POLITYKA ENERGETYCZNA – ENERGY POLICY JOURNAL

2025 ✦ Volume 28 ✦ Issue 2 ✦ 211–230

DOI: 10.33223/epj/207197

Rafał Sowiński[1], Aleksandra Komorowska[2]

Forecasting electricity prices in the Polish Day-Ahead Market using machine learning models

Abstract: Given the constantly changing market situation for electricity prices, driven by shifts in the energy mix, regulatory reforms, and broader socio-economic factors, it is necessary to reassess the understanding of price forecasting periodically. Traditional statistical methods may struggle when faced with heightened volatility, nonlinear dependencies, and rapidly changing input features. In contrast, machine learning models, particularly Artificial Neural Networks (ANNs), can adapt more effectively to complex, non-stationary patterns in price time series. In this study, six distinct artificial neural network (ANN) architectures were developed and trained using eight years of historical Polish Day-Ahead Market electricity price data (2016–2024). Four of these were plain deep learning models: a Multilayer Perceptron (MLP), a Convolutional Neural Network (CNN), a Long Short-Term Memory (LSTM) model, and a Gated Recurrent Unit (GRU) model. Two others were hybrid models combining convolutional layers with recurrent layers. The hybrid architectures, namely CNN+LSTM and CNN+GRU, were designed to leverage the capacity of CNN to automatically extract features from narrower sliding windows of past prices and the LSTM/GRU layers' ability to capture long-term temporal dependencies. The models' performances were evaluated using three metrics: Mean Absolute Error (MAE), Root Mean Squared Error

✉ Corresponding Author: Rafał Sowiński; e-mail: sowinski@min-pan.krakow.pl

[1] The Department of Energy Policy and Markets, Mineral and Energy Economy Research Institute, Polish Academy of Sciences, Poland; ORCID iD: 0009-0002-3836-1816; e-mail: sowinski@min-pan.krakow.pl

[2] The Department of Energy Policy and Markets, Mineral and Energy Economy Research Institute, Polish Academy of Sciences, Poland; ORCID iD: 0000-0002-9604-1071; e-mail: komorowska@min-pan.krakow.pl

(RMSE), and the coefficient of determination (R^2). The top-performing CNN+LSTM achieved an MAE of 75.21 PLN/MWh, an RMSE of 103.64 PLN/MWh, and an R^2 of 0.59. Results were also compared against several models previously reported in the literature. These results may be used to improve price forecasting by indicating the optimal pathways for building forecasting models and, in extension, lead to more efficient power system planning.

Keywords: electricity, Day-Ahead Market, artificial neural networks, prices, forecasting

Introduction

The ongoing global decarbonization process is reshaping the structure of electricity generation in power systems (Obiora et al. 2024). In 2024, renewable energy sources (RES) accounted for 30% of global electricity production, representing a 23% increase in solar energy and a 10% increase in wind energy generation compared to the previous year (Ember 2024). As environmentally friendly, renewable sources of electricity increasingly replace solid fuels, electricity generation becomes more variable and less predictable due to the high dependency on weather conditions, which directly affect wind and solar output (Pepłowska 2025). These conditions vary across countries and regions, depending on their climatic zones. Moreover, the pace of changes in the energy mix is often regulated or supported by national legislation and policy frameworks (Liu and Feng 2023). Consequently, countries that began their RES transition earlier have gained more experience and are actively addressing the related challenges (Dong et al. 2019).

As the electricity market becomes increasingly volatile, conventional statistical models may become overly complex and less accurate (Tschora et al. 2022). Given that a wide range of public and private stakeholders rely on price forecasts, inaccuracies can lead to financial losses, coordination challenges, and complications in system planning (Alboom 2025). Instead of using conventional statistical models, these challenges may be addressed by implementing machine learning (ML) models, which do not require any predefined functional forms and can better capture non-linear patterns in the data (Ejdys et al. 2015; Miller and Bućko 2014). Although such models have been successfully applied in various contexts, their effectiveness depends on the relevance of the training data. As electricity markets evolve, older models trained on outdated datasets may lose predictive accuracy.

In recent years, a growing number of studies have employed machine learning techniques to forecast electricity prices. These approaches vary in terms of model architecture, input data, and geographical focus. In 2010, Halicka developed and trained a Multilayer Perceptron (MLP) machine learning model with two hidden layers, consisting of 30 and 5 neurons, respectively, and seven input variables (price, volume, wind speed, cloud cover, temperature, hour of the day and day of the week) in 24 samples sliding windows (Halicka 2010). The model achieved a mean absolute error (MAE) of 2.93 PLN/MWh, using data from the 2008 Polish Day-Ahead Market for training (once with backpropagation and once using the conjugate gradient method) and from January 2009 for testing.

Miller and Bućko built an MLP model with a single hidden layer, consisting of five neurons, using the backpropagation method on 2010 Polish electricity market data, that was able to predict the price with a 4.05% mean absolute percentage error (MAPE) (Miller and Bućko 2014). Ejdys et al. claim to have developed and trained an MLP model in 2015, which reached a MAPE of 4.71% (or an MAE of 7.93 PLN/MWh), but did not specify the input data time range, nor the training method. As for architecture, the name "MLP 7-2-1" suggests seven input neurons, two neurons in a hidden layer, and one output neuron (Ejdys et al. 2015). As these studies relied on older data, the resulting models would probably prove unreliable in today's day-ahead electricity market. Other, more recent studies have focused on foreign electricity markets and have not considered the specific characteristics of the Polish power system (Belenguer et al. 2025; Mao et al. 2025).

Although the abovementioned studies provide important early contributions to electricity price forecasting using MLP architectures, their current relevance is limited. All models were trained on datasets from over a decade ago and, therefore, do not accurately reflect the structural changes that have occurred in the Polish electricity market since then, such as increased generation from renewable energy sources, shifts in demand patterns, and regulatory developments. Consequently, there is a need for an updated, context-specific model that incorporates the latest market dynamics and leverages recent advances in machine learning techniques.

This study contributes to the existing body of research by developing a machine-learning model explicitly tailored to the Polish Day-Ahead electricity market. In contrast to earlier models, which were based on outdated data or failed to account for local market conditions, the proposed approach integrates recent datasets and accurately represents the current structure of the Polish power system. This study utilizes the most recent available meteorological and generation data from Poland, combined with various artificial neural network (ANN) architectures. To further enhance the accuracy of electricity price forecasting, various combinations of input variables, layer types, and training methods are compared.

The remainder of this paper is organized as follows: Section 2 presents the data sets used and describes the applied methodology. Section 3 discusses the results, while Section 4 concludes the study.

