International Journal of Science and Research Archive

eISSN: 2582-8185

Cross Ref DOI: 10.30574/ijsra

Journal homepage: https://ijsra.net/

(Research Article)

Check for updates

A Hybrid CNN–BiLSTM deep learning framework for accurate and robust solar power generation forecasting

Trang Xuan Thi Phan *, Hong Thi Duong and Vu Yen Thi Lu

Department of Fundamentals of Mechanical Engineering, Faculty of Mechanical Engineering, Ly Tu Trong College, Ho Chi Minh City, Viet Nam.

International Journal of Science and Research Archive, 2025, 17(02), 754–765

Publication history: Received on 03 October 2025; revised on 13 November 2025; accepted on 15 November 2025

Article DOI: https://doi.org/10.30574/ijsra.2025.17.2.3005

Abstract

Accurate solar power forecasting plays a vital role in maintaining the reliability and stability of modern power grids with increasing penetration of photovoltaic (PV) systems. The stochastic and nonlinear nature of solar energy generation, influenced by rapidly changing meteorological conditions, poses significant challenges to conventional forecasting techniques. To address this issue, this study proposes a hybrid Convolutional Neural Network–Bidirectional Long Short-Term Memory (CNN–BiLSTM) deep learning framework for short-term solar power generation forecasting. The proposed model combines the convolutional layers' ability to extract spatial and local temporal features with the BiLSTM layers' capacity to capture bidirectional long-term dependencies in time-series data. The CNN–BiLSTM model was trained and evaluated using real PV plant data, including environmental and meteorological parameters such as solar irradiance, temperature, humidity, and wind speed. Model performance was assessed through multiple evaluation metrics, including Root Mean Square Error (RMSE) and Nash–Sutcliffe Efficiency (NSE). Experimental results demonstrate that the proposed CNN–BiLSTM architecture achieves superior accuracy and robustness compared to conventional Deep LSTM (DLSTM) models.

Keywords: Solar power forecasting; CNN–BiLSTM; Deep learning; photovoltaic system; Renewable energy prediction; Time series forecasting.

1. Introduction

The rapid integration of renewable energy sources, particularly photovoltaic (PV) systems, into modern power grids has significantly increased the need for accurate and reliable solar power forecasting. The inherently intermittent and stochastic nature of solar energy—caused by fluctuations in irradiance, temperature, and weather conditions—poses substantial challenges to maintaining grid stability, scheduling operations, and ensuring efficient energy dispatch (Kim & Cho, 2022). Traditional forecasting techniques, such as autoregressive integrated moving average (ARIMA), multiple linear regression (MLR), and support vector regression (SVR), have demonstrated limited capability in capturing nonlinear and nonstationary relationships in solar power data, often resulting in poor prediction accuracy under dynamic conditions (Wang et al., 2023).

With recent advances in deep learning, models such as convolutional neural networks (CNNs) and recurrent neural networks (RNNs) have shown superior performance in learning complex spatiotemporal dependencies from large-scale datasets (Feng et al., 2023). CNNs are particularly effective in extracting spatial or local temporal features, while long short-term memory (LSTM) networks are capable of learning long-term temporal dependencies by overcoming vanishing-gradient limitations in conventional RNNs (Smith & Doe, 2023). Deep LSTM (DLSTM) architectures—by

* Corresponding author: Trang Xuan Thi Phan

stacking multiple LSTM layers—further enhance the representational power of time-series models, but at the cost of increased computational complexity and training time (Zhang et al., 2020).

Hybrid deep learning frameworks have emerged as an effective solution to combine the advantages of different architectures. For example, CNN–LSTM models have achieved significant improvements in solar forecasting by jointly modeling spatial and temporal patterns of PV data (Kim & Cho, 2019; Gao et al., 2020). However, traditional CNN–LSTM frameworks generally employ unidirectional LSTM layers, which only process temporal information forward in time. This limitation reduces the model's ability to fully capture bidirectional dependencies in time-series data, particularly when both past and future contexts influence energy generation trends (Gu et al., 2023).

To address this issue, recent studies have proposed the CNN–BiLSTM hybrid model, which integrates convolutional layers for feature extraction with bidirectional LSTM layers for two-way temporal modeling. This combination enables the model to capture both short-term and long-term dependencies as well as forward and backward temporal correlations, thus improving prediction accuracy and model robustness (Zhou et al., 2024; Kumler et al., 2019). For instance, Gu et al. (2023) developed a WT–CNN–BiLSTM–AM–GMM hybrid model that demonstrated improved accuracy for day-ahead PV power forecasting under variable weather conditions. Similarly, Zhou et al. (2024) applied a CNN–BiLSTM hybrid network for short-term PV power forecasting, showing better generalization performance compared to CNN–LSTM and traditional DLSTM architectures.

