ENHANCING SHORT-TERM SOLAR GENERATION FORECASTING ACCURACY USING A CNN-TRANSFORMER MODEL

ANH THO – VU THI[1*], TUAN ANH NGUYEN[1], MANH HAI PHAM[1]
[1] *Electric Power University, Hanoi, Vietnam;*
** Corresponding author: thovta@epu.edu.vn*
DOIs: https://www.doi.org/ 10.46242/jstiuh.v83i5.6178

Abstract. The increasing penetration of solar energy into power systems poses significant challenges to grid dispatch and operation due to its inherent variability and strong dependence on meteorological conditions. Recently, Transformer-based deep learning models have demonstrated great potential in time-series analysis owing to their self-attention mechanism. However, the vanilla Transformer architecture tends to prioritize long-term dependencies, which may result in an insufficient capture of local, short-term fluctuations in meteorological data. This study proposes a hybrid CNN-Transformer model for short-term solar power forecasting. A one-dimensional convolutional neural network (1D-CNN) is employed to extract local spatio-temporal features, which are subsequently fed into the Transformer block to learn long-term dependencies. Experimental results conducted on a real-world dataset from a 50MW solar power plant indicate that the CNN-Transformer model achieves superior performance. Specifically, the model significantly reduces the error metrics, with an RMSE of 2,919.42 kW and an NMAPE of 3.45%, consistently outperforming the vanilla Transformer model.

Keywords. Power generation forecasting, solar power, time series forecasting, CNN-Transformer.

1 INTRODUCTION

The transition to renewable energy sources, particularly solar power, has become a major global trend. However, the power output of solar power plants depends directly on meteorological conditions [1]. Sudden fluctuations in these conditions pose significant risks to grid stability [2]. Consequently, short-term solar power forecasting (ranging from minutes to hours) plays a pivotal role in energy management systems (EMSs) [3], [4]. Comprehensive reviews have further shown that photovoltaic forecasting methods increasingly combine meteorological variables with data-driven learning algorithms to improve robustness across different forecast horizons [5]-[7].

Traditional forecasting methods, such as ARIMA and SVR, often struggle to represent highly nonlinear relationships [8]. Although deep learning models such as RNNs, LSTMs, and GRUs can model nonlinear temporal patterns, their sequential processing limits parallelization and can make long-sequence learning computationally demanding [9]. The Transformer, based on multi-head self-attention [10], processes all time steps in parallel and learns the relative importance of observations across a sequence. Recent studies have reported improved PV forecasting accuracy using support-vector, tree-based, convolutional, and recurrent deep-learning architectures [11]-[15].

Nevertheless, solar power data contain localized fluctuations, such as rapid irradiance changes caused by transient cloud cover. A vanilla Transformer may not explicitly emphasize these short-range patterns. To address this limitation, this study integrates a one-dimensional convolutional neural network (1D-CNN) [16] into the Transformer architecture. The CNN serves as a local feature extractor before the subsequent Transformer layers learn longer-range dependencies. Hybrid CNN-based models have shown that convolutional layers can strengthen local feature extraction, while CNN-LSTM and CNN-LSTM-Transformer architectures can jointly represent short-term variations and longer temporal dependencies [17]-[19]. More recently, Transformer and recurrent architectures have also been investigated for multi-step PV power forecasting [20].

The remainder of this paper is organized as follows. Section II describes the methodology, Section III presents the results and discussion, and Section IV concludes the study.

© 2026 Industrial University of Ho Chi Minh City

2 METHODOLOGY

2.1 Dataset and Data Preprocessing

This study utilizes a historical operational dataset from a solar power plant with a rated capacity of 50 MW. The data was collected over a one-year period, from January 1, 2022, to December 31, 2022, with a recording resolution of 5 minutes. The input features comprise micro-meteorological parameters, specifically ambient temperature, PV panel temperature, and solar irradiance, while the solar power output serves as the target variable.