1. Materials and methods

This section provides a description of the datasets used in the study, the methods applied for their processing and analysis, and the architectures of the machine learning implemented. The datasets include meteorological and electricity generation data relevant to the Polish Day-Ahead electricity market. The overall workflow of the study is presented in Figure 1, outlining the main

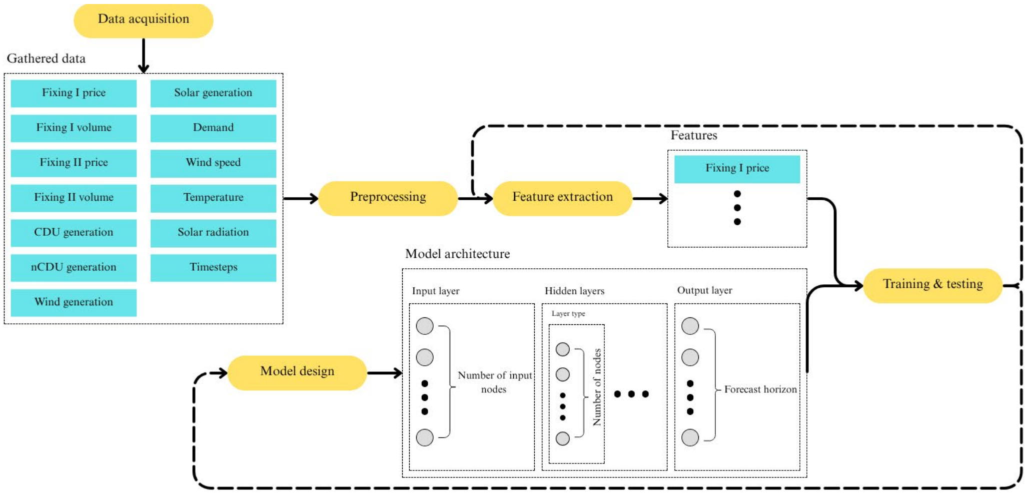

Fig. 1. Workflow of the study

Rys. 1. Przebieg pracy na rzecz badania

steps from data acquisition through preprocessing, feature extraction, and model design to model training, testing, and evaluation.

1.1. Data assumptions

Initially, the relevant datasets were acquired from open-access and publicly available sources. The Polish Transmission System Operator (TSO) shares hourly electricity generation data, including outputs from Centrally Dispatched Units (CDU), non-centrally Dispatched Units (nCDU), wind and solar generation, as well as electricity demand, up to June 14, 2024 (PSE SA 2025). Day-ahead market data, comprising prices and volumes for fixing I and II, were obtained from Energy Instrat, an open-access platform offering electricity and fuels market data (Energy Instrat 2025). The earliest available records date back to January 1, 2016. These two dates define the overall temporal scope of the datasets from January 1, 2016, to June 14, 2024. To provide a general overview of the dataset's characteristics, Figure 2 presents the weekly averages of fixing I price and traded volume over the entire period. The sudden increase in electricity price growth is closely related to the COVID-19 pandemic, the ongoing Russo-Ukrainian armed conflict, and the resulting surge in global fuel prices (Adolfsen et al. 2025; Zhong et al. 2020).

Temperature, solar radiation, and wind speed data were obtained from the ERA5 Climate Data Store (CDS) in the form of simulated reanalysis (ERA5 2025). All datasets were available at hourly resolution. A comparison of the simulated wind speed data with real-world measurements from the Polish Institute of Meteorology and Water Management (IMGW) revealed that the ERA5 wind data exhibited an average relative error of 12.99% (IMGW 2025). This discrepancy

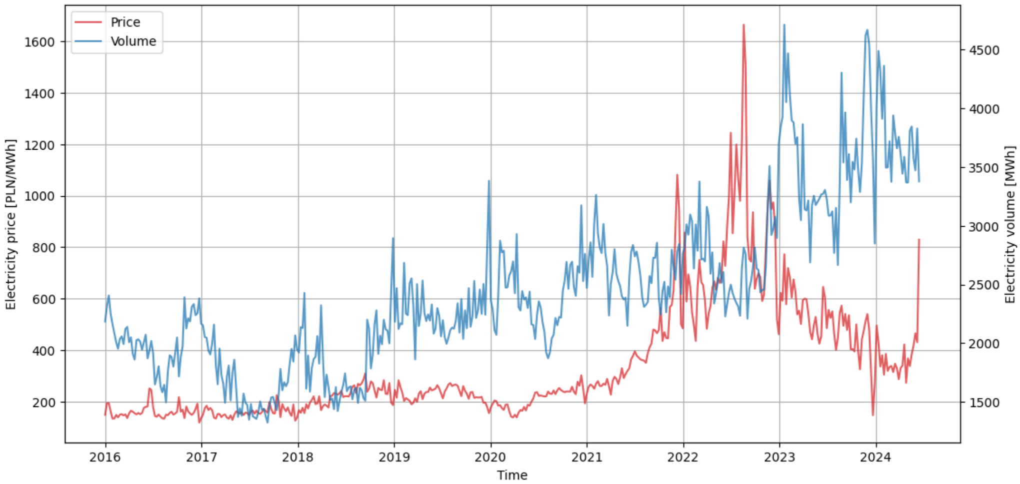

Fig. 2. Average weekly electricity price and traded volume (fixing I) in the Polish market in 2016–2024 (Energy Instrat)

Rys. 2. Średnie tygodniowe ceny i wolumenu energii elektrycznej (fixing I) na Rynku Dnia Następnego między 2016–2024 (Energy Instrat)

is likely due to the spatial difference between the locations of the meteorology station and the nearest ERA5 datapoint, located approximately 8 km southeast of the station. Additionally, the station's proximity to the shoreline may contribute to local wind patterns not fully captured by the ERA5 dataset. Nevertheless, this level of error is generally considered acceptable for large--scale energy modeling and site screening purposes (Dubus et al. 2023; Tsai et al. 2022)

As shown in Figure 1, the gathered data were subject to preprocessing, which eliminated empty values by interpolating adjacent values. The datasets were then analyzed both visually and statistically to identify the most suitable features for training the machine learning model. At the same time, a variety of artificial neural network (ANN) models were being developed. This involved iterating on the number of layers, the number of neurons per layer, and the types of layers to achieve the highest possible accuracy.

Most of the steps were performed using the open-source programming language Python and a set of specialized libraries. The *pandas* library was employed for data analysis and manipulation, while *NumPy* supported matrix operations and other mathematical functions. *SciPy* provided algorithms for scientific computing and was primarily used for scaling the data. *TensorFlow* provided API interfaces for building ML models, including the open-source deep learning library *Keras*, which was primarily used in this study. Every custom model was trained using the *Adam* (Adaptive Moment Estimation) optimizer provided by *Keras*.

All collected data were subjected to statistical analysis to identify the most relevant features for training the machine learning models. The analysis included, among other things, visual assessment of the dependency plots, calculation of correlation coefficients between input variables, and the construction of linear regression models to evaluate the coefficients of determination. The analysis was conducted in Python, using the *statsmodels* library alongside *NumPy*, which provided the necessary tools and functions.

1.2. Methods – Artificial Neural Network models

To capture the complex dependencies influencing electricity prices, data-driven forecasting techniques were employed in this study. Among them, Artificial Neural Networks (ANNs) were selected due to their proven effectiveness in modeling nonlinear time series data. Artificial Neural Networks discern complex, nonlinear patterns and relationships from historical training data, which can then be applied to forecast future values. ANN is a type of machine learning model that mimics the inner workings of a biological nervous system (Benalcazar et al. 2017).