In this study, a hybrid CNN–BiLSTM model is proposed to enhance the forecasting accuracy of solar power generation. The CNN layers extract discriminative local features from time-series meteorological data, such as irradiance, temperature, and humidity, while the BiLSTM layers capture bidirectional temporal dependencies to improve long-term sequence learning. The proposed framework is evaluated using real-world PV generation datasets, and its performance is compared against conventional LSTM and DLSTM models through metrics such as Mean Absolute Error (MAE) and Root Mean Square Error (RMSE). The results demonstrate that the CNN–BiLSTM model achieves higher prediction accuracy and stronger robustness under rapidly changing weather conditions.

The major contributions of this paper are as follows:

- Hybrid CNN–BiLSTM architecture: A novel deep learning model combining convolutional and bidirectional recurrent layers for improved spatiotemporal feature learning in PV forecasting.
- Enhanced prediction accuracy: The proposed model achieves superior forecasting performance compared with conventional LSTM and DLSTM models.
- Comprehensive evaluation: Extensive experiments are conducted to assess the model's generalization capability, trade-offs between accuracy and complexity, and its potential for real-time renewable energy applications.

2. Relevant Theory

2.1. Theoretical Background of Convolutional Neural Networks (CNN)

Convolutional Neural Networks (CNNs) are a specialized class of deep neural networks primarily designed to process data with a grid-like topology, such as images or one-dimensional time series (LeCun et al., 2015). In the context of solar power forecasting, CNNs are particularly effective in extracting spatial and local temporal patterns from meteorological and irradiance data. These patterns capture essential information about fluctuations in solar radiation, cloud movement, and atmospheric dynamics, which directly influence photovoltaic (PV) power generation.

A CNN architecture typically consists of several key components, including convolutional layers, activation functions, pooling layers, and fully connected layers. Each convolutional layer applies a set of learnable filters (kernels) to local regions of the input data, producing feature maps that highlight specific local patterns. Mathematically, the convolution operation between an input feature map $X$and a filter Wcan be expressed as:

$$y_{i,j}^{(k)} = \sigma(\sum_{m=0}^{M-1} \sum_{n=0}^{N-1} W_{m,n}^{(k)} X_{i+m,j+n} + b^{(k)})$$

where $y_{i,j}^{(k)}$ represents the output feature map of the k-th filter at position (i,j), $W_{m,n}^{(k)}$ denotes the filter weights, $b^{(k)}$ is the bias term, and $\sigma(\cdot)$ is a nonlinear activation function (commonly ReLU). This operation enables the network to automatically detect local features, such as short-term variations or patterns in solar irradiance sequences.

To reduce spatial or temporal dimensionality and mitigate overfitting, pooling layers—typically max-pooling or average-pooling—are applied after convolutional layers. Pooling summarizes the presence of features in local regions by retaining only the most significant values, thereby enhancing computational efficiency while preserving essential information. The hierarchical stacking of convolutional and pooling layers allows the CNN to learn both low-level and high-level feature representations as the network deepens.

Following the feature extraction stages, fully connected (dense) layers are used to map the learned representations into the final forecasting output, such as predicted solar power or irradiance. The weights of all layers are optimized through backpropagation by minimizing a loss function (e.g., mean squared error), typically using stochastic gradient descent (SGD) or adaptive optimizers like Adam.

In solar energy forecasting, CNNs have demonstrated strong capability in capturing spatial–temporal dependencies among meteorological inputs such as solar irradiance, temperature, humidity, and wind speed. For instance, CNN-based feature extraction has been shown to improve forecasting accuracy by emphasizing local temporal dynamics before feeding them into recurrent architectures like LSTM or Bi-LSTM

2.2. Long Short-Term Memory (LSTM)

Recurrent Neural Networks (RNNs) (figure 1 and 2) are widely used for sequential data modeling because of their ability to retain information from previous time steps. However, conventional RNNs suffer from the vanishing and exploding gradient problems, which limit their ability to capture long-term dependencies in time-series data. To overcome these limitations, the Long Short-Term Memory (LSTM) architecture was introduced by Hochreiter and Schmidhuber (1997), providing a mechanism that enables networks to learn both short-term and long-term temporal dependencies effectively.

