019) by a network with one ConvLSTM encoder.

When the prediction horizon was extended to 3-4 weeks, the quality of the result still remained quite good. It has to be noted, that we performed more experiments for these cases, when the output of the network was feeded back as its new input. Thus, more than half a year incertainity was introduced into the system. Still, the average of the square root of MSE on the test set did not fail dramatically (23.784 and 22.626).

Finally, the elimination of autocorrelation prooved to give precise forecasts for the first 100-150 hours – its reason can be the size of the time window –, however, the further predicated values start to converge quickly to zero.

5.4 The quality of our results compared to other works

The literature on forecasts for energy market is significant. However, these studies differ not only in the applied model but the considered input variables, the input data itself, the prediction horizon and the performance measure that was used, for example. Since we have found no other study on the same data source and approach that we used, our comparison is more qualitative than numerical. Our experiments covered several settings, however, we separate the comparison into two parts based on the prediction horizon: 1 hour and 24 hours.

Singhal and Swarup [10] also applied artificial neural network to forecast electricity price for a one hour prediction horizon. Next to the historical electricity price, they involved date- and time-related data, forecasted demand and fuel price, too, in their forecast as input variables. The MAE values of their forecasts were between 0.682 and 9.282. 8 years later, in 2019, Cheng et al. [2] published their hybrid model (using Empirical Wavelet Transform, SVM and BiLSTM) that reached 0.1116 MAE on electricity price forecasting for 1 hour prediction horizon, based on historical electricity price data. Khalid et al. [4] published a model in 2020, that was based on Jaya-LSTM (JLSTM) and resulted in 0.0047 MAE in electricity price forecast. They also applied historical electricity price as input variable and used hourly prediction horizon. Considering them, the resulted 0.01-0.02 MAE that we got for our forecasts with 1 hour prediction horizon (see Table 1) is competitive.

Comparing the results of our forecasts with 24 hours prediction horizon to other published results, MSE was used as performance measure. As most of the papers on electricity price forecast uses shorter (mainly hourly) prediction horizon, the papers that were selected for the comparison with daily prediction horizon deal with carbon, gas or oil price forecast. Rahim et al. [7] used a fuzzy rule-based time series method to forecast crude palm oil prices for the next 24 hours. Their method used historical oil price and their result was qualified by 0.02 MSE. Su et al. [11] published a least squares regression boosting algorithm for forecasting natural gas spot price based on historical gas price. They got 0.4376 MSE for the case of daily prediction horizon. Also for natural gas spot price forecasting, Siddiqui [9] used an artificial neural network. Using historical gas price, for a daily prediction horizon he reached 0.026 MSE. Yahsi et al. [14] used both ANN and decision tree-based models to forecast carbon price with daily prediction horizon. As input variables they used the historical prices of the different energy sources, and other energy market-based indexes. The performance of their models' forecast was characterized by MSE between 0.3416 and 1.7065. Taking into account these results, the 0.002-0.004 minimum MSE values that we got for our forecasts with 24 hours or longer prediction period (see Table 2 and Table 3) are promising.

6 Conclusion

In our research we made an attempt to forecast the electricity price for the future using deep neural networks. Several experiments were performed, investigating the effect of different types of input data and the network structure on the quality of the forecast. We have experienced that – opposite to the international scientific literature on the electricity price forecast in the western countries – meteorologic data has only very small impact on the Hungarian electricity price. Its reason can be the low level spread of the renewable energy sources in the Hungarian energy industry. Moreover, in Hungary, only a small percentage of the population applies electricity for heating, and the usage of air conditioners in the households is only in starting phase.

Based on the performed experiments, deep neural network with one ConvLSTM encoder prooved to give the most precise forecast. Though, the sequence of electricity prices may contain huge offset values, noise filtering was not able to improve the forecast significantly. Its reason can be the flexibility of neural networks that makes handling of input offsets easy by default.

We have experienced that short-term forecasts were mostly precise, in constrast to the long-term forecasts. However, because the contracts of the electricity market are concluded mainly at least one year earlier than the transactions, long-term forecasts have much higher importance than short-term forecasts.