The raw data underwent cleaning to handle missing values and was normalized to the range [0, 1] using MinMaxScaler. A sliding-window technique with a 24-hour look-back period (288 five-minute time steps) was then used to generate time-series samples. The 2022 dataset was divided chronologically, with the first 80% used for training and the remaining 20% retained as a held-out test set for model assessment and selection. In addition, an entirely separate forecasting dataset collected from June 10 to June 12, 2023 was used for independent out-of-sample evaluation. To generate forecasts for all 864 observations in this three-day period, the 288 observations immediately preceding the forecasting period were used as the initial input sequence. Specifically, measurements from June 9, 2023 provided the 24-hour historical context required to forecast the first observation on June 10. These preceding observations were used only as model inputs and were excluded from the calculation of RMSE and NMAPE. All final results reported in Table 1 and Figures 2-3 were obtained from the 864 forecast-measurement pairs between June 10 and June 12, 2023.

2.2 Model Architecture: CNN-Transformer

The proposed CNN-Transformer architecture combines the ability of convolutional neural networks (CNNs) to extract local spatio-temporal features with the ability of self-attention to capture longer-range dependencies. The processing pipeline comprises three functional blocks: local feature extraction, global representation learning, and forecasting. This design is consistent with recent hybrid forecasting studies in which convolutional layers are used for local representation learning and recurrent or attention-based modules are used for longer-range temporal modeling [15], [17]-[19].

2.2.1 Local Feature Extraction Block (1D-CNN Block)

The model input is a time-series matrix $X \in \mathbb{R}^{(T \times D)}$, where T denotes the sliding-window length (T = 288 time steps, corresponding to the previous 24 hours at five-minute resolution) and D denotes the number of micro-meteorological features (D = 3: solar irradiance, ambient temperature, and PV panel temperature). Rather than feeding the raw data directly into the Transformer, X is first processed by a one-dimensional convolutional neural network (1D-CNN) block that functions as a local feature extractor.

- **Conv1D Layer**

This layer utilizes kernels that slide along the temporal axis to detect abrupt short-term fluctuations, such as sudden drops in solar irradiance caused by cloud cover. The convolutional layer is designed to extract local features from the input sequence. Similar to traditional CNN architectures [16], the convolution operation is performed between the kernels and the input data. However, in this study, the Rectified Linear Unit (ReLU) activation function is employed instead of the Sigmoid function to enhance model convergence and mitigate the gradient saturation problem. The operation of the 1D-CNN layer is mathematically formulated as follows:

$$H_{conv} = \mathrm{ReLU}(W_c * X + b_c) \quad (1)$$

Where W_c and b_c denote the weight matrix and the bias vector of the convolutional layer, respectively, and ***** represents the convolution operator.

- **MaxPooling1D Layer**

Following the Conv1D layer, the data pass through a MaxPooling1D layer for down-sampling. This operation retains the most prominent local features and halves the temporal sequence length. The reduction lowers the quadratic computational cost, $O(T^2)$, of the attention matrix in the subsequent Transformer block and may also help reduce overfitting. The output of this block is denoted by H_{pool}, in which the temporal dimension is reduced from T to T/2.

2.2.2 Global Representation Learning Block (Transformer Encoder Block)

The feature matrix H_{pool} is subsequently fed into the Transformer Encoder to analyze long-term temporal patterns (e.g., diurnal power cycles). This block comprises two core components:

- **Multi-Head Self-Attention (MHSA) Mechanism**

Unlike Recurrent Neural Networks (RNNs) that process data sequentially, the Self-Attention mechanism enables the model to consider all time steps simultaneously, computing the mutual influence between them. The input data is multiplied by learnable weight matrices to generate three distinct matrices: Query (Q), Key (K), and Value (V). The Attention Score is determined using the following formula [10]:

$$\text{Attention}(Q, K, V) = \text{softmax}(\frac{QK^T}{\sqrt{d_k}})\ V \quad (2)$$

where d_k denotes the dimension of the query and key vectors. The scaling factor $1/\sqrt{d_k}$ prevents excessively large dot products from driving the softmax function into regions with very small gradients during training.

To capture diverse aspects of the data (for instance, one head may specialize in learning temperature patterns while another focuses on solar irradiance), the model employs multiple parallel attention mechanisms, known as Multi-Head Attention. The outputs from these individual heads are subsequently concatenated and linearly transformed to produce the final representation.