For this study, four different ANN architectures were employed:

- Multilayer Perceptron (MLP),
- Long-Short Term Memory (LSTM),
- Gated Recurrent Unit (GRU),
- Convolutional Neural Network (CNN),

and some combinations of the aforementioned.

An ANN works by activating nodes, called neurons, which are grouped in layers. Each node is connected to others, and each connection is assigned a weight. This leads to a simple neuron model function as shown in Equation 1, where h is the output, σ is the activation function (sigmoid in this example), w is the weight, x is the input, and b is the bias (Haykin 1999).

$$h(x) = \sigma(w \cdot x + b) \tag{1}$$

1.2.1. Feedforward Neural Network

The most basic type of artificial neural network (ANN) is a Feedforward Neural Network (FFN), in which the flow of information is strictly unidirectional: from the input layer through the hidden layers to the output layer. This structure can be represented as a sequence of neurons feeding into subsequent layers. By extending Equation 1, the output of a neuron can be represented as Equation 2, where σ is the nonlinear activation function of the nth node, w is the weight matrix, x is the input matrix, and b is the bias (Haykin 1999). A graphical representation of an FFN is shown in Figure 3. Due to the simplicity of the model, it requires significantly less computational power for training, but it may be less effective for more complex or highly temporal data (Jain and Medsker 1999). An MLP is a name given to modern FFN with fully connected neurons activated nonlinearly.

$$h(x) = \sigma_n\left(w_n \cdot \sigma_{n-1}\left(w_{n-1} \cdots \sigma_1\left(w_1 \cdot x_1 + b_1\right) \cdots + b_{n-1}\right) + b_n\right) \tag{2}$$

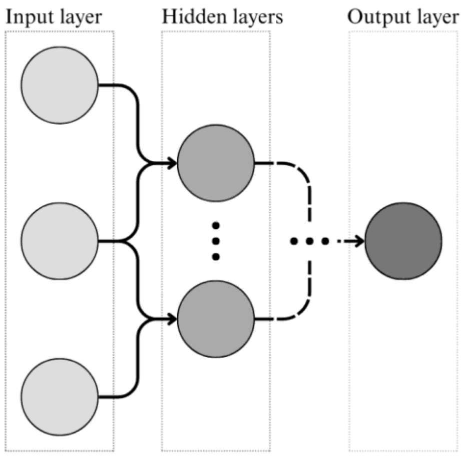

Fig. 3. Visual representation of a Feedforward Neural Network (based on: Jain and Medsker 1999)

Rys. 3. Reprezentacja wizualna jednokierunkowej sieci neuronowej

1.2.2. Recurrent Neural Network

A Recurrent Neural Network (RNN) derives from the FFN. The main difference between these two is that an RNN allows a connection between a neuron and another from a previous layer. This creates a feedback loop. As a result, an RNN is expected to perform better with sequential data than an FFN (Jain and Medsker 1999). Figure 4 shows the visual representation of an exemplary RNN. Equation 3 is a mathematical representation of an RNN neuron, where σ is the nonlinear activation function of a neuron in the moment t, w_a and w_x are the recurrent and input weight vectors, a stands for the recurrent information, x is the input, and b represents the bias (Jain and Medsker 1999).

$$h(a,x) = \sigma(w_a \cdot a_{t-1} + w_x \cdot x_t + b) \tag{3}$$

1.2.3. Long-Short Term Memory

However, the feedback loop introduces a challenge: the recurrent information tends to either vanish or explode during training, hence the *vanishing/exploding gradient problem*. To address this issue, Hochreiter and Schmidhuber proposed an improved type of recurrent neuron called the Long-Short Term Memory (LSTM) cell (Hochreiter and Schmidhuber 1997). The LSTM cell contains additional gates that regulate the inclusion or exclusion of information. This allowed the LSTM networks, unlike regular RNNs, to selectively forget

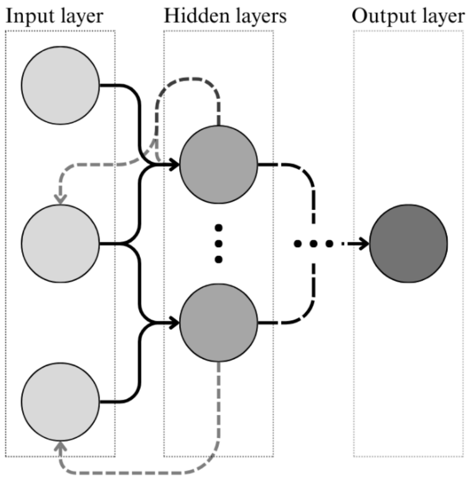

Fig. 4. The representation of a Recurrent Neural Network (based on: Jain and Medsker 1999)

Rys. 4. Reprezentacja rekurencyjnej sieci neuronowej

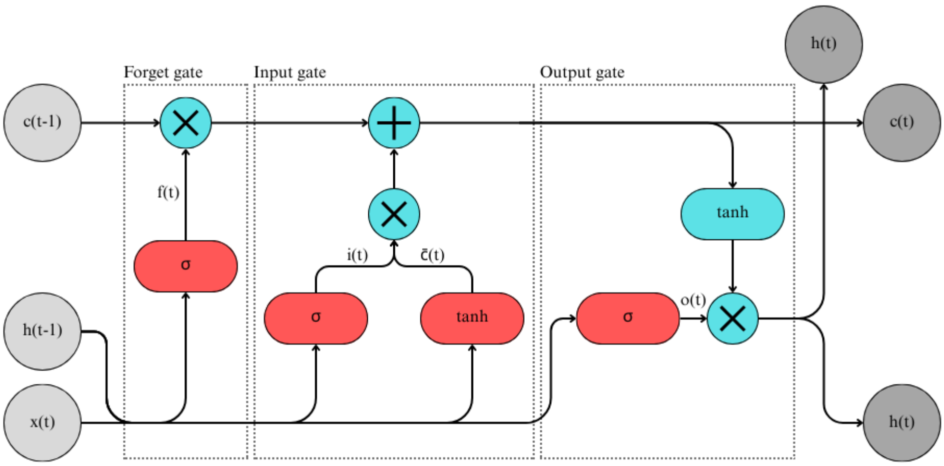

Fig. 5. Single LSTM cell representation (based on: Yu et al. 2019).

