Energy Proceedings
Vol 64, 2026
ISSN 2004-2965

Adaptive and Interpretable Electricity Price Forecasting via Sliding-Update BiLSTM with Attention#

Qichen Wang[1,2,3], Zhengmeng Hou[1,2,3*], Yilin Guo[1,2,3**], Liangchao Huang[1,2,3], Tianle Shi[1,2,3], Tian Zhang[1,2], Nan Cai[1,2]

1 Institute of Subsurface Energy Systems, Clausthal University of Technology, 38678 Clausthal-Zellerfeld, Germany

2 Research Centre of Energy Storage Technologies, Clausthal University of Technology, 38640 Goslar, Germany

3 Sino-German Research Institute of Carbon Neutralization and Green Development, Zhengzhou University, Zhengzhou 450001, China

(Corresponding Author: hou@tu-clausthal.de; yilin.guo@tu-clausthal.de)

ABSTRACT

Electricity price forecasting is highly challenging due to market volatility, abrupt spikes, and recurrent seasonal patterns. This study proposes an adaptive and interpretable deep learning framework, termed Sliding-Update Bidirectional LSTM with Attention (SU-BiLSTM-AM), which combines a dynamic update mechanism, bidirectional sequence modeling, and attention-based feature weighting to improve accuracy and robustness under volatile market conditions. Using day-ahead prices from the European Energy Exchange (EEX), SU-BiLSTM-AM consistently outperforms recurrent and sequence-learning baselines, reducing forecast errors by 14–23% compared with LSTM. The model demonstrates robustness under extreme regimes and maintains reasonable computational cost, making it suitable for intelligent energy systems, bidding strategies, and risk management in volatile electricity markets.

Keywords: electricity price forecasting, bidirectional LSTM, attention mechanism, intelligent energy systems, renewable integration, time series

NONMENCLATURE

Abbreviations	
LSTM	Long Short-Term Memory
SU	Sliding-Update Window
BiLSTM	Bidirectional LSTM
AM	Attention Mechanism
RNN	Recurrent Neural Network
EEX	European Energy Exchange
RNN	Recurrent Neural Network
GRU	Gated Recurrent Unit
TCN	Temporal Convolutional Network
Symbols	
x_t	Input vector at time step t
h_t	Hidden state at time step t
c_t	Cell state at time step t
f_t	Forget gate vector at time step t
i_t	Input gate vector at time step t
o_t	Output gate vector at time step t
W_f, W_i, W_o	Weight matrices for gates
b_f, b_i, b_o	Bias vectors for gates
e_t	Attention score for time step t
α_t	Normalized attention weight at t
v, W_h, b_h	Trainable parameters in AM
H	Final context vector
p_t	Electricity price at time step t
X_t	Input vector at time step t using SU
T	Length of sliding window

1. INTRODUCTION

Electricity price forecasting plays a vital role in modern energy systems, directly affecting bidding strategies, renewable energy integration, and risk management. However, electricity markets are characterized by nonlinearity, strong seasonality, and abrupt price spikes, making accurate forecasting particularly challenging [1]. Price fluctuations arise from diverse factors such as transmission constraints, weather conditions, and sudden changes in demand, which together produce complex temporal dependencies that traditional statistical models often fail to capture [2].

Classical time series methods such as ARIMA, and machine learning techniques like support vector regression, have achieved limited success, largely due to their assumptions of stationarity and their inability to model long-range nonlinear dependencies [3]. Deep learning has emerged as a promising alternative, with recurrent neural networks, LSTM networks, and their variants demonstrating significant improvements in modeling temporal dynamics [4]. Nevertheless,

This is a paper for the 17th International Conference on Applied Energy (ICAE2025), December 8-12, 2025, Bangkok, Thailand.

conventional LSTM and GRU models may underestimate extreme events, while Transformer-based models require large datasets and high computational cost, limiting their applicability in volatile electricity markets [5].

Hybrid approaches have therefore gained attention, combining different architectures or introducing mechanisms such as attention and residual correction to enhance responsiveness to volatility [6]. These methods highlight the importance of model designs that are specifically tailored to the unique dynamics of electricity markets.

Building on this perspective, the present study develops SU-BiLSTM-AM, a hybrid forecasting framework designed to improve accuracy and robustness under conditions of high volatility. The framework integrates three components:

- a sliding update mechanism to ensure adaptability to evolving market conditions;
- a bidirectional LSTM backbone to capture dependencies from both past and future contexts;
- an attention layer to emphasize critical short-term lags relevant for near-term prediction.