- **Add & Norm Mechanism and Feed-Forward Network (FFN)**

To prevent information loss as the neural network increases in depth, the architecture incorporates Residual Connections (skip connections) combined with Layer Normalization [10]. This mechanism ensures stable gradient flow and facilitates the training of deeper architectures. Subsequently, the normalized output is passed through a Position-wise Feed-Forward Network, which consists of two linear transformations separated by a non-linear activation function, to further project the features into a high-dimensional representation space.

$$H_{attn} = \text{LayerNorm}(H_{pool} + \text{MHSA}(H_{pool})) \quad (3)$$

Subsequently, the data is processed through a Position-wise Feed-Forward Network (FFN). In this architecture, the FFN is constructed using 1D-convolutional layers with a kernel size of 1 to enhance the non-linearity of the feature representation [10]:

$$H_{out} = \text{LayerNorm}(H_{attn} + \text{FFN}(H_{attn})) \quad (4)$$

2.2.3 Output Forecasting Block

The output from the Transformer Encoder is a high-level representation matrix containing complex spatio-temporal information. To produce the final power forecast, the data undergoes the following stages:

- Global Average Pooling 1D (GAP): This layer computes the mean value along the temporal axis, compressing the entire sequence into a single feature vector. This vector encapsulates the essential features derived from both local and global learning processes.
- Output Layer: Finally, a Fully Connected (Dense) layer utilizing a linear activation function maps this feature vector to a single real scalar value, $\hat{y}$**.** This value represents the predicted solar power output (kW) for the subsequent time step.

The seamless integration of these three structural blocks ensures that the model remains robust against transient microscopic noise (mitigated by the CNN) while preserving the macroscopic cyclical patterns of the solar power system (captured by the Transformer).

2.3 Forecasting Procedure

Based on the proposed architecture, the forecasting procedure distinguishes between the historical 2022 dataset used for model development and the independent dataset from June 10 to June 12, 2023 used for final forecasting evaluation, as illustrated in Figure 1.

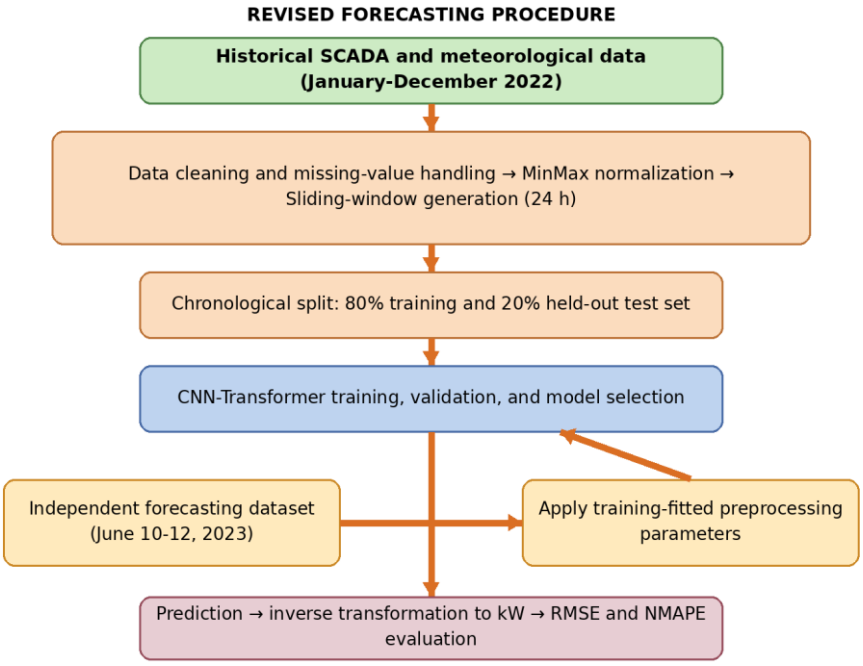


Figure 1: Revised forecasting procedure distinguishing model-development data from the independent forecasting dataset.

Step 1: Data Initialization and Preprocessing

Historical SCADA data from 2022 are first cleaned to address missing or abnormal values. MinMaxScaler is fitted using the training portion of the historical dataset and is then used to normalize the corresponding variables. The normalized observations are converted into time-series samples using a sliding window with a 24-hour look-back period (288 five-minute time steps).

Step 2: Data Splitting

After sliding-window generation, the 2022 samples are divided chronologically. The first 80% are used to train the models, while the remaining 20% form a held-out test set used for preliminary generalization assessment and model selection. The independent dataset from June 10 to June 12, 2023 is not included in this split.

