JCMM Journal of Computers, Mechanical and Management
Journal homepage: www.jcmm.co.in
e-ISSN: 3009-075X
AAN PUBLISHING

Volume 5 Issue 2
Article Number: 25394

Humidity-Aware Hybrid Transformer-LSTM Framework for IoT-Enabled Photovoltaic Power Prediction

Ahmed Mohammed * [1,2], Ranjit Singh Sarban Singh[3], and Saad Aslam[1]

[1]Department of Smart Computing and Cyber Resilience, Faculty of Engineering and Technology, Sunway University, Selangor, Malaysia 47500
[2]Department of Computer Science, Faculty of Science, Gombe State University, Gombe, Nigeria 760253
[3]Research Center for Human-Machine Collaboration (HUMAC), Faculty of Engineering and Technology, Sunway University, Selangor, Malaysia 47500

Abstract

Accurate short-term photovoltaic (PV) power forecasting is essential for Internet of Things (IoT)-enabled monitoring, control, and performance assessment of small-scale solar installations. While electrical variables and temperature are widely used in data-driven PV forecasting models, the contribution of ambient humidity remains insufficiently characterized, particularly in persistently humid environments. This study investigates the role of ambient humidity as a contextual environmental feature and evaluates a humidity-aware Hybrid Transformer-LSTM framework for short-term PV power prediction using real-world IoT data collected from multiple photovoltaic panels over a 34-day monitoring period. The proposed hybrid architecture integrates a Transformer-based self-attention mechanism for cross-feature interaction modeling with LSTM-based recurrent learning to capture temporal persistence. Model performance is evaluated against LSTM-only, Transformer-only, Random Forest, and Linear Regression baselines using a strictly time-ordered train-test split, complemented by architectural and feature ablation studies, rolling time-based validation, cross-panel testing, and robustness analysis under input perturbation. Experimental results show that LSTM-based models achieve the highest predictive accuracy on the short-duration dataset, while ambient humidity provides only marginal and context-dependent benefit as a supplementary environmental feature. Transformer-only models perform poorly under data-limited conditions, while the Hybrid Transformer-LSTM achieves competitive accuracy and demonstrates stable behavior under temporal validation, spatial generalization, and sensor noise. These findings highlight that the primary contribution of this study lies in rigorous evaluation and deployment-aware validation rather than absolute accuracy gains, positioning hybrid attention-recurrent architectures as robustness-oriented solutions for IoT-enabled solar PV monitoring systems.

Keywords: Photovoltaic Power Prediction; IoT-Based PV Monitoring; Humidity-Aware Modeling; Hybrid Transformer-LSTM; Smart Energy Systems

*Corresponding Author: Ahmed Mohammed (ahmed.m26@imail.sunway.edu.my)
Received: 28 Dec 2025; **Revised:** 20 Jan 2026; **Accepted:** 23 Jan 2026; **Published:** 30 Apr 2026

DOI: 10.57159/jcmm.5.2.25394.

1. Introduction

The rapid deployment of photovoltaic (PV) systems in distributed and residential environments has increased the need for accurate short-term power forecasting and reliable monitoring infrastructures [1]. Unlike utility-scale solar plants, small-scale PV installations operate under localized microclimatic conditions and limited instrumentation, making data-driven forecasting approaches essential for operational optimization, energy management, and fault detection [2].

PV power output is primarily governed by solar irradiance and cell temperature [3, 4]. Elevated panel temperatures reduce conversion efficiency, while irradiance variability drives short-term power fluctuations. Ambient humidity has traditionally been regarded as a secondary environmental factor and is therefore often excluded from short-term PV forecasting models. Nonetheless, several studies suggest that humidity can indirectly influence PV performance. These effects are commonly attributed to reduced convective cooling, surface moisture accumulation, and coupled thermal dynamics, particularly in tropical and subtropical climates [5–7]. Despite these physical insights, the explicit contribution of humidity to data-driven PV power prediction remains insufficiently validated using real-world operational data.

Recent advances in Internet of Things technologies have enabled continuous, high-resolution acquisition of electrical and environmental measurements from PV installations using low-cost sensors, embedded devices, and lightweight communication protocols [8, 9]. However, many PV forecasting studies still rely on randomized data splitting or isolated evaluation windows [10], which can introduce temporal leakage and optimistic performance estimates, limiting their applicability in real deployments.

To address these gaps, this study investigates humidity-aware PV power forecasting using a Hybrid Transformer-LSTM framework trained on real-world IoT-collected data. Unlike many existing studies, this work adopts strictly time-ordered data splitting, rolling (walk-forward) validation, cross-panel testing, and robustness analysis to avoid temporal leakage and optimistic bias. This evaluation-first design prioritizes deployment realism and generalization reliability over headline accuracy metrics. In addition, the study presents a deployment-realistic IoT data pipeline incorporating edge-level persistence and an application-layer MQTT acknowledgment mechanism, ensuring reliable end-to-end data delivery under intermittent connectivity.

The main contributions of this work are summarized as follows:

- Empirical characterization of humidity-power relationships using real-world IoT PV data, clarifying the contextual and regime-dependent role of ambient humidity in short-term solar PV power prediction.
- Systematic architectural and feature ablation comparing LSTM-only, Transformer-only, and Hybrid Transformer-LSTM models under identical training and validation conditions.
- A rigorously validated IoT-oriented forecasting framework emphasizing leakage-free preprocessing, time-ordered and rolling validation, cross-panel generalization analysis, and robustness testing, thereby prioritizing reliability and methodological rigor over headline accuracy metrics.

The remainder of this paper is organized as follows. Section 2 reviews related work on PV power forecasting, environmental effects, and IoT-based monitoring systems. Section 3 presents the system architecture, dataset, preprocessing, and the proposed Hybrid Transformer-LSTM model. Section 4 reports the experimental results, including ablation studies and validation outcomes. Section 5 discusses the findings and their practical implications. Section 6 concludes the study and outlines limitations and future research directions.

2. Related Work

Photovoltaic performance modeling has traditionally relied on physical and semi-empirical formulations derived from equivalent circuit representations, thermodynamic principles, and empirical loss models [4, 11]. While these approaches offer physical interpretability, they require detailed system parameters that are often unavailable in small-scale deployments.

To address these limitations, data-driven machine learning methods have been widely adopted. Techniques such as support vector regression, random forests, and artificial neural networks have demonstrated strong predictive capability by learning nonlinear relationships directly from historical measurements [12, 13]. However, many classical machine learning models operate on fixed input vectors and lack explicit mechanisms to model temporal dependencies, limiting their effectiveness for PV power forecasting, where sequential dynamics are prominent. Recurrent neural networks, particularly Long Short-Term Memory (LSTM) architectures, address this limitation and have demonstrated strong performance in PV power forecasting [14]. More recent studies have explored hybrid architectures combining attention mechanisms or convolutional layers with recurrent models [15, 16]. While such hybrids may improve robustness, the benefits are often evaluated under limited validation settings.

Ambient humidity has received comparatively little attention as an explicit input feature in short-term PV forecasting. Despite some literature discussing indirect physical effects, data-driven prediction models rarely include humidity, even in humid climates [6, 7]. Consequently, the empirical value of humidity as a predictive variable remains poorly characterized.

