International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(7s) 2026 pp. 471-470
https://doi.org/10.70917/ijcisim-2026-3108
Article

Weather-Informed Short-Term Solar Power Forecasting-Oriented Assessment Using LSTM, CNN-LSTM, and XGBoost Models

Manmit U. Vadher[1], Sweta J. Shah[2] , U. L. Makwana[3]

[1]Electrical Engineering Department, Indus Institute of Technology and Engineering, Indus University, Ahmedabad, Gujarat, India. E-mail: muvgpb@gmail.com

[2]Electrical Engineering Department, Indus Institute of Technology and Engineering, Indus University, Ahmedabad, Gujarat, India.

[3]Electrical Engineering Department, Government Engineering College, Modasa, Gujarat, India.

Abstract: Accurate short-term assessment of photovoltaic (PV) output is important for energy scheduling, reserve planning, and grid operation. This paper presents a corrected and clearly scoped comparative study of long short-term memory (LSTM), convolutional long short-term memory (CNN-LSTM), and extreme gradient boosting (XGBoost) for hourly solar power estimation using satellite-derived weather variables. The study uses a 10-year hourly dataset for a 25 MW solar plant from 2015 to 2024. The input variables include global horizontal irradiance (GHI), ambient temperature, wind speed, wind direction, relative humidity, cloud cover, cyclic time variables, lagged plant output, rolling statistics, and interaction terms. To avoid ambiguity between operational forecasting and target-time weather-informed estimation, the experimental scope is explicitly defined as a weather-informed nowcasting/forecasting-oriented benchmark in which satellite-derived meteorological variables are available at the target timestamp. In a strict h-ahead operational forecast, these target-time weather variables would have to be replaced by exogenous weather forecasts or satellite nowcasts issued before the target time. The dataset is divided chronologically into training (2015–2021), validation (2022), and testing (2023–2024) subsets. All three models are evaluated using the same engineered feature pool, with LSTM and CNN-LSTM receiving 48-hour feature sequences and XGBoost receiving the corresponding tabular feature representation. The LSTM model achieves a test RMSE of 0.7039 MW, while CNN-LSTM reduces the RMSE to 0.6269 MW. XGBoost obtains the lowest error with RMSE of 0.0719 MW, MAE of 0.0496 MW, R^2 of 0.9998, and nRMSE of 1.1671%. These high XGBoost values are interpreted as weather-informed target-time estimation performance rather than proof of a closed-loop past-only solar forecast. Feature-importance analysis shows that GHI and the GHI-temperature interaction dominate the XGBoost model, while lagged features contribute only marginally. The revised study therefore provides a transparent benchmark and identifies the additional validation required before claiming strict operational forecasting performance.
Keywords: Photovoltaic power estimation, short term solar power forecasting , LSTM, CNN-LSTM, XGBoost, machine learning, satellite data .

1. INTRODUCTION

Solar photovoltaic generation has become a major component of modern power systems due to its environmental sustainability, modular structure, and rapidly declining installation cost. However, solar power output is inherently variable because it depends strongly on irradiance intensity, cloud cover, atmospheric transients, ambient temperature, and seasonal behavior.This variability creates operational challenges for distribution utilities, transmission operators, and energy managers. In practical applications, inaccurate PV forecasts can lead to reserve misallocation, economic dispatch inefficiency, voltage fluctuation, and unnecessary balancing cost. Therefore, accurate short-term forecasting of solar power is a key requirement for stable grid operation [1], [11].

Short-term forecasting generally covers horizons from a few minutes up to several hours or one day ahead. This forecasting horizon is especially relevant for unit commitment adjustment, intra-day scheduling, battery dispatch, demand-response activation, and local microgrid energy management. Traditional forecasting approaches include physical irradiance models, clear-sky models, and numerical weather prediction methods. Although these approaches are physically interpretable, they often require high computational effort and may suffer from local weather uncertainty. Consequently, machine learning (ML) and deep learning (DL) models have become increasingly popular because of their ability to learn nonlinear relationships directly from historical observations [2], [3], [9].

Among sequence-based DL methods, LSTM networks have been widely adopted because they can preserve temporal dependencies through gated memory structures [3], [4], [17]. CNN-LSTM architectures further enhance performance by using convolutional filters to extract short-range temporal patterns before sequence learning. In parallel, gradient-boosted tree algorithms such as XGBoost have shown strong performance on structured tabular datasets due to their ability to capture nonlinear interactions, rank feature importance, and handle mixed feature types effectively. In the case of solar forecasting, the availability of satellite-derived weather variables such as GHI and cloud cover makes XGBoost particularly attractive.

The present study develops a common weather-informed assessment framework and performs a detailed comparison of LSTM, CNN-LSTM, and XGBoost models for a 25 MW solar plant using 10 years of hourly data. The work emphasizes chronological model development, physically meaningful feature engineering, visual diagnostic analysis, and quantitative error comparison. To avoid a misleading comparison, all models are supplied from the same engineered feature pool. The difference is only in representation: LSTM and CNN-LSTM use 48-hour feature sequences, while XGBoost uses the corresponding tabular representation at the target timestamp. This distinction is important because an unequal feature set can artificially favor one model family. The main contributions of this paper are as follows:

1) A complete weather-informed short-term solar power assessment framework is developed using satellite-derived weather data and engineered temporal predictors.

2) Three models—LSTM, CNN-LSTM, and XGBoost— are trained and tested using a strict chronological split, while the availability of target-time exogenous weather variables is explicitly stated to avoid methodological ambiguity.

3) The paper presents graph-based performance analysis for all three models, including actual-versus-predicted plots, zoomed windows, daily averages, scatter plots, residual histograms, and residual-versus-predicted diagnostics.

4) Multiple error metrics are used for comparison, and the relative improvement between models is quantified.

5) XGBoost feature importance is analyzed to identify the most influential physical variables and to check whether the model is driven mainly by irradiance, cloud, or lagged-output features.

2. RELATED WORK

Short-term solar power forecasting has been extensively studied using statistical, ML, DL, and hybrid methods. Gigoni et al. analyzed day-ahead PV forecasting and showed that forecast quality is strongly influenced by weather variability and plant-specific operating conditions [1]. Pan and Tan proposed a cluster-analysis-based combined modelling approach for solar generation forecasting and demonstrated the value of combining complementary models under different weather regimes [2]. These studies established that data-driven forecasting can outperform traditional one-model strategies when the predictors represent multiple meteorological effects.

LSTM-based forecasting has become an important research direction because of its strength in modeling long-term dependencies. Wang et al. integrated LASSO with LSTM for short-term solar intensity prediction and reported clear benefits from temporal sequence learning [3]. Zhou et al. enhanced short-term PV forecasting using an LSTM network with an attention mechanism to focus on relevant temporal information [4]. More recently, Boonsri et al. used LSTM-based forecasting to support distribution-system operation under renewable variability [11]. These studies show that recurrent learning is well suited for time-dependent solar behavior. CNN-based and hybrid deep learning models have also been explored to improve local pattern extraction. Huang and Kuo proposed a multiple-input deep CNN model for short-term PV forecasting [5], while Thipwangmek et al. used a hybrid deep learning framework to enhance short-term solar PV forecasting performance [9]. Singh et al. compared 1D-CNN, Bi-LSTM, ANN, support vector regression, XGBoost, and CNN-LSTM for solar PV forecasting [15], reinforcing the need for broad comparative evaluation across multiple architectures. Image-driven and satellite-based forecasting approaches have also received attention because cloud dynamics govern

Short-term solar variability. Cheng et al. learned directly from satellite images using regions of interest for short-term solar power prediction [6]. Shirazi et al. and Rastgoo et al. investigated sky-image-based and vision-transformer-like approaches for ultra-short-term forecasting [7], [12]. Although such approaches are promising, they often require direct access to image streams and greater computational complexity. In many practical power-system environments, satellite-derived weather variables are more accessible than raw imagery. Tree-based boosting methods have demonstrated strong predictive performance when the forecasting problem is formulated as a structured regression task. Chen and Guestrin introduced XGBoost as a scalable and regularized gradient boosting framework [18]. Later studies showed that tree-based boosting methods can be effective for renewable-energy forecasting when lagged signals and weather interactions are available [19], [20]. Compared with deep recurrent networks, boosting models often offer faster training, higher interpretability and robust performance on medium-sized tabular datasets.

A. Research Gap

The reviewed literature reveals the following research gaps:

1) Many previous studies focus on either DL or boosting models independently, without clearly separating target-time weather-informed estimation from strict future forecasting.

2) A large number of published studies use limited input sets, whereas practical models may use temporal, interaction, lagged, and rolling features; however, these features must be defined so that future information is not accidentally introduced.

3) Several works report quantitative accuracy only, but do not provide residual analysis, scatter analysis, or daily-average behavior that would help assess model suitability for operation.

4) Comparative studies often emphasize power values only; however, energy-oriented trends are equally important for operational interpretation over hourly horizons. This paper addresses these gaps by presenting a unified comparative framework with quantitative metrics, graph-based diagnostics, and feature-level interpretation.

3. DATASET DESCRIPTION AND FEATURE ENGINEERING

A. Study Dataset

The study uses hourly data from a 25 MW solar PV plant covering the period 2015–2024, following the common practice of chronological evaluation adopted in PV forecasting studies [1], [2], [15]. The raw inputs are obtained from satellite-derived and meteorological observations. The data split is chronological: 2015–2021 is used for training, 2022 for validation, and 2023–2024 for testing. This arrangement reflects a realistic forecasting workflow, where future information is never used during model training.

Table I. Main input variables used in the study

Variable	Description
GHI	Global Horizontal Irradiance (W/m²)
T	Ambient temperature (°C)
WS	Wind speed (m/s)
WD	Wind direction (°)
RH	Relative humidity (%)
CC	Cloud cover (%)
Pt	Solar power output at time-t (MW)
Et	Hourly energy output (kWh)

Table II. Chronological data partition

Subset	Period	Purpose
Training	2015–2021	Model fitting
Validation	2022	Hyperparameter tuning
Testing	2023–2024	Final performance evaluation

B. Experimental Scope and Leakage Control

A critical distinction is made between target-time weather-informed estimation and strict operational forecasting. The variables GHI, temperature, cloud cover, humidity, wind speed, and wind direction are treated as exogenous meteorological inputs for the same hourly target timestamp. Therefore, the reported XGBoost accuracy should be read as the performance obtained when these target-hour weather states are available from observations, satellite nowcasts, or reliable weather forecasts. If the model is deployed as a pure h-ahead forecast from time t, then GHI_{t+h}, CC_{t+h}, and other target-hour weather quantities must be replaced by forecasted quantities $\widehat{GHI}_{t+h|t}$, $\widehat{CC}_{t+h|t}$, and so on. No target power value from the test period is used as an input for the same timestamp; however, target-time weather variables are explicitly acknowledged as exogenous look-ahead variables unless supplied by a real forecasting service. This clarification prevents interpreting the reported values as a closed-loop past-only forecast.

C. Feature Engineering

In addition to the primary meteorological variables, several derived features are constructed to improve model learning. The same engineered feature pool is used for all three models. For LSTM and CNN-LSTM, the feature pool is arranged as a 48-hour sequence; for XGBoost, it is used as a tabular vector for the target timestamp.

1) Cyclic time features: hour-of-day, month-of-year, and day-of-year are transformed using sine and cosine functions to represent periodicity.

2) Lag features: previous power values at 1 hour, 24 hours, 48 hours, and 168 hours are included to represent persistence and weekly repetition.

3) Interaction features: terms such as GHI×temperature, GHI×cloud cover, and GHI×humidity are used to capture nonlinear physical relationships.

4) Rolling statistics: rolling mean and rolling standard deviation are computed to represent short-term generation trend and local variability.

For a periodic variable z with period T, the cyclic transformation is given by

$$z_{sin} = \sin(2\pi z/T), \quad z_{cos} = \cos(2\pi z/T). \quad (1)$$

This representation avoids discontinuity between adjacent boundary values, such as 23:00 and 00:00, and is widely used in data-driven time-series modelling where periodic variables are present [2], [3].