The objective of this work is to provide a forecasting model that combines accuracy, stability, and computational efficiency. Validation on European Energy Exchange data demonstrates significant improvements over recurrent baselines and state-of-the-art sequence models. In particular, SU-BiLSTM-AM reduces forecast errors by more than 20% compared with standard LSTM and exhibits enhanced robustness under extreme price regimes. These features underline its potential value for real-time electricity market applications.

2. THEORY AND CALCULATION

2.1 Recurrent modeling background

Recurrent neural networks (RNNs) capture sequential dependencies by updating hidden states iteratively [7]. LSTM networks improve on RNNs with gating mechanisms:

$$f_t = \sigma(W_f[h_{t-1}, x_t] + b_f) \tag{1}$$

$$i_t = \sigma(W_i[h_{t-1}, x_t] + b_i) \tag{2}$$

$$c_t = f_t \odot c_{t-1} + i_t \odot tanh(W_c[h_{t-1}, x_t] + b_c) \tag{3}$$

$$h_t = o_t \odot tanh(c_t) \tag{4}$$

where f_t, i_t, o_t denote the gates, c_t the cell state, and $\mathrm{h_t}$ the hidden output. While effective, conventional LSTMs rely only on past information and often fail to capture extreme volatility.

2.2 Sliding update mechanism

To address distributional drift, a sliding update strategy was introduced. At each new time step, the latest observation is appended to the window and the oldest value is discarded, maintaining a constant input length [8]:

$$X_t = \{p_{t-T+1}, \dots, p_t\} \Rightarrow X_{t+1} = \{p_{t-T+2}, \dots, p_{t+1}\} \tag{5}$$

where p_t denotes the electricity price at time t. This ensures that the input sequence continuously reflects the most recent market conditions.

2.3 Bidirectional LSTM

A BiLSTM layer extends LSTM by processing the sequence in both forward and backward directions [9]. The two hidden states are concatenated:

$$h_t = \overrightarrow{h_t} \oplus \overleftarrow{h_t} \tag{6}$$

This bidirectional representation captures dependencies from both past and near-future contexts within the input window, which is particularly beneficial for cyclical energy market data.

2.4 Attention mechanism

To further enhance predictive power, an attention layer was incorporated. The attention score for each hidden state $\mathrm{h_t}$ is computed as:

$$e_t = v^\top tanh(W_h h_t + b_h) \tag{7}$$

$$\alpha_t = \frac{\exp(e_t)}{\sum_{k=1}^{T} \exp(e_k)} \tag{8}$$

The final sequence representation is obtained by:

$$H = \sum_{t=1}^{T} \alpha_t h_t \tag{9}$$

This mechanism highlights short-term lags with strong predictive relevance, improving both accuracy and interpretability.

2.5 Model integration

The proposed SU-BiLSTM-AM integrates three modules into a unified framework. The sliding-update strategy adapts inputs to evolving market conditions, the bidirectional LSTM captures temporal dependencies in both directions, and the attention mechanism emphasizes critical lags for accurate and interpretable forecasts. The combined representation is fed into a dense output layer to generate predictions, as illustrated in Fig. 1.

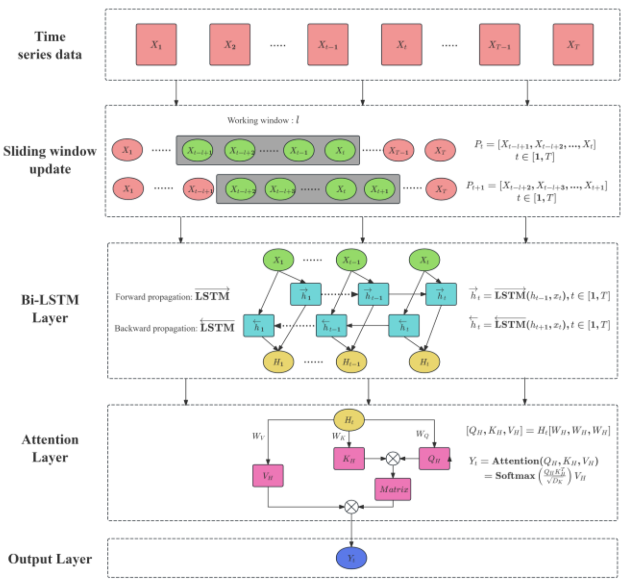

Fig. 1 Framework of SU-BiLSTM-AM model

3. MATERIAL AND METHODS

3.1 Data collection

The dataset used in this study was obtained from the European Energy Exchange (EEX) [10], covering the period from January 2018 to June 2023, as shown in Fig. 2 and Fig. 3. Day-ahead electricity prices were recorded at 15-minute intervals, resulting in 96 observations per day. A total of 175,200 data points from 2018-2022 were used for training and validation, while 17,184 data points from January-June 2023 were reserved for testing. The price series exhibits clear daily, weekly, and seasonal cycles, as well as significant volatility in 2022, with peaks above €300/MWh. These characteristics highlight the nonlinear and nonstationary nature of electricity prices and justify the use of advanced deep learning approaches.