Developments in IoT-based PV monitoring emphasize the importance of reliable data pipelines and edge resilience for reliable data analytics [17–19]. These insights motivate forecasting studies that prioritize validation rigor and deployment realism alongside predictive performance.

3. Methodology

To ensure a leakage-free, deployment-realistic evaluation, all preprocessing and model assessments in this study strictly adhere to temporal ordering. A strictly time-ordered train-test split is employed, where earlier observations are used exclusively for training and later observations for testing, without random shuffling. Data imputation, normalization, and sequence construction are performed after time-ordered splitting, with all learned statistics derived exclusively from the training subset. Model performance is evaluated using this fixed split, complemented by rolling (walk-forward) validation to assess temporal generalization under conditions representative of real IoT deployments.

3.1. IoT-Based Photovoltaic Monitoring Architecture

An end-to-end IoT-based monitoring architecture was deployed to enable reliable acquisition, fault-tolerant transmission, and analytics-driven utilization of electrical and environmental data from small-scale photovoltaic (PV) panels under realistic network conditions. All sensors were sampled synchronously at fixed 60-second intervals to ensure temporal alignment across electrical and environmental measurements. The architecture, illustrated in Figure 1, is organized into six functional layers: sensing, edge processing, local persistence, communication, backend analytics and machine learning, and visualization.

At the sensing layer, each rooftop PV panel is instrumented with multiple sensors to capture synchronized electrical and environmental parameters. An INA226 sensor measures panel voltage and current, from which instantaneous electrical power is computed. Panel surface temperature is measured using a DS18B20 sensor, while ambient temperature and relative humidity are recorded using an SHT45 sensor. Together, these measurements provide the necessary inputs to analyze thermal behavior, humidity–temperature interactions, and their combined influence on PV power output.

The edge processing layer is implemented on a Raspberry Pi Zero W, which performs periodic data acquisition at 60-second intervals, timestamps all measurements, and executes lightweight preprocessing tasks, such as data formatting and integrity checks. To ensure robustness under intermittent network connectivity, each acquired data record is first written to a local publisher database residing on the edge device. This local persistence mechanism enables fault tolerance by preventing data loss during temporary communication outages.

The communication layer employs the Message Queuing Telemetry Transport (MQTT) protocol over Wi-Fi to support lightweight and efficient data transmission. The Raspberry Pi operates as an MQTT publisher, transmitting sensor data to a remote broker. An acknowledgment (ACK) mechanism is implemented so that, once data are successfully received and stored in the server-side subscriber database, an ACK is returned to the edge node. Upon receipt of the ACK, the corresponding local record is marked as synchronized, ensuring eventual consistency between edge and server storage while avoiding redundant retransmissions.

At the backend analytics and machine learning layer, the MQTT subscriber service stores incoming data in a centralized database that serves as the authoritative repository for long-term storage, statistical analysis, and model development. This database is used to train, validate, and evaluate machine learning models for short-term PV power prediction. In the present study, the collected time-series data are used to train recurrent, attention-based, and hybrid deep learning models, enabling systematic architectural and feature ablation, temporal validation, and robustness analysis.

Finally, the visualization layer provides a monitoring dashboard for real-time visualization of PV power output, environmental conditions, and historical trends. Model outputs can be integrated into the dashboard to support predictive insights, performance assessment, and decision-making. The proposed architecture enables reliable, high-resolution data acquisition and analytics, forming a practical foundation for IoT-enabled PV forecasting and monitoring applications.

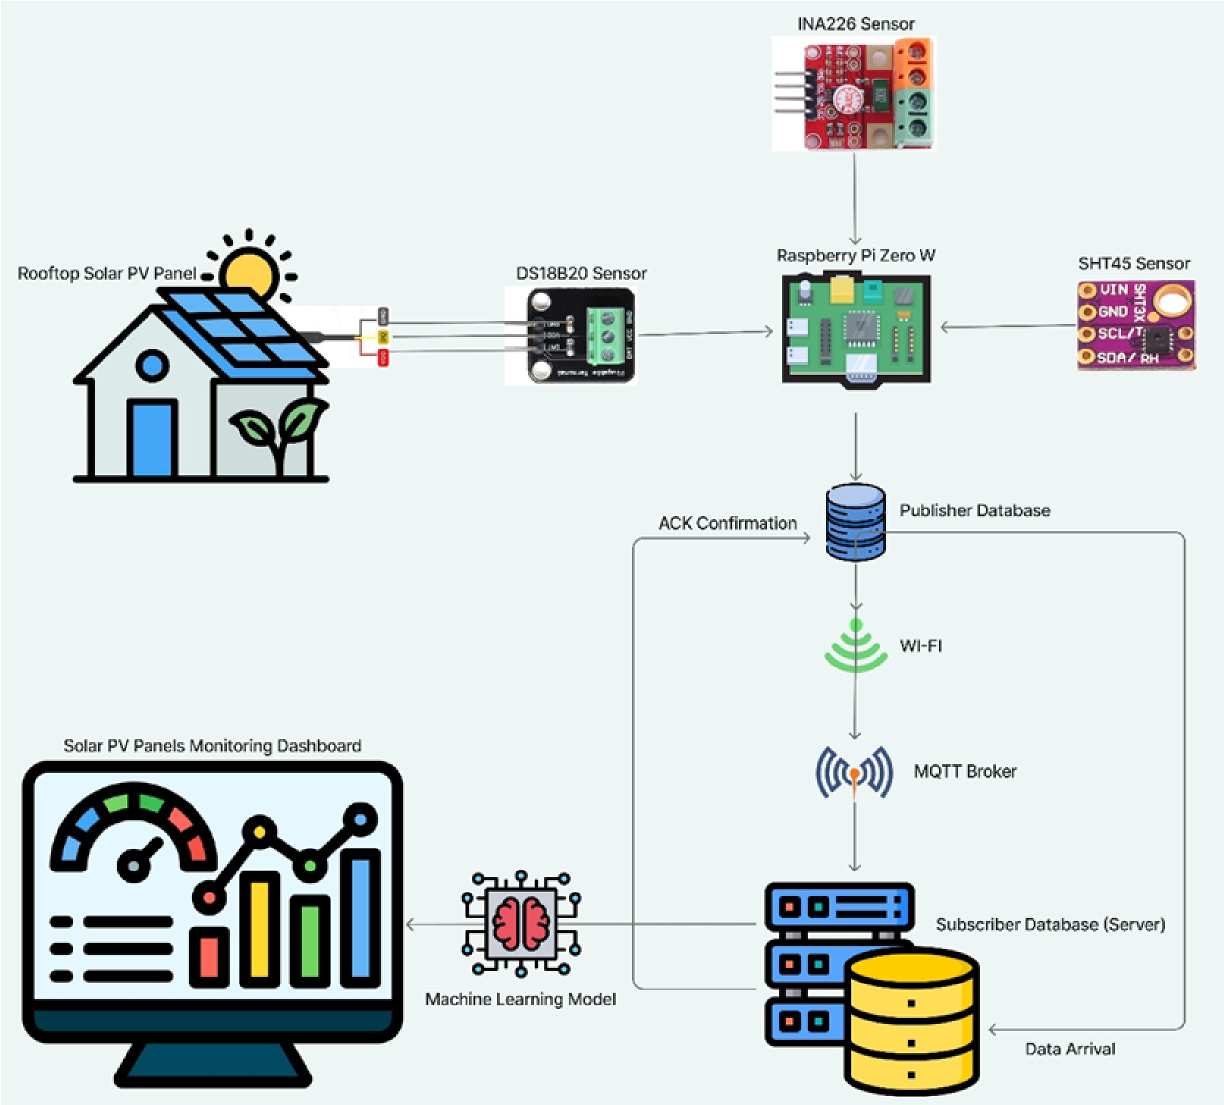

Figure 1: End-to-end IoT-based photovoltaic monitoring and data transmission architecture.