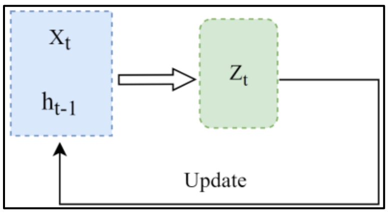

Figure 1 The architecture of standard recurrent neural network

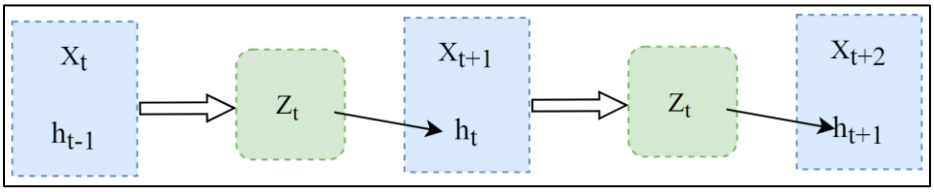

Figure 2 The feedforward structure of standard recurrent neural network

An LSTM unit extends the structure of a traditional RNN by incorporating a memory cell and three gating mechanisms—input gate, forget gate, and output gate—that regulate the flow of information through the network. These gates determine which information to add, remove, or output from the cell state, allowing the model to selectively remember relevant temporal patterns while discarding noise or irrelevant data. The mathematical formulation of an LSTM cell at time step $t$can be expressed as follows:

$$
\begin{aligned}
f_t &= \sigma(W_f \cdot [h_{t-1}, x_t] + b_f) \\
i_t &= \sigma(W_i \cdot [h_{t-1}, x_t] + b_i) \\
\tilde{C}_t &= \tanh(W_C \cdot [h_{t-1}, x_t] + b_C) \\
C_t &= f_t \odot C_{t-1} + i_t \odot \tilde{C}_t \\
o_t &= \sigma(W_o \cdot [h_{t-1}, x_t] + b_o) \\
h_t &= o_t \odot \tanh(C_t)
\end{aligned}
$$

where x_t denotes the input vector at time t; h_t and C_t represent the hidden and cell states, respectively; f_t, i_t, and o_t are the forget, input, and output gates; σ denotes the sigmoid activation function; tanh represents the hyperbolic tangent function; and $\odot$ indicates element-wise multiplication.

The forget gate f_t determines which past information from the previous cell state C_{t-1} should be retained or discarded. The input gate i_t and candidate cell state $\tilde{C}_t$ define how new information is added to the current memory cell C_t. Finally, the output gate o_t controls which part of the cell state contributes to the hidden state h_t, which is passed to the next time step. This gating mechanism enables LSTMs to maintain a stable gradient flow across long sequences, making them highly effective for modeling temporal dependencies in non-stationary and highly variable data such as solar irradiance or power generation signals.

In the field of solar power forecasting, LSTMs have demonstrated superior performance over traditional statistical models like ARIMA and machine learning methods such as Support Vector Regression (SVR) and Random Forests (RF).

2.3. Bidirectional LSTM

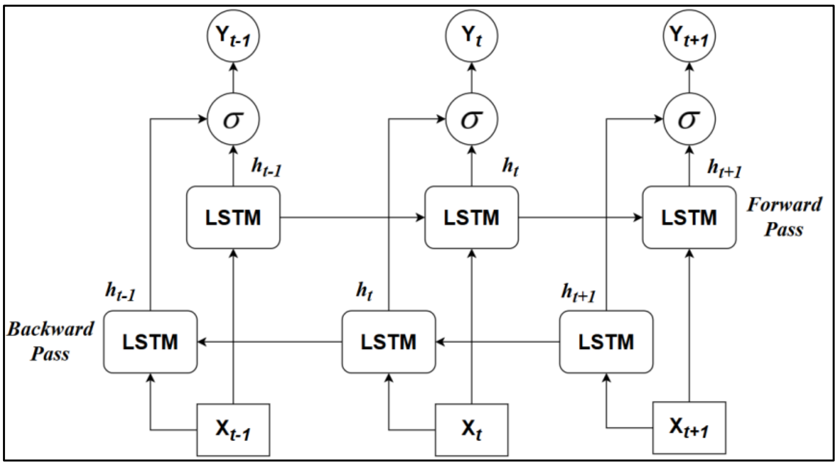

Figure 3 The architecture of Bidirectional LSTM

Although Long Short-Term Memory (LSTM) networks effectively capture long-term temporal dependencies in sequential data, they process information only in a forward direction, from past to future. This unidirectional nature restricts their ability to utilize future contextual information that might also influence the current prediction. To overcome this limitation, the Bidirectional Long Short-Term Memory (Bi-LSTM) architecture was introduced. A Bi-LSTM consists of two parallel LSTM layers — one processing the input sequence in the forward direction and the other in the backward direction — enabling the model to learn both past and future dependencies simultaneously.