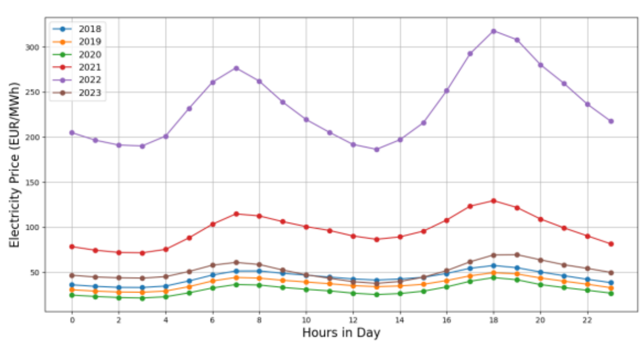

Fig. 2 Electricity price patterns at daily scales in *in the German energy market*

window lengths

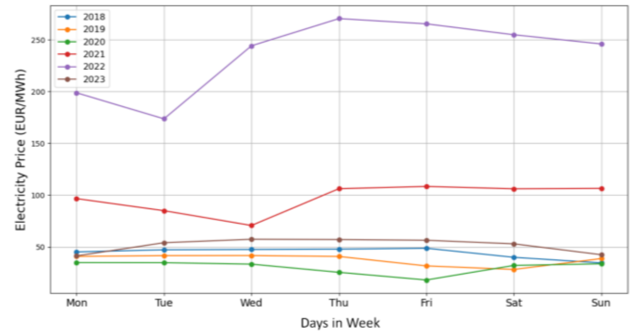

Fig. 3 Electricity price patterns at weekly scales in the German energy market

3.2 Feature construction

To improve forecasting accuracy, ten exogenous features were incorporated in addition to historical prices. These variables capture demand, renewable generation, weather, and calendar effects, ensuring physical and economic relevance. All series were resampled at 15-minute intervals and standardized using z-score normalization. The resulting input tensor has dimensions (T,F), where T is the sliding window length and F=10 the number of features. Fig. 4 displays the attention heatmap across time steps (T1-T10) and features (F1-F10).

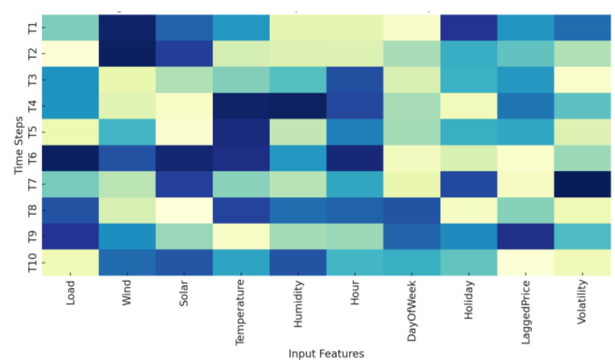

Fig. 4 Attention heat map of multivariate features

3.3 Experimental setup

The dataset was split into 80% training, 10% validation, and 10% testing, following chronological order. All models were implemented under identical conditions. Training employed the Adam optimizer with an initial learning rate of 0.002, batch size of 512, and early stopping based on validation loss. Each run was repeated with multiple random seeds to ensure robustness. Baseline models included RNN, GRU, LSTM, BiLSTM, Transformer, and TCN, trained with selected features suggested in prior studies. Experiments were executed on GPUs with mixed-precision training to accelerate computation.

3.4 Evaluation metrics

Root Mean Square Error (RMSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE) are used as evaluation metrics, defined as:

$$RMSE = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(P_i - \hat{P}_i)^2} \tag{10}$$

$$MAE = \frac{1}{N}\sum_{i=1}^{N}\left|P_i - \hat{P}_i\right| \tag{11}$$

$$MAPE = \frac{1}{N}\sum_{i=1}^{N}\left|\frac{P_i - \hat{P}_i}{P_i}\right| \times 100\% \tag{12}$$

where, P_i represents the actual electricity price at time point i, $\hat{P}_i$ is the predicted price at the same time point, and *N* is the number of time points in the test period. Together, these metrics provide a comprehensive evaluation of model accuracy, stability, and practical significance.

