Deep Learning Approach for Solar Radiation Forecasting in a Tropical Region Using LSTM Networks

Manuel Ospina[1], Gabriel Chanchí[2], Álvaro Realpe[3]
Faculty of Engineering-Systems Engineering Program, University of Cartagena, Cartagena, Colombia[1, 2]
Faculty of Engineering-Chemical Engineering Program, University of Cartagena, Cartagena, Colombia[3]

Abstract*—Solar radiation forecasting is a key task for energy planning, grid management, and photovoltaic deployment, especially in tropical regions where weather variability reduces operational reliability. This work applies deep learning techniques to forecast hourly solar radiation in Mompox, Colombia, using Long Short-Term Memory (LSTM) neural networks. Three temporal windows were studied (5, 24, and 720 hours) to examine how sequence length affects prediction accuracy and model behavior. Hourly radiation data from 2021 to 2022 were used for training, and independent datasets from 2023 to 2024 were used for external validation to ensure long-term assessment and reproducibility. Most existing studies use short input windows designed for mid-latitude environments (5–24 hours), which do not capture multi-day tropical cloud persistence or sub-seasonal radiation variability. This gap limits forecasting accuracy and restricts practical use in tropical energy planning. To address this issue, this study introduces a long temporal input design that allows the model to learn month-scale variability more effectively. The three network configurations were trained under the same settings, allowing a direct comparison between short, daily, and long input memories. The LSTM-720 model performed best, achieving the lowest RMSE and the most stable predictions across all validation years, showing its ability to reconstruct both diurnal cycles and broader seasonal dynamics. Unlike most solar forecasting work, which treats window size as a tuning parameter, this study introduces a long-context LSTM design based on a 720-hour sequence. This allowed the model to learn intra-month atmospheric persistence—an essential tropical feature that short windows cannot represent—positioning the approach as a methodological contribution that expands the temporal learning paradigm rather than a configuration adjustment. Time-series comparisons revealed close agreement between measured and predicted radiation, particularly during stable climate periods. The proposed framework can support practical applications in solar plant design, renewable energy scheduling, and operational grid strategies in tropical regions. Future work will integrate satellite information and hybrid deep learning architectures to enhance spatial transferability and long-term forecasting accuracy.

***Keywords*—Deep learning; LSTM networks; renewable energy; solar radiation forecasting; time series prediction**

This study is considered a product within the framework of the project: “Desarrollo Sostenible Mediante la Implementación de Energía Solar Fotovoltaica en Zonas Rurales no Conectadas a la Red Eléctrica Nacional en Mompox, Colombia (022-2024)”, which is part of an internal research initiative at the University of Cartagena, Cartagena, Colombia.

I. Introduction

Given the growing constraints on fossil fuel generation [1], integrating clean energy systems (such as solar) with conventional power generation [2], [3] has become a key strategy to enhance efficiency and reduce emissions [4], [5]. At the solar energy level, solar radiation is a fundamental variable for energy management, particularly in the planning and design of photovoltaic (PV) systems [6], [7], agricultural applications [8], [9], and environmental modelling [10]–[12].

The use of solar energy significantly reduces greenhouse gas emissions [13], as its generation produces no CO_2 or other pollutants [14], thereby lowering dependence on non-renewable resources and promoting ecosystem conservation as a clean and sustainable electricity source [15]–[17]. In tropical regions, accurate forecasting is even more crucial for the effective planning, construction, and operation of PV power plants because sunlight availability changes rapidly with cloud cover and seasonal variation [18], [19].

The town of Mompox, located in northern Colombia, offers an interesting setting for solar radiation prediction. It combines a humid tropical climate with high solar potential, but has limited meteorological monitoring infrastructure. Although Colombia’s Caribbean region receives abundant solar energy, few predictive models have been developed specifically for this area. This lack of localized modeling limits the efficient deployment of solar technologies and constrains opportunities for grid optimization and energy independence.

Traditional statistical approaches, such as ARIMA and Support Vector Machines (SVM), often fall short when modeling the nonlinear and time-dependent behavior typical of solar radiation [20], [21]. Recent progress in artificial intelligence has changed this landscape [22]–[25]. Recurrent Neural Networks (RNNs) [26], and particularly Long Short-Term Memory (LSTM) architectures [27], can capture long-term dependencies and nonlinear trends in meteorological time series [4], [28]. Several studies have demonstrated their effectiveness in predicting solar irradiance across different regions [29]–[31], however, few have explored their use in small tropical towns like Mompox, where climatic variability and data scarcity pose additional challenges.