Step 3: Model Configuration and Construction

The proposed CNN-Transformer architecture is implemented with the specific hyperparameters detailed in the following sections. These configurations are selected to balance the model's complexity and its ability to capture both local and global features. Model Architecture Parameters as below:

1D-CNN Block:

- Number of Filters: 32
- Kernel Size: 3
- Activation Function: ReLU
- Max Pooling Size: 2

Transformer Block:

- Embedding Dimension d_{model}: 32
- Number of Attention Heads: 2 (To capture independent data perspectives simultaneously)
- Feed-Forward Network (FFN) Units: 32
- Dropout Rate: 0.1 (To mitigate overfitting by randomly deactivating 10% of neurons)

Training Configurations:

- Batch Size: 128 (The model updates its weights after processing every 128 samples).
- Epochs: 50 (The model iterates through the entire training dataset 50 times).
- Optimizer: Adam (Adaptive Moment Estimation).

- Loss Function: Mean Squared Error (MSE).

Step 4: Model Training

The training set is fed into the CNN-Transformer network in batches. During each epoch, the model performs forward propagation to generate the predicted power output, which is then compared against the actual labels to calculate the loss. Subsequently, the backpropagation algorithm is employed to update the network weights. This iterative process continues until the loss function (MSE) converges or the maximum number of predefined epochs is reached.

Step 5: Forecasting and Performance Evaluation

After training and model selection, the trained vanilla Transformer and CNN-Transformer models are applied to the independent forecasting period from June 10 to June 12, 2023. To forecast the first five-minute observation on June 10, each model is initialized using the immediately preceding 288 observations, corresponding to the 24-hour period on June 9, 2023. A rolling one-step-ahead forecasting procedure is then used throughout the three-day evaluation period. At each forecasting instant, the input window is updated with the most recent measured observations, rather than with previous model forecasts, and the latest 288 observations are used to predict the subsequent five-minute power output. This procedure avoids recursive accumulation of forecast errors. The June 9 observations serve only as historical input context and are not included in the evaluation. The same preprocessing parameters fitted on the training data are applied to the independent dataset. Because the model outputs are normalized, an inverse transformation is subsequently performed to restore the predictions to the original power scale in kW. The final RMSE and NMAPE values reported in Table 1 and Figures 2-3 are calculated exclusively from the 864 forecast-measurement pairs obtained between June 10 and June 12, 2023.

2.4 Evaluation Metrics

The performance of the proposed model is quantitatively evaluated using two widely adopted statistical indicators:

- **Root Mean Square Error (RMSE)**

The RMSE measures the deviation between predicted values and actual measurements and gives greater weight to large errors. A lower RMSE indicates better forecasting accuracy. It is defined as follows [21]:

$$\mathrm{RMSE} = \sqrt{\frac{1}{\mathrm{n}}\sum_{\mathrm{i}=1}^{\mathrm{n}}\left(\hat{\mathrm{y}}_{\mathrm{i}} - \mathrm{y}_{\mathrm{i}}\right)^2} \tag{5}$$

Where:

- $\hat{y}_i$ is the predicted power output (kW);
- y_i is the actual power output (kW);
- n is the total number of data points (samples).

- **Normalized Mean Absolute Percentage Error (NMAPE)**

In PV power systems, NMAPE is useful because output power commonly falls to zero at night. The metric normalizes the absolute forecasting error by the plant's rated capacity of 50 MW (50,000 kW), rather than by the instantaneous measured output, thereby avoiding division by zero and enabling evaluation over the full operating range. NMAPE is defined as follows [22]:

$$NMAPE = \frac{100}{n}\sum_{i=1}^{n}\frac{\left|\hat{\mathrm{y}}_{\mathrm{i}} - \mathrm{y}_{\mathrm{i}}\right|}{P_{rate}} \tag{6}$$

where:

- $\hat{y}_i$ is the predicted solar power output at sample i (kW);
- y_i is the measured solar power output at sample i (kW);
- P_{rate}d is the rated capacity of the solar power plant (50,000 kW);
- n is the total number of evaluated forecasting samples.

3 RESULTS AND DISCUSSION

The forecasting performance of the vanilla Transformer and the proposed CNN-Transformer was evaluated on the complete independent dataset collected from June 10 to June 12, 2023. With a five-minute sampling interval, this three-day forecasting dataset contains 864 observations. The RMSE and NMAPE

values reported in Table 1 and Figure 2 were calculated using all 864 samples.

Table 1: Performance comparison on all 864 samples of the independent forecasting dataset.