3.2. Dataset Description

3.2.1 Data Collection System

The dataset comprises 66,982 observations collected from three PV panels (PV000, PV001, PV002) over a 34-day period (17 November–21 December 2025), sampled at 60-second intervals.

Solar irradiance measurements were not available in the deployed system due to hardware and cost constraints typical of small-scale and residential IoT installations. Instead, the dataset reflects a minimal-sensor configuration commonly encountered in real-world deployments, where electrical and basic environmental measurements are prioritized over specialized radiometric instrumentation. While the absence of irradiance limits physical interpretability, it enables evaluation of forecasting performance under realistic sensing constraints and motivates the emphasis on temporal learning and validation rigor. The database schema for data storage is given in Table 1.

Table 1: Database schema for IoT-based PV monitoring dataset

Column Name	Data Type	Description
id	BIGINT (PK)	Unique auto-increment identifier for each sensor record
panel_id	VARCHAR	Logical identifier of the PV panel (e.g., PV000, PV001, PV002)
device_id	VARCHAR	Unique hardware identifier of the IoT edge device
voltage	FLOAT	Measured panel output voltage (V)
current	FLOAT	Measured panel output current (A)
power	FLOAT	Instantaneous electrical power output (W)
energy	FLOAT	Accumulated energy produced by the panel (Wh)
panel_temperature	FLOAT	Surface temperature of the solar PV panel (°C)
ambient_temperature	FLOAT	Surrounding environmental temperature (°C)
ambient_humidity	FLOAT	Ambient relative humidity (%)
sent_data_count	INTEGER	Counter for data packets transmitted
sent_at	BIGINT	Timestamp (UNIX epoch in ms) when sent
received_at	BIGINT	Timestamp when received by server
created_at	DATETIME	Server-side record creation timestamp
is_synced	TINYINT	Synchronization flag (0 = not synced, 1 = synced)

3.3. Data Preprocessing

3.3.1 Missing Value Handling

An initial data quality assessment revealed a small proportion of missing values across several attributes. Specifically, voltage, current, and power measurements contained 13 missing records (0.019%), ambient temperature and humidity contained 454 missing records (0.678%), and panel temperature contained a single missing record (0.001%).

To prevent temporal information leakage, missing values were handled after time-ordered data splitting. Panel-wise median imputation and time-based interpolation were applied separately to the training and test subsets, ensuring that future information did not influence model training. This approach preserves panel-specific temporal continuity while maintaining strict separation between training and evaluation data.

Given the low proportion of missing values and strong temporal smoothness of the signals, this preprocessing step minimizes distortion while ensuring leakage-free model evaluation:

$$x_{i,t}^{\text{imputed}} = \text{median}\left(\{x_{i,s} : s \in T_i, x_{i,s} \neq \text{null}\}\right) \tag{1}$$

where $x_{i,t}$ represents the value of attribute i at time t, and T_i denotes all timestamps for the same panel where attribute i is non-null.

3.3.2 Feature Scaling

Feature normalization was performed using Min-Max scaling. Scaling parameters were learned exclusively from the training data and subsequently applied to the test data. Although models were trained using normalized inputs for numerical stability, all predictions were inverse transformed prior to evaluation, and all reported metrics and figures are presented in physical units (watts).

3.4. Correlation Analysis

To quantify relationships between environmental variables and PV power output, Pearson correlation coefficients were computed for all numerical features. The Pearson correlation coefficient between variables X and Y is defined as:

$$r_{XY} = \frac{\sum_{i=1}^{n}(x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^{n}(x_i - \bar{x})^2}\sqrt{\sum_{i=1}^{n}(y_i - \bar{y})^2}} \tag{2}$$

To empirically assess the relevance of environmental variables prior to model training, Pearson correlation coefficients were computed between PV power output and key electrical and environmental parameters, as illustrated in Figure 2. As expected, PV power shows strong positive correlations with current ($r = 0.96$) and voltage ($r = 0.71$), reflecting their direct electrical relationship. Panel temperature shows a moderate positive correlation with power ($r = 0.41$), consistent with thermal effects observed in small-scale PV systems.

Ambient humidity demonstrates a weak but consistent positive correlation with PV power output ($r = 0.15$). Although the magnitude of this relationship is modest, its stability across the dataset supports including humidity as a contextual feature rather than a primary driver. Notably, ambient humidity is strongly negatively correlated with ambient temperature ($r = -0.85$), indicating coupled environmental dynamics that may indirectly influence PV performance. These observations motivate the inclusion of humidity in the multivariate learning framework while avoiding assumptions of direct causality. The cross-panel validation results presented in Section 4.4 further complement this analysis by explicitly assessing panel-dependent generalization behavior.

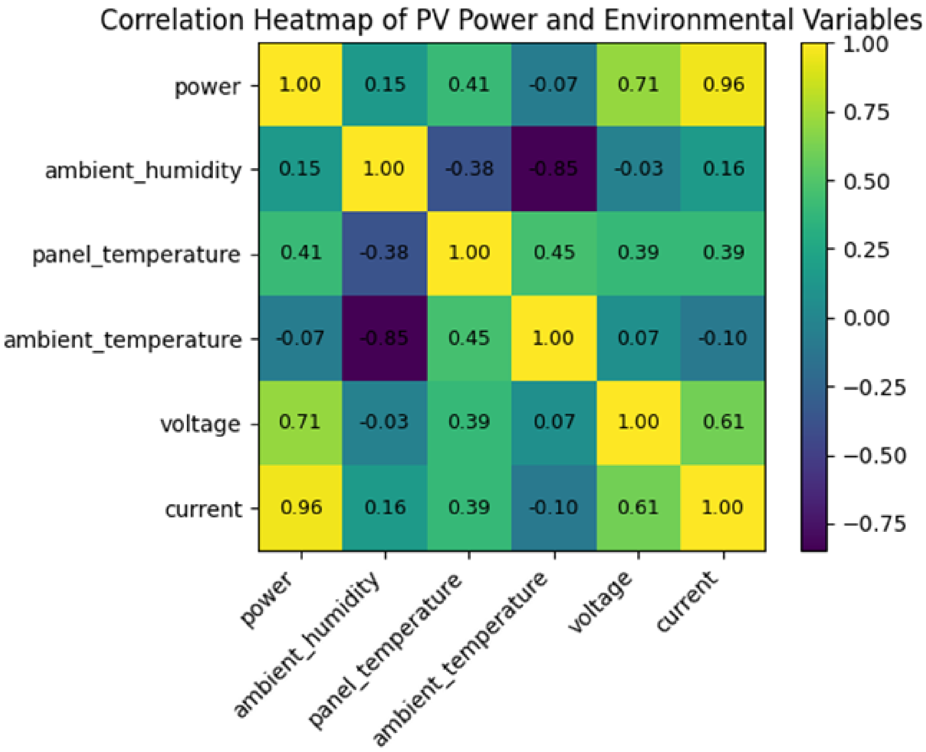

Figure 2: PV panels correlation heatmap.