In this context, the research evaluates the ability of LSTM models to forecast solar radiation at multiple temporal resolutions (5, 24, and 720 hours) using real measurements from

Mompox. The model was trained with data from 2021 to 2022 and validated with independent datasets from 2023 to 2024, ensuring reproducibility and robustness. This approach can be adapted to other data-limited regions seeking to improve renewable energy forecasting. Predicting solar radiation in tropical climates remains a challenging task. In these environments, cloud systems can evolve within minutes, humidity levels fluctuate abruptly, and short-lived storms frequently interrupt solar exposure [32]. Such variability introduces nonlinear and nonstationary behavior that traditional statistical models struggle to capture effectively [33].

Deep learning methods, particularly Long Short-Term Memory (LSTM) networks, offer a promising alternative because they can learn how these rapid changes in cloud cover and humidity interact over time. By retaining information from previous states, LSTMs can represent the temporal continuity and sudden transitions that characterize tropical weather, resulting in more accurate and resilient radiation forecasts [34]. Yet, few studies have examined their effectiveness in humid equatorial zones like Mompox, where diurnal and seasonal cycles interact in distinctive ways.

Solar radiation forecasting in tropical regions is strategically important for national energy planning, particularly in areas where rural electrification, PV expansion, and microgrid deployment depend on accurate generation forecasting. Improving prediction stability in data-limited locations, such as Mompox, supports lower operational uncertainty, reduces reserve margin dependence, enables more reliable agricultural scheduling, and improves climate adaptation strategies. This context shows that the research is not only a methodological exercise, but also responds to a regional need for energy autonomy and system planning, making the development of long-memory forecasting approaches an urgent technical requirement for tropical countries.

Tropical solar radiation dynamics differ fundamentally from conditions addressed in most LSTM-based forecasting studies. Prior works overwhelmingly employ short input memories ranging from several hours to one day, an assumption that aligns with mid-latitude irradiance behavior but not with tropical environments. In the tropics, radiation variability is governed by persistent multi-day cloud systems, moisture accumulation cycles, and sub-seasonal atmospheric structures that extend far beyond diurnal periodicity. These characteristics produce long temporal dependencies that short window LSTMs cannot capture, resulting in phase shift errors, amplitude smoothing, and loss of climatic context. This gap in the literature motivated the exploration of long-memory recurrent models. By incorporating a 720-hour input window, equivalent to a full intra-month sequence, we aimed to enable the network to learn climatic persistence patterns rather than only daily oscillations, addressing a structural limitation of existing approaches and positioning this work beyond parameter tuning studies.

The novelty of this study lies in its methodological approach rather than parameter scaling. While previous LSTM works overwhelmingly rely on short input windows (5–24 hours), we demonstrate that tropical radiation dynamics demand long-context modelling to reconstruct sub-seasonal variability. By implementing a 720-hour sequence, the model learns persistent atmospheric structures that cannot be captured by traditional diurnal frameworks, establishing long-memory LSTM forecasting as a required paradigm for tropical environments rather than a configuration variant.

The main contributions of this research are summarized as follows: 1) a long-context LSTM architecture was introduced that uses a 720-hour input sequence tailored to represent seasonal tropical radiation dynamics, a configuration not previously explored in related forecasting literature. 2) The model's robustness was evaluated through a multi-year external validation over unseen 2023 and 2024 datasets, demonstrating long-term predictive stability rarely addressed in prior deep learning work for tropical environments. 3) It was shown that long-context learning improves forecasting accuracy over traditional short-window and daily configurations, providing evidence that tropical radiation behavior requires long-memory modelling rather than parameter scaling, and 4) a reproducible and scalable modelling framework was constructed based on raw ground solar radiation measurements, which can be transferred to other tropical locations with limited climatological infrastructure.

To address this gap, this study proposes an LSTM-based approach optimized for tropical solar radiation forecasting using high-resolution ground data. It systematically analyzes how different window sizes (from 5 to 720 hours) influence prediction accuracy and identifies the most suitable configuration for practical energy management in tropical environments.

II. Related Works

Forecasting solar irradiance has evolved rapidly in recent years with the introduction of deep learning models, especially Long Short-Term Memory (LSTM) networks and their hybrid variants. Several studies have shown that LSTM architectures consistently outperform classical methods such as ARIMA, Artificial Neural Networks (ANN), and Support Vector Machines (SVM) when handling the nonlinear dynamics of solar radiation. For instance, [35] employed an LSTM network for short-term irradiance forecasting and found that combining multiple input sequences helped the model generalize better while reducing both variance and bias. In a similar vein, [23] compared ARIMA, feed-forward neural networks, and LSTM using Colombian datasets and reported that the LSTM model delivered the highest accuracy under cloudy conditions, one of the most challenging scenarios for solar prediction.