4. RESULTS

Electricity price forecasting is characterized by high volatility and recurrent seasonality. To evaluate the proposed SU-BiLSTM-AM, extensive experiments were conducted against recurrent baselines (RNN, GRU, LSTM, BiLSTM) and sequence models (Transformer, TCN) [11-12]. All models were trained under identical settings and repeated across multiple runs. Performance was assessed using RMSE, MAE, and MAPE as defined in Section 3.4.

4.1 Predictive accuracy

The averaged error metrics across repeated runs and multiple sliding window settings are summarized in Fig. 5 and detailed values for the optimal performance configuration for each model are provided in Table 1. All results represent mean values across independent runs, with error bars denoting one standard deviation. LSTM consistently outperforms RNN and GRU, confirming the advantage of gated recurrence in electricity price modeling. Building on this, SU-BiLSTM-AM further reduces errors: compared with LSTM, the average RMSE decreases by 14.0%, MAE by 22.8%, and MAPE by 23.0%. Moreover, the standard deviations (sd) are substantially smaller, suggesting more stable performance across independent runs. In contrast, Transformer and TCN do not surpass recurrent structures in this domain, despite their strong performance in natural language and generic sequence modeling. The discrepancy likely reflects the strong periodicity, high noise, and relatively limited data volume in electricity markets, which favor inductive biases inherent in recurrent models combined with attention.

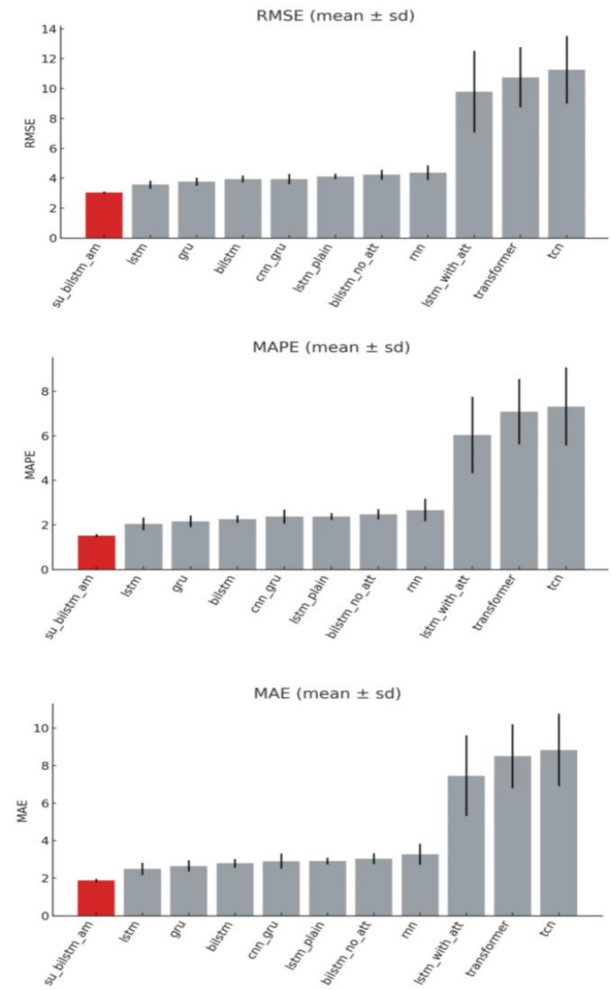

Fig. 5 Comparative error performance of baseline

Table 1 Comparative error metrics (best-performing configuration per model)

Model	Window length	RMSE mean	MAE mean	MAPE mean
su_bilstm_am	24	2.98165	1.81070	1.47529
lstm	48	3.46635	2.34410	1.91207
gru	96	3.51428	2.37874	1.93139
bilstm	24	3.71897	2.53876	2.04752
cnn_gru	48	3.78261	2.71377	2.24536
rnn	48	3.90442	2.70987	2.14785
lstm_plain	72	4.05886	2.85935	2.36090
bilstm_no_att	24	4.14360	2.93067	2.38089
lstm_with_att	24	5.44086	4.00331	3.30366
transformer	24	8.55742	6.73984	5.61337
tcn	24	8.58719	6.61131	5.33687

4.2 Sensitivity to input window length

Given the importance of sliding window design in time series forecasting, we performed a sensitivity analysis by varying the input length (L = 24, 48, 72, 96, 168). The results (Fig. 6) show that SU-BiLSTM-AM yields stable accuracy across all window lengths, with RMSE varying only within 2.68-3.01. In contrast, several baselines demonstrated larger fluctuations, particularly Transformer and TCN at longer horizons. This indicates that the proposed model does not critically depend on manual tuning of the input window, supporting its applicability in real-world forecasting where the optimal horizon length is unknown.