3.5. Hybrid Transformer-LSTM Model

The proposed Hybrid Transformer-LSTM architecture, shown in Figure 3, integrates a lightweight self-attention mechanism with a recurrent neural network to jointly model cross-feature interactions and temporal dependencies in PV power time series. The model applies a multi-head self-attention layer with 4 attention heads, followed by layer normalization, and then feeds the resulting representations into an LSTM layer with 64 hidden units. A dropout rate of 0.2 is applied prior to the final linear output layer.

The selection of four attention heads reflects a trade-off between representational capacity and data availability. Preliminary experiments with higher headcounts did not yield performance improvements and led to less stable convergence due to the limited dataset duration. Similarly, the 60-minute input window was chosen empirically to balance short-term temporal persistence with computational efficiency; shorter windows failed to capture gradual power transitions, while longer windows introduced redundancy without measurable accuracy gains.

For the Transformer-only baseline, a lightweight self-attention block was intentionally employed without positional encoding or stacked feed-forward sublayers. This design reflects the short, smooth, and highly autocorrelated nature of PV power time series, where deep Transformer stacks were empirically ineffective. The poor performance of Transformer-only models, therefore, reflects the limitations of attention mechanisms under data-limited conditions rather than an implementation deficiency.

Early stopping was intentionally avoided to maintain comparability across architectural and feature-ablation experiments. All models were trained for a fixed number of epochs, and convergence behavior was monitored to confirm stable loss trajectories. While this approach may increase the risk of overfitting, rolling time-based validation and cross-panel testing mitigate this concern by evaluating generalization under unseen temporal and spatial conditions.

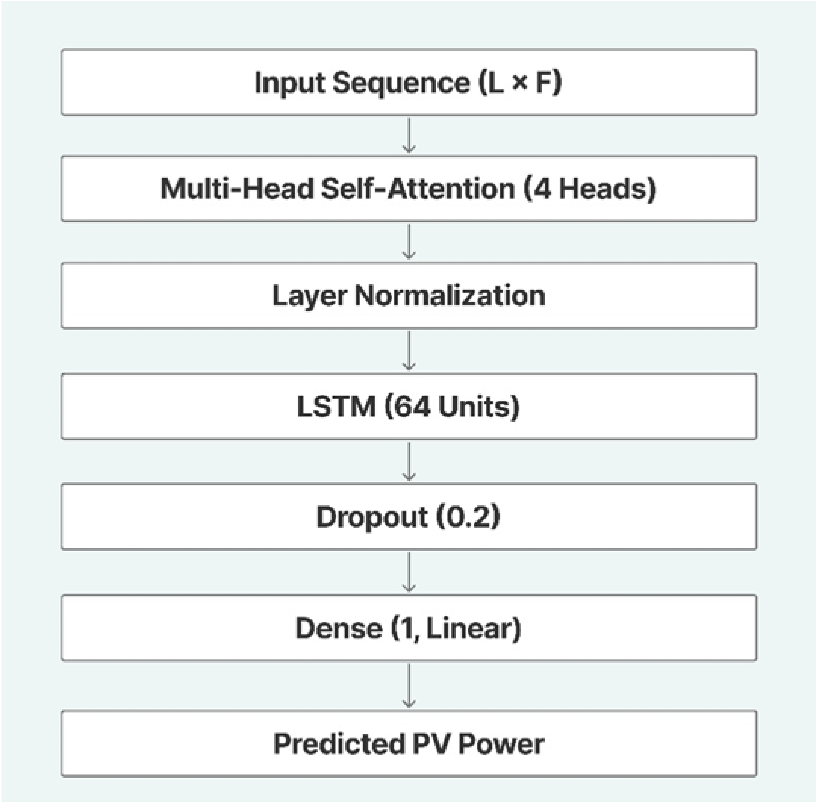

Figure 3: Hybrid Transformer-LSTM architecture.

3.6. Baseline Models and Evaluation Metrics

Baseline models include Linear Regression, Random Forest, Transformer-only, and LSTM-only models. The Random Forest was configured with 200 trees, maximum depth of 20, minimum samples per leaf of 2, and bootstrap sampling enabled, with all other parameters set to their default values.

Performance metrics used include MAE, RMSE, nRMSE, and R^2:

Mean Absolute Error (MAE):

$$\text{MAE} = \frac{1}{n}\sum_{i=1}^{n} |y_i - \hat{y}_i| \tag{3}$$

Root Mean Squared Error (RMSE):

$$\text{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2} \tag{4}$$

Coefficient of Determination (R^2):

$$R^2 = 1 - \frac{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{n}(y_i - \bar{y})^2} \tag{5}$$

The normalized RMSE (nRMSE) is computed by dividing RMSE by the mean observed PV power of the corresponding test set, enabling scale-independent comparison across models and panels. These metrics are reported in Table 2.

4. Results

This section presents the experimental results from the proposed humidity-aware Hybrid Transformer-LSTM framework and the baseline models. All experiments were carried out with strictly time-ordered data splitting, leakage-free preprocessing, and the same training protocols for all models to ensure fair evaluation. Along with single-split evaluation, rolling time-based validation, cross-panel testing, and noise-robustness analysis, these methods were used to test generalization and deployment robustness.

4.1. Time-Ordered Test Split Performance

Table 2 summarizes the predictive performance of all evaluated models under a strictly time-ordered train-test split. Among all evaluated approaches, the LSTM model without ambient humidity achieved the best overall performance, with an RMSE of approximately 0.029 W and an R^2 exceeding 0.99, indicating strong temporal learning for short-term photovoltaic power prediction. Incorporating ambient humidity did not improve predictive accuracy for LSTM or Hybrid models and, in some cases, slightly degraded performance. The Hybrid Transformer-LSTM achieved competitive accuracy but did not outperform the simpler LSTM architecture. Transformer-only models exhibited substantially higher error (RMSE $\approx$ 0.28 W), confirming that attention mechanisms alone are ineffective for short, smooth PV time series under data-limited conditions. Classical baselines (Random Forest and Linear Regression) underperformed sequence-based models, highlighting the importance of explicit temporal modeling.