Hybrid models have also gained attention for their enhanced representational power. A review by [36] noted that CNN–LSTM architectures, which merge spatial and temporal learning, consistently outperform single-model approaches in irradiance forecasting, albeit with higher computational requirements. Recent developments even include BiLSTM–Transformer hybrids that achieved the lowest RMSE in seven-day forecast experiments, outperforming both standalone LSTM and GRU networks [24]. Comprehensive literature reviews further confirm that deep learning techniques, including LSTM, GRU, CNN–LSTM, and attention-based networks, yield significant accuracy improvements over traditional approaches, though at the expense of longer training times [37]. For example, [19] compared LightGBM, LSTM, and GRU using hourly radiation data. While LSTM achieved an RMSE of about 59 W/m², LightGBM provided slightly better accuracy (≈54.8 W/m²) and faster computation, suggesting that tree-based ensembles can still compete when computational resources are constrained.

The field has also begun shifting toward probabilistic and multivariate forecasting. Hybrid frameworks that integrate BiLSTM and GRU layers, combined with Bayesian hyperparameter optimization and dropout regularization, have achieved strong performance in multivariate irradiance prediction [25]. Additionally, several studies that merge global climate models with ground-based measurements through hybrid deep learning and machine learning pipelines have produced reliable daily radiation forecasts [22]. Despite these advancements, certain research gaps remain clear. Few studies have focused on tropical towns such as Mompox, where meteorological data are scarce, and climate variability is high. Equally uncommon are works that compare multiple forecasting horizons (short-term, daily, and monthly) within a unified framework, or that evaluate pure LSTM configurations at extended horizons of up to 720 hours. The present study directly addresses these gaps by developing an LSTM model tailored to solar radiation forecasting in Mompox, Colombia. Using datasets from 2021 to 2022 for training and from 2023 to 2024 for validation, this work provides a comprehensive, horizon-sensitive modeling framework specifically designed for tropical regions with limited historical observations.

III. METHODOLOGY

This section presents the methodological framework designed to forecast solar radiation in Mompox using Long Short-Term Memory (LSTM) neural networks. It describes the dataset collected from IDEAM [38], the preprocessing steps applied, the configuration of the neural architectures, and the training and validation strategy adopted. Together, these components ensure transparency, reproducibility, and methodological rigor for the subsequent analyses.

Hourly solar radiation data were collected for Mompox (9°14'N, 74°26'W; elevation 10 m a.s.l.), a town located in the tropical Magdalena River basin of northern Colombia. The dataset covers the period from January 2021 to December 2022 and corresponds to global horizontal irradiance (Wh/m²). The information was provided in CSV format by the Instituto de Hidrología, Meteorología y Estudios Ambientales (IDEAM), Colombia's national meteorological authority [38].

The original file also included ancillary variables such as air temperature and wind speed. However, these were excluded to focus the analysis exclusively on solar radiation prediction. Additional datasets spanning January–December 2023 and January–July 2024 were used solely for external validation. In total, the dataset contains about 17,500 hourly records. Missing values, representing less than 1 % of the data, were interpolated linearly to maintain the temporal continuity of the series. Extreme outliers, defined as values exceeding ±3 standard deviations from the mean, were replaced using a rolling median filter. A summary of the dataset, including temporal coverage and partitioning strategy, is shown in Table I.

The time series of global horizontal irradiance (Wh/m²) was used directly, without normalization or scaling, in order to preserve the physical meaning of the radiation values. Input–output pairs were created using a sliding window method, in which a sequence of past radiation observations was used to predict the next value. Three different input window sizes were evaluated to capture varying temporal behaviors: 5 hours: short-term fluctuations, 24 hours: daily cycles, and 720 hours: monthly or intra-seasonal patterns. For each configuration, the data were chronologically divided into training (67 %), testing (22 %), and validation (11 %) subsets (see Table I). This temporal split preserved the natural sequence of the observations and avoided data leakage between model phases.

TABLE I. SUMMARY OF THE DATASET

Dataset	Data Features		
	Period	***Frequency***	***Samples***
Training	2021-2022	Hourly	11,910
Testing	2021-2022	Hourly	3,853
Validation	2021-2022	Hourly	1,752
External validation	2023	Hourly	8,760
External validation	Jan-Jul 2024	Hourly	5,040

All forecasting models were implemented in Python 3.10 using TensorFlow 2.15 with the Keras API. Three LSTM architectures were designed, each corresponding to one of the input window sizes.

- LSTM-5: Input (5 × 1), one LSTM layer with 64 units, followed by a dense layer with 8 neurons (ReLU activation) and a linear output layer.
- LSTM-24: Input (24 × 1), one LSTM layer with 64 units, followed by a dense layer with 8 neurons (ReLU activation) and a linear output layer.
- LSTM-720: Input (720 × 1), one LSTM layer with 720 units, two dense layers with 16 and 8 neurons (both ReLU activation), and a final linear output layer.

The increase in the number of LSTM units for the 720-hour configuration was intentional and proportional to the longer input sequence length. A higher number of recurrent units provides the model with greater representational capacity to retain and process extended temporal dependencies across a full monthly cycle. Preliminary tests indicated that using fewer than 720 units led to underfitting, particularly in reproducing smooth transitions between consecutive days, while the chosen configuration achieved lower RMSE and faster convergence without overfitting. Thus, the increase in hidden units was empirically justified as it improved the model's ability to capture long-term patterns inherent to tropical solar radiation.