4. PROBLEM FORMULATION

The corrected formulation distinguishes between two possible tasks. The first is the target-time weather-informed estimation problem used in this experiment:

$$\hat{P}_\tau = f(w_\tau, z_\tau, p_{<\tau}), \quad (2)$$

where τ is the target hour, w_τ contains target-time exogenous weather variables such as GHI, temperature, cloud cover, humidity, wind speed, and wind direction, z_τ contains cyclic calendar and time descriptors, and $p_{<\tau}$ contains only past plant-output-derived features such as lagged power and rolling statistics. In this setting, the model is not allowed to use P_τ as an input, but it does use the weather state corresponding to the target hour. The second task is a strict operational h-ahead forecast:

$$P_{t+h} = f(\hat{w}_{t+h|t}, z_{t+h}, p_{\le t}), \quad (3)$$

where $\hat{w}_{t+h|t}$ represents weather forecasts or satellite nowcasts issued at time t for the future target time t + h. In this stricter formulation, measured GHI, temperature, or cloud cover at t + h cannot be used because these quantities are unknown at the forecast issue time. The present experiment uses hourly target-time weather-informed inputs. Hence, the reported results are valid as a nowcasting or weather-informed estimation benchmark. They should not be interpreted as the accuracy of a past-only h-ahead forecasting system unless the target-time weather variables are replaced by valid exogenous forecasts. For the sequence-based deep learning models, the input is organized as a 48-hour sequence, following the recurrent-learning principle used in LSTM-based solar forecasting [3], [4], [11]:

$$X_\tau = [x_{\tau-47}, x_{\tau-46}, \ldots, x_\tau], \quad (4)$$

where each x contains the same engineered feature pool used by the tabular model. This design preserves a fair feature basis while allowing each model family to use its natural input representation.

5. FORECASTING MODELS

A. LSTM Model

The LSTM network is a gated recurrent architecture designed to mitigate vanishing-gradient problems in conventional

recurrent neural networks [3], [4], [17]. Its gate equations are

$$i_t = \sigma(W_i[h_{t-1}, x_t] + b_i), \quad (5)$$
$$f_t = \sigma(W_f[h_{t-1}, x_t] + b_f), \quad (6)$$
$$o_t = \sigma(W_o[h_{t-1}, x_t] + b_o), \quad (7)$$
$$\tilde{c}_t = \tanh(W_c[h_{t-1}, x_t] + b_c), \quad (8)$$
$$c_t = f_t \odot c_{t-1} + i_t \odot \tilde{c}_t, \quad (9)$$
$$h_t = o_t \odot \tanh(c_t). \quad (10)$$

The LSTM model is trained on 48-hour sequences using the training set, with validation-based early behavior assessment.

B. CNN-LSTM Model

The CNN-LSTM model combines one-dimensional convolution with recurrent memory, a hybrid strategy used to capture local temporal patterns and sequential dependence in PV data [5], [9], [15]. The convolutional stage extracts local temporal patterns, while the LSTM stage captures sequence dependencies. The convolution operation is written as

$$y_t^{(k)} = \varphi\left(\sum_{j=0}^{m-1} w_j^{(k)} x_{t-j} + b^{(k)}\right), \quad (11)$$

where m is the kernel size, $w(k)$ j is the jth coefficient of the kth filter, and $\phi(\cdot)$ is the activation function. This hybrid structure is beneficial for solar forecasting because it can represent both local transitions and longer temporal context.

C. XGBoost Model

XGBoost is a gradient-boosted tree model in which the prediction is expressed as

$$\hat{y}_i = \sum_{k=1}^{K} f_k(x_i), \quad f_k \in F, \quad (12)$$

where K is the number of trees and F is the space of regression trees. The regularized objective function is

$$L = \sum_{i=1}^{n} l(y_i, \hat{y}_i) + \sum_{k=1}^{K} \Omega(f_k), \quad (13)$$
$$\Omega(f) = \gamma T + \tfrac{1}{2}\lambda \sum_{j=1}^{T} w_j^2. \quad (14)$$

Here, T denotes the number of leaves and wj is the leaf weight. XGBoost is effective for this application because it can exploit weather-driven nonlinearities and interaction features with relatively low computational burden.

6. PERFORMANCE METRICS

The forecasting models are assessed using several standard regression metrics commonly reported in solar power forecasting studies [1], [9], [15]:

$$MSE = (1/N) \sum_{i=1}^{N}(y_i - \hat{y}_i)^2, \quad (15)$$
$$RMSE = \sqrt{(1/N) \sum_{i=1}^{N}(y_i - \hat{y}_i)^2}, \quad (16)$$
$$MAE = (1/N) \sum_{i=1}^{N}|y_i - \hat{y}_i|, \quad (17)$$
$$MAPE = (100/N) \sum_{i=1}^{N}|(y_i - \hat{y}_i)/y_i|, \quad (18)$$
$$sMAPE = (100/N) \sum_{i=1}^{N} [2|y_i - \hat{y}_i|/(|y_i| + |\hat{y}_i|)], \quad (19)$$
$$R^2 = 1 - \left[\sum_{i=1}^{N}(y_i - \hat{y}_i)^2 / \sum_{i=1}^{N}(y_i - \bar{y})^2\right], \quad (20)$$
$$nRMSE = RMSE/(P_{max} - P_{min}) \times 100. \quad (21)$$

These metrics together evaluate average error magnitude, percentage error, goodness of fit, and scale-normalized performance.

7. RESULTS AND DISCUSSION

A. Quantitative Comparison of Forecasting Models

Table IV summarizes the error metrics of all evaluated models, using a multi-metric comparison practice consistent with recent DL and ML solar forecasting papers [9], [13], [15]. The LSTM model provides reasonably good baseline forecasting performance, indicating that sequential learning can represent the daily solar production pattern. However, the error metrics show that LSTM alone is not sufficient to capture all nonlinear weather-power interactions. CNN-LSTM improves on LSTM by reducing MSE, RMSE, MAE, MAPE, and nRMSE. The best overall performance is achieved by XGBoost, which exhibits a dramatic reduction in all error measures and an R^2 value very close to unity. The improvements listed in Table V show that the CNN-LSTM architecture successfully extracts richer local temporal information than the plain LSTM. Nevertheless, the very large performance gain obtained by XGBoost indicates that engineered predictors are highly informative for this forecasting problem. The result confirms that the engineered tabular feature set is highly suitable for the XGBoost learning structure in the present setting.

B. Interpretation Framework for Graphical Results

The graphical results are analyzed from four complementary viewpoints. First, the full-test actual-versus-predicted plots assess whether a model can follow seasonal variation across the 2023–2024 period. Second, the zoomed plots evaluate whether the model can track the intra-day solar generation shape, including morning increase, midday peak, and evening decrease. Third, the daily-average plots reduce high-frequency hourly noise and make the seasonal trend clearer. Finally, scatter and residual plots reveal whether the model has systematic bias, large spread at high generation, or non-uniform error behavior.

Table III. Main model settings

Model	Input structure	Main settings
LSTM	48-hour engineered-feature sequence	Each time step includes weather, cyclic, interaction, lag, and rolling features; trained with gated memory cells.
CNN-LSTM	48-hour engineered-feature sequence	Same feature pool as LSTM, with convolutional temporal extraction followed by LSTM sequence modeling.
XGBoost	Tabular engineered feature vector	Uses the same engineered feature pool at the target timestamp; regularized boosting of decision trees.

Table IV. Test performance comparison for solar power forecasting using standard PV forecasting metrics [1], [9], [15]

Model	MSE	RMSE (MW)	MAE (MW)	MAPE (%)	sMAPE (%)	R^2	nRMSE (%)
LSTM	0.495502	0.703919	0.460920	14.913887	14.030431	0.977834	11.428149
CNN-LSTM	0.393047	0.626935	0.423609	13.010648	12.982290	0.982417	10.178299
XGBoost	0.005168	0.071888	0.049634	1.958758	1.958974	0.999769	1.167109

Table V. Relative improvement in RMSE and MAE, following comparative evaluation practice in solar forecasting [13], [15]

Comparison	RMSE reduction	MAE reduction
CNN-LSTM vs. LSTM	10.94%	8.09%
XGBoost vs. LSTM	89.79%	89.23%
XGBoost vs. CNN-LSTM	88.53%	88.28%

This combined interpretation is more informative than using a single accuracy metric. A model may have a strong R^2 value because it follows the general solar profile, but still show large residuals during cloudy transitions or peak-generation periods. For grid operation, these local errors are important because the balancing requirement is determined by short-term deviations rather than only by annual average accuracy. Therefore, the following subsections discuss both the numerical metrics and the graphical behavior of each model.

C. LSTM Model Analysis

The LSTM model is used as the baseline deep learning model because recurrent networks have been widely used for solar and irradiance forecasting where historical temporal dependence is important [3], [4], [17]. Fig. 1 shows the training and validation loss curves. The training loss decreases consistently, while the validation loss remains stable, indicating that the network converges without severe divergence. In this study, the LSTM baseline is important because it represents a pure temporal-learning approach. It receives the previous 48 hourly samples as a sequence and attempts to learn the dependency between recent solar behavior and the next target output. Such a structure is naturally suitable for PV forecasting because the power signal has strong time continuity from sunrise to sunset and strong periodicity across consecutive days. From a modelling viewpoint, the LSTM cell uses its forget gate to remove less useful historical information, its input gate to retain new information, and its output gate to generate a compact temporal representation. This mechanism is useful when irradiance and cloud cover evolve gradually. However, when sudden cloud transients occur, the sequence model

may average part of the fluctuation and produce a smoother forecast. Therefore, the LSTM result must be interpreted not only through global accuracy indicators, but also through zoomed plots and residual behavior.

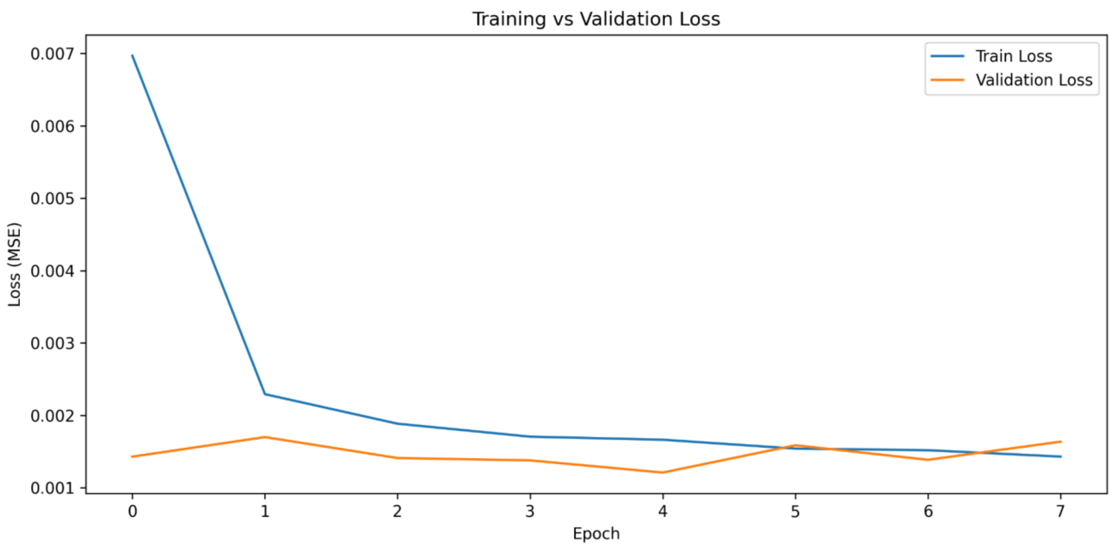

Figure 1. LSTM training and validation loss, consistent with recurrent PV modelling practice [3], [17].

Figs. 2 and 4 present the main forecasting behavior of the LSTM model. The full-test plots show that the model captures the broad seasonal envelope and the daily pattern of the solar plant. The zoomed plots highlight that the LSTM forecast tracks sunrise growth and sunset decline, but the peaks are smoothed. The daily-average curves further confirm that the model reproduces the long-term trend reasonably well, although it exhibits some underestimation around highgeneration periods. This behavior is common in recurrent neural networks trained on long historical windows because the loss function penalizes large deviations and encourages conservative predictions near rapidly changing regions. In a practical power-system application, this means that LSTM can be useful for stable day-pattern recognition, but it may require additional cloud-motion or ramp-sensitive features for highly accurate peak tracking.