Table 2: Model performance under a strictly time-ordered test split

Model	**MAE**	**RMSE**	**nRMSE**	R^2
LSTM (Without Humidity)	0.0109	0.0294	0.0114	0.9916
LSTM (With Humidity)	0.0130	0.0303	0.0117	0.9911
Hybrid (Without Humidity)	0.0121	0.0321	0.0124	0.9900
Hybrid (With Humidity)	0.0103	0.0324	0.0126	0.9898
LR (With Humidity)	0.0458	0.0631	0.0244	0.9614
LR (Without Humidity)	0.0478	0.0640	0.0248	0.9603
RF (With Humidity)	0.0275	0.0999	0.0387	0.9031
RF (Without Humidity)	0.0234	0.1028	0.0398	0.8974
Transformer (With Humidity)	0.1485	0.2802	0.1085	0.2381
Transformer (Without Humidity)	0.1511	0.2803	0.1085	0.2375

4.1.1 RMSE-Only Comparison

For improved readability, comparative performance across models is visualized using RMSE-only bar charts, separated into configurations with and without ambient humidity as shown in Figures 4 and 5, respectively.

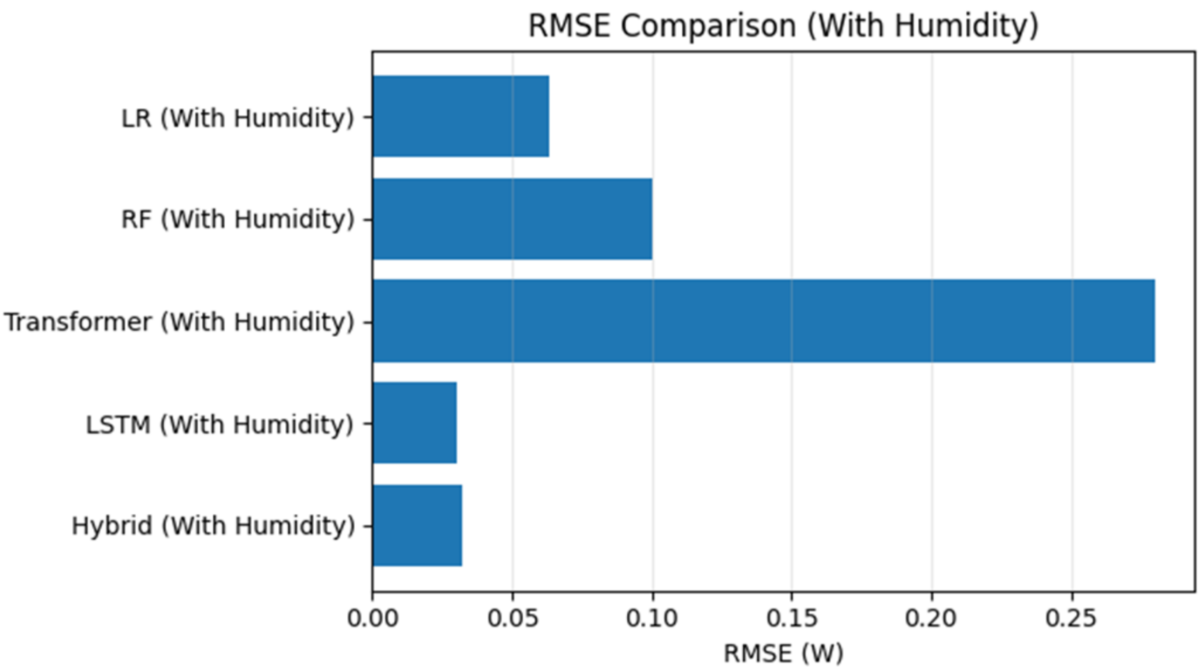

Figure 4: Performance comparison across models and feature ablations (with humidity).

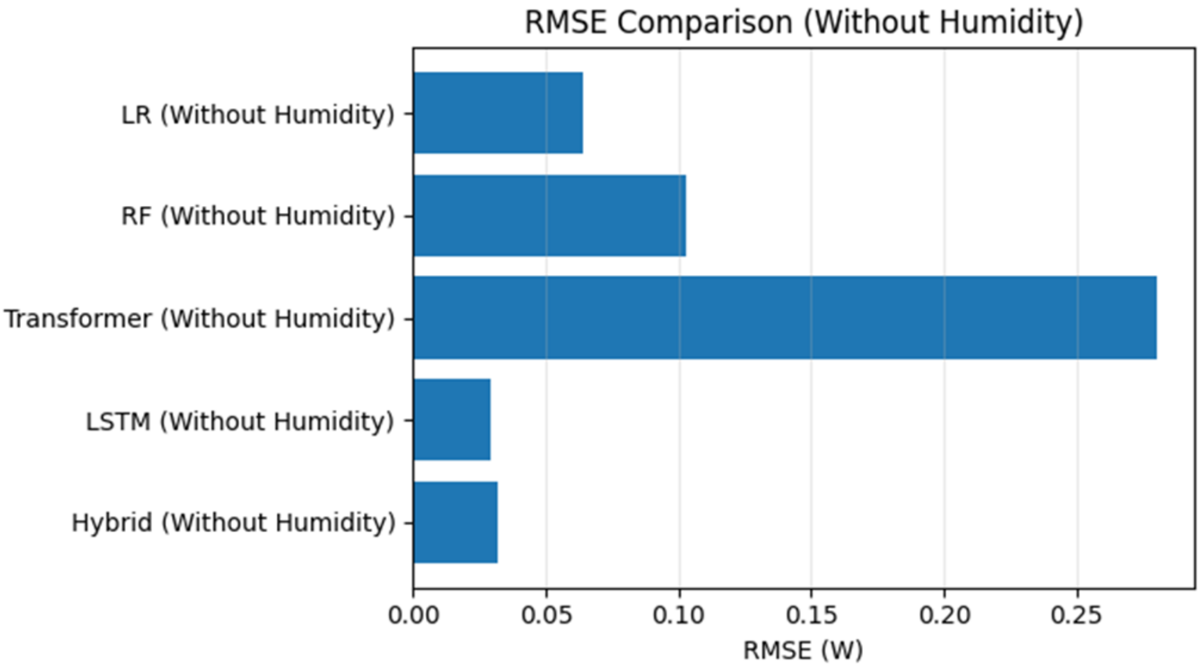

Figure 5: Performance comparison across models and feature ablations (without humidity).

4.2. Feature Ablation: Effect of Ambient Humidity

The impact of ambient humidity as an input feature is evaluated through a controlled ablation study, with performance differences directly observable in Figures 4 and 5. Across all model classes, including humidity results in marginal and inconsistent changes in prediction error. While minor improvements are observed in certain configurations, humidity does not consistently enhance predictive accuracy.

These findings suggest that ambient humidity acts as a contextual auxiliary feature rather than a dominant predictor of PV power output. Its contribution appears to depend on local operating conditions and temporal regimes, underscoring the need for empirical validation before incorporating additional environmental variables into short-term forecasting models.

4.3. Qualitative Prediction Analysis

To complement the quantitative metrics, Figure 6 presents the actual versus predicted PV power for the best-performing model (LSTM without humidity) over a representative segment of the test set. Predictions closely follow measured power, accurately capturing both smooth temporal trends and short-term fluctuations.