The choice of the 720-hour input configuration was not defined as a tuning exercise but as a methodological decision derived from climatic characteristics of the study region. The thirty-day temporal span was selected to encode full intra-month variability related to cloud persistence, humidity accumulation, and low-frequency atmospheric patterns typical of tropical environments. Unlike short window configurations (5, 24 hours), the 720 hour structure enables the model to learn sub-seasonal energy dynamics and long term irradiance memory, which are not recoverable through repeated diurnal cycles. This design transforms the LSTM into a long context recurrent learner, allowing it to internalize radiation persistence behavior and avoid phase misalignment effects commonly reported in tropical solar forecasting. Therefore, the 720-hour model represents a conceptual modelling framework rooted in regional climate physics rather than a parametric extension of existing approaches.

All networks used the Adam optimizer (learning rate = 0.001) and the Mean Absolute Error (MAE) as the loss function. Model performance was tracked using the Root Mean Squared Error (RMSE) metric. Each model was trained for 30 epochs, with a ModelCheckpoint callback to save the best-performing weights according to validation loss, thus minimizing the risk of overfitting.

Fig. 1 illustrates the general LSTM configuration applied for solar radiation forecasting, showing the sequential structure from the input window to the output layer.

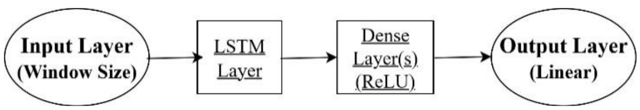

Fig. 1. General LSTM architecture for solar radiation forecasting.

Model training was performed in a Google Colab Pro environment equipped with an NVIDIA Tesla T4 GPU. The same optimizer, learning rate, and loss function were maintained for consistency across all experiments. Each training session lasted 30 epochs, and the ModelCheckpoint callback was again used to retain the best weights based on validation loss.

Data splitting followed a chronological order to preserve temporal dependencies: 67 % for training, 22 % for testing, and 11 % for validation, corresponding to 11,910, 3,853, and 1,752 hourly samples, respectively. Performance was evaluated using RMSE, MAE, and coefficient of determination (R2) for each subset. Additionally, the trained networks were tested on the external datasets from 2023 and 2024 to assess their ability to generalize to unseen data. Training progress was monitored through loss and validation-loss curves to detect possible underfitting or overfitting behaviors. The main parameters used during training are summarized in Table II.

TABLE II. TRAINING PARAMETERS

Parameter	Description
Framework / API	TensorFlow 2.15 / Keras
Programming language	Python 3.10
Environment	Google Colab (GPU: NVIDIA Tesla T4)
Optimizer	Adam
Learning rate	0.001
Loss function	Mean absolute Error (MAE)
Evaluation metric	Root Mean Squared Error (RMSE)
Epochs	30
Batch size	Default (TensorFlow)
Validation strategy	Chronological split (67% / 22% / 11%)

All procedures in this section were designed to ensure reproducibility and consistency across model configurations. Combining historical hourly radiation data with structured input windows and systematic model training established a solid foundation for evaluating forecast performance. This methodological design can be easily extended in future research by incorporating additional meteorological variables or developing hybrid architectures that combine LSTM with other deep learning models.

IV. RESULTS

This section presents and interprets the results obtained from the three LSTM configurations developed for solar radiation forecasting in Mompox, Colombia. The analysis integrates both quantitative and visual assessments of model performance, highlighting how different temporal window sizes influence prediction accuracy and generalization capacity. The following subsections examine, in detail, how each model performed in reproducing observed solar radiation patterns, comparing predicted and measured values across multiple time horizons.

A. Model Comparison by Input Window Size

To examine how the temporal context affects forecasting accuracy, three LSTM configurations were tested using input window sizes of 5, 24, and 720 hours. All models were trained under the same conditions (30 epochs, Adam optimizer, learning rate = 0.001) using hourly solar radiation data. The goal was to identify the sequence length that best balances short-term fluctuations with longer-term seasonal patterns.

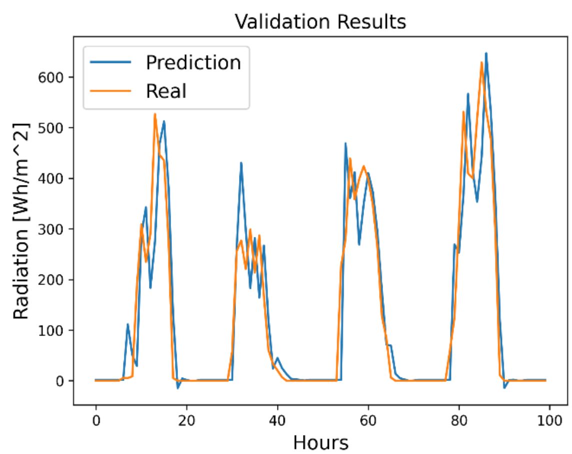

Fig. 2. Predicted vs. observed radiation using LSTM-5 (*window = 5 hours*).