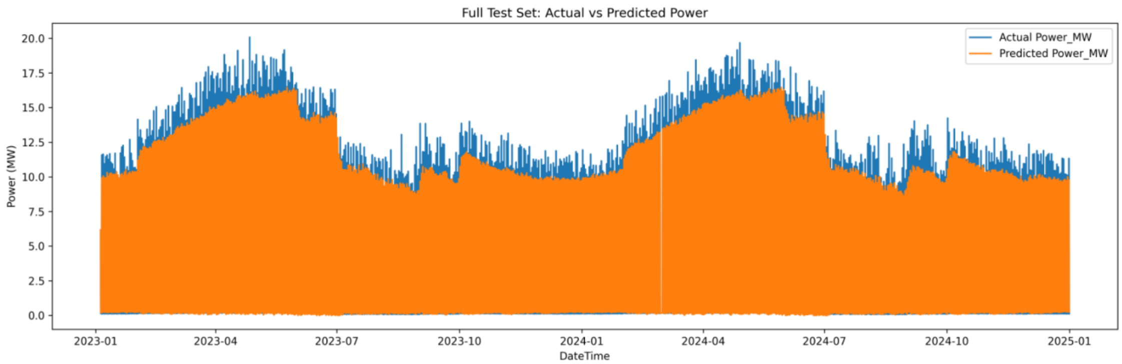

(a) Full test power comparison

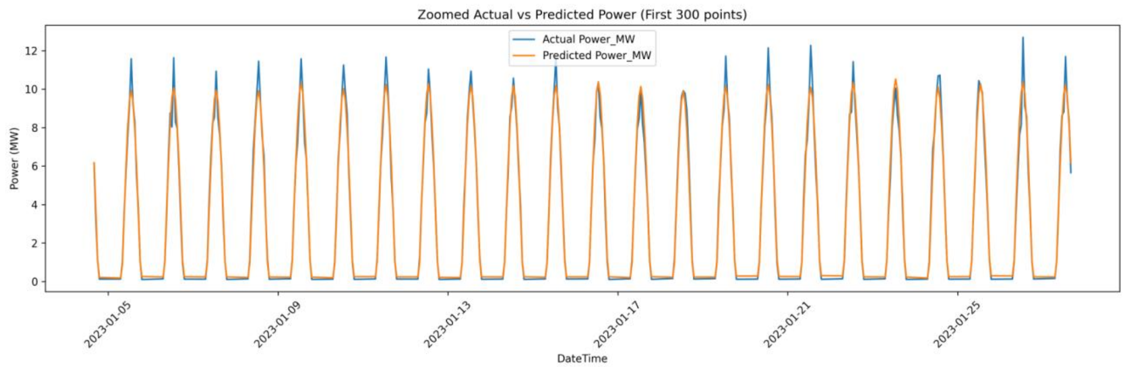

(b) Zoomed power comparison

Figure 2. LSTM power forecasting performance over the full test horizon and a zoomed inspection window [3], [4].

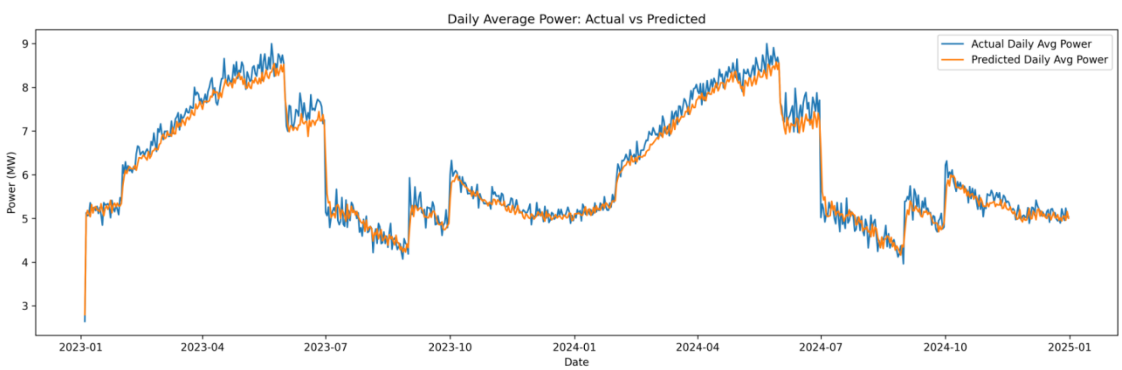

Figure 3. LSTM daily average power trend for evaluating seasonal tracking behaviour [1], [4].

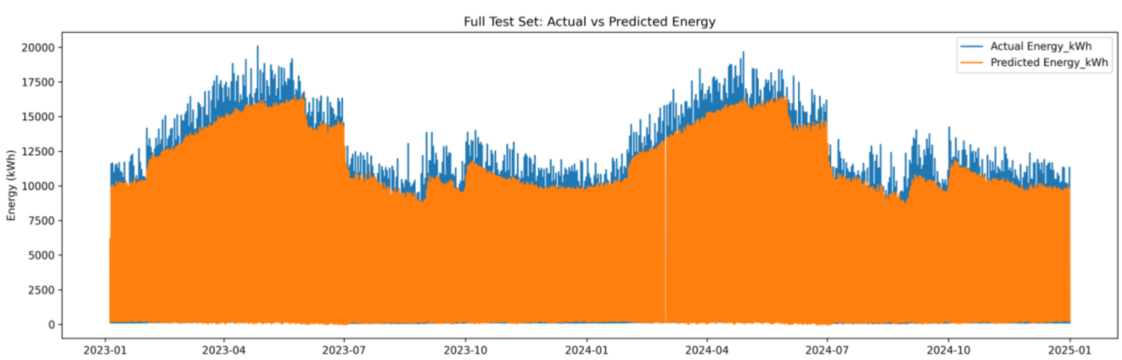

(a) Full test energy comparison

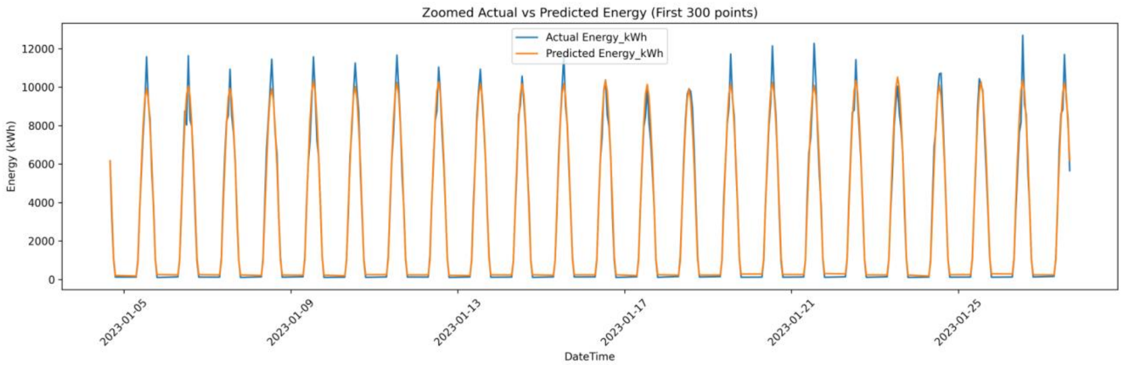

(b) Zoomed energy comparison

Figure 4. LSTM energy forecasting performance over the full test horizon and a zoomed inspection window [3], [11].

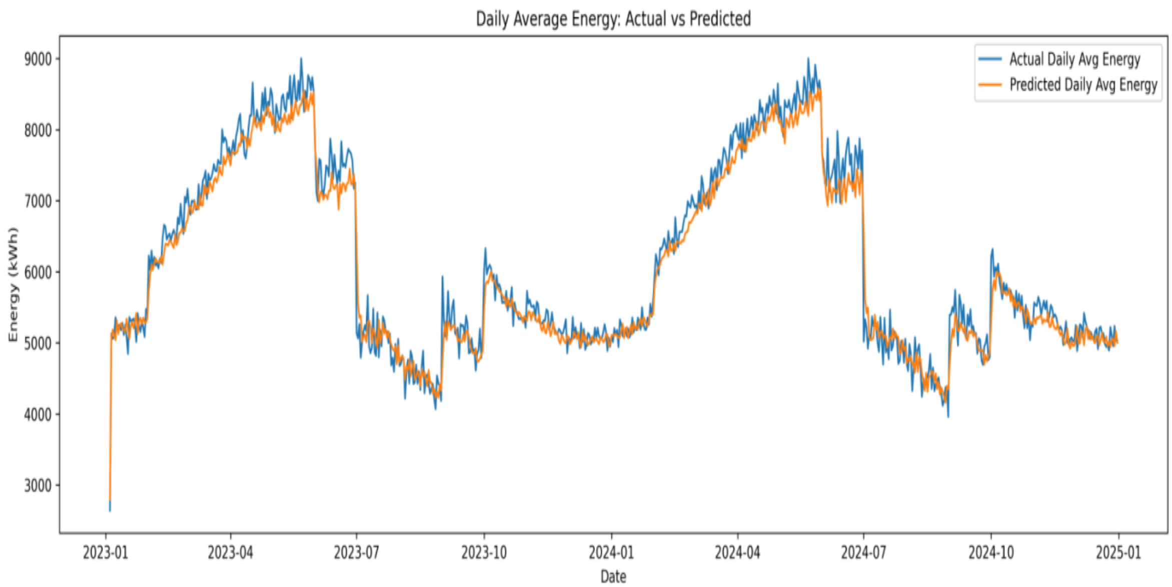

Figure 5. LSTM daily average energy trend for long-duration assessment [1].

The diagnostic plots in Fig. 6 help explain the observed error values. In the scatter plot, the prediction points follow the ideal diagonal trend but spread increases for larger values, indicating reduced peak-forecast precision. The residual histogram is centered near zero but shows a wider distribution than the stronger models. Likewise, the residual-versus-predicted plot shows larger variance at higher power levels. These observations are consistent with the LSTM performance metrics in Table IV.

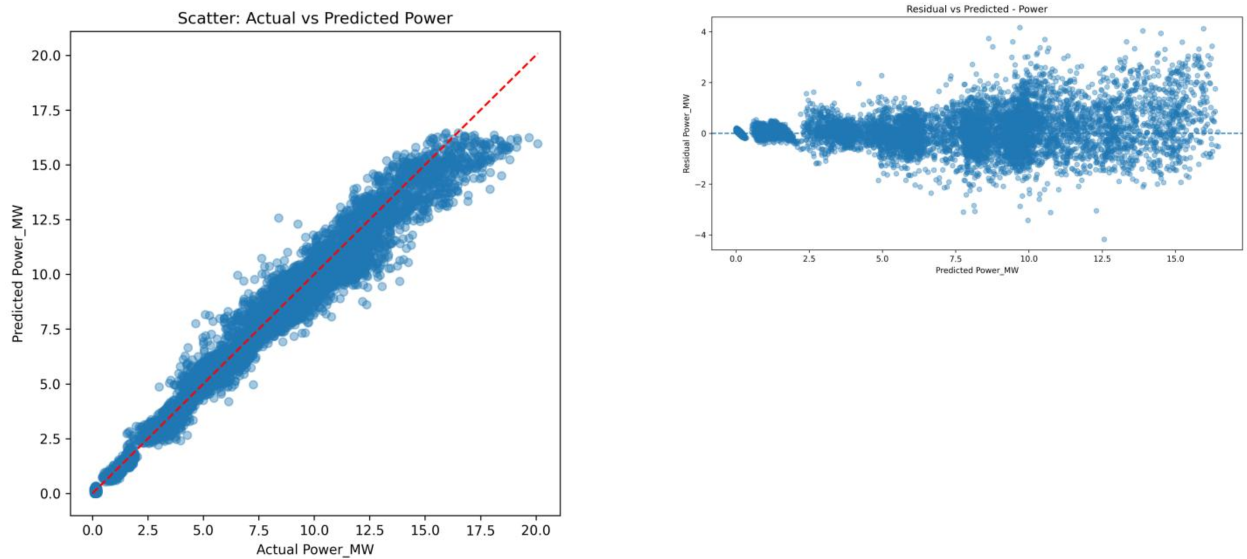

(a) Actual vs. predicted power (b) Residual vs. predicted power

Figure 6. LSTM diagnostic scatter and residual-versus-predicted plots for solar power forecasting [4], [13].