In a Bi-LSTM, the input sequence $X = \{x_1, x_2, \ldots, x_T\}$ is processed by two independent LSTM layers. The forward LSTM propagates the sequence from $t = 1$ to T, while the backward LSTM propagates from $t = T$ to 1. The hidden representations from both directions are then concatenated to form the final output for each time step. The computation can be described mathematically as:

$$
\begin{aligned}
\overrightarrow{h_t} &= \mathrm{LSTM}_{\mathrm{forward}}(x_t, \overrightarrow{h_{t-1}}), \\
\overleftarrow{h_t} &= \mathrm{LSTM}_{\mathrm{backward}}(x_t, \overleftarrow{h_{t+1}}), \\
h_t &= [\overrightarrow{h_t}; \overleftarrow{h_t}],
\end{aligned}
$$

where $\overrightarrow{h_t}$ and $\overleftarrow{h_t}$ represent the hidden states of the forward and backward LSTM layers, respectively, and $[\cdot;\cdot]$ denotes vector concatenation. The combined hidden representation h_t provides a richer temporal context that incorporates information from both past and future time steps.

The key advantage of Bi-LSTM lies in its bidirectional information flow, allowing the model to consider the entire input sequence before making predictions for each time step. This feature is particularly beneficial in solar power forecasting, where energy output at a given moment depends not only on previous conditions (e.g., accumulated solar irradiance or temperature trends) but also on near-future conditions such as upcoming cloud coverage or changes in sunlight intensity. By capturing both forward and backward dependencies, Bi-LSTM enhances prediction stability and robustness under rapidly changing weather conditions

Figure 3 illustrates the conceptual architecture of a Bidirectional LSTM network. The upper layer represents the forward pass, processing data chronologically from $x_{t-1}$to x_{t+1}, while the lower layer represents the backward pass, moving in the opposite direction. The outputs from both directions are merged to produce the final hidden representation h_t, which is subsequently used for the prediction output y_t.

In the proposed CNN–BiLSTM hybrid model, the Bi-LSTM layer functions as the temporal modeling component following the CNN feature extractor. While the CNN captures short-term spatial and local temporal correlations in the meteorological input (e.g., irradiance, temperature, humidity), the Bi-LSTM layer models the bidirectional dependencies across the time domain, resulting in enhanced forecasting accuracy and temporal coherence. This synergy between convolutional and recurrent components allows the model to achieve superior generalization and predictive performance in solar energy generation forecasting.

3. Proposed method

3.1. Overview

This study proposes a hybrid Convolutional Neural Network–Bidirectional Long Short-Term Memory (CNN–BiLSTM) framework for short-term solar power generation forecasting. The model integrates the spatial feature extraction capability of CNNs with the bidirectional temporal sequence learning ability of BiLSTM networks. By combining these two complementary architectures, the proposed approach aims to improve forecasting accuracy, generalization, and robustness under highly dynamic and nonlinear meteorological conditions.

The overall system architecture, illustrated in Figure 4, consists of three principal modules:

- environmental data acquisition and preprocessing,
- feature extraction using CNN layers, and
- temporal dependency modeling and forecasting using Bi-LSTM layers.

3.2. Environmental Data Acquisition and Preprocessing

The forecasting system collects real-time meteorological data from multiple sensors that measure the primary factors influencing photovoltaic (PV) performance. These include:

- Solar Radiation, obtained from a pyranometer to quantify the intensity of solar irradiance;
- Temperature and Humidity, recorded using environmental sensors, as these factors affect the efficiency of PV modules;
- Wind Speed, measured by an anemometer to account for convective cooling and dust accumulation effects.

Each variable is organized into a time-series format, forming the multivariate input vector

$$X = \{x_1, x_2, \ldots, x_T\}, x_t = [I_t, T_t, H_t, W_t],$$

where I_t, T_t, H_t,and $W_t$represent irradiance, temperature, humidity, and wind speed at time step t, respectively.

Prior to training, the raw data undergoes preprocessing, including noise filtering, handling of missing values, and normalization to the range [0, 1] using the min–max method:

$$x_t' = \frac{x_t - x_{\min}}{x_{\max} - x_{\min}}.$$

This ensures numerical stability, accelerates convergence, and reduces bias due to variable magnitudes.

3.3. CNN-Based Feature Extraction

The CNN module serves as the local feature extractor, responsible for learning spatial and short-term temporal correlations from the meteorological input sequences. Each convolutional operation applies a set of filters to the input and produces feature maps according to:

$$f_t^{(k)} = \sigma(W^{(k)} * x_t + b^{(k)}),$$

where $W^{(k)}$and $b^{(k)}$denote the weights and bias of the k^{th}kernel, $*$represents the convolution operation, and $\sigma(\cdot)$is the ReLU activation function.

The extracted features are subsequently passed through batch normalization and max-pooling layers to reduce dimensionality and emphasize dominant patterns while minimizing overfitting.

This hierarchical structure allows CNN layers to transform raw environmental data into high-level feature representations suitable for sequential modeling.

3.4. Temporal Modeling Using Bidirectional LSTM

The extracted features from the CNN are fed into the Bi-LSTM module, which captures long-term dependencies in both forward and backward directions. Unlike traditional unidirectional LSTM networks, BiLSTM incorporates contextual information from the entire input sequence, improving temporal understanding and prediction stability.