The LSTM-5 model reproduced rapid fluctuations in solar radiation, particularly on clear days, but it struggled to represent complete diurnal transitions. As shown in Fig. 2, the predicted values follow the general pattern of observed radiation but show amplitude mismatches during peak irradiance hours. According to Fig. 2, this limitation is mainly due to the network's short memory window, which restricts its ability to capture the full daily energy cycle.

When the input window was expanded to 24 hours (Fig. 3), the model captured the daily periodicity of solar radiation. Fig. 3 shows that the predicted curves align more closely with observed diurnal patterns, reducing phase shifts and errors at sunrise and sunset. However, the model still tended to underestimate radiation peaks on cloudy days, suggesting that a one-day context is insufficient to describe multi-day atmospheric variability.

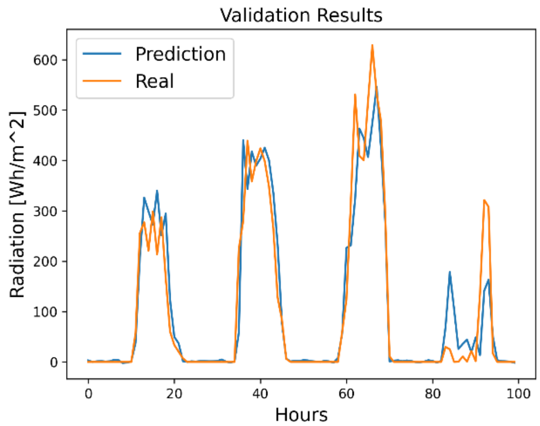

Fig. 3. Predicted vs. observed radiation using LSTM-24 (*window = 24 hours*).

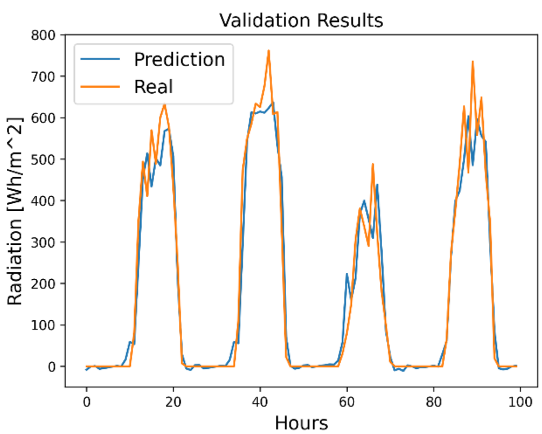

Fig. 4. Predicted vs. observed radiation using LSTM-720 (window = 720 hours).

The LSTM-720 model, trained on a monthly (30-day) input window, produced the most accurate and stable forecasts (Fig. 4). By learning over a broader temporal context, the network captured both daily and intra-monthly variations, resulting in smoother transitions and reduced noise. This improvement confirms that long-range dependencies are critical for reliable radiation prediction in tropical regions, where multi-day weather persistence is common.

In accordance with the above, extending the temporal window from 5 to 720 hours significantly enhanced model performance. Although computational cost increased with larger input sequences, the gains in accuracy justified using the LSTM-720 configuration for subsequent analyses. This setup successfully captures the natural temporal structure of solar radiation and provides the foundation for the evaluations presented in Section B and C.

B. Model Training Behavior

During training, the LSTM-720 model showed fast and stable convergence. Both training and validation losses decreased sharply within the first ten epochs and then plateaued, as illustrated in Fig. 5.

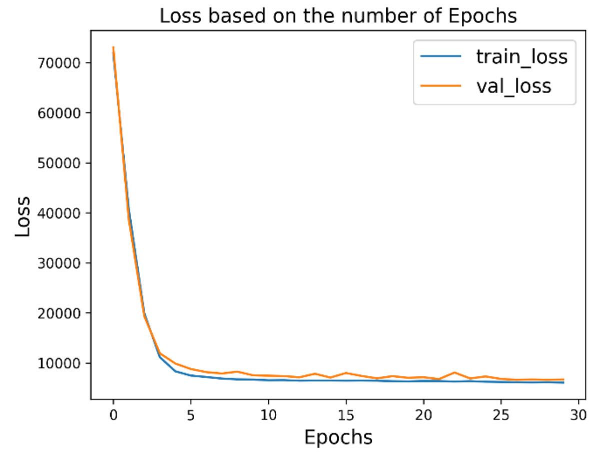

Fig. 5. Training and validation loss curves of LSTM-720 model.