Based on our research, it turned out that neural networks learnt mostly the seasonality of the data. It can be a reason for the limitation that forecasts cannot be more precise. On the other hand, the price of electricity can be influenced by other factors that are hard to model and forecast: the political environment, natural disasters, deployment or breakdown of power plants. Because of these circumstances, electricity price forecast remains a challenging problem, especially for a long-term prediction horizon.

Acknowledgements

This research was supported by the National Research, Development and Innovation Office in the project 'Research and development of an innovative risk assessment expert system for the renewable energy market' (contract id: KFI_16-1-2017-0489).

References

[1] M. Castelli, A. Groznik, A. Popovič, Forecasting Electricity Prices: A Machine Learning Approach, *Algorithms* **13,** 5 (2020) 1–16. ⇒209

[2] H. Cheng, X. Ding, W. Zhou, R. Ding, A hybrid electricity price forecasting model with Bayesian optimization for German energy exchange, *International Journal of Electrical Power & Energy Systems* **110** (2019) 653-666. ⇒219

[3] P. Gabrielli, M. W utrich, S. Blume, G. Sansavini, Data-driven modeling for long-term electricity price forecasting, *Energy* **244** (2022) 123107. ⇒209

[4] R. Khalid, N. Javaid, F.A. Al-Zahrani, K. Aurangzeb, E.U. Qazi, T. Ashfaq, Electricity Load and Price Forecasting Using Jaya-Long Short Term Memory (JLSTM) in Smart Grids, *Entropy (Basel)* **22** (2019) 10. ⇒219

[5] H. Lu, X. Ma, M. Ma, S. Zhu, Energy price prediction using data-driven models: A decade review, *Computer Science Review* **39** (2021) 100356. ⇒209, 210

[6] T. Miseta, A. Fodor, Á. Vathy-Fogarassy, Energy trading strategy for storage-based renewable power plants, *Energy* **250** (2022) 123788. ⇒209

[7] N.F. Rahim, M. Othman, R. Sokkalingam, E. Abdul Kadir, Forecasting Crude Palm Oil Prices Using Fuzzy Rule-Based Time Series Method, *IEEE Access* **6** (2018) 32216–32224. ⇒219

[8] Á. Romero, J.R. Dorronsoro, J. Diáz, Day-Ahead Price Forecasting for the Spanish Electricity Market, *Int. J. Interact Multimedia Artif. Intell.* **5,** 4 (2019) 42–50. ⇒210

[9] A.W. Siddiqui, Predicting Natural Gas Spot Prices Using Artificial Neural Network, *2nd International Conference on Computer Applications & Information Security (ICCAIS)* (2019) 1–6. ⇒220

[10] D. Singhal, K.S. Swarup, Electricity price forecasting using artificial neural networks, *International Journal of Electrical Power & Energy Systems* **33,** 3 (2011) 550–555. ⇒219

[11] M. Su, Z. Zhang, Y. Zhu, D. Zha, Data-Driven Natural Gas Spot Price Forecasting with Least Squares Regression Boosting Algorithm, *Energies* **12** (2019) 1094. ⇒220

[12] B. Uniejewski, R. Weron, F. Ziel, Variance Stabilizing Transformations for Electricity Spot Price Forecasting, *IEEE Trans. Power Syst.* **33,** 2 (2017) 2219–2229. ⇒210

[13] T. Windler, J. Busse, J. Rieck, One month-ahead electricity price forecasting in the context of production planning, *J. Cleaner Production* **238** (2019) 117910. ⇒210

[14] M. Yahsi, E. Canakoglu, S. Agrali, Carbon price forecasting models based on big data analytics, *Carbon Management* **10** (2019) 1–13. ⇒220

[15] X. Zhang, J. Wang, Y. Gao, A hybrid short-term electricity price forecasting framework: Cuckoo search-based feature selection with singular spectrum analysis and SVM, *Energy Economics* **81** (2019) 899–913. ⇒210

Received: November 18, 2022 • Revised: December 8, 2022