Model	RMSE (kW)	NMAPE (%)
Vanilla Transformer	4,693.75	5.63
CNN-Transformer	2,919.42	3.45

Table 1 shows a clear performance advantage for the proposed architecture on the independent forecasting dataset. Compared with the vanilla Transformer, the CNN-Transformer reduces RMSE from 4,693.75 kW to 2,919.42 kW, corresponding to a 37.80% reduction. NMAPE decreases from 5.63% to 3.45%, corresponding to a 38.72% reduction.

To emphasize these distinctions visually, Figure 2 compares the absolute RMSE and NMAPE values, while the discussion reports the corresponding relative error reductions.

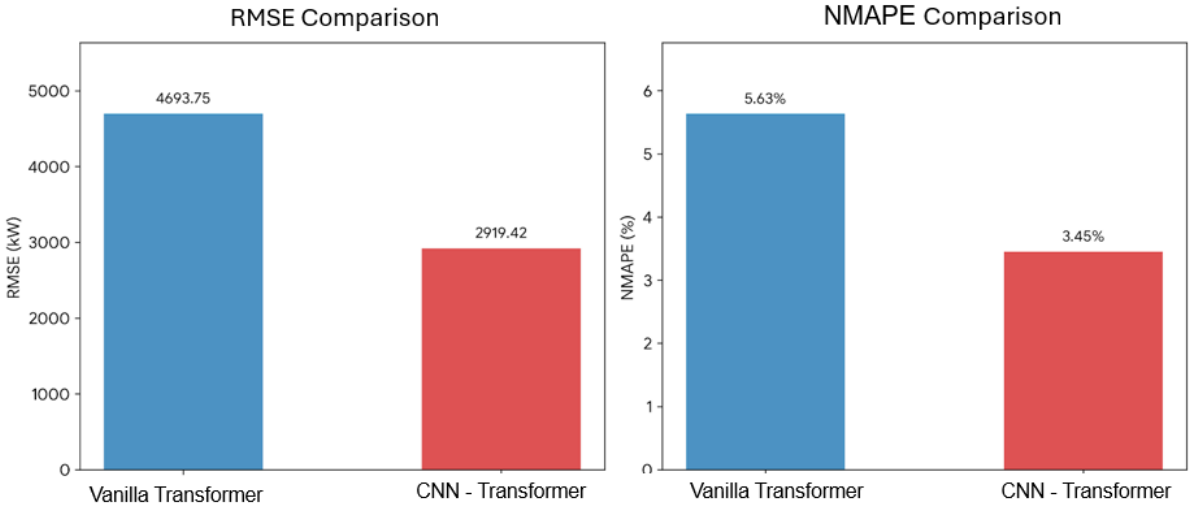


Figure 2: Comparison of RMSE and NMAPE on the independent forecasting dataset (June 10-12, 2023). The CNN-Transformer achieves reductions of 37.80% in RMSE and 38.72% in NMAPE relative to the vanilla Transformer.

As illustrated in Figure 2, RMSE decreases from 4,693.75 kW to 2,919.42 kW, and NMAPE decreases from 5.63% to 3.45%. These reductions quantitatively demonstrate the clearer advantage of the CNN-Transformer over the vanilla Transformer on the independent forecasting period.

Beyond the aggregate error metrics, Figure 3 presents 500 consecutive samples selected from the complete independent forecasting dataset of 864 observations. Only 500 samples are displayed to maintain graphical clarity and make the differences among the measured power, vanilla Transformer forecast, and CNN-Transformer forecast easier to observe.

Figure 3 provides a detailed visual comparison of the forecasting behavior of the two models over 500 consecutive samples selected from the complete independent forecasting dataset. The vanilla Transformer produces a relatively smooth forecasting curve and captures the general daily variation in power output. However, it responds slowly to several rapid changes in the measured signal, leading to temporal lag and substantial deviations during steep rising and falling intervals. This limitation is particularly evident near major power transitions, where the measured output changes abruptly while the vanilla Transformer continues to follow a smoother trajectory. In comparison, the CNN-Transformer reacts more strongly to short-term fluctuations and captures several local variations more effectively, indicating that the convolutional feature-extraction block improves sensitivity to rapid changes in the meteorological input sequence. Nevertheless, the proposed model does not eliminate all forecasting errors. It occasionally produces sharp overestimated peaks, suggesting that some rapid irradiance variations may be amplified. In addition, during several low-output and zero-output periods, both models continue to predict positive power values, revealing a systematic overestimation under nighttime or near-zero irradiance conditions. The largest discrepancies are therefore concentrated around abrupt power ramps and periods in which the measured output approaches zero. Although the CNN-Transformer does not completely remove these errors, it provides a closer overall representation of the measured temporal variations than the vanilla Transformer. This visual assessment is supported by the quantitative results calculated from all 864