The close alignment of the two curves indicates good generalization and minimal overfitting, confirming that the model effectively learned the underlying temporal dependencies. The smooth decline in loss values reflects the appropriate choice of hyperparameters, particularly the moderate learning rate (0.001) and batch size (32), which ensured stable gradient descent and reproducible convergence. Overall, these results demonstrate that the selected architecture and optimization settings are well suited for modeling the complexity of tropical solar radiation time series.

C. Validation and External Testing Performance

The predictive ability of the final LSTM-720 model was assessed in three validation stages: 1) internal validation using data from 2021 to 2022, 2) external validation with unseen data from 2023, and 3) an additional external validation using 2024 data. Each stage tested how well the model generalized under distinct climatic conditions.

Fig. 6 compares the observed and predicted hourly radiation during the internal validation period. The model accurately reproduced both clear and cloudy day dynamics, closely matching the amplitude and timing of daily cycles.

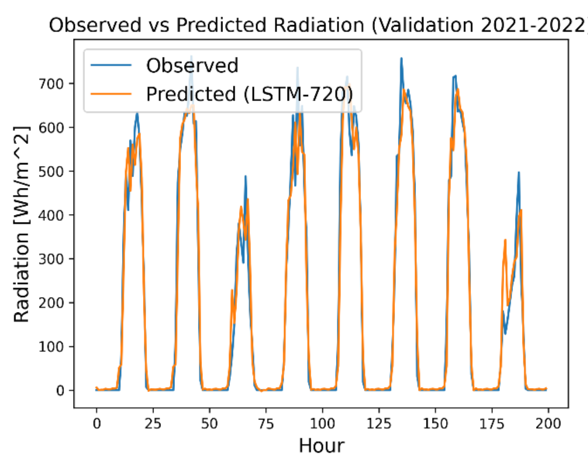

Fig. 6. Internal validation for the years 2021 to 2022 of LSTM-720 model.

Minor underestimations were observed at midday peaks, a common feature of networks trained without data normalization, but overall performance remained excellent, with RMSE = 63.88 Wh/m² and R^2 = 0.9535, explaining over 95 % of the observed variance. These metrics demonstrate strong calibration and confirm that the model learned the key temporal structures of solar radiation.

The first external test, shown in Fig. 7, evaluates the model on independent data from 2023. Predictions closely tracked the observed series throughout the year, with small amplitude differences during peak irradiance. Despite slightly higher RMSE (71.93 Wh/m²) and a marginally lower R^2 (0.913), the model maintained high accuracy, demonstrating robust generalization across years. This performance decay is modest and expected when forecasting beyond the training period. It also confirms that the temporal patterns learned from 2021 to 2022 remain valid under new atmospheric conditions.

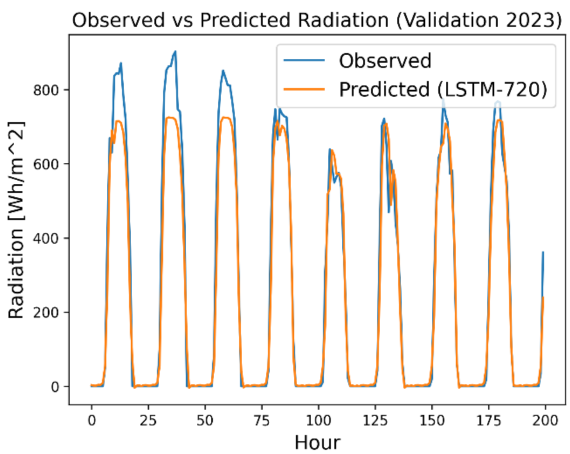

Fig. 7. External validation for the year 2023 of LSTM-720 model.

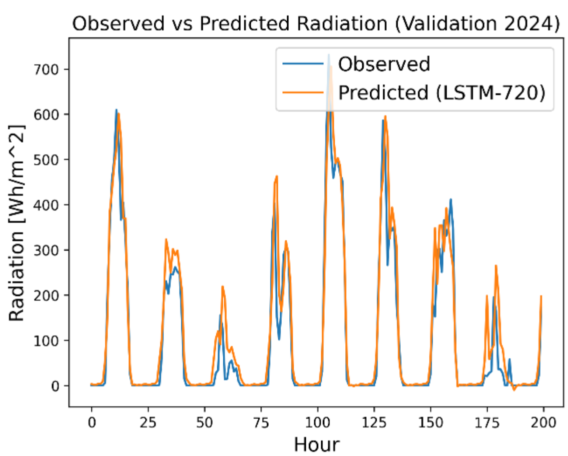

Fig. 8. External validation for the year 2024 of LSTM-720 model.