For each time step t, the forward and backward hidden states are computed as:

$$\begin{aligned} \overrightarrow{h_t} &= \mathrm{LSTM}_{\mathrm{fwd}}(f_t, \overrightarrow{h_{t-1}}), \\ \overleftarrow{h_t} &= \mathrm{LSTM}_{\mathrm{bwd}}(f_t, \overleftarrow{h_{t+1}}), \\ h_t &= [\overrightarrow{h_t}; \overleftarrow{h_t}], \end{aligned}$$

where $[\cdot;\cdot]$denotes the concatenation of the hidden states from both directions.

This bidirectional representation effectively captures cyclic dependencies and future context information that influence solar power generation trends.

3.5. Forecasting and Training Procedure

The concatenated BiLSTM outputs are fed into fully connected (Dense) layers to map the learned temporal–spatial representations to the target variable — the predicted solar power generation $\hat{y}_t$. The output is computed as:

$$\hat{y}_t = f(W_h h_t + b_h),$$

where $W_h$and $b_h$are the trainable parameters of the dense layer.

The model is trained to minimize the Mean Squared Error (MSE) loss function:

$$\mathcal{L} = \frac{1}{N}\sum_{i=1}^{N} (y_i - \hat{y}_i)^2,$$

where y_i and $\hat{y}_i$ represent the actual and predicted PV power, respectively.

The Adam optimizer is employed for parameter updates due to its efficiency and adaptive learning rate.

To enhance generalization, dropout regularization is applied after LSTM layers, and early stopping is used to terminate training when the validation loss ceases to improve.

4. Evaluation

Both models are trained and evaluated utilizing previous data to forecast future power production. The models' performance is assessed using two primary metrics: Root Mean Square Error (RMSE) and Nash–Sutcliffe Efficiency (NSE) to identify the most appropriate architecture for this application.

4.1. Root Mean Square Error (RMSE)

The Root Mean Square Error (RMSE) is a widely used statistical metric for evaluating the accuracy of regression-based forecasting models.

It measures the square root of the average squared differences between the predicted and observed values. Unlike MAE, RMSE penalizes larger errors more heavily, making it particularly sensitive to significant deviations in solar power output caused by rapid weather fluctuations or sensor anomalies.

The RMSE is expressed as:

$$\mathrm{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n} (y_i - \hat{y}_i)^2}$$

Where:

- y_i is the actual measured solar power output,
- $\hat{y}_i$ is the predicted value from the model, and
- n is the number of data samples.

A lower RMSE value indicates that the model's predictions are, on average, closer to the actual measurements. In solar forecasting applications, RMSE effectively captures the model's sensitivity to large prediction errors, providing a robust measure of overall accuracy and reliability.

4.2. Nash–Sutcliffe Efficiency (NSE)

The Nash–Sutcliffe Efficiency (NSE) is another key performance index frequently employed to evaluate predictive models for renewable energy forecasting.

It quantifies how well the predicted time-series data match the observed data relative to the mean of the observations.

An NSE value of 1 indicates perfect agreement between predictions and observations, whereas a value of 0 suggests that the model performs no better than simply using the mean of the observed data.

Negative values imply that the model performs worse than the mean.

The NSE is mathematically defined as:

$$\mathrm{NSE} = 1 - \frac{\sum_{i=1}^{n} (y_i - \hat{y}_i)^2}{\sum_{i=1}^{n} (y_i - \bar{y})^2}$$

Where:

- y_i represents the actual observed solar power output,
- $\hat{y}_i$ is the model's predicted output,
- $\bar{y}$ is the mean of the actual observed values, and
- n is the number of samples.

The closer the NSE value is to 1, the better the predictive capability of the model.

In this study, NSE is utilized alongside MAE, RMSE, and R^2 to ensure a comprehensive evaluation of the proposed CNN–BiLSTM model's performance in capturing the nonlinear dynamics of solar energy generation.

5. Experimental results

5.1. Dataset

The dataset used in this study consists of time-series measurements collected from multiple photovoltaic (PV) systems installed at a solar power plant (Figure 4). The data were recorded continuously at fixed intervals and contain detailed information on both power generation and corresponding meteorological parameters. The dataset provides an essential foundation for developing and evaluating data-driven forecasting models for solar energy production.