observations in the independent forecasting period. As reported in Table 1, the CNN-Transformer reduces RMSE from 4,693.75 kW to 2,919.42 kW, corresponding to a 37.80% improvement, while NMAPE decreases from 5.63% to 3.45%, representing a 38.72% improvement. These results confirm that the proposed architecture improves overall forecasting accuracy, although further work is required to reduce overestimation during low-power periods and improve robustness under abrupt irradiance fluctuations.

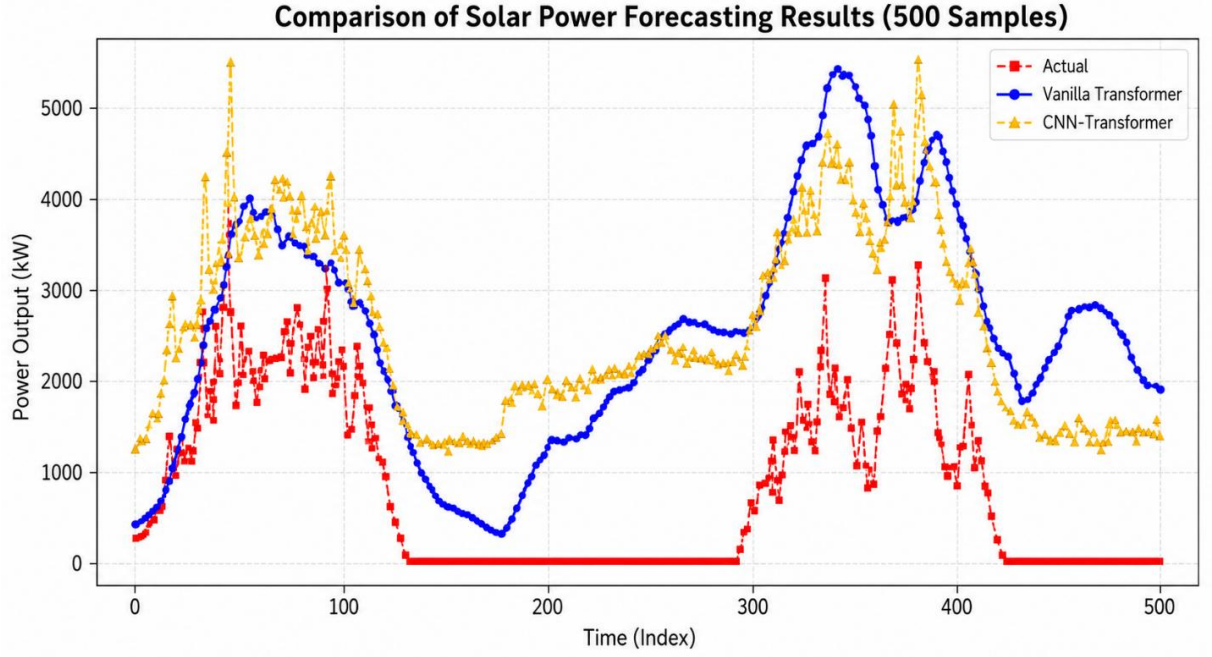


Figure 3: Comparison of actual and forecasted power over 500 consecutive samples selected from the complete independent forecasting dataset of 864 five-minute observations.

4 CONCLUSION

The results show that integrating a 1D Convolutional Neural Network as a front-end to the Transformer Encoder improves the extraction of local temporal features. On the independent forecasting dataset collected from June 10 to June 12, 2023, the proposed model achieved an RMSE of 2,919.42 kW and an NMAPE of 3.45%, outperforming the vanilla Transformer.

These results indicate that the CNN-Transformer has potential for use in EMS- and SCADA-assisted forecasting applications. Further validation across additional seasons, operating conditions, and solar power plants is nevertheless required before direct operational deployment.