Fig. 8 presents the results for the 2024 dataset. The predicted time series aligned closely with the observed data, accurately reproducing diurnal cycles and the effects of intermittent cloud cover. Model metrics remained strong (RMSE = 65.53 Wh/m², MAE = 33.97 Wh/m², R^2 = 0.902), confirming stable performance and adaptability across consecutive years. The slightly narrower error range compared with 2023 suggests that weather patterns in 2024 were more similar to those seen during training.

A summary of all validation results is shown in Table III. R^2 values consistently above 0.90 confirm the model's ability to generalize well across different years and atmospheric conditions. The lowest RMSE corresponds to the internal validation, while external datasets remain within ±8 Wh/m² of that baseline, evidence of minimal performance decay. The MAE values (33–36 Wh/m²) indicate that average hourly errors are small relative to typical radiation levels in Mompox (> 600 Wh/m²). Altogether, these results emphasize the temporal stability and predictive robustness of the LSTM-720 architecture for operational solar radiation forecasting.

TABLE III. PERFORMANCE METRICS OF THE LSTM-720 MODEL ACROSS DATASETS

Dataset	Metrics		
	RMSE (Wh/m²)	*MAE (Wh/m²)*	*R^2*
2021–2022	63.88	35.66	0.954
2023	71.93	36.43	0.913
2024	65.53	33.97	0.902

To further evaluate model robustness, residual distributions were analyzed across the three validation periods (Fig. 9). Residuals were calculated as the difference between observed and predicted radiation, providing direct insight into systematic bias and variability. Across all years, residuals remained centered around zero with narrow interquartile ranges, indicating consistent performance and no major bias.

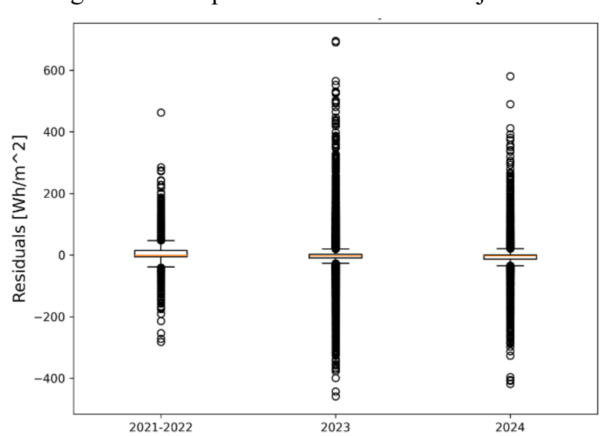

Fig. 9. Residual error distributions by year.

The statistical indicators summarized in Table IV show small negative mean values, suggesting a slight underestimation of peak irradiance.

Standard deviations stayed below 70 Wh/m², matching the RMSE magnitudes reported earlier. Skewness values near zero confirm symmetric error distributions, a desirable trait in predictive modeling. Slightly broader residual spreads in 2023 and 2024 correspond to stronger local variability, likely due to convective cloud formation and intermittent tropical rainfall. Even so, the overall bias remains low, confirming that the LSTM-720 model captures dominant radiation dynamics across changing weather conditions. Future research could enhance accuracy further by incorporating auxiliary variables such as temperature, humidity, or cloud fraction, or by employing ensemble architectures to better capture transient atmospheric behaviors.

TABLE IV. STATISTICAL ANALYSIS OF MODEL RESIDUALS FOR THE VALIDATION DATASETS

Dataset	Metrics		
	Mean Residual (Wh/m^2)	*Std. Deviation (Wh/m^2)*	*Skewness*
2021–2022	-2.47	59.82	-0.06
2023	-3.85	67.49	0.12
2024	-1.91	63.07	-0.09

V. DISCUSSION AND IMPLICATIONS

The discussion of the broader implications of these findings in the context of previous research focused on three main aspects: 1) the effect of input window size on temporal learning, 2) the model's stability when applied to unseen years (2023–2024), and 3) the statistical characteristics of forecast residuals. Together, these analyses provide a comprehensive understanding of how LSTM networks can effectively model solar radiation dynamics under tropical climatic conditions and characteristics.

The findings confirm that deep recurrent neural networks, when trained on sufficiently long sequences, can model the nonlinear and seasonal variability of tropical solar radiation. Compared with conventional methods such as ARIMA or SVM [18], [20], [26], [28], LSTM models provide superior capacity to represent temporal dependencies without extensive feature engineering. This agrees with previous work on deep learning for solar forecasting [31], [35]–[37], further validating the effectiveness of recurrent architectures for meteorological time-series prediction.

Beyond outperforming traditional statistical baselines, the proposed 720-hour LSTM model offers a structural advantage over existing short-window and hybrid machine-learning approaches. Studies using ARIMA, LightGBM, CNN–RNN and BiLSTM–GRU configurations have shown accuracy improvements under standard mid-latitude conditions [19], [20], [24], [25], yet these models rely on short temporal dependencies and lose stability when climatic variability increases. In contrast, the long-context design used in this work was able to reconstruct multi-day and intra-month persistence patterns that those architectures cannot represent, demonstrating that forecasting improvement resulted from modelling strategy rather than parameter tuning. This positions the LSTM-720 framework as a distinct methodological alternative capable of recovering tropical atmospheric structure, rather than as another competing model variant. As such, the findings highlight a clear modelling advantage: long-memory recurrent networks enable tropical radiation learning behavior that existing methods structurally fail to reproduce.