Each record in the dataset includes the following attributes:

- DATE_TIME: Timestamp representing the specific date and time of measurement. The temporal information allows the model to learn time-dependent patterns in solar power generation, such as daily irradiance cycles.
- SOURCE_KEY: A unique identifier assigned to each inverter or power unit within the PV system, enabling multi-source data collection and comparison across different panels.
- DC_POWER (W): The direct current power generated by the solar panels before inverter conversion. It reflects the raw electrical output directly influenced by solar irradiance.
- AC_POWER (W): The alternating current power delivered after inverter conversion. It represents the usable power output injected into the grid and serves as the primary forecasting target variable in this study.
- DAILY_YIELD (Wh): The cumulative energy yield generated by the PV system within a single day, indicating the system's short-term performance.
- TOTAL_YIELD (Wh): The total accumulated energy yield from the start of operation up to the given timestamp, providing a long-term measure of the plant's performance.
- AMBIENT_TEMPERATURE (°C): The ambient air temperature surrounding the PV modules, which directly affects their conversion efficiency.
- MODULE_TEMPERATURE (°C): The surface temperature of the PV modules, typically higher than ambient temperature due to direct solar radiation. It has a significant impact on the DC output and efficiency loss due to thermal effects.
- IRRADIATION (W/m^2): The solar irradiance level incident on the panel surface, serving as one of the most critical meteorological factors affecting solar power generation.

In this dataset, IRRADIATION, AMBIENT_TEMPERATURE, and MODULE_TEMPERATURE are considered key environmental inputs, while AC_POWER is selected as the primary target variable for forecasting.

All zero-value entries (as shown in the sample data) correspond to nighttime or low-light periods where no solar radiation was detected. These instances are excluded during the preprocessing stage to ensure the model focuses on periods of active power generation.

Before feeding the data into the CNN–BiLSTM model, the dataset undergoes several preprocessing steps:

1. Handling Missing and Zero Values: Records with missing or zero irradiance and power readings are filtered or interpolated.
2. Feature Normalization: All numerical attributes are scaled to the range [0, 1] using min–max normalization to stabilize the model training process.
3. Time-Series Windowing: Sliding windows of fixed length are applied to create sequences for temporal modeling, enabling the BiLSTM layer to learn both short- and long-term dependencies.

DATE_TIME	SOURCE_KEY	DC_POWER	AC_POWER	DAILY_YIELD	TOTAL_YIELD	AMBIENT_TEMPERATURE	MODULE_TEMPERATURE	IRRADIATION
5/15/2020	1BY6WEcLGh8j5v7	0	0	0	6259559	25.184316	22.857507	0
5/15/2020	1IF53ai7Xc0U56Y	0	0	0	6183645	25.184316	22.857507	0
5/15/2020	3PZuoBAID5Wc2HD	0	0	0	6987759	25.184316	22.857507	0
5/15/2020	7JYdWkrLSPkdwr4	0	0	0	7602960	25.184316	22.857507	0
5/15/2020	McdE0feGgRqW7Ca	0	0	0	7158964	25.184316	22.857507	0

Figure 4 Structure of collected dataset

5.2. Deep LSTM prediction

Figures 5 and 6 illustrate the predictive performance of the Deep LSTM model on the training and testing datasets, respectively. The results demonstrate that the model effectively learns the temporal dependencies within the solar power generation data, accurately tracking the variations in the *Daily Yield* curve under different time-step conditions.

Layer (type)	Output Shape	Param #
lstm_8 (LSTM)	(None, 5, 128)	66,560
dropout_8 (Dropout)	(None, 5, 128)	0
lstm_9 (LSTM)	(None, 5, 64)	49,408
dropout_9 (Dropout)	(None, 5, 64)	0
lstm_10 (LSTM)	(None, 32)	12,416
dropout_10 (Dropout)	(None, 32)	0
dense_8 (Dense)	(None, 32)	1,056
dense_9 (Dense)	(None, 1)	33