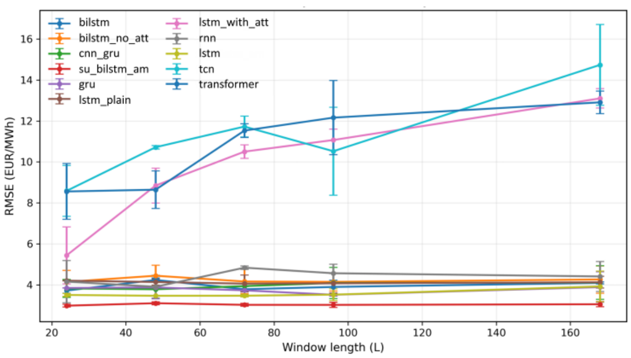

Fig. 6 Sensitivity of forecasting error to different input window lengths

4.3 Trade-off between accuracy and computational cost

To examine efficiency, we compared RMSE against per-run training time across all models (Fig. 7). SU-BiLSTM-AM achieves the lowest RMSE among all models while maintaining moderate training times. Transformer and TCN models demand substantially higher computational resources yet deliver higher errors. Conventional recurrent models such as GRU and LSTM are computationally efficient but consistently less accurate. These results indicate that SU-BiLSTM-AM achieves favorable accuracy-efficiency balance, making it practical for large-scale forecasting.

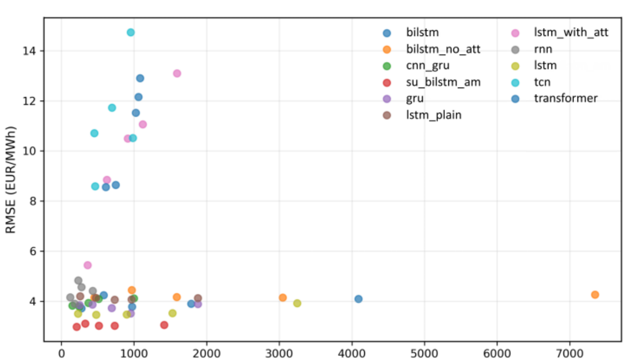

Fig. 7 Computation cost vs. accuracy for different models

4.4 Attention-based interpretability

To examine how temporal dependencies are exploited, we analyzed attention distributions. Fig. 8 overlays the average attention weight across input lags for varying window lengths *L*. A consistent pattern emerges: the model assigns disproportionately high weight to the most recent observations, reflecting the dominant role of short-term dynamics. For shorter windows (*L=24, 48*), attention peaks sharply at the latest hours, whereas for longer windows (*L=72, 96, 168*), the distribution flattens but still privileges the near past. This suggests that predictive power is primarily driven by short-term dependencies, aligning with domain knowledge that electricity prices are shaped by immediate supply-demand shifts.