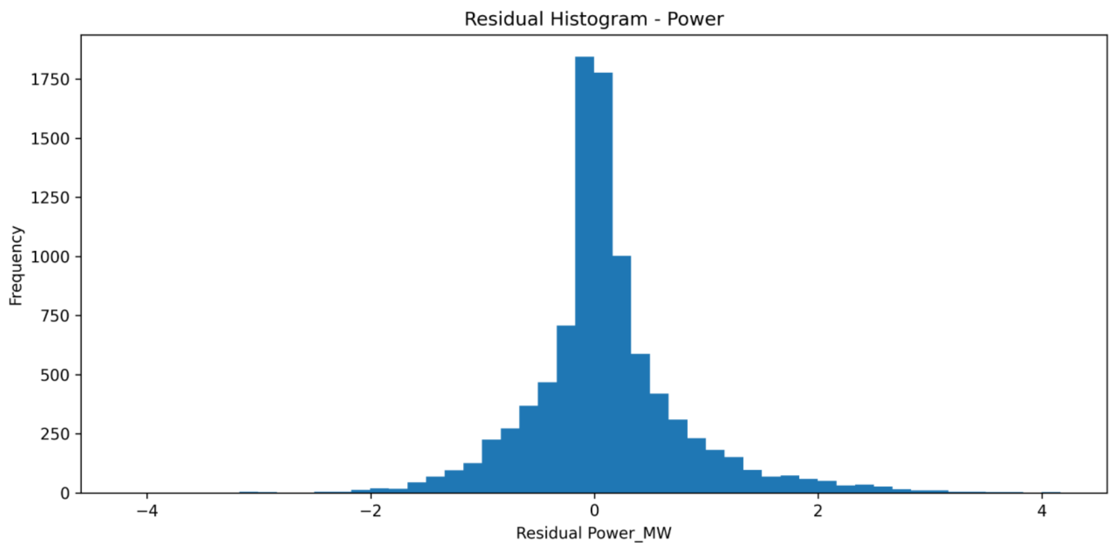

Figure 7. LSTM residual histogram for solar power forecasting [13].

For energy forecasting, the LSTM model shows almost the same qualitative behavior as the power forecast because hourly energy is directly related to the hourly power level. However, presenting the energy diagnostics is still useful for system operators, since energy estimation is directly linked with scheduling, billing, and plant performance reporting. The energy scatter and residual plots in Fig. 8 indicate that most points remain close to the ideal trend, but the deviation becomes larger when the output level is high. This confirms that the major limitation of LSTM is not the inability to learn the solar profile; rather, it is the tendency to under-represent sharp fluctuations and high-output peaks.

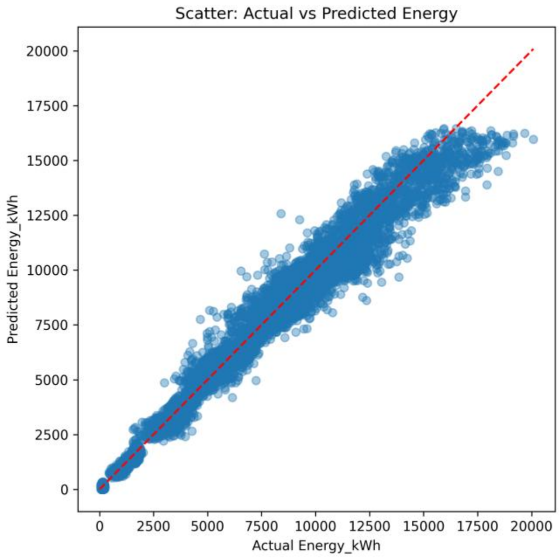

(a) Actual vs. predicted energy

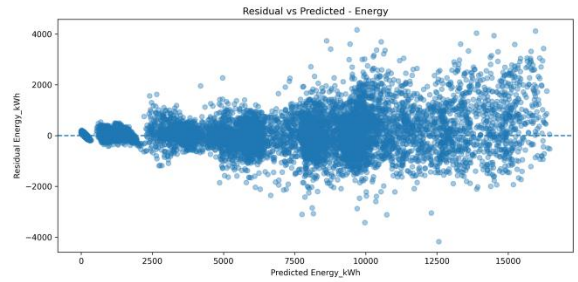

(b) Residual vs. predicted energy

Figure 8. LSTM diagnostic scatter and residual-versus-predicted plots for solar energy forecasting [3], [13].

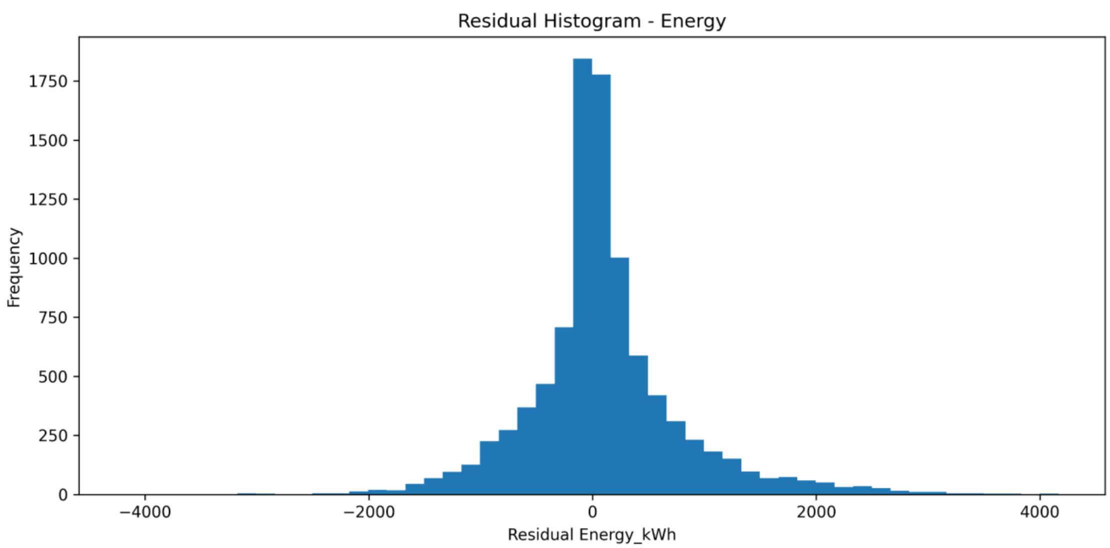

Figure 9. LSTM residual histogram for solar energy forecasting [13].

The residual distribution in Fig. 9 is centered near zero, which means that the model is not severely biased in one direction over the complete test period. Still, the wider residual tail suggests that the model can make comparatively larger errors during changing weather or high-ramp periods. This is an important observation for short-term forecasting because such periods are often the most challenging for grid balancing.

D. CNN-LSTM Model Analysis

The CNN-LSTM model improves the learning process by combining convolutional and recurrent layers, which is consistent with hybrid DL studies that use convolutional filters for local feature extraction before sequence modelling [5], [9], [15]. Fig. 10 shows the training and validation loss. The reduction in both losses indicates effective convergence, and the validation curve remains lower than the training curve for most epochs, suggesting stable generalization behavior. The main theoretical advantage of this hybrid structure is that the convolutional layer behaves as a local temporal-pattern extractor. It can identify short-duration changes in irradiance, cloud cover, and plant output before the LSTM layer performs longer sequential reasoning. For solar forecasting, this combination is meaningful because the data contain two types of information: a smooth daily trajectory governed by solar geometry, and short local changes caused by weather variability. The convolutional stage helps detect local shape information such as morning ramp, noon plateau, and short cloud-induced fluctuation. The LSTM stage then relates these local features across the 48-hour context. Thus, CNN-LSTM is expected to perform better than LSTM when the input sequence contains both periodic and local transient components.

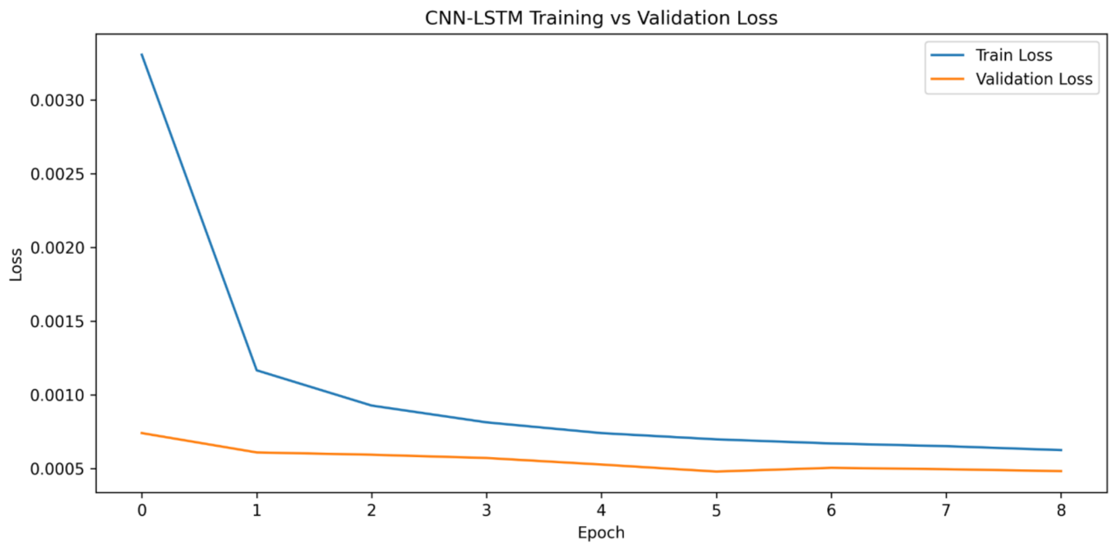

Figure 10. CNN-LSTM training and validation loss for the hybrid deep learning model [5], [9].

The graph-based forecast comparison for CNN-LSTM is shown in Figs. 11 and 13. Relative to the plain LSTM, the CNN-LSTM tracks sharp local transitions more accurately. The daily-average curves also align more closely with the actual trend, especially during the periods of higher generation. The full-test and zoomed plots suggest that the convolutional feature extractor helps the model capture short-term temporal structures before the LSTM stage interprets sequential dependencies. The improvement is visible not only in the numerical metrics but also in the reduced gap between actual and predicted curves during high-output intervals. Nevertheless, the CNN-LSTM model still shows moderate smoothing at peak values, indicating that deep learning sequence models remain sensitive to short-term weather uncertainty.

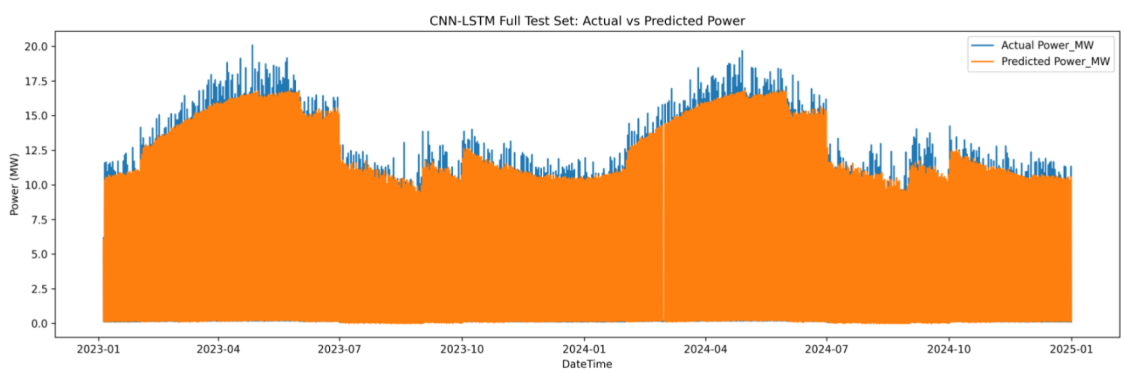

(a) Full test power comparison

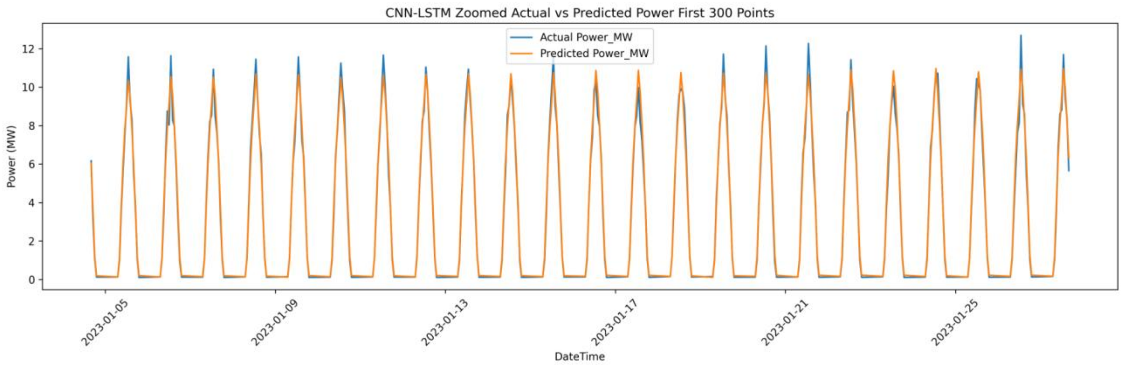

(b) Zoomed power comparison

Figure 11. CNN-LSTM power forecasting performance over the full test horizon and a zoomed inspection window [9], [15].