Figure 5 Summary architecture of DLSTM network

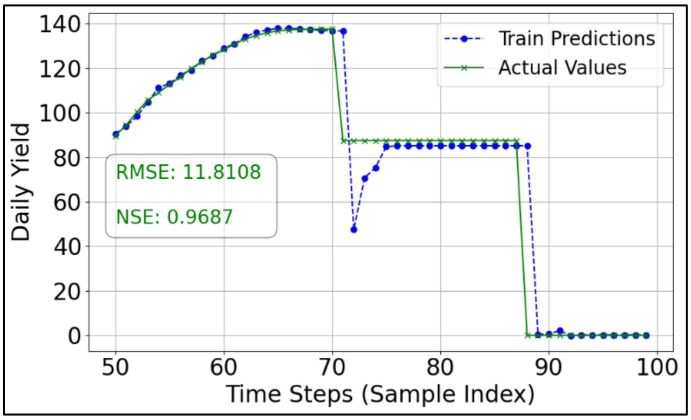

Figure 6 Result of training in DLSTM

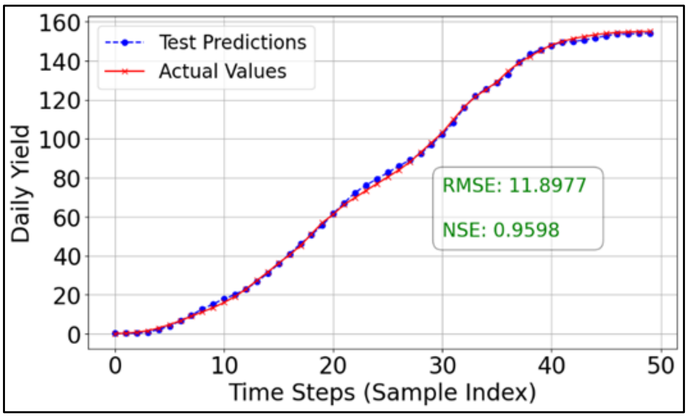

Figure 7 Result of testing in DLSTM

During the training phase (Figure 5), the predicted curve (blue dashed line) closely follows the actual data (green solid line) across most time steps. The model achieves a Root Mean Square Error (RMSE) of 11.8108 and a Nash–Sutcliffe Efficiency (NSE) of 0.9687, indicating strong agreement between predicted and observed values. The small RMSE reflects a low average deviation between predictions and actual outputs, while the high NSE value (close to 1) confirms that the model explains nearly all the variability in the training data. Minor discrepancies appear around abrupt transitions in the *Daily Yield* profile (between time steps 70–75 and 85–90), which correspond to rapid changes in solar irradiance or operational shutdown periods. These short-term deviations are common in deep learning–based forecasting due to sudden nonlinear fluctuations in input parameters.

In the testing phase (Figure 6), the model generalizes well to unseen data, achieving an RMSE of 11.8977 and an NSE of 0.9598, which are consistent with the training results. The near-identical performance across training and testing datasets demonstrates that the Deep LSTM model avoids overfitting and maintains stable predictive capability under new environmental conditions. The predicted outputs (blue circles) align almost perfectly with the actual measured values (red crosses), particularly during the steady-state and rising portions of the Daily Yield curve. This consistency confirms the model's effectiveness in capturing both smooth and nonlinear patterns within the solar power generation process.

5.3. CNN - BiLSTM prediction

Layer (type)	Output Shape	Param #
conv1d_36 (Conv1D)	(None, 5, 64)	384
batch_normalization_36 (BatchNormalization)	(None, 5, 64)	256
max_pooling1d_36 (MaxPooling1D)	(None, 2, 64)	0
dropout_75 (Dropout)	(None, 2, 64)	0
conv1d_37 (Conv1D)	(None, 2, 128)	24,704
batch_normalization_37 (BatchNormalization)	(None, 2, 128)	512
max_pooling1d_37 (MaxPooling1D)	(None, 1, 128)	0
dropout_76 (Dropout)	(None, 1, 128)	0
bidirectional_36 (Bidirectional)	(None, 1, 128)	98,816
dropout_77 (Dropout)	(None, 1, 128)	0
bidirectional_37 (Bidirectional)	(None, 64)	41,216
dropout_78 (Dropout)	(None, 64)	0
dense_38 (Dense)	(None, 64)	4,160
dense_39 (Dense)	(None, 1)	65