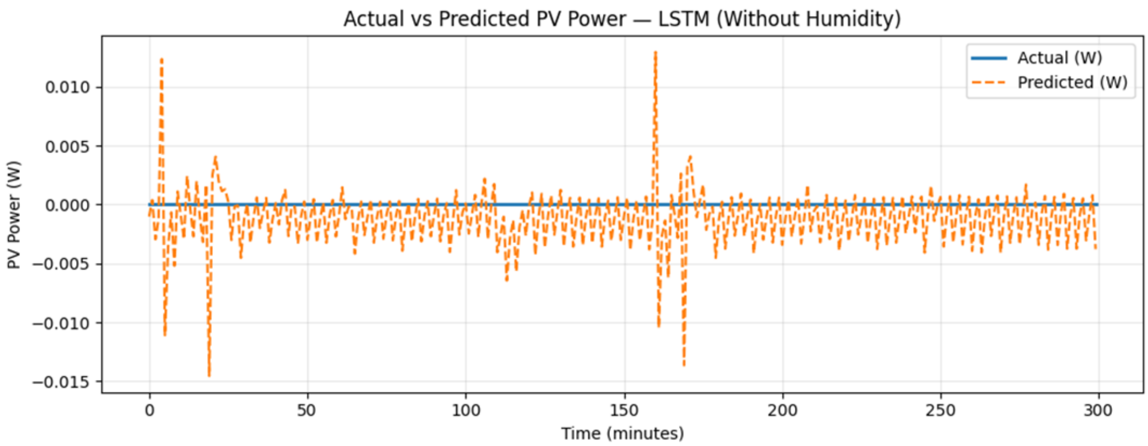

Figure 6: Actual versus predicted PV power using the LSTM model without humidity.

To further assess prediction fidelity, a scatter plot of predicted versus actual values with a 45-degree reference line is shown in Figure 7. The tight clustering of points around the diagonal confirms strong agreement between predictions and ground truth, supporting the high R^2 values reported in Table 2.

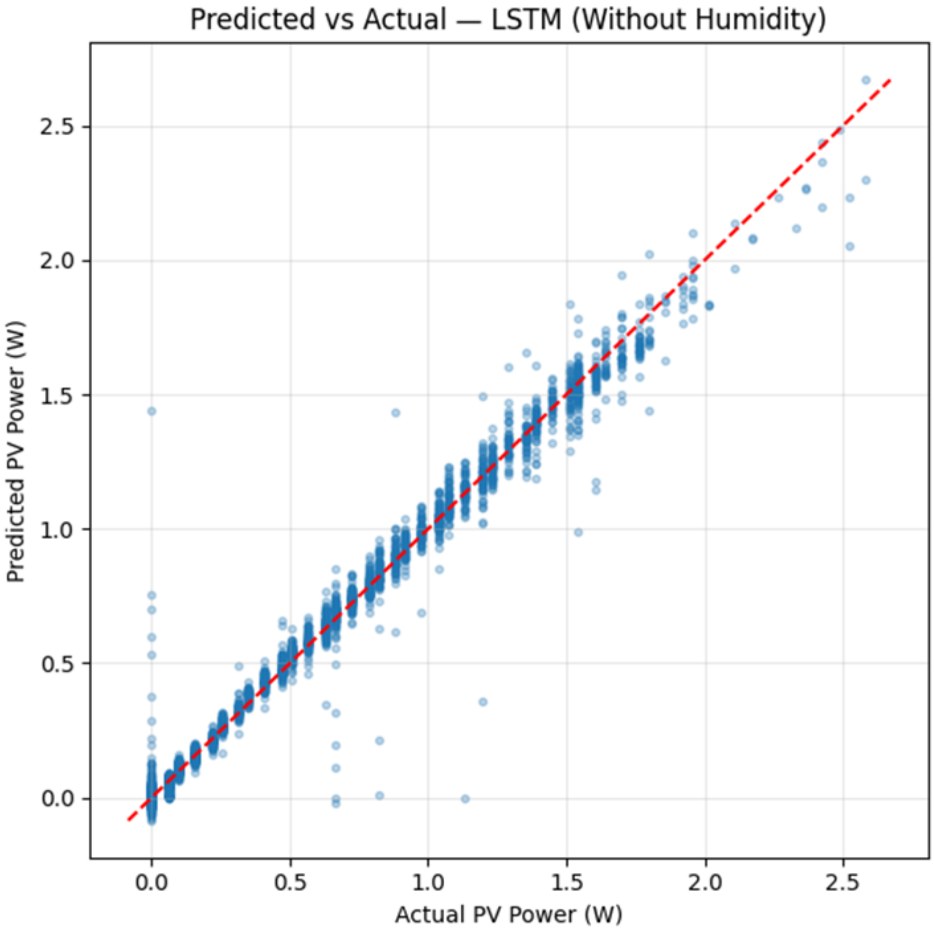

Figure 7: Scatter plot of the actual versus predicted PV power using the LSTM model without humidity.

The residual distribution for the LSTM model, shown in Figure 8, is tightly centered around zero with a standard deviation of approximately 0.03 W, indicating the absence of systematic bias or heavy-tailed error behavior.

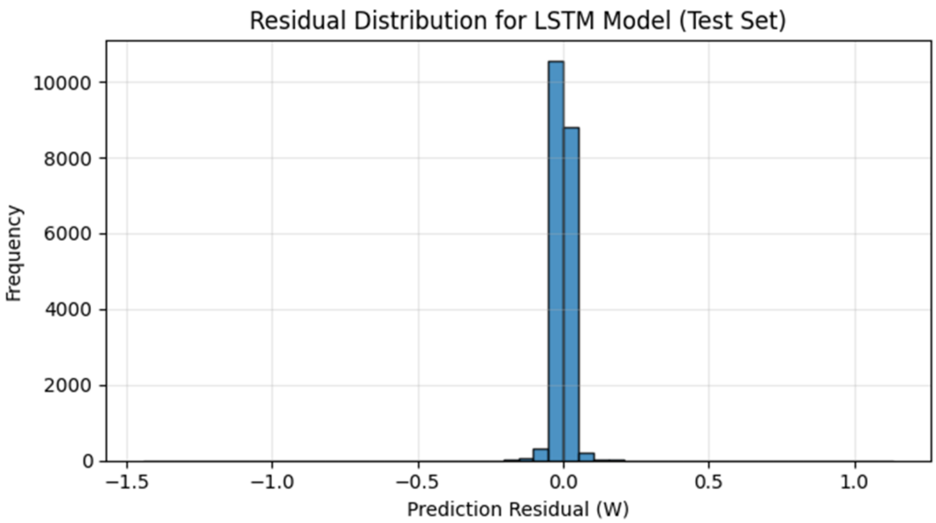

Figure 8: Distribution of prediction residuals for the LSTM model (test set).

4.4. Cross-Panel Generalization

Cross-panel validation was conducted using a leave-one-panel-out strategy to assess spatial generalization. As shown in Figure 9, prediction error varies across panels. PV000 exhibits a substantially higher RMSE than PV001 and PV002.

This disparity is attributed to panel-specific operating conditions rather than model instability. Differences in panel orientation, localized shading, load coupling, or sensor calibration can introduce distinct power dynamics that are not fully captured when training on heterogeneous panels. The observed variability, therefore, reflects realistic inter-panel differences in distributed PV deployments, underscoring the importance of cross-panel evaluation rather than indicating a limitation of the modeling approach.