Rys. 5. Reprezentacja pojedynczej komórki sieci typu długa pamięć krótkotrwała

information over time (Hochreiter and Schmidhuber 1997). Figure 5 represents a single LSTM cell. Equation 4 explains each step of the LSTM cell's data processing, where σ and *tanh* are nonlinear activation functions of a neuron in the moment *t*, *w* are appropriate weights, *x* is the input, a is the output, *c* is the cell state, and *b* the bias (Hochreiter and Schmidhuber 1997).

$$
\begin{aligned}
f_t &= \sigma\left(w_{fa} \cdot a_{t-1} + w_{fx} \cdot x_t + b_f\right) \\
i_t &= \sigma\left(w_{ia} \cdot a_{t-1} + w_{ix} \cdot x_t + b_i\right) \\
\overline{c}_t &= tanh\left(w_{\overline{c}a} \cdot a_{t-1} + w_{\overline{c}x} \cdot x_t + b_{\overline{c}}\right) \\
c_t &= f_t \cdot c_{t-1} + i_t \cdot \overline{c}_t \\
o_t &= \sigma\left(w_{oa} \cdot a_{t-1} + w_{ox} \cdot x_t + b_o\right) \\
h_t &= o_t \cdot \tanh\left(c_t\right)
\end{aligned}
\tag{4}
$$

1.2.4. Gated Recurrent Unit

In 2014, in order to reduce the number of parameters in an LSTM cell, Chu et al. proposed the RNN Encoder-Decoder model. Commonly known as the Gated Recurrent Unit (GRU), this hidden unit implements the update and the reset gates (Cho et al. 2025). The first determines whether to update the hidden state and the second whether the previous hidden state should be ignored, as shown in Figure 6. Equation 5 shows calculations during every intermediate step of the process, where u (*update gate*), r (*reset gate*) and $\tilde{a}$ are intermediate states of information in the moment t, w are appropriate weights, x is the input, a is the output, and b is the bias (Agarap 2018).

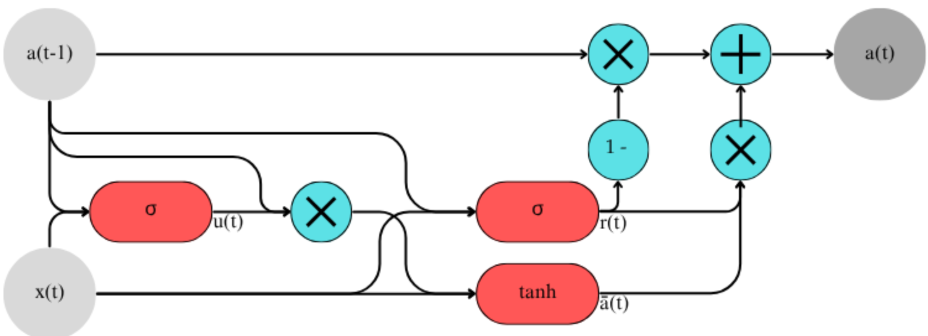

Fig. 6. Gated Recurrent Unit representation (based on: Agarap 2018)

Rys. 6. Reprezentacja bramkowej jednostki rekurencyjnej

$$
\begin{aligned}
u_t &= \sigma\left(w_{ua} \cdot a_{t-1} + w_{ux} \cdot x_t + b_u\right) \\
r_t &= \sigma\left(w_{ra} \cdot a_{t-1} + w_{rx} \cdot x_t + b_r\right) \\
\widetilde{a_t} &= \tanh\left(w_{ar}\left(r_t \otimes a_{t-1}\right) + w_{ax} \cdot x_t + b_a\right) \\
a_t &= \left(1 - u_t\right) \otimes a_{t-1} + u_t \otimes \widetilde{a_t}
\end{aligned}
\tag{5}
$$

1.2.5. Convolutional Neural Network

A Convolutional Neural Network (CNN) is a type of Artificial Neural Network (ANN) architecture composed of filters that produce feature maps by sliding over the input data, thereby capturing local patterns. Commonly used in tandem, pooling layers reduce the dimensionality of these feature maps while retaining crucial information (O'Shea and Nash 2015). Figure 7 presents the CNN architecture.

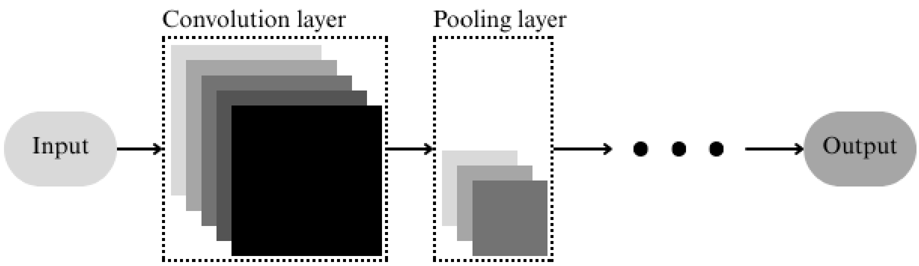

Fig. 7. Convolutional Neural Network representation (based on: O'Shea and Nash 2015)

Rys. 7. Reprezentacja konwolucyjnej sieci neuronowej

Each architecture was run through hyperparameter tuning using *KerasTuner*. The tuner trains models in a loop, adjusting their hyperparameters, namely the number of layers, the number of neurons per layer, the activation function, and the dropout rate between iterations. Models were trained using the Adam optimizer provided by the *keras* library. The most efficient model of each type was saved for further testing and comparison. Electricity prices, along with other highly related time series, were split into training, validation, and testing datasets in a ratio of 0.7:0.15:0.15. The final models allow for electricity price forecasting with varying levels of accuracy, depending on the specific model architecture. Other models, namely the Temporal Convolution Network (TCN) and Transformer, did not show promising results during the initial analysis and were excluded from further tuning.

2. Results and discussion

The training and testing results of the machine learning models are presented in this section. The following models were analyzed: Multilayer Perceptron (MLP), Long-Short Term Memory (LSTM), Convolutional Neural Network (CNN), Gated Recurrent Unit (GRU), and hybrid models built from combinations of these architectures. The characteristics of the best-performing models from each architecture are presented in Table 1. The number and type of layers, number

Table 1. Results of hyperparameter optimisation

Tabela 1. Wyniki optymalizacji hiperparametrów

Model	Number and types of layers	Number of neurons or filters	Activation functions
MLP	Input, Dense, Dense (output)	168, 5, 24	None, Linear, Linear
LSTM	Input, LSTM, Dropout, Dense (output)	7, 4, (rate 0.2), 24	None, Hyperbolic tangent, None, Linear
CNN	Input, 1D convolution, Max pooling, Flatten, Dense (output)	168, 5 (kernel size 3), (pool size 2), None, 24	None, Rectified Linear Unit (ReLU), None, None, Linear
GRU	Input, GRU, Dense (output)	7, 3, 24	None, Hyperbolic tangent, Linear
CNN + LSTM	Input, 1D Convolution, Max pooling, LSTM, Dense (output)	7, 9 (kernel size 3), (pool size 2), 6, 24	None, ReLU, None, Hyperbolic tangent, Linear
CNN + GRU	Input, 1D Convolution, Max pooling, GRU, Dense (output)	7, 8 (kernel size 3, (pool size 2), 6, 24	None, ReLU, None, Hyperbolic tangent, Linear

of neurons, and activation functions are used to describe them. It is worth noting that since all activation functions were linear, the MLP model is not a multilayer perceptron sensu stricto. The name was not changed, however, for the sake of convention.