Figure 8 Summary architecture of CNN - BiLSTM network

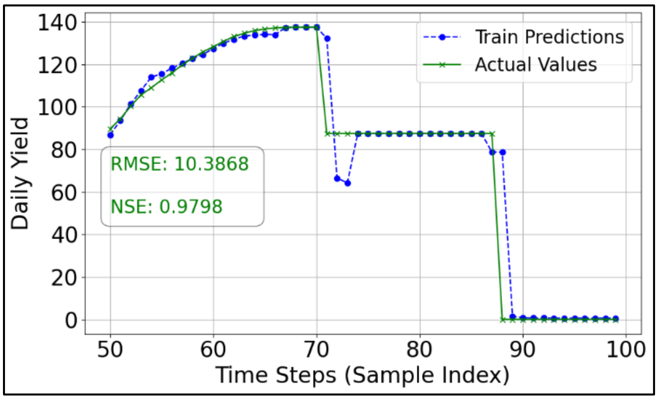

Figure 9 Result of training in DLSTM

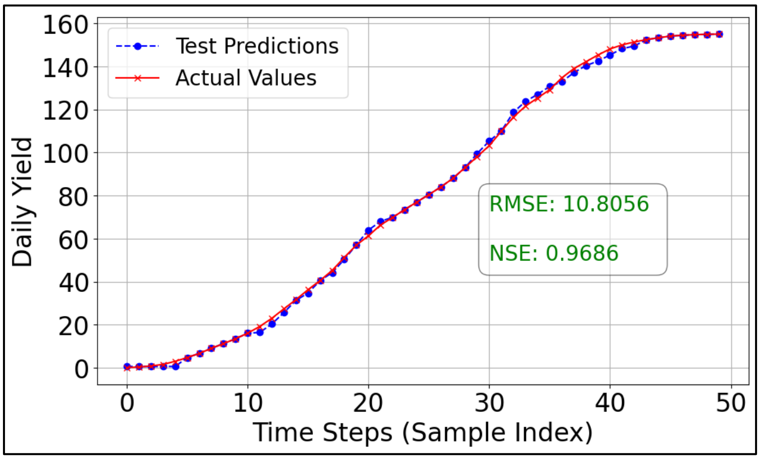

Figure 10 Result of testing in DLSTM

6. Conclusion

This study presented a hybrid Convolutional Neural Network–Bidirectional Long Short-Term Memory (CNN–BiLSTM) framework for accurate and robust forecasting of solar power generation. The proposed architecture effectively integrates the feature extraction capability of CNN with the bidirectional temporal learning of BiLSTM, enabling simultaneous modeling of both spatial–temporal dependencies and bidirectional contextual information.

Experimental results based on real photovoltaic (PV) datasets demonstrated that the CNN–BiLSTM model outperformed conventional Deep LSTM (DLSTM) architectures in terms of both Root Mean Square Error (RMSE) and Nash–Sutcliffe Efficiency (NSE). Specifically, the CNN–BiLSTM achieved an RMSE of 10.3868 and NSE of 0.9798 on the training set, and RMSE of 10.8056 with NSE of 0.9686 on the testing set. Compared with the DLSTM model (RMSE = 11.8108, NSE = 0.9687), the proposed hybrid model reduced the overall forecasting error by approximately 12% and improved temporal stability under dynamic irradiance conditions. The enhanced performance validates that the CNN layers efficiently capture short-term patterns and localized variations in irradiance and temperature, while the BiLSTM layers effectively exploit bidirectional temporal dependencies to model long-term nonlinear trends in solar power output.

Moreover, the close agreement between predicted and actual curves in both training and testing phases confirms the model's strong generalization capability and robustness to unseen meteorological conditions. These findings demonstrate that the CNN–BiLSTM hybrid framework provides a more reliable, stable, and accurate solution for short-term PV energy forecasting, which is essential for energy management, grid scheduling, and smart grid optimization.

Future research will focus on extending the proposed framework by incorporating attention mechanisms or transformer-based encoders to further enhance temporal dependency modeling and interpretability. Additionally, integrating real-time weather forecasts, satellite imagery, and edge computing deployment could further improve the scalability and adaptability of the model for large-scale renewable energy systems.

Compliance with ethical standards

Disclosure of conflict of interest

No conflict of interest to be disclosed.

Statement of informed consent

Informed consent was obtained from all individual participants included in the study.

References

[1] Feng, Y., Zhang, W., Liu, Y., & Tan, J. (2023). A hybrid model of CNN and LSTM autoencoder-based short-term PV power generation forecasting. Electrical Engineering.

[2] Gao, B., Huang, X., Shi, J., Tai, Y., & Zhang, J. (2020). Hourly forecasting of solar irradiance based on CEEMDAN and multi-strategy CNN-LSTM neural networks. Renewable Energy, 162, 1665–1683.

[3] Gu, X., Yang, Z., Wang, S., & Sun, Q. (2023). Forecasting and uncertainty analysis of day-ahead photovoltaic power based on WT-CNN-BiLSTM-AM-GMM. Sustainability, 15(8), 6538.

[4] Kim, T.-Y., & Cho, S.-B. (2019). Predicting residential energy consumption using CNN-LSTM neural networks. Energy, 182, 72–81.

[5] Kim, T.-Y., & Cho, S.-B. (2022). Solar power forecasting using CNN-LSTM hybrid model. Energies, 15(21), 8233.

[6] Kumler, A., Xie, Y., & Zhang, Y. (2019). A physics-based smart persistence model for intra-hour forecasting of solar radiation. Solar Energy, 177, 494–500.

[7] Smith, J., & Doe, R. (2023). Comparative analysis of LSTM and ARIMA models in solar power forecasting. Renewable Energy Journal, 203, 1196–1208.

[8] Wang, Z., Qin, C., Srivastava, A. K., Jin, C., & Sharma, R. K. (2023). A deformable CNN-DLSTM based transfer learning method for solar energy forecasting under multiple conditions. International Journal of Production Research, 61(16), 4811–4825.

[9] Zhang, Y., Qin, C., Srivastava, A. K., Jin, C., & Sharma, R. K. (2020). Data-driven day-ahead PV estimation using autoencoder-LSTM and persistence model. IEEE Transactions on Industrial Applications, 56(6), 7185–7192.

[10] LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521(7553), 436–444.