Future Research Directions: Despite the high performance achieved, the current model is limited to point forecasting based on surface meteorological data. Future studies will aim to expand the feature space by integrating satellite cloud imagery to track cloud movement. Furthermore, developing the model toward probabilistic forecasting will be prioritized to provide confidence intervals for power output, thereby better supporting risk assessment tasks in smart grid operations.

REFERENCES

[1] M. Q. Raza, M. Nadarajah, and C. Ekanayake, On recent advances in PV output power forecast, *Sol. Energy*, vol. 136, no. July, pp. 125–144, 2016, doi: 10.1016/j.solener.2016.06.073.

[2] S. Shivashankar, S. Mekhilef, H. Mokhlis, and M. Karimi, Mitigating methods of power fluctuation of photovoltaic (PV) sources - A review, *Renew. Sustain. Energy Rev.*, vol. 59, pp. 1170–1184, 2016, doi: 10.1016/j.rser.2016.01.059.

[3] M. Diagne, M. David, P. Lauret, J. Boland, and N. Schmutz, Review of solar irradiance forecasting methods and a proposition for small-scale insular grids, *Renew. Sustain. Energy Rev.*, vol. 27, pp. 65–76, 2013, doi: 10.1016/j.rser.2013.06.042.

[4] C. Voyant et al., Machine learning methods for solar radiation forecasting: A review, *Renewable Energy*, vol. 105, pp. 569–582, 2017, doi: 10.1016/j.renene.2016.12.095.

[5] U. K. Das et al., Forecasting of photovoltaic power generation and model optimization: A review, *Renew. Sustain. Energy Rev.*, vol. 81, pp. 912–928, 2018, doi: 10.1016/j.rser.2017.08.017.

[6] G. Li, S. Xie, B. Wang, J. Xin, Y. Li, and S. Du, Photovoltaic power forecasting with a hybrid deep learning approach, *IEEE Access*, vol. 8, pp. 175871–175880, 2020, doi: 10.1109/ACCESS.2020.3025860.
[7] M. J. Mayer and G. Gróf, Extensive comparison of physical models for photovoltaic power forecasting, *Appl. Energy*, vol. 283, Art. no. 116239, 2021, doi: 10.1016/j.apenergy.2020.116239.
[8] U. K. Das et al., SVR-based model to forecast PV power generation under different weather conditions, *Energies*, vol. 10, no. 7, Art. no. 876, 2017, doi: 10.3390/en10070876.
[9] T.-A. Nguyen, M.-H. Pham, T.-K. Duong, and M.-P. Vu, A Recent Invasion Wave of Deep Learning in Solar Power Forecasting Techniques Using ANN, *in 2021 IEEE International Future Energy Electronics Conference (IFEEC)*, Nov. 2021, pp. 1–6, doi: 10.1109/IFEEC53238.2021.9661747.
[10] A. Vaswani et al., Attention Is All You Need, *in Advances in Neural Information Processing Systems 30 (NeurIPS 2017)*, pp. 5998–6008, 2017.
[11] M. Wang, J. Peng, Y. Luo, Z. Shen, and H. Yang, Comparison of different simplistic prediction models for forecasting PV power output: Assessment with experimental measurements, *Energy*, vol. 224, Art. no. 120162, 2021, doi: 10.1016/j.energy.2021.120162.
[12] F. Wang, Z. Zhen, B. Wang, and Z. Mi, Comparative study on KNN and SVM based weather classification models for day ahead short term solar PV power forecasting, *Appl. Sci.*, vol. 8, no. 1, Art. no. 28, 2018, doi: 10.3390/app8010028.
[13] J. Wang, P. Li, R. Ran, Y. Che, and Y. Zhou, A short-term photovoltaic power prediction model based on the gradient boost decision tree, *Appl. Sci.*, vol. 8, no. 5, Art. no. 689, 2018, doi: 10.3390/app8050689.
[14] M. Pan et al., Photovoltaic power forecasting based on a support vector machine with improved ant colony optimization, *J. Clean. Prod.*, vol. 277, Art. no. 123948, 2020, doi: 10.1016/j.jclepro.2020.123948.
[15] K. Wang, X. Qi, and H. Liu, Photovoltaic power forecasting based LSTM-Convolutional Network, *Energy*, vol. 189, Art. no. 116225, 2019, doi: 10.1016/j.energy.2019.116225.
[16] H. Yi, H. Wang, R. Zhang, and J. Peng, Forecasting Photovoltaic Power Using a Convolutional Neural Network and the Wavelet Transformation, *DEStech Trans. Comput. Sci. Eng.*, no. CSAE, pp. 350–357, 2018, doi: 10.12783/dtcse/csae2017/17503.
[17] S.-C. Lim, J.-H. Huh, S.-H. Hong, C.-Y. Park, and J.-C. Kim, Solar power forecasting using CNN-LSTM hybrid model, *Energies*, vol. 15, no. 21, Art. no. 8233, 2022, doi: 10.3390/en15218233.
[18] H. Aprillia, H.-T. Yang, and C.-M. Huang, Short-term photovoltaic power forecasting using a convolutional neural network-salp swarm algorithm, *Energies*, vol. 13, no. 8, Art. no. 1879, 2020, doi: 10.3390/en13081879.
[19] E. M. Al-Ali, M. A. T. Alsaeedi, S. H. Al-Salman, and A. A. Al-Saadi, Solar energy production forecasting based on a hybrid CNN-LSTM-Transformer model, *Mathematics*, vol. 11, no. 3, Art. no. 676, 2023, doi: 10.3390/math11030676.
[20] J. Kim, J. Obregon, H. Park, and J.-Y. Jung, Multi-step photovoltaic power forecasting using transformer and recurrent neural networks, *Renew. Sustain. Energy Rev.*, vol. 200, Art. no. 114479, 2024, doi: 10.1016/j.rser.2024.114479.
[21] E. Nziyumva, R. Hu, Y. Xu, J. Niyogisubizo, Z. Li, and P. C. Nshimyumukiza, Solar Radiation Forecasting Using Ensemble-Based Hybrid LGBM-GB-MLP Model: A Novel Stacked Generalization Method, *SSRN Electron. J.*, 2022, doi: 10.2139/ssrn.4076358.
[22] Electricity Regulatory Authority of Vietnam, Decision No. 67/QD-DTDL promulgating the procedure for forecasting the generating capacity and electricity output of renewable energy sources, Aug. 10, 2021.