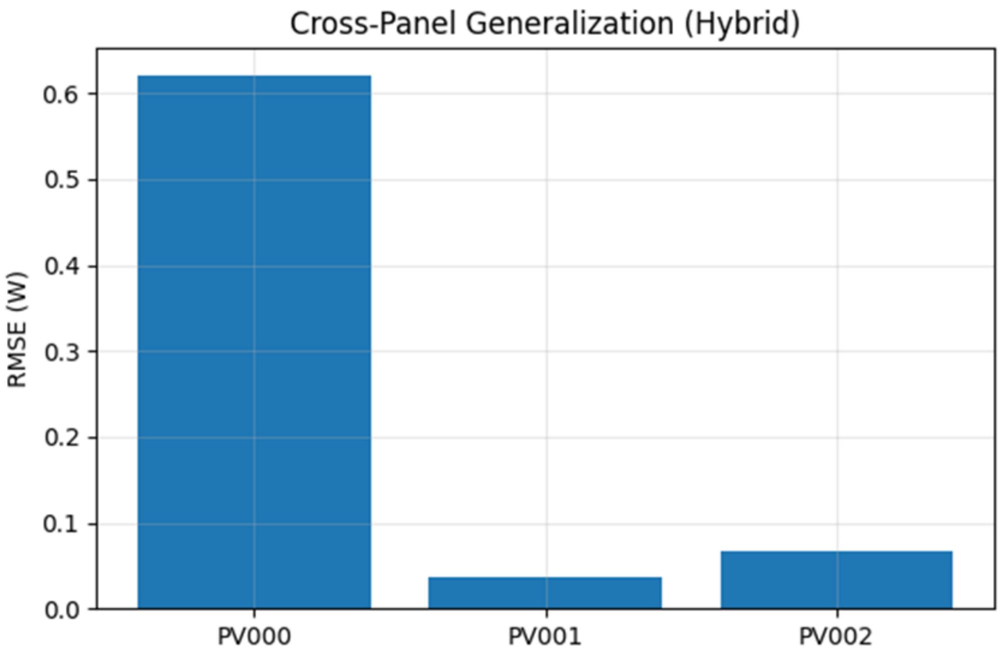

Figure 9: Cross-panel generalization results.

4.5. Rolling Time-Based Validation and Robustness Analysis

To further evaluate temporal generalization, rolling (walk-forward) validation was performed using expanding training windows and unseen validation segments. Rolling RMSE values were consistently higher than those obtained from the single test split, indicating that isolated evaluation windows may yield optimistic estimates for short-duration datasets and confirming that the reported pointwise performance metrics are optimistic but bounded.

In addition, robustness to sensor noise was evaluated by injecting zero-mean Gaussian noise corresponding to approximately 2% of the input feature variance into the test data. Under noisy conditions, prediction error increased moderately (RMSE $\approx$ 0.05–0.06 W) but remained bounded, indicating graceful degradation rather than catastrophic failure. This behavior is consistent with realistic measurement uncertainty in low-cost IoT sensors and supports the suitability of recurrent models for deployment in practical monitoring environments.

5. Discussion

The experimental results in Section 4 provide several important insights into short-term photovoltaic power forecasting using IoT data, particularly regarding feature relevance, model inductive bias, and robustness for deployment.

5.1. When and Why Humidity Helps

Although ambient humidity shows a weak linear association with PV power output, including it as an input feature does not consistently improve predictive accuracy. These findings suggest that humidity functions as a contextual auxiliary variable, whose utility depends on operating regime and interaction with temperature rather than serving as a dominant predictor. Its role is therefore best assessed empirically through ablation rather than assumed a priori.

5.2. Why LSTM Outperforms the Hybrid Model

The consistent superiority of the LSTM-only model indicates that temporal persistence dominates predictive structure in short-duration PV datasets. The additional representational capacity of attention mechanisms is not fully exploited under these conditions, highlighting that increased architectural complexity does not guarantee improved accuracy in data-limited regimes.

5.3. Practical Benefit of the Hybrid Model in IoT Settings

Although not the most accurate model on this dataset, the Hybrid Transformer-LSTM demonstrates advantages in robustness and generalization. Its stable behavior under rolling validation and noise perturbation indicates resilience to temporal drift and sensor uncertainty, key considerations for real-world IoT deployments. In contrast, Random Forest achieves reasonable pointwise accuracy but lacks explicit temporal modeling and robustness to evolving operating regimes.

5.4. Implications for PV Forecasting Research

These findings underscore the limitations of relying solely on single-split performance metrics. High R^2 values may reflect favorable evaluation windows rather than genuine generalization. By combining architectural and feature ablation with temporal and spatial validation, this study provides a balanced evaluation framework that prioritizes reliability and methodological rigor over headline accuracy.

6. Conclusion

This study investigated humidity-aware short-term photovoltaic (PV) power forecasting using real-world Internet of Things (IoT) data and a Hybrid Transformer-LSTM modeling framework. Rather than assuming the benefits of additional environmental features or architectural complexity, the work adopted a rigorous, validation-driven approach to examine the role of ambient humidity and the effectiveness of hybrid attention-recurrent models under data-limited conditions. Through systematic feature and architectural ablation, time-ordered evaluation, rolling validation, cross-panel testing, and robustness analysis, the results demonstrate that rigorous validation and deployment realism are more critical to reliable PV forecasting than marginal gains in accuracy or increased model complexity.

The experimental results demonstrate that LSTM-based models achieve the highest predictive accuracy on the 34-day dataset, highlighting the dominant role of temporal persistence in short-term PV power generation. Ambient humidity was found to provide a marginal, context-dependent benefit, confirming its role as a supplementary environmental feature rather than a universally influential predictor. The Hybrid Transformer-LSTM achieved competitive performance and demonstrated stable behavior under temporal drift, spatial generalization, and sensor noise perturbations, though it did not consistently outperform simpler recurrent architectures. Transformer-only models proved unsuitable for short-duration datasets, whereas classical machine learning baselines lacked the temporal modeling and robustness required for continuous IoT deployment.

Several limitations of the present study should be acknowledged. First, the experimental evaluation is based on data collected from a single deployment site over a relatively short 34-day monitoring period. Although the dataset reflects realistic operating conditions and high temporal resolution, the findings may not fully generalize to PV systems operating in different climatic zones, with seasonal variations, or with different system configurations. Second, the set of input variables is limited to electrical measurements and a small number of environmental parameters. While the ablation analysis clarifies the contextual role of humidity, other influential factors, such as solar irradiance, wind speed, soiling, and shading dynamics, were not available and may further improve model performance and interpretability. Third, although the Hybrid Transformer-LSTM demonstrates robustness advantages, its computational complexity exceeds that of simpler models such as Linear Regression and Random Forest, which may pose challenges for large-scale or resource-constrained edge deployments without further optimization. Finally, the absence of direct irradiance measurements further limits physical interpretability and may reduce predictive performance under rapidly changing sky conditions; however, this constraint reflects realistic IoT deployments and does not affect the validity of the comparative and methodological findings. Future work will address these limitations in several directions. First, multisite and multi-climate studies will be conducted to evaluate the generalizability of humidity-aware modeling across diverse environmental conditions and longer temporal horizons. Second, explainable learning techniques, including attention-based analysis and SHAP-driven feature attribution, will be explored to quantify the dynamic contribution of humidity and other variables under varying operating regimes.