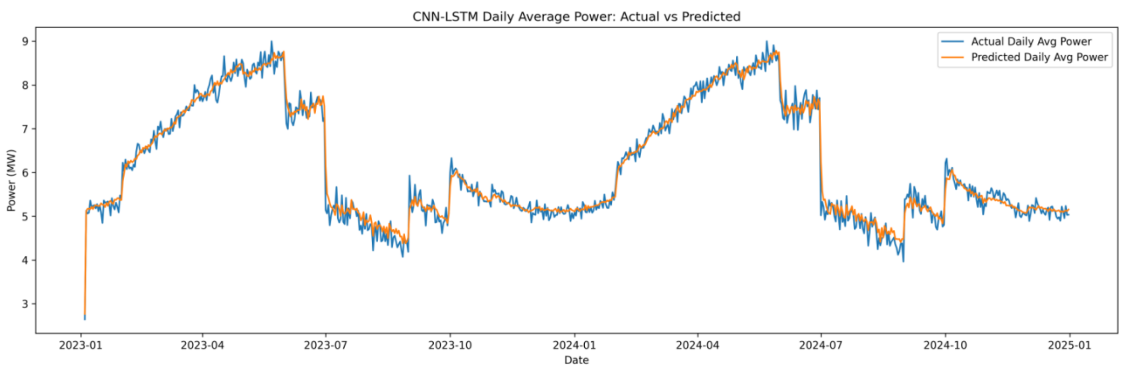

Figure 12. CNN-LSTM daily average power trend [9].

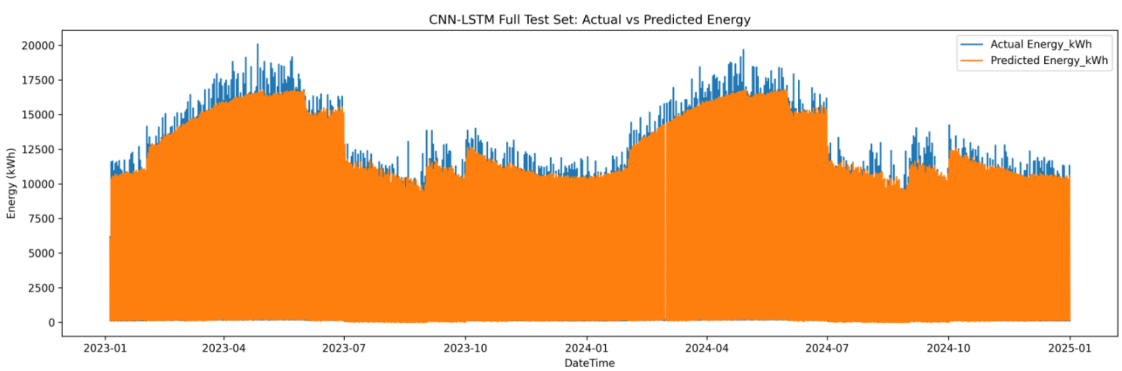

(a) Full test energy comparison

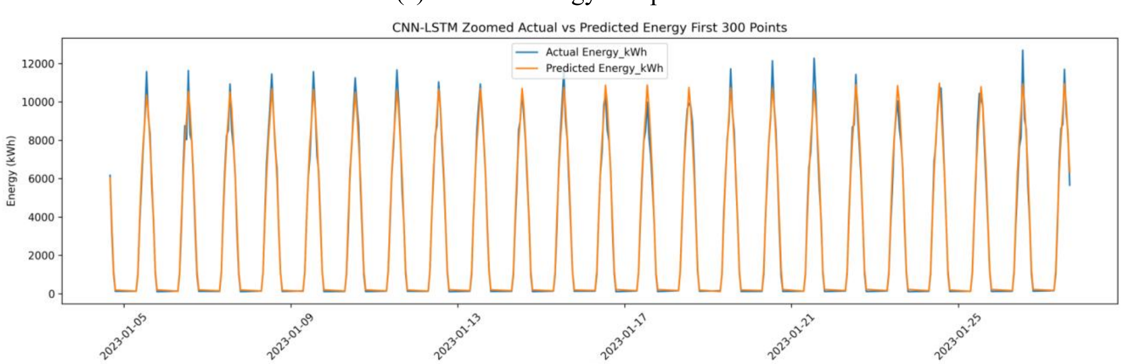

(b) Zoomed energy comparison

Figure 13. CNN-LSTM energy forecasting performance over the full test horizon and a zoomed inspection window [5], [15].

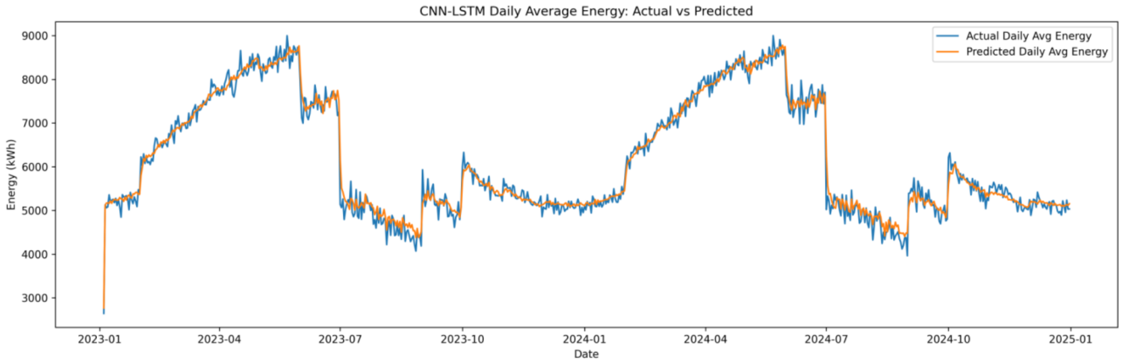

Figure 14. CNN-LSTM daily average energy trend [15].

The diagnostic plots of CNN-LSTM are provided in Fig. 15. The scatter plot is tighter around the identity line than that of the LSTM model, and the residual histogram is more concentrated around zero. The residual-versus-predicted plot still reveals some spread at high power levels, but the overall distribution is narrower than the plain LSTM case. These graphical observations explain the improvement of RMSE from 0.7039 MW to 0.6269 MW.

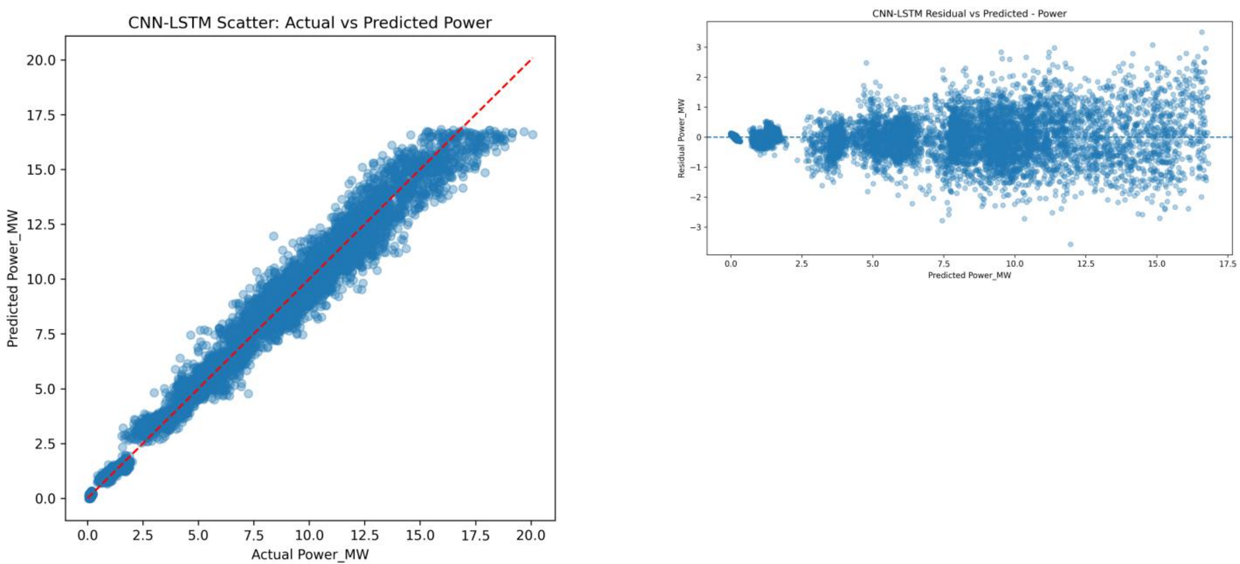

(a) Actual vs. predicted power

(b) Residual vs. predicted power

Figure 15. CNN-LSTM diagnostic scatter and residual-versus-predicted plots for solar power forecasting [9], [13].

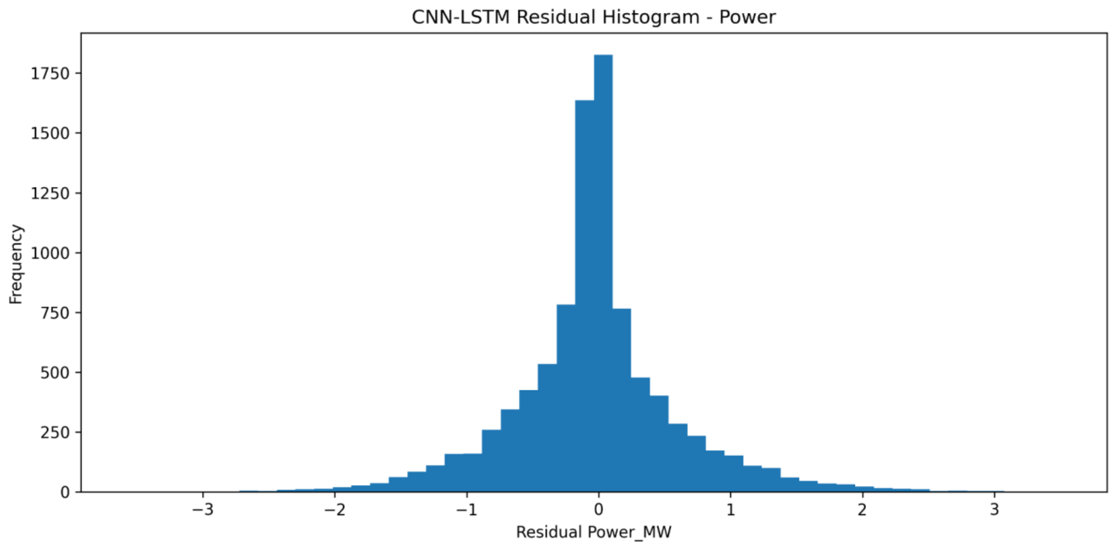

Figure 16. CNN-LSTM residual histogram for solar power forecasting [13].

The energy-based diagnostics of CNN-LSTM are shown in Fig. 17. Compared with the LSTM energy diagnostics, the point cloud is more compact and the residual spread is lower. This supports the conclusion that adding convolutional feature extraction improves the ability of the model to follow local changes in the time series. In solar forecasting, this advantage is meaningful because the same daily solar envelope may contain different local shapes depending on cloud cover and atmospheric conditions.