NÂNG CAO ĐỘ CHÍNH XÁC DỰ BÁO NGẮN HẠN CÔNG SUẤT PHÁT ĐIỆN MẶT TRỜI SỬ DỤNG MÔ HÌNH CNN-TRANSFORMER

VŨ THỊ ANH THƠ[1]*, NGUYỄN TUẤN ANH[1], PHẠM MẠNH HẢI[2]
[1] *Khoa Kỹ thuật điện - Trường Đại học Điện lực*
[2] *Khoa Năng lượng và Phát triển bền vững - Trường Đại học Điện lực*
** Tác giả liên hệ: thovta@epu.edu.vn*

Tóm tắt. Sự gia tăng nhanh chóng tỷ trọng nguồn điện mặt trời trong hệ thống điện đặt ra nhiều thách thức cho công tác điều độ và vận hành lưới điện do tính biến động cao và sự phụ thuộc chặt chẽ vào điều kiện khí tượng. Gần đây, các mô hình học sâu dựa trên kiến trúc Transformer đã cho thấy tiềm năng lớn trong phân tích chuỗi thời gian nhờ cơ chế tự chú ý (self-attention). Tuy nhiên, kiến trúc Transformer nguyên bản có xu hướng ưu tiên các phụ thuộc dài hạn, dẫn tới việc nắm bắt chưa đầy đủ các dao động cục bộ, ngắn hạn của dữ liệu khí tượng. Nghiên cứu này đề xuất mô hình lai CNN-Transformer cho bài toán dự báo ngắn hạn công suất phát điện mặt trời. Mạng nơ-ron tích chập một chiều (1D-CNN) được sử dụng để trích xuất các đặc trưng không gian - thời gian cục bộ, sau đó đưa vào khối Transformer để học các phụ thuộc dài hạn. Kết quả thực nghiệm trên bộ dữ liệu vận hành thực tế của một nhà máy điện mặt trời công suất 50 MW cho thấy mô hình CNN-Transformer đạt hiệu quả vượt trội, với RMSE là 2.919,42 kW và NMAPE là 3,45%, tốt hơn đáng kể so với mô hình Transformer nguyên bản.

Từ khóa. Dự báo công suất phát điện, điện mặt trời, dự báo chuỗi thời gian, CNN-Transformer.

Received on August 28 – 2026
Revised on October 08 – 2026