The developed models were compared with one another and with those identified during the literature review, using standard metrics such as Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and the coefficient of determination (R^2). Mean absolute error is a measure of the average absolute difference between the predicted and observed historical testing data. The exact mathematical formula for MAE is presented in Equation 6, where i is the current sample out of n samples, x is the original sample, and y is the predicted value. Root Mean Squared Error is the square root of the average square of the difference between the predicted and observed historical testing data, as shown in Equation 7, where i is the current sample out of n samples, x is the original sample and y is the predicted value. The coefficient of determination is usually explained as the measure of how much of the variance the model accounts for. Computationally, R^2 is the proportion of the residual sum of squares to the total sum of squares subtracted from one, as in Equation 8, where i is the current sample out of n samples, x is the original sample, y is the predicted value, and $^{-}$ is the average sample value. An R^2 value below zero means that the model's prediction is less accurate than a constant function predicting the mean of the data ($\overline{x}$).

$$MAE = \frac{1}{n} \cdot \sum_{i=1}^{n} \left| x_i - y_i \right| \tag{6}$$

$$RMSE = \sqrt[2]{\frac{1}{n} \cdot \sum_{i=1}^{n} \left(x_i - y_i \right)^2} \tag{7}$$

$$\begin{gathered} \overline{x} = \frac{1}{n} \cdot \sum_{i=1}^{n} x_i \\ R^2 = 1 - \frac{\sum_{i=1}^{n} \left(x_i - y_i \right)^2}{\sum_{i=1}^{n} \left(x_i - \overline{x} \right)^2} \end{gathered} \tag{8}$$

To compare the models, each one was trained and tested on the same dataset, using the same set of features selected as best suited for the analysis, and evaluated through *time series cross-validation*. The original performance claims reported by the authors of the compared studies could not be replicated due to the lack of specifications, differences in implementation environments, and inaccessible data. After the initial tuning had demonstrated the advantage of simpler feedforward models, another tuning process was conducted, focusing solely on the MLP model to identify the optimal combination of hyperparameters. Since this architecture is less complex, a larger number of combinations could be tested within a reasonable time frame.

As shown in Table 2, the hybrid CNN+LSTM model outperformed the others. Flattened CNN+LSTM forecasting results are shown in Figure 8. The last 15% of the time series was used as the testing data. The lowest MAE and RMSE values suggest that this architecture allows for the most accurate electricity price forecasting. A high coefficient of determination also ensures that the model can explain a significant percentage of the market variability. The convolutional layers learn to detect local patterns, such as peaks or short-term sequences, while the recurrent layers allow for effective temporal modeling. The MLP model performed only slightly worse in terms of MAE, and the CNN model was not far behind in terms of MAE, RMSE, and R^2. This shows that appropriately adjusted simple models with carefully adapted features may still outperform plain RNNs.

Models based on compared studies showed lower accuracy, with the Halicka (2010) model performing comparably to, though slightly below, a constant mean function, which may reflect

Table 2. Performance comparison of machine learning models trained on 8 years of data using time series cross-validation

Tabela 2. Porównanie wydajności modeli uczenia maszynowego trenowanych na danych z 8 lat, po walidacji krzyżowej szeregów czasowych

Metric	Model								
	Halicka (2010)	Miller and Bućko (2014)	Ejdys et al. (2015)	MLP	CNN	LSTM	GRU	CNN + + LSTM	CNN + + GRU
MAE	209.61	216.27	100.87	77.28	76.48	96.02	89.90	75.21	79.34
RMSE	336.32	229.40	133.14	107.85	105.77	129.18	123.17	103.64	107.75
R^2	–0.59	0.75	0.62	0.56	0.58	0.37	0.42	0.59	0.56

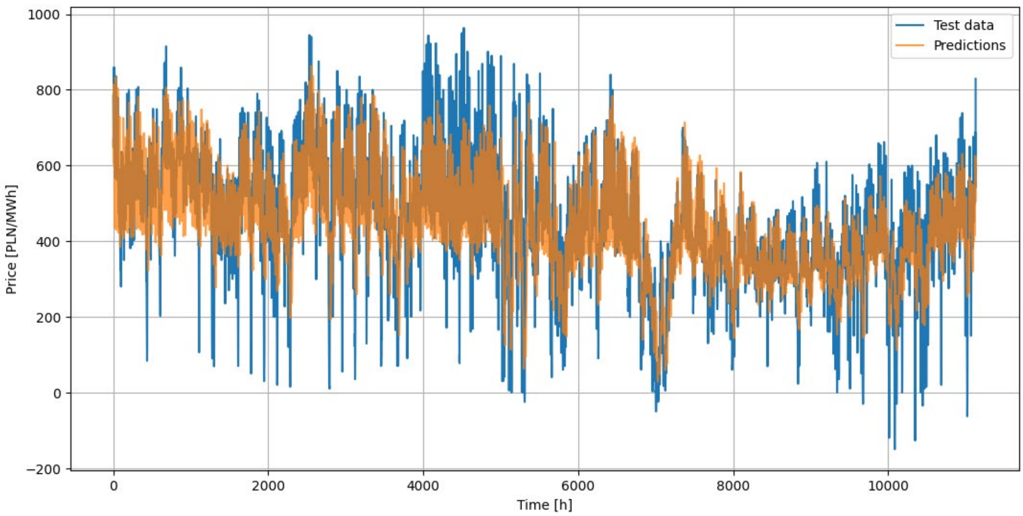

Fig. 8. CNN+LSTM forecasting results on top of test data. Forecasts were flattened for better visibility

Rys. 8. Wyniki prognozy modelu CNN+LSTM naniesione na dane testowe. Wyniki spłaszczono dla lepszej widoczności

differences in dataset characteristics or modeling assumptions. The models were trained using the methods specified in the papers, except for Ejdys et al. (2015), for which the method was not reported, and the *Keras Adam* optimizer was used instead. A full comparison is presented in Figure 9.

Limiting the input data to shorter time ranges improves the model's accuracy. By narrowing it down to just one month (March 2016), the CNN+LSTM MAE decreased to 13.64 PLN/MWh (0.1% Mean Absolute Percentage Error). This, however, sacrifices the long-range dependencies of learning.