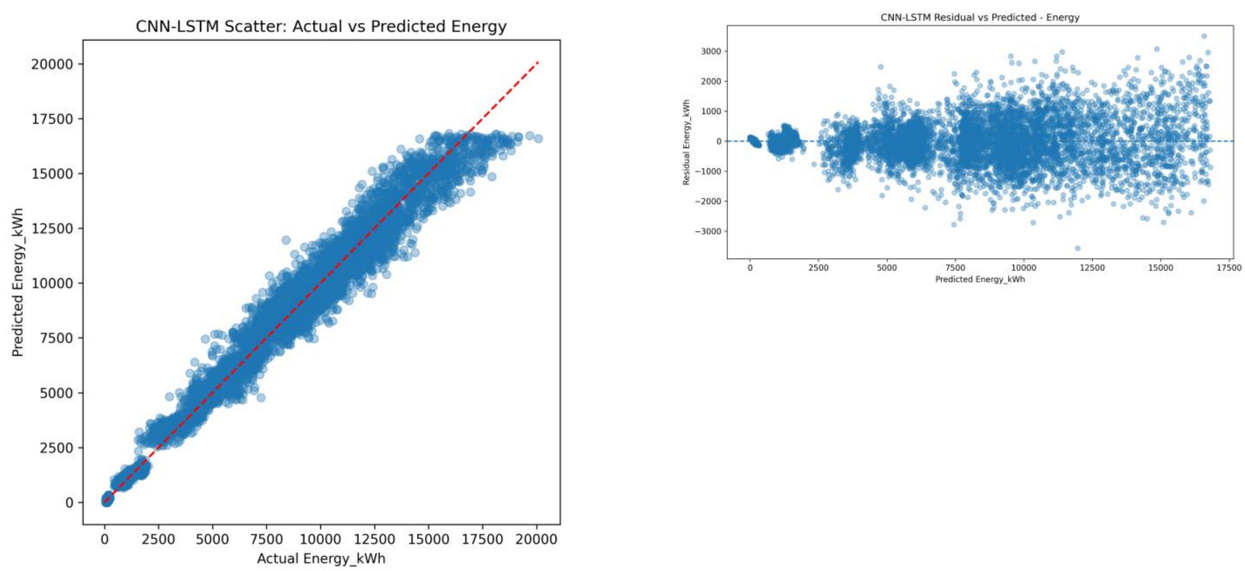

(a) Actual vs. predicted energy

(b) Residual vs. predicted energy

Figure 17. CNN-LSTM diagnostic scatter and residual-versus-predicted plots for solar energy forecasting [9], [13].

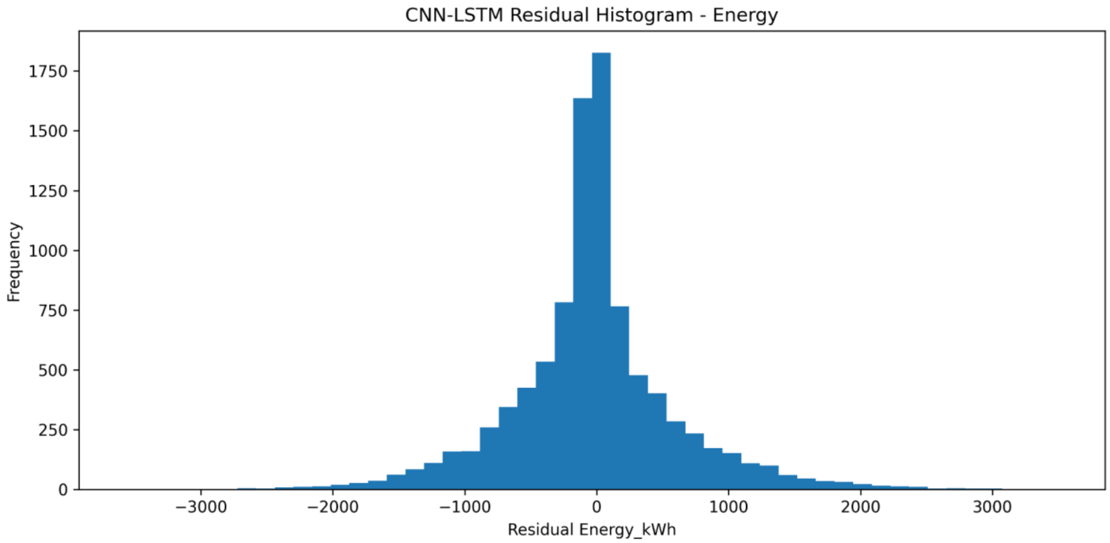

Figure 18. CNN-LSTM residual histogram for solar energy forecasting [13].

Although CNN-LSTM improves the prediction accuracy, its residuals still expand during medium and high generation intervals. This indicates that the sequence model learns the deterministic component of the solar profile more easily than the stochastic component introduced by cloud movement. Therefore, if CNN-LSTM is selected for real-time use, additional short-horizon cloud indicators or probabilistic correction layers may further improve peak and ramp prediction.

E. XGBoost Model Analysis

The XGBoost model gives the best forecasting performance among the three evaluated models [18]–[20]. Unlike the sequence-based deep learning models, XGBoost learns directly from the engineered tabular feature space, including GHI, cloud cover, temporal variables, lagged power, rolling statistics, and interaction features. Fig. 19 shows the full testperiod comparison and the zoomed comparison of the first 300 points. The predicted curves almost coincide with the actual series, which reflects the very low RMSE and MAE values. The superior result of XGBoost can be explained by the nature of the available input data. The dataset is not only a raw time series; it also contains physically meaningful tabular features such as irradiance, cloud cover, humidity, temperature, cyclic time indicators, and lagged power. Gradient-boosted trees are very effective at dividing such a feature space into nonlinear operating regions. For example, the model can learn different mappings for high-GHI clear-sky periods, cloudy low-GHI conditions, morning ramping intervals, and seasonal transitions. This local partitioning ability is one reason why XGBoost can outperform a sequence model when rich engineered features are available [15], [18]. The feature-importance analysis shows that the strong XGBoost result is driven mainly by irradiance-related information rather than by lagged-output variables. Lagged and rolling variables remain useful as stabilizing temporal descriptors, but their importance is much smaller than that of GHI and the GHI-temperature interaction. Therefore, the XGBoost result should be understood primarily as the outcome of target-time weather-informed modeling with a powerful nonlinear tabular learner, not simply as an effect of persistence [18], [19].

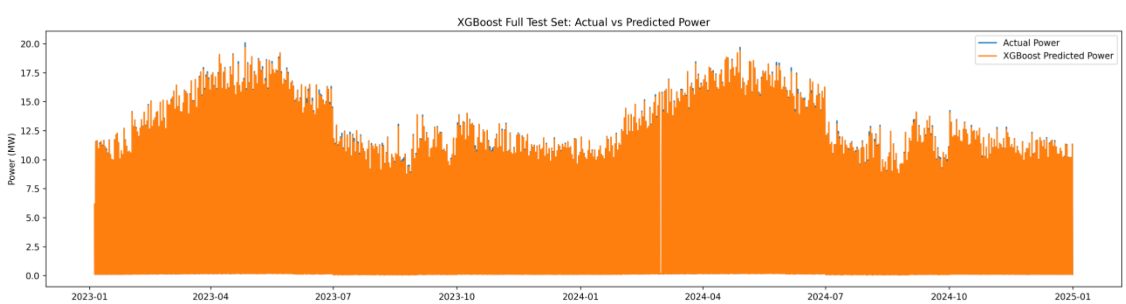

(a) Full test power comparison

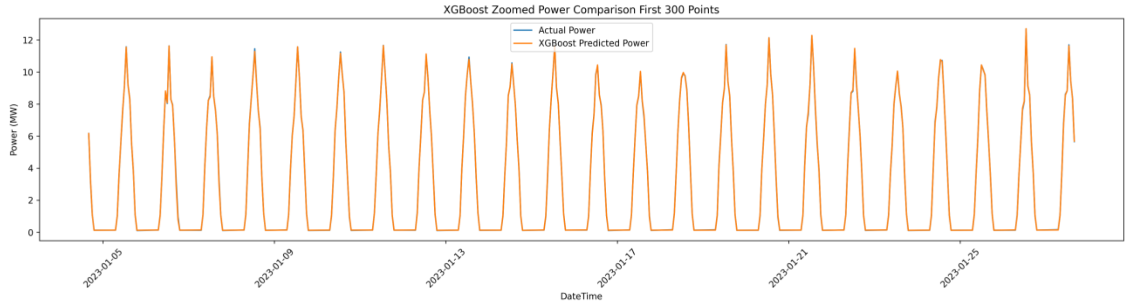

(b) Zoomed power comparison

Figure 19. XGBoost actual versus predicted power comparison under weather-informed assessment [18]–[20].

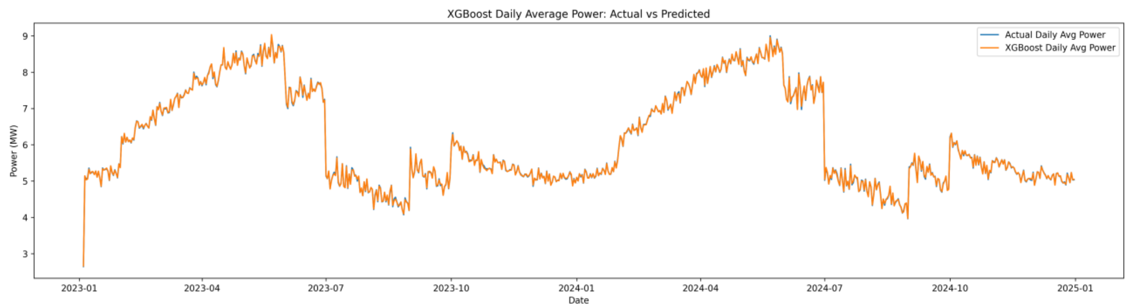

Figure 20. XGBoost daily average power comparison [18], [19].

The corresponding diagnostic plots are shown in Fig. 21. The actual-versus-predicted scatter points are very tightly concentrated along the ideal diagonal line, the residual histogram is strongly centered near zero, and the residual-versuspredicted plot exhibits very limited dispersion. Compared with the DL models, the error band is significantly narrower over the entire operating range. The residual pattern also indicates that the XGBoost model does not suffer from a strong systematic bias across low, medium, and high generation levels. This is important for operational forecasting because underprediction at high generation may affect reserve planning, while overprediction during cloudy periods may create dispatch imbalance [1], [11], [13].

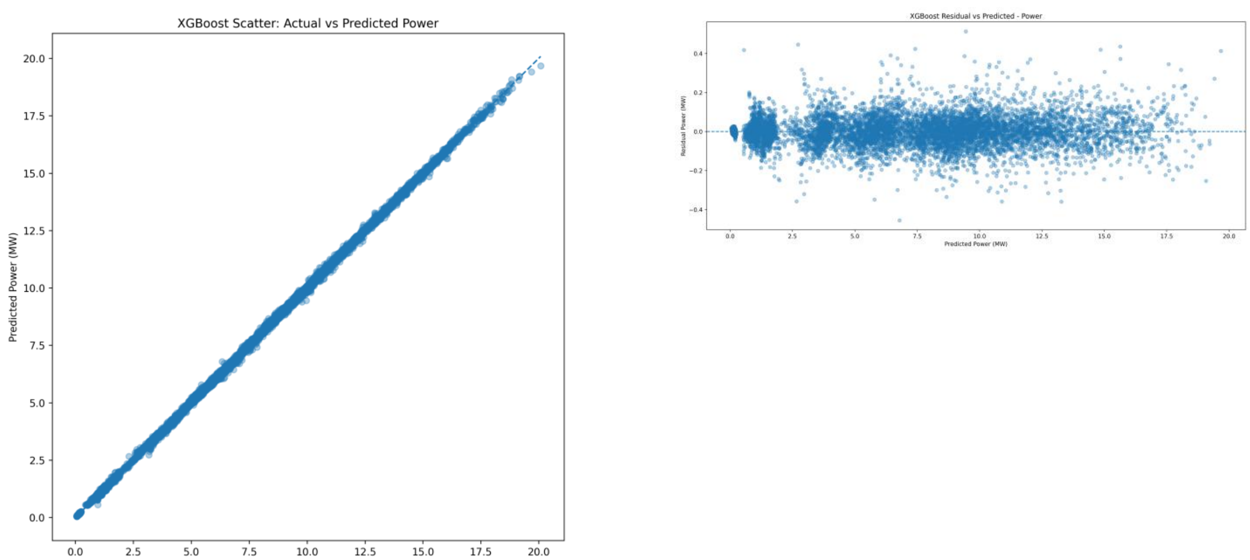

(a) Actual vs. predicted power

(b) Residual vs. predicted power

Figure 21. XGBoost diagnostic scatter and residual-versus-predicted plots for solar power forecasting [13], [18].