Third, communication-efficient and deployment-aware learning strategies, such as edge-based inference and federated learning, will be investigated to reduce data transmission overhead while maintaining predictive reliability. Finally, future extensions will integrate additional environmental variables and fault indicators to enable joint power forecasting and anomaly detection within a unified IoT-based predictive maintenance framework.

Author Contributions

Ahmed Mohammed: Conceptualization, Methodology, Software, Formal analysis, Writing: original draft; **Ranjit Singh Sarban Singh**: Supervision, Writing: review and editing; **Saad Aslam**: Supervision, Writing: review and editing.

Declaration of Competing Interests

The authors declare that they have no known competing financial interests or personal relationships that could have influenced the work reported in this paper.

Data Availability Statement

The dataset used in this study is available from the corresponding author upon request.

Use of Generative AI

The authors used artificial intelligence tools (ChatGPT) for code debugging and grammar checking during manuscript preparation. All scientific content, analysis, model development, and conclusions were independently designed, executed, and validated by the authors. No AI-generated text was included in the final manuscript without author review and verification.

Acknowledgment

The authors would like to acknowledge Sunway University for providing research grant support and financial assistance to conduct this research.

Funding Declaration

This work was supported under the research grant GRTIN-RAG-DCIS-04-2024, titled “Real-Time IoT-Based Solar Photovoltaic Panel Parameters Analysis and Fault Notification Hardware System Device Development.”

Ethics Approval and Consent

This study does not involve human participants, animals, or sensitive personal data. No ethics approval was required.

References

[1] International Energy Agency, “Renewables 2023: Analysis and forecast to 2028,” tech. rep., IEA Publications, 2022.

[2] S. Mekhilef, R. Saidur, and M. Kamalisarvestani, “Effect of dust, humidity and air velocity on efficiency of photovoltaic cells,” *Renewable and Sustainable Energy Reviews*, vol. 16, no. 5, pp. 2920–2925, 2012.

[3] E. H. M. Ndiaye, A. Ndiaye, M. Faye, D. Gueye, A. Ba, and M. Traore, “Analysis of the impact of irradiance and temperature on photovoltaic production: A statistical and machine learning approach,” *MethodsX*, vol. 15, p. 103716, 2025.

[4] A. H. Arab, B. Taghezouit, K. Abdeladim, and S. Semaoui, “Maximum power output performance modeling of solar photovoltaic modules,” *Energy Reports*, vol. 6, pp. 680–686, 2020.

[5] M. I. Islam, M. S. Jadin, A. A. Mansur, and T. Alharbi, “Electrical performance and degradation analysis of field-aged pv modules in tropical climates: A comparative experimental study,” *Energy Conversion and Management: X*, vol. 24, p. 100719, 2024.

[6] A. Awasthi, A. K. Shukla, M. M. S.R., C. Dondariya, K. Shukla, D. Porwal, and G. Richhariya, “Review on sun tracking technology in solar pv system,” *Energy Reports*, vol. 6, pp. 392–405, 2020.

[7] R. J. Mustafa, M. R. Gomaa, M. Al-Dhaifallah, and H. Rezk, “Environmental impacts on the performance of solar photovoltaic systems,” *Sustainability*, vol. 12, no. 2, p. 608, 2020.

[8] N. B. Sushmi and D. Subbulekshmi, “Real-time ultra short-term irradiance forecasting using a novel r-gru model for optimizing pv controller dynamics,” *Results in Engineering*, vol. 26, p. 105046, 2025.

[9] N. H. Alombah, J. N. Mungwe, A. Harrison, W. F. Mbasso, and H. B. Fotsin, “Advanced iot-based monitoring system for real-time photovoltaic performance evaluation: Conception, development and experimental validation,” *Scientific African*, vol. 28, p. e02763, 2025.

[10] F. Aksan, A. Pawlica, V. Suresh, and P. Janik, “A comparative study of machine learning models for pv energy prediction in an energy community,” *Energies*, vol. 18, no. 22, p. 5980, 2025.

[11] R. Asghar, F. R. Fulginei, M. Quercio, and A. Mahrouch, “Artificial neural networks for photovoltaic power forecasting: A review of five promising models,” *IEEE Access*, vol. 12, pp. 90461–90485, 2024.

[12] K. Ferkous, S. Menakh, M. Guermoui, A. Bellaour, B. Bekkar, A. Rabehi, T. F. Agajie, and M. Benghanem, “Optimized solar power forecasting: A multi-decomposition framework using vmd and swarm techniques,” *AIP Advances*, vol. 15, no. 9, 2025.

[13] M. Tradacete-Ágreda, E. Santiso-Gómez, F. J. Rodríguez-Sánchez, P. J. Hueros-Barrios, J. A. Jiménez-Calvo, and C. Santos-Pérez, “High-performance iot module for real-time control and self-diagnose pv panels under working daylight and dark electroluminescence conditions,” *Internet of Things*, vol. 25, p. 101006, 2023.

[14] R. H. Casanova and A. Conde, “Enhancement of lstm models based on data pre-processing and optimization of bayesian hyperparameters for day-ahead photovoltaic generation prediction,” *Computers & Electrical Engineering*, vol. 116, p. 109162, 2024.

[15] E. Lodhi, X. Liu, G. Xiong, M. A. Khan, Z. Lodhi, T. Nawaz, A. Dilawar, S. Tarkoma, and F. Wang, “Smartpv-aiot: an aiot-integrated framework for fault diagnosis and remote monitoring in photovoltaic systems,” *Energy Conversion and Management: X*, vol. 27, p. 101117, 2025.

[16] M. E. M. Ismail, “The impact of cooling systems on the efficiency of solar panels across different climates: An analytical study based on climate variability,” *SSRN Electronic Journal*, 2025.

[17] H. A. Kazem and M. T. Chaichan, “The effect of dust accumulation and cleaning methods on pv panels’ outcomes based on an experimental study of six locations in northern oman,” *Solar Energy*, vol. 187, pp. 30–38, 2019.

[18] N. Jannah, M. S. A. Hanifah, T. S. Gunawan, S. A. Zabidi, S. H. Yusoff, and S. N. M. Sapihie, “Comparative analysis of mlp and cnn-lstm models for solar power generation forecasting,” in *Proceedings of the IEEE International Conference on Smart Instrumentation, Measurement and Applications (ICSIMA)*, pp. 284–289, 2023.

[19] C. Chauhan, G. K. Saxena, C. A. Kshirsagar, R. K. Solanki, G. Kumar, and S. B. Goyal, “Optimizing predictive maintenance in industrial iot networks using machine learning: A comparative study of svm, dt and ann,” *Journal of Computers, Mechanical and Management*, vol. 4, no. 5, pp. 8–16, 2025.