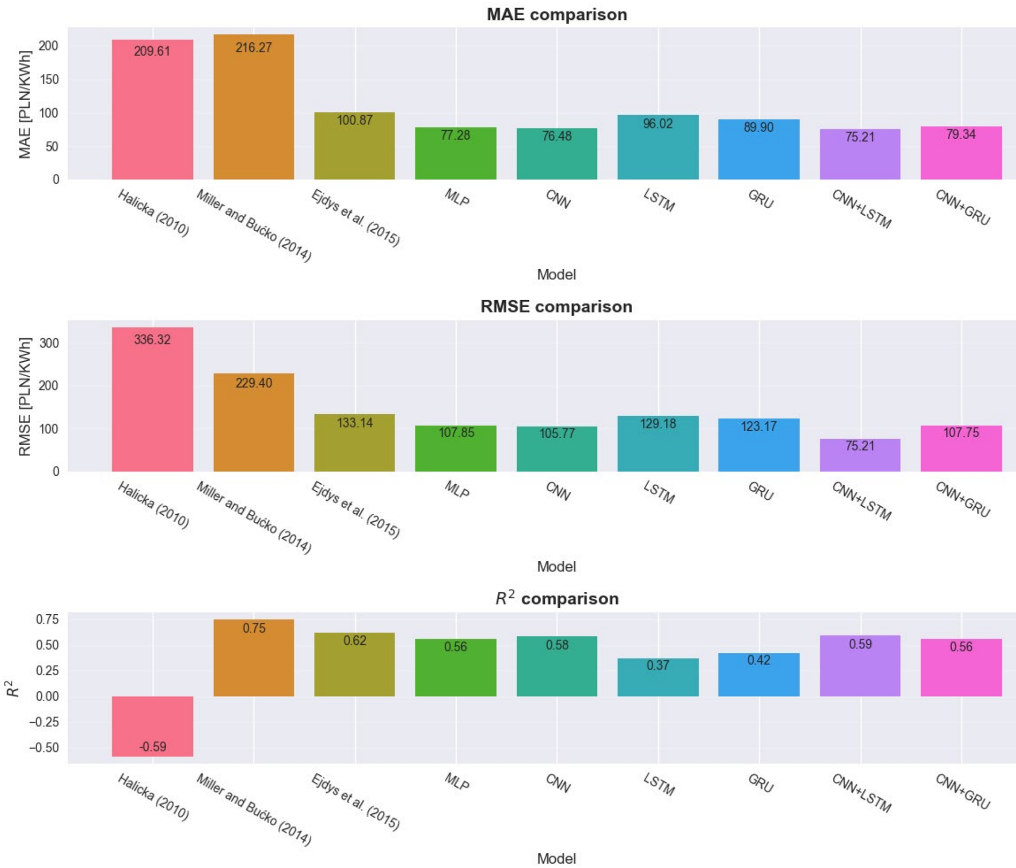

Fig. 9. Comparison of machine learning models according to different metrics

Rys. 9. Porównanie modeli uczenia maszynowego na podstawie różnych wskaźników

To investigate the impact of seasonality on the model's accuracy, an analysis was conducted for each season separately. Figure 10 presents the first 500 forecasting results, organized by season for better visibility. Figure 11 shows a comparison of performance metrics across seasons. The model performed best in forecasting data from spring, achieving an MAE of 50.43 PLN/MWh and RMSE of 94.39 PLN/MWh. In contrast, the worst performance was observed for summer data forecasting, reaching an MAE of 78.10 PLN/MWh and an RMSE of 200.64 PLN/MWh. The model achieves a lower MAE per season than overall, except for summer, while also exhibiting higher RMSE, except in spring. This may be attributed to the increased volatility of input variables within individual seasons. Since RMSE is more sensitive to outliers, this could also indicate that although average error decreased, a greater number of high-magnitude prediction errors may have occured.