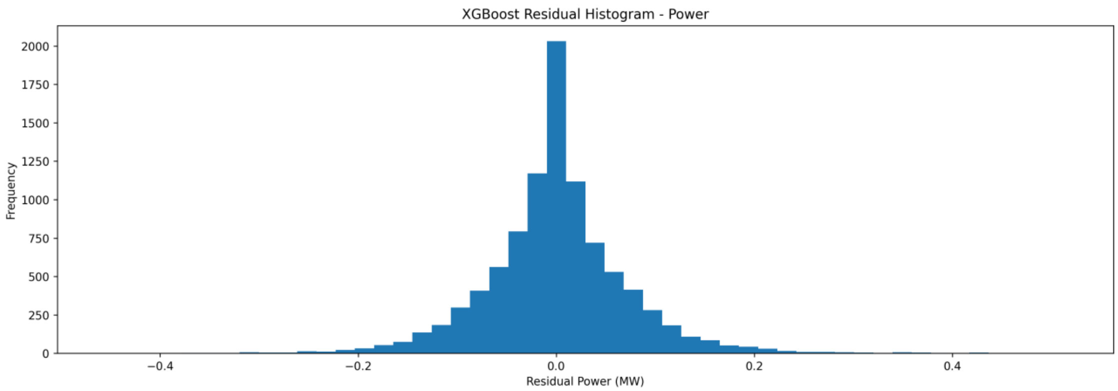

Figure 22. XGBoost residual histogram for solar power forecasting [18].

The predictive quality of XGBoost can be attributed mainly to the target-time irradiance and weather variables, together with physically meaningful interaction terms. The gradient boosting framework captures nonlinear thresholds and operating regimes effectively for this structured dataset. However, the result should not be generalized to strict future forecasting unless the same target-time meteorological variables are replaced by valid exogenous forecasts or nowcasts. From an operational point of view, this behaviour is important because PV power forecasting is not only a curve-fitting task. A plant operator requires a model that performs well during different operating regimes, including low-generation hours, fast morning ramp, high irradiance peak,

Table VI. Top XGBoost features and their physical interpretation [18], [19]

Feature	Importance
GHI (W/m²)	0.558412
GHI×Temperature	0.291878
Cloud cover (%)	0.044623
Hour cosine	0.038023
Month cosine	0.015352
Month sine	0.010647
Relative humidity (%)	0.009162
GHI×Cloud cover	0.007191
Lag power (168 h)	0.004441
GHI×Humidity	0.004320

cloudy transition, and evening decline. XGBoost is able to form conditional rules for these regimes through decision-tree splits. For example, the same GHI level may lead to different power output under different temperature and cloud-cover conditions. A boosted-tree model can represent such conditional dependence without requiring a manually defined physical equation. This makes XGBoost particularly suitable when satellite-derived tabular weather variables are available with recent plant measurements.

F. Feature-Importance Interpretation

The XGBoost feature-importance plot is presented in Fig. 23, while the top features are summarized in Table VI, following the interpretability advantage typically associated with tree-based boosting models [18], [19]. GHI is the dominant feature by a large margin, which is physically expected because solar radiation drives PV production directly. The GHI-temperature interaction is the second most important feature, confirming that PV output depends not only on irradiance magnitude but also on temperature-dependent conversion efficiency. Cloud cover, hour cosine, and month-based cyclic features also contribute meaningfully, as they describe sky obstruction and regular diurnal/seasonal patterns.

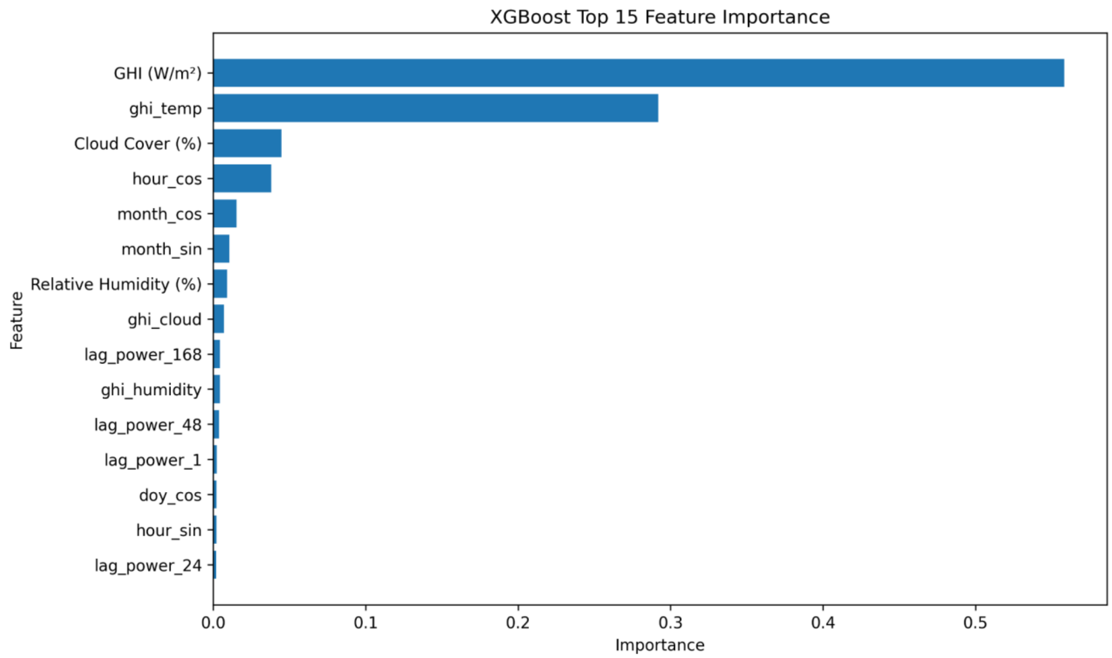

Figure 23. Top 15 XGBoost feature-importance values, showing the interpretability advantage of tree boosting [18], [19].

The importance ranking has practical significance. If a forecasting system must be simplified for online implementation, the plot suggests that irradiance, cloud cover, and a few engineered interactions can preserve most of the predictive

power. This also indicates that feature engineering is a major contributor to the performance advantage of XGBoost.

G. Error Comparison and Practical Interpretation

From a target-time estimation perspective, the error reduction achieved by XGBoost is substantial. The difference between an RMSE of 0.7039 MW and 0.0719 MW shows that the tabular boosting model reproduces the measured power curve very closely when target-hour weather inputs are available. For operational reserve planning, however, these numbers should be treated as an upper-bound benchmark. A deployable forecast would require the same experiment to be repeated with weather forecasts available at the forecast issue time [6], [13]. The comparison between LSTM and CNN-LSTM is also meaningful. Although both belong to the deep learning family, the CNN-LSTM architecture yields improved prediction quality because the convolution layer performs local temporal pattern extraction before long-term sequence modeling. This is particularly useful for solar data, where clear-sky behavior follows smooth trajectories but intermittent cloud effects may produce local disturbances [9], [15].

H. Detailed Error Behaviour Across Operating Regions

The error behaviour of a PV forecasting model is not uniform over the full operating range. At very low generation, especially during early morning, evening, and night periods, the absolute error is naturally small because the plant output is close to zero. During medium generation, the model must follow the morning ramp and afternoon decline, where the output changes with both solar angle and cloud conditions. During high generation, even small irradiance variations can produce large changes in power, and therefore the residual spread usually increases [1], [7], [12]. The LSTM residual plots show that the model has broader dispersion at medium and high predicted values. This occurs because the LSTM emphasizes sequential continuity and tends to smooth rapid local variation. CNN-LSTM reduces this spread by detecting local patterns before recurrent processing. XGBoost provides the narrowest residual band because it can directly split the feature space into different operating regimes. For example, the relationship between GHI and power is different under high cloud cover and low cloud cover, and tree-based models can represent such conditional relationships efficiently [18], [19].

I. Operational Significance of Forecasting Accuracy

For a utility or PV plant operator, the practical value of a forecast depends on how much uncertainty it removes from scheduling decisions. A reduction in RMSE from 0.7039 MW for LSTM to 0.0719 MW for XGBoost is significant for a 25 MW plant. Lower error reduces the reserve margin needed to compensate for forecast deviation. It also improves the accuracy of day-ahead and intra-day energy planning, especially when the PV plant is combined with battery storage or contractual power delivery requirements [1], [11].

MAPE and sMAPE provide additional interpretation only when computed carefully. Since PV output becomes zero at night, percentage errors are reported for the positivegeneration subset rather than all 24-hour samples. RMSE and MAE are therefore the primary scale-dependent error measures, while daytime MAPE and sMAPE are supporting indicators. With this definition, the metrics consistently show that CNN-LSTM improves over LSTM, and XGBoost gives the strongest weather-informed estimation performance [13], [18].

J. Role of Feature Engineering in Model Performance

The results also show that feature engineering plays a decisive role in short-term solar forecasting. GHI directly represents available solar radiation, cloud cover reflects atmospheric obstruction, temperature affects module conversion efficiency, and humidity can represent atmospheric moisture conditions. Cyclic features allow the model to distinguish between morning and evening even when some weather values are similar. Lagged power features represent persistence, which is one of the strongest baselines for short-horizon solar forecasting [2], [13]. Interaction features are especially useful because the effect of one variable may depend on another variable. For example, a high GHI value under high temperature may not produce the same output as the same GHI value under moderate temperature due to module efficiency effects. Similarly, GHI combined with cloud cover can represent whether irradiance is stable or influenced by variable sky conditions. XGBoost benefits strongly from such interactions because decision trees can split the feature space according to nonlinear thresholds [18], [19].

K. Model Selection Guidelines

The comparative results suggest that model selection should depend on the forecasting environment. If the available data are mainly sequential and contain limited engineered weather features, LSTM or CNN-LSTM may be suitable because they can learn temporal dependencies directly. If the dataset contains structured weather variables, lagged power values, and interaction features, XGBoost can be a stronger practical choice because it offers high accuracy, faster training, and better interpretability. CNN-LSTM remains useful when the goal is to develop a deep learning framework capable of extension to image sequences, sky-camera inputs, or multi-step temporal prediction. XGBoost is more suitable when the operator requires a reliable tabular model that can be interpreted through feature importance. Therefore, the proposed comparison does not indicate that one model family is universally superior; it shows that, for the given 10-year structured satellite-derived weather dataset, XGBoost is the most effective model [13], [15].

L. Relevance of Satellite-Derived Weather Inputs

Satellite-derived weather variables are valuable for PV forecasting because they provide spatially consistent information

about irradiance and cloud conditions. In conventional plantlevel monitoring, only local measurements may be available, and these measurements may not fully describe approaching cloud movement. Satellite-derived GHI and cloud cover provide an external view of atmospheric conditions, which helps the forecasting model distinguish between clear-sky generation, partially cloudy generation, and heavily attenuated irradiance conditions. In the present dataset, the dominance of GHI and cloud cover in the XGBoost feature-importance result confirms this physical relationship [6], [7]. Another advantage of satellite-derived data is that they can be integrated with historical plant output without requiring expensive ground-based sensor networks at every location. This is particularly useful for distributed solar plants, remote PV parks, and utility-scale plants located in regions where dense meteorological stations are not available. When such data are combined with temporal and lagged features, the forecasting model receives both weather-state information and plant-response information. This combination explains why the tree-based model performs strongly in the present study [18], [19].

M. Comparison of Deterministic and Practical Forecast Requirements

The forecasts developed in this paper are deterministic, meaning that each model produces a single expected power or energy value for each time step. Deterministic forecasts are useful for dispatch planning, inverter-level monitoring, and plant performance assessment. However, in real power-system operation, forecast uncertainty is also important. A forecast with low average error but unknown uncertainty may still lead to suboptimal reserve scheduling during highly variable cloud conditions. Therefore, the residual distributions shown in this paper provide an initial indication of uncertainty, even though a formal probabilistic forecasting method is not implemented [10], [13], [16]. The residual histogram and residual-versus-predicted plots can be used by operators to identify where uncertainty is concentrated. For example, if the residual spread is high at large predicted power values, the operator may keep additional reserves during high-irradiance but cloud-variable hours. If the residuals are tightly concentrated around zero, the forecast can be treated as more reliable. In this respect, the XGBoost model is more attractive than the two deep learning models because its residual band is much narrower across the operating range [13], [18].