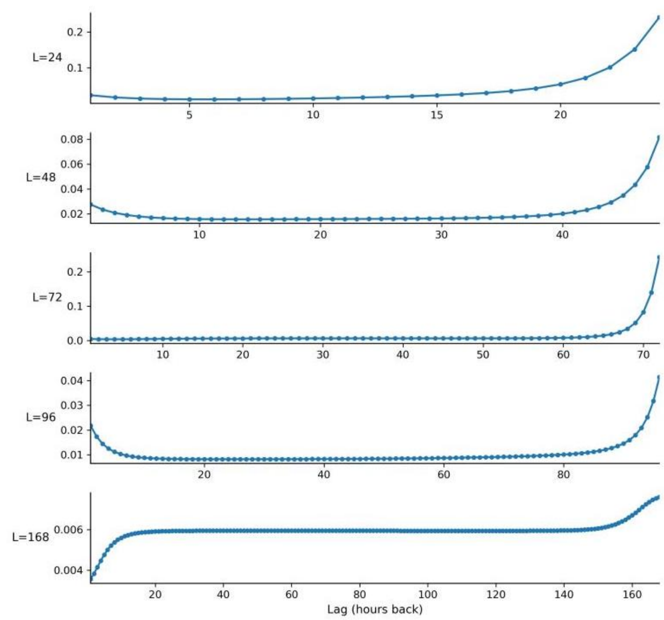

Fig. 8 Attention profiles at different input lengths

Across all experiments, SU-BiLSTM-AM consistently demonstrates superior accuracy, stability, and interpretability compared with baseline models. The improvements are statistically significant, robust under extreme price regimes, and achieved without prohibitive computational cost. Sensitivity analyses reveal that performance is stable across varying input horizons, while attention visualizations provide interpretable evidence that the model captures the dominant influence of short-term market dynamics. Together, these results suggest that the proposed architecture provides a reliable and practical framework for electricity price forecasting, with potential applicability to other energy time series characterized by high volatility and recurrent seasonality.

5. CONCLUSIONS

A hybrid forecasting framework, SU-BiLSTM-AM, was developed to address volatility and nonstationarity in electricity markets. The model integrates three components: a sliding update mechanism for adaptability, a bidirectional LSTM backbone for capturing cross-temporal dependencies, and an attention layer for amplifying short-term dynamics. Experiments on EEX data from 2018-2023 demonstrated that SU-BiLSTM-AM consistently improves forecasting accuracy compared with recurrent and sequence baselines. On average, errors were reduced by more than 20% relative to LSTM, with enhanced robustness under extreme market regimes and faster convergence than conventional recurrent models. The results indicate that the synergy of sliding updates, bidirectional modeling, and attention is essential for sustaining predictive accuracy in volatile energy markets. Remaining challenges include the reliable forecasting of rare and abrupt price spikes, suggesting that future work should explore anomaly-aware and regime-switching extensions. Overall, the proposed framework provides a practical and interpretable approach to electricity price forecasting, with potential applications in renewable integration, demand management, and market risk assessment.

ACKNOWLEDGEMENT

This work was supported by European Union's "Horizon Europe programme"-LOC3G (Grant No. 101129729) and the Henan Centre for Outstanding Overseas Scientists in Green Low-Carbon Technology and Energy Transition (Grant No. GZS2024001). We would like to thank the anonymous reviewers for their invaluable comments and suggestions.

REFERENCE

[1] De Gooijer JG, Hyndman RJ. 25 years of time series forecasting. Int J Forecast 2006;22(3):443-73.

[2] Beck HP, Hou Z, Wang Q, et al. The German energy transition after the Russia-Ukraine war–challenges and opportunities. Carbon Neutral Syst 2025;1:4.

[3] Qiu D, Dong Z, Ruan G, Zhong H, Strbac G, Kang C. Strategic retail pricing and demand bidding of retailers in electricity market: a data-driven chance-constrained programming. Adv Appl Energy 2022;7:100100.

[4] Wang JQ, Du Y, Wang J. LSTM based long-term energy consumption prediction with periodicity. Energy 2020;197:117197.

[5] Wang Q, Hou Z, Guo Y, Huang L, Fang Y, Sun W, Ge Y. Enhancing energy transition through sector coupling: a review of technologies and models. Energies 2023;16:5226.

[6] Eren Y, Küçükdemiral İ. A comprehensive review on deep learning approaches for short-term load forecasting. Renew Sustain Energy Rev 2024;189:114031.

[7] Yu Y, Si X, Hu C, Zhang J. A review of recurrent neural networks: LSTM cells and network architectures. Neural Comput 2019;31(7):1235-70.

[8] Beck HP, Hou MZ, Wang Q, Shi T, Guo Y, Huang L, et al. A comprehensive review of sustainable energy systems in the context of the German energy transition—part 1: development status and system foundations. Carbon Neutral Syst 2025;1(1):12.

[9] Siami-Namini S, Tavakoli N, Namin AS. The performance of LSTM and BiLSTM in forecasting time series. In: Proc IEEE Int Conf Big Data (Big Data); 2019 Dec; Los Angeles. p. 3285-92.

[10] European Energy Exchange AG (EEX). https://www.eex.com/en/. Accessed 06 Aug 2025.

[11] Lago J, De Ridder F, De Schutter B. Forecasting spot electricity prices: deep learning approaches and empirical comparison of traditional algorithms. Appl Energy 2018;221:386-405.

[12] Hewamalage H, Bergmeir C, Bandara K. Recurrent neural networks for time series forecasting: current status and future directions. Int J Forecast 2021;37(1):388-427.