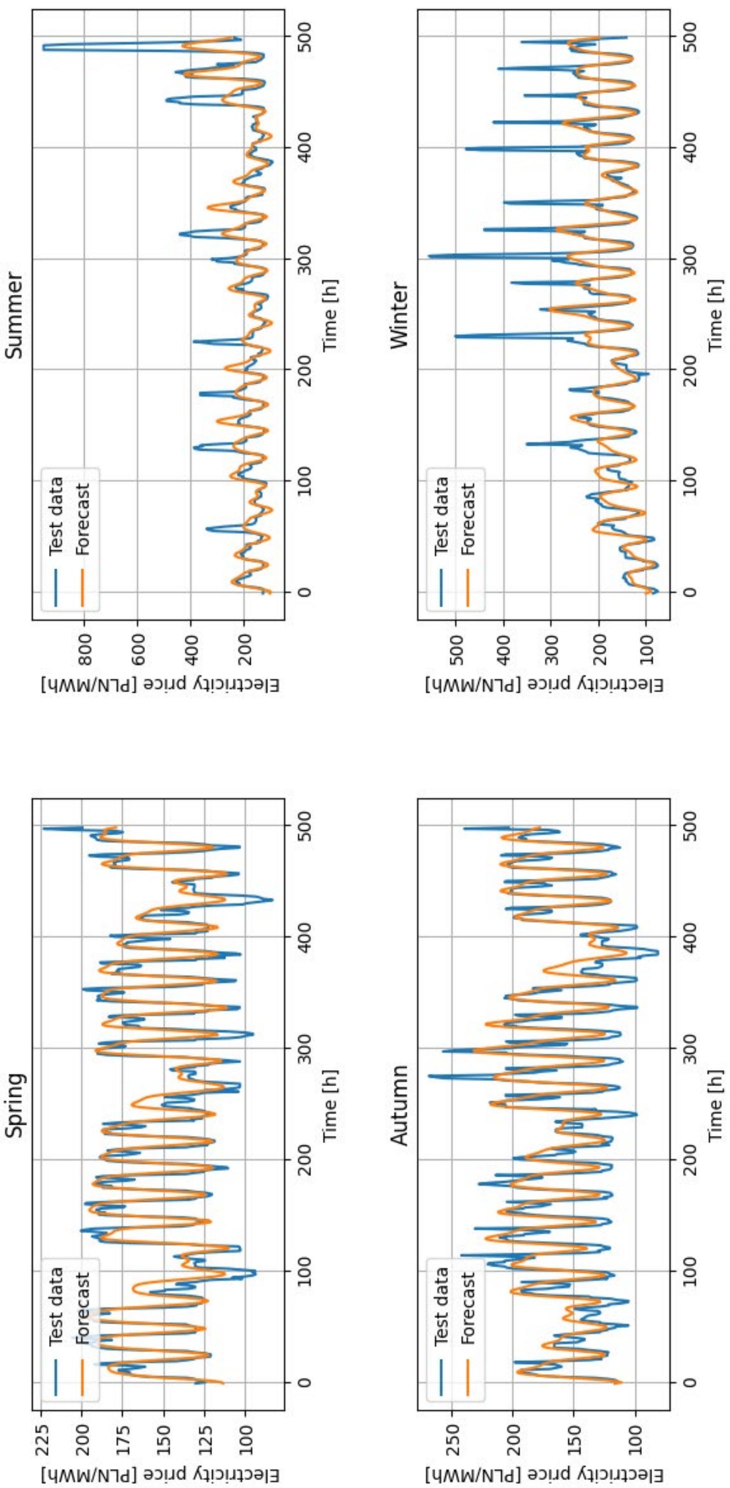

Fig. 10. Forecasting results per season

Rys. 10. Wyniki predykcji dla każdej z pór roku

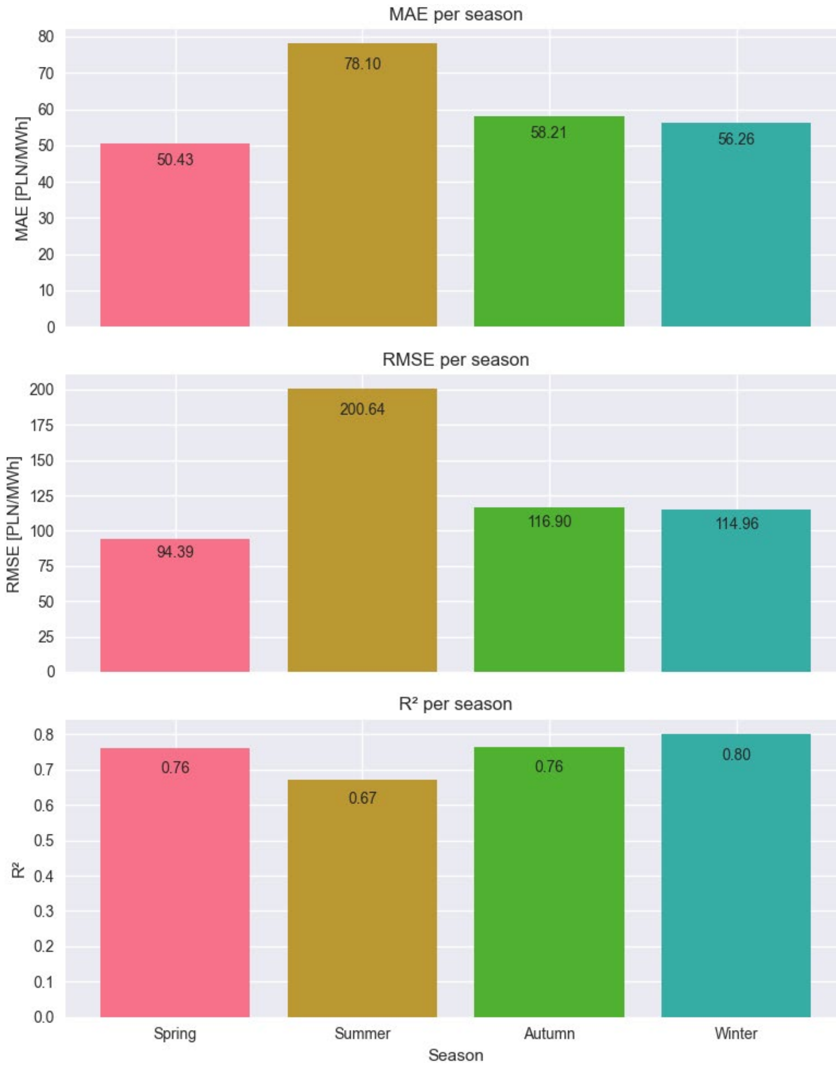

Fig. 11. Comparison of the model seasonal performance

Rys. 11. Porównanie dokładności modelu w różnych porach roku

Conclusions

The study focuses on electricity price forecasting in the Polish Day-Ahead Market using machine learning models trained on current market data and employing newer and more optimized methods than those used in the compared studies. The need for accurate and timely price forecasts is becoming increasingly important due to the growing share of intermittent renewable energy sources, which contribute to higher market volatility and uncertainty.

Initial feature extraction, supported by the literature, revealed that historical electricity prices and volumes, as well as wind speed, temperature, cloud cover, and temporal data, are the most relevant input data. This indicates that both market-related and weather-dependent variables play a crucial role in shaping short-term price dynamics and should be prioritized when developing forecasting models.

While it is generally assumed that RNNs perform better for time series forecasting, their complexity might lead to overfitting. They are also less tolerant to noise and external factors than MLPs. The sequential nature of RNNs may cause the model to expect patterns from random fluctuations or one-time events, leading to incorrect predictions. Explicitly feeding an MLP with temporal features may improve its long-term dependency learning, further diminishing the RNN's advantage. CNNs cover comparably small time windows in order to extract patterns. The higher accuracy of CNNs over RNNs may indicate that local patterns have a greater impact on electricity prices than long-term trends. The study showed that simpler models with fewer layers and neurons perform better in the context of electricity price than more complex architectures.

Combining a CNN for local pattern extraction with an RNN responsible for long-term temporal modeling appears to be an optimal approach, as it leads to lower forecasting errors and a higher coefficient of determination than either of its components. The accuracy of the models could be further improved by narrowing the input data to shorter time ranges. This would, however, limit their applicability since the long-term temporal patterns could not be observed.

Known machine learning-based electricity price forecasting techniques require continuous updating, as the power system is subject to ongoing changes due to a range of internal and external factors, such as shifts in the energy mix, fuel price fluctuations, and other socio-economic factors. Since machine learning has gained significant attention in recent years, the architectures and methods used to train machine learning models are being continually improved. This raises the need to update current forecasting models even further.

Future research should focus on further developing the hybrid models or exploring entirely different architectures. Another way to improve these findings would be to include the socio-economic factors that were outside of the scope of this study.

This paper was partially carried out as part of the statutory research activity of the Mineral and Energy Economy Research Institute of the Polish Academy of Sciences.

The Authors have no conflicts of interest to declare.

References

Adolfsen et al. 2022 – Adolfsen, J.F., Kuik, F., Schuler, T. and Lis, E.M. 2022. The impact of the war in Ukraine on euro area energy markets. *ECB Economic Bulletin* 4.

Agarap, A.F.M. 2018. A Neural Network Architecture Combining Gated Recurrent Unit (GRU) and Support Vector Machine (SVM) for Intrusion Detection in Network Traffic Data. *10th International Conference on Machine Learning and Computing*, pp. 26–30, DOI: 10.1145/3195106.3195117.

Alboom, R.K. 2025. *Energy Forecasting Inaccuracies and Their Direct Impact on Grid Performancee*. Dubai: Rochester Institute of Technology.

Belenguer et al. 2025 – Belenguer, E., Segarra-Tamarit, J., Pérez, E. and Vidal-Albalate, R. 2025. Short-term electricity price forecasting through demand and renewable generation prediction. *Mathematics and Computers in Simulation* 229, pp. 350–361, DOI: 10.1016/j.matcom.2024.10.004.

Benalcazar et al. 2017 – Benalcazar, P., Krawczyk, M. and Kamiński, J. 2017. Forecasting global coal consumption: An artificial neural network approach. *Gospodarka Surowcami Mineralnymi – Mineral Resources Management* 33(4), pp. 29–44, DOI: 10.1515/gospo-2017-0042.

Chreng et al. 2022 – Chreng, K., Lee, H.S., and Tuy, S. 2022. A Hybrid Model for Electricity Demand Forecast Using Improved Ensemble Empirical Mode Decomposition and Recurrent Neural Networks with ERA5 Climate Variables. *Energies* 15(19), DOI: 10.3390/en15197434.

Cho et al. 2014 – Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H. and Bengio, Y. 2014. Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation. Proceedings of the *2014 Conference on Empirical Methods in Natural Language Processing* (*EMNLP*), pp. 1724–1734, DOI: 10.3115/v1/D14-1179.

Dong et al. 2019 – Dong, S., Li, H., Wallin, F., Avelin, A., Zhang, Q. and Yu, Z. 2019. Volatility of electricity price in Denmark and Sweden. *Energy Procedia* 158, pp. 4331–4337, DOI: 10.1016/j.egypro.2019.01.788.

Dubus et al. 2023 – Dubus, L., Saint-Drenan, Y.-M., Troccoli, A., De Felice, M., Moreau, Y., Ho-Tran, L., Goodess, C., Amaro Silva, R. and Sanger, L. 2023. C3S Energy: A climate service for the provision of power supply and demand indicators for Europe based on the ERA5 reanalysis and ENTSO-E data. *Meteorological Applications* 30(5), DOI: 10.1002/met.2145.