N. Implications for Energy Management and Grid Operation

Accurate solar forecasting supports several downstream applications. In a grid-connected PV plant, the forecast helps estimate expected injection into the grid and reduces mismatch between scheduled and actual generation. In a microgrid, the forecast can be used for battery charge-discharge scheduling, diesel generator commitment, and demand-response planning. In a hybrid renewable system, accurate PV forecasts help coordinate solar generation with wind generation and storage [1], [11].

For a 25 MW solar plant, even a small percentage improvement in forecasting accuracy can have operational value. Lower RMSE reduces balancing energy requirements, while lower MAE improves the expected accuracy of hourly generation planning. The reduction in nRMSE from 11.4281% for LSTM to 1.1671% for XGBoost indicates that the proposed feature-engineered boosting approach can support more confident operational decisions. This is especially important for short-term horizons where grid operators need rapid and reliable estimates rather than only long-term energy averages [1], [13].

O. Discussion With Respect to Literature The outcome of this study is consistent with the broader literature showing that no single forecasting model is universally dominant for all datasets. LSTM-based models are strong when temporal dependencies are the main information source, while CNN-based hybrid models improve performance by extracting local patterns. Satellite-image and sky-image methods are valuable when visual cloud dynamics are available. In contrast, boosting-based methods become very competitive when the dataset contains structured weather variables and carefully designed lag features. The present results therefore add a practical observation: for satellite-derived weather data converted into a tabular forecasting dataset, XGBoost may offer stronger accuracy and interpretability than deep sequence models. This does not reduce the value of LSTM or CNN-LSTM; rather, it highlights the importance of matching the model type to the data form. When raw images or high-dimensional spatio-temporal data are used, deep learning may become more advantageous. When compact weather and plant-output features are used, regularized boosting can be highly effective [18]–[20].

P. Implementation Considerations for Practical Deployment For practical deployment, the forecasting model should be evaluated not only by accuracy but also by data availability, retraining effort, computational time, and interpretability. LSTM and CNN-LSTM models require sequence preparation, scaling consistency, and careful monitoring of training convergence. They can be computationally heavier than tree-based methods, especially when the model architecture is expanded or when multi-step forecasting is added. However, once trained, they can be integrated into an online forecasting pipeline if input sequences are updated hourly [3], [4]. XGBoost is generally easier to deploy in a tabular environment because it accepts engineered input features directly and provides feature-importance information. This makes it suitable for plant operators who need to understand why a forecast changes. For example, if the predicted power decreases due to a drop in GHI and an increase in cloud cover, the feature-importance interpretation provides a physically meaningful explanation. Such interpretability is useful in operational control rooms where forecasts must be trusted by engineers and dispatchers [18], [19]. Another practical consideration is retraining frequency. Solar plant behavior may change over time due to panel degradation, cleaning cycles, inverter maintenance, seasonal soiling,

or sensor recalibration. Therefore, any deployed model should be periodically retrained or recalibrated using recent data. In this study, the chronological split already tests the ability of the models to generalize from past years to future years. The strong test performance of XGBoost suggests that the learned mapping remains stable for the 2023–2024 period, but long-term deployment should still include periodic validation [13].

Q. Recommended Forecasting Pipeline

Based on the findings of this work, a practical deployment pipeline should be arranged in four stages. In the first stage, satellite-derived weather variables, weather forecasts or satellite nowcasts, and plant SCADA measurements are collected and time-aligned. In the second stage, missing or inconsistent values are treated, and cyclic, lagged, rolling, and interaction features are created without using future target power. In the third stage, the model is trained using only past data and validated on a separate chronological period. In the fourth stage, operational testing must use only information available at the forecast issue time. This final stage is essential for separating genuine forecasting performance from target-time weather-informed estimation accuracy [6], [13]. This pipeline is important because solar forecasting models are highly sensitive to preprocessing. Random splitting of data may produce over-optimistic results because samples from the same season or adjacent days can appear in both training and testing subsets. The chronological split used in this paper avoids that issue and provides a more realistic estimate of future forecasting ability. Similarly, scaling parameters should be learned only from the training data and then applied to validation and testing data. This prevents future information from entering the model-development process. For operational use, the forecast output can be generated every hour as new weather and plant measurements become available. The model output may then be passed to energy management, reserve scheduling, or battery dispatch modules. If the forecast error exceeds a predefined threshold, the system can trigger a retraining check or request updated satelliteweather inputs. Such a closed-loop forecasting arrangement would make the proposed modelling framework more useful for real-time renewable-energy management.

R. Limitations and Future Work

Even though the forecasting performance is strong, several limitations remain. First, the comparison is based on a single plant dataset and a specific forecasting horizon. Second, the study uses satellite-derived weather variables rather than raw satellite imagery or numerical weather prediction fields. Third, uncertainty quantification is not explicitly modeled; the present study focuses on deterministic forecasting. Future work can extend the analysis to probabilistic forecasting, confidence intervals, transfer learning across plants, and hybrid frameworks that integrate satellite images with structured weather features.

8. CONCLUSION

This paper presented a detailed comparative study of weather-informed short-term solar power estimation and

forecasting-oriented assessment using LSTM, CNN-LSTM, and XGBoost models with satellite-derived weather data. A consistent chronological training-validation-testing framework was adopted for a 25 MW solar PV plant dataset spanning 2015–2024. The results show that all three models can capture the fundamental solar generation pattern, but their predictive accuracy differs substantially. The LSTM model provides a useful deep learning baseline with RMSE of 0.7039 MW. The CNN-LSTM model improves performance to 0.6269 MW, demonstrating the benefit of combining convolutional feature extraction with recurrent memory. Under the target-time weather-informed setting, XGBoost obtains the lowest error with RMSE of 0.0719 MW, MAE of 0.0496 MW, daytime MAPE of 1.9588%, R^2 of 0.9998, and nRMSE of 1.1671%. Graph-based analysis across all three models shows that XGBoost offers the closest match to actual power, the tightest scatter distribution, and the narrowest residual spread. Feature-importance analysis identifies GHI, GHI-temperature interaction, and cloud cover as the most influential inputs, while lagged features have comparatively small importance values. Overall, the study concludes that XGBoost is highly effective for weather-informed short-term PV power estimation. A separate strict h-ahead experiment using forecasted weather inputs is required before claiming the same accuracy for operational future forecasting.

ACKNOWLEDGMENT

The authors acknowledge the support of the Electrical Engineering Department, Indus Institute of Technology and Engineering, Indus University, Ahmedabad for academic guidance and research support.

References

1. L. Gigoni, A. Betti, E. Crisostomi, A. Franco, M. Tucci, F. Bizzarri, and D. Mucci, “Day-ahead hourly forecasting of power generation from photovoltaic plants,” IEEE Transactions on Sustainable Energy, vol. 9, no. 2, pp. 831–842, Apr. 2018, doi: 10.1109/TSTE.2017.2762435.
2. C. Pan and J. Tan, “Day-ahead hourly forecasting of solar generation based on cluster analysis and model,” IEEE Access, vol. 7, pp. 112921– 112930, 2019, doi: 10.1109/ACCESS.2019.2935273.
3. Y. Wang, Y. Shen, S. Mao, X. Chen, and H. Zou, “LASSO and LSTM integrated temporal model for short-term solar intensity forecasting,” IEEE Internet of Things Journal, vol. 6, no. 2, pp. 2933–2944, Apr. 2019, doi: 10.1109/JIOT.2018.2877510.
4. H. Zhou, Y. Zhang, L. Yang, Q. Liu, K. Yan, and Y. Du, “Short-term photovoltaic power forecasting based on long short term memory neural network and attention mechanism,” IEEE Access, vol. 7, pp. 78063– 78074, 2019, doi: 10.1109/ACCESS.2019.2923006.
5. C.-J. Huang and P.-H. Kuo, “Multiple-input deep convolutional neural network model for short-term photovoltaic power forecasting,” IEEE Access, vol. 7, pp. 74822–74834, 2019, doi: 10.1109/ACCESS.2019.2921238.
6. L. Cheng, H. Zang, Z. Wei, T. Ding, R. Xu, and G. Sun, “Shortterm solar power prediction learning directly from satellite images with regions of interest,” IEEE Transactions on Sustainable Energy, vol. 13, no. 1, pp. 629–639, Jan. 2022, doi: 10.1109/TSTE.2021.3123476.

7. E. Shirazi, I. Gordon, A. Reinders, and F. Catthoor, “Sky images for short-term solar irradiance forecast: A comparative study of linear machine learning models,” IEEE Journal of Photovoltaics, vol. 14, no. 4, pp. 691–702, Jul. 2024, doi: 10.1109/JPHOTOV.2024.3398365.
8. J. Lee, J. Kang, S. Lee, and H.-M. Oh, “Ultra-short term photovoltaic generation forecasting based on data decomposition and customized hybrid model architecture,” IEEE Access, vol. 12, pp. 21768–21782, 2024, doi: 10.1109/ACCESS.2024.3362234.
9. N. Thipwangmek, N. Suetrong, A. Taparugssanagorn, S. Tangparitkul, and N. Promsuk, “Enhancing short-term solar photovoltaic power forecasting using a hybrid deep learning approach,” IEEE Access, vol. 12, pp. 113279–113296, 2024, doi: 10.1109/ACCESS.2024.3440035.
10. I. K. Bazionis, M. A. Kousounadis-Knousen, V. E. Katsigiannis, F. Catthoor, and P. S. Georgilakis, “An advanced hybrid Boot-LSTM-ICSO-PP approach for day-ahead probabilistic PV power yield forecasting and intra-hour power fluctuation estimation,” IEEE Access, vol. 12, pp. 44334–44353, 2024, doi: 10.1109/ACCESS.2024.3381049.
11. T. Boonsri, N. Kaewdornhan, S. Somchit, and R. Chatthaworn, “Enhancing distribution system management through LSTM-based short-term forecasting of solar irradiation and load,” IEEE Access, vol. 13, pp. 199799–199814, 2025, doi: 10.1109/ACCESS.2025.3634681.
12. R. Rastgoo, N. Amjady, S. Lin, and S. M. Muyeen, “Ultra-short-term solar power prediction using sky image sequences by a residual vision reformer,” IEEE Transactions on Sustainable Energy, vol. 16, no. 4, pp. 2972–2985, Oct. 2025, doi: 10.1109/TSTE.2025.3575520.
13. N. B. Nam, E. Ogliari, S. Leva, E. Pafumi, D. Alberti, and M. Q. Duong, “Comparative analysis of conformal prediction techniques and machine learning models for very short-term solar power forecasting,” Energy and AI, vol. 21, Art. no. 100573, 2025, doi: 10.1016/j.egyai.2025.100573.
14. Z. Song, F. Xiao, Z. Chen, and H. Madsen, “Probabilistic ultra-short-term solar photovoltaic power forecasting using natural gradient boosting with attention-enhanced neural networks,” Energy and AI, vol. 20, Art. no. 100496, 2025.
15. P. K. Singh, A. Saraswat, and Y. Gupta, “Deep learning prediction models for short-term solar photovoltaic power generation forecasting,” Next Energy, vol. 11, Art. no. 100531, 2026.
16. R. Rastgoo, N. Amjady, S. Lin, R. Shah, and S. M. Muyeen, “Short-term probabilistic solar power forecasting using a deep stochastic framework,” IEEE Access, vol. 14, 2026, doi: 10.1109/ACCESS.2026.3693255.
17. S. Hochreiter and J. Schmidhuber, “Long short-term memory,” Neural Computation, vol. 9, no. 8, pp. 1735–1780, 1997, doi: 10.1162/neco.1997.9.8.1735.
18. T. Chen and C. Guestrin, “XGBoost: A scalable tree boosting system,” in Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, 2016, pp. 785–794, doi: 10.1145/2939672.2939785.
19. R. S. Kumar, P. S. Meera, V. Lavanya, and S. Hemamalini, “Brown bear optimized random forest model for short term solar power forecasting,” Results in Engineering, vol. 25, Art. no. 104583, 2025.
20. X. Liu, J. Zhan, and X. Liu, “Short-term photovoltaic (PV) power forecasting: Leveraging the capabilities of AdaBoost and decision trees,” Journal of Electrical Engineering & Technology, 2026, doi: 10.1007/s42835-026-02791-z.