Ejdys et al. 2015 – Ejdys, J., Halicka, K. and Godlewska, J. 2015. Forecasting the price of electricity on the energy exchange (*Prognozowanie cen energii elektrycznej na giełdzie energii*). *Zeszyty Naukowe Politechniki Śląskiej* 77, pp. 53–61 (*in Polish*).

Ember 2024. Yearly Electricity Data. [Online] https://ember-energy.org/data/yearly-electricity-data/ [Accessed: 2025-02-19].

Energy Instrat Hourly electricity prices (Day Ahead Market). [Online] https://energy.instrat.pl/ceny/energia-rdn-godzinowe/ [Accessed: 2025-03-24] (*in Polish*).

ERA5 Hourly data on single levels from 1940 to present. [Online] https://cds.climate.copernicus.eu/datasets/reanalysis-era5-single-levels?tab=overview [Accessed: 2025-03-24].

Halicka, K. 2010. Application of artificial neural networks in forecastng on energy market (*Wykorzystanie sztucznych sieci neuronowych do prognozowania cen na giełdzie energii*). *Rynek Energii* 1 (*in Polish*).

Haykin, S. 1999. Neural networks: a comprehensive foundation by Simon Haykin. *The Knowledge Engineering Review* 13(4) pp. 409–412.

Hochreiter, S. and Schmidhuber, J. 1997. Long Short-Term Memory. *Neural Computation* 9(8), pp. 1735–1780.

IMGW Index of /data/dane_pomiarowo_obserwacyjne/dane_meteorologiczne/dobowe. [Online] https://danepubliczne.imgw.pl/data/dane_pomiarowo_obserwacyjne/dane_meteorologiczne/dobowe/ [Accessed: 2025-03-24] (*in Polish*).

Jain, L.C. and Medsker, L.R. 1999. *Recurrent Neural Networks: Design and Applications*. Boca Raton: CRC Press.

Liu, Y. and Feng, C. 2023. Promoting renewable energy through national energy legislation. *Energy Economics* 118, DOI: 10.1016/j.eneco.2023.106504.

Mao et al. 2025 – Mao, X., Chen, S., Yu, H., Duan, L., He, Y. and Chu, Y. 2025. Simplicity in dynamic and competitive electricity markets: A case study on enhanced linear models versus complex deep-learning models for day-ahead electricity price forecasting. *Applied Energy* 383, DOI: 10.1016/j.apenergy.2024.125201.

Miller, A. and Bućko, P. 2014. Using artificial neural networks to forecasting the energy exchange price (*Zastosowanie sztucznych sieci neuronowych do prognozowania cen na giełdzie energii*). *Zeszyty Naukowe Wydziału Elektrotechniki i Automatyki Politechniki Gdańskiej* 40 (*in Polish*).

Obiora et al. 2024 – Obiora, S.C., Bamisile, O., Hu, Y., Ozsahin, D.U. and Adun, H. 2024. Assessing the decarbonization of electricity generation in major emitting countries by 2030 and 2050: Transition to a high share renewable energy mix. *Heliyon* 10(8), DOI: 10.1016/j.heliyon.2024.e28770.

O'Shea, K. and Nash, R. 2015. An Introduction to Convolutional Neural Networks. *International Journal for Research in Applied Science and Engineering Technology* 10(12), pp. 943–947, DOI: 10.48550/arXiv.1511.08458.

Pepłowska, M. 2025. Climate neutrality in Poland: the role of the coal sector in achieving the 2050 goal. *Polityka Energetyczna – Energy Policy Journal* 28(1), pp. 113–125, DOI: 10.33223/epj/202307 .

PSE SA TSO Reports / Market data. [Online] https://raporty.pse.pl/ [Accessed: 2025-03-24] (*in Polish*).

Tschora et al. 2022 – Tschora, L., Pierre, E., Plantevit, M. and Robardet, C. 2022. Electricity price forecasting on the day-ahead market using machine learning. *Applied Energy* 313, DOI: 10.1016/j.apenergy.2022.118752.

Yu et al. 2019 – Yu, Y., Si, X., Hu, C., and Zhang, J. 2019. A Review of Recurrent Neural Networks: LSTM Cells and Network Architectures. *Neural Computation* 31(7), pp. 1235–1270, DOI: 10.1162/neco_a_01199.

Zhong et al. 2020 – Zhong, H., Member, S., Tan, Z., Member, S., He, Y., Xie, L. and Kang, C. 2020. Implications of COVID-19 for the electricity industry: A comprehensive review. *CSEE Journal of Power and Energy Systems* 6(3), pp. 489–495, DOI: 10.17775/CSEEJPES.2020.02500.

Rafał Sowiński, Aleksandra Komorowska

Przewidywanie cen energii elektrycznej na Polskim Rynku Dnia Następnego z wykorzystaniem modeli uczenia maszynowego

Streszczenie

Ze względu na stale zmieniające się ceny energii elektrycznej, spowodowane zmianami w miksie energetycznym, regulacyjnymi i innymi czynnikami społeczno-ekonomicznymi, konieczne staje się okresowe weryfikowanie podejścia do prognozowania cen. Tradycyjne metody statystyczne mogą zawodzić w warunkach nasilonej zmienności, nieliniowych zależności i często zmieniających się cech wejściowych. Modele uczenia maszynowego, a zwłaszcza Sztuczne Sieci Neuronowe (SSN), potrafią skutecznie dostosowywać się do złożonych, niestacjonarnych wzorców w szeregach czasowych. W niniejszym badaniu opracowano i wytrenowano sześć różnych modeli SSN, korzystając z danych historycznych z Polskiego Rynku Dnia Następnego z lat 2016–2024. Cztery z tych modeli to czyste modele głębokiego uczenia: wielowarstwowy perceptron (MLP), sieć konwolucyjna (CNN), długa pamięć krótkotrwała (LSTM) oraz bramkowa jednostka rekurencyjna (GRU). Dwa pozostałe to architektury hybrydowe, oznaczone jako CNN+LSTM i CNN+GRU, łączą zdolność CNN do wychwytywania cech z węższych okien czasowych i umiejętność warstw rekurencyjnych do uczenia się zależności długoterminowych. Wydajność modeli oceniano na podstawie trzech miar: średniego błędu bezwzględnego (MAE), pierwiastka ze średniego błędu kwadratowego (RMSE) i współczynnika determinacji (R^2). Najlepsze wyniki osiągnęła architektura CNN+LSTM, uzyskując MAE na poziomie 75,21 zł/MWh, RSME równe 103,64 zł/MWh i R^2 wynoszące 0,59. Wyniki te mogą zostać wykorzystane do usprawnienia procesów prognozowania cen energii elektrycznej poprzez wskazanie wytycznych dotyczących projektowania modeli prognostycznych opartych na uczeniu maszynowym, co z kolei może wiązać się z wydajniejszym planowaniem działania systemu energetycznego.

Słowa kluczowe: energia elektryczna, Rynek Dnia Następnego, sztuczne sieci neuronowe, ceny, prognoza