Away from accuracy improvements, these findings translate into general modelling insights that apply to broader forecasting tasks. The results show that tropical radiation prediction requires architectures capable of learning sub-seasonal patterns rather than isolated diurnal signals, establishing a transferable design principle: input horizons must extend into the multi-day domain to reconstruct persistence dynamics. This principle aligns with recent studies reporting that deep recurrent models outperform statistical and short-window baselines when climatic variability increases [22], [23], [24], [25], [31]. In particular, works comparing ARIMA, LightGBM, hybrid CNN–RNN and BiLSTM–GRU approaches show that when atmospheric conditions are highly unstable, conventional models lose capacity to retain temporal structure [19], [20], [24], [25]. These links support the interpretation that forecasting improvements observed here are not dataset-dependent but reflect an underlying modelling requirement for tropical climates. Thus, the contribution extends beyond a single architecture: it provides guidance for window selection, long-context design, and application to other subtropical and equatorial solar forecasting problems.

A key strength of the LSTM-720 configuration is its durability over time. Beyond outperforming shorter input configurations, the results indicate that the 720 hour model delivers a fundamentally different learning behavior. Rather than improving accuracy through parameter scaling, the long-window architecture captured low frequency variability, radiation persistence, and intra month climatic structures that short-memory networks failed to represent. This shift reveals an underlying methodological implication: tropical radiation forecasting requires long context recurrent models to reconstruct non periodic atmospheric processes, not merely diurnal oscillations. The model's stability across 2021–2024 external validation years further demonstrates that this behavior is structural rather than dataset dependent. Consequently, this work contributes a conceptual advance to the field by demonstrating that long memory design is a necessary modelling component for tropical solar forecasting, positioning the approach as a generalizable framework rather than a sensitivity study.

In addition to demonstrating performance improvements, the findings from this study offer practical modelling guidance. The results show that tropical irradiance prediction benefits structurally from extended temporal contexts, establishing a clear design principle: forecasting windows must exceed diurnal or daily horizons to reconstruct persistence dynamics. Short-memory configurations captured rapid fluctuations but systematically failed to reproduce multi-day energy cycles, confirming that traditional daily window assumptions are unsuitable for humid tropical climates. Moreover, the stability of the 720-hour model across independent yearly datasets suggests that long-context LSTMs can retain climatic structure without retraining, providing a blueprint for robust tropical forecasting architectures. These elements sharpen the contribution of this work beyond experimentation results and translate into modelling guidelines applicable to future solar forecasting studies in similar environments.

Further to these findings, the results open several structured research directions. First, incorporating meteorological predictors such as cloud fraction, humidity or temperature may reveal whether tropical persistence originates from atmospheric drivers rather than radiation memory alone, helping to explain why monthly windows outperform daily horizons. Second, evaluating GRU, BiLSTM–GRU and Transformer-based models could determine whether attention mechanisms improve multi-day pattern extraction beyond recurrent gating, clarifying if the benefit of long-context learning is architectural or temporal. Third, transferring and retraining the 720-hour framework in other tropical and subtropical climates would test whether the monthly memory requirement is a regional phenomenon or a general forecasting principle. Finally, integrating satellite imagery or reanalysis data may allow the model to learn cloud-system morphology, enabling spatial forecasting and operational deployment. Together, these directions illustrate how the present findings extend beyond model testing and establish a research pathway toward scalable tropical solar forecasting systems.

Even without retraining, the model retained high accuracy when applied to new years, confirming its potential for long-term operational use in solar energy management. Its consistent performance across varying atmospheric conditions demonstrates its adaptability to the complex climate patterns typical of tropical zones. The LSTM-720 model proved capable of capturing both rapid fluctuations and seasonal trends in solar radiation. Its ability to generalize with minimal degradation highlights the reliability of deep learning methods for renewable energy forecasting in tropical regions. Beyond Mompox, this approach offers a scalable framework adaptable to other data-limited locations that experience strong climatic variability.

Although the long-context design produced strong predictive performance, this study presents several limitations that should be acknowledged. First, the model was trained and evaluated using ground-based measurements from a single tropical location. This restricts the generalizability of the results to other geographic regions, where climatological dynamics and atmospheric behavior may differ significantly. Second, the network relied exclusively on previous radiation values, without integrating meteorological predictors such as cloud cover, humidity, temperature, or satellite observations. The absence of these external drivers may limit model sensitivity during extreme or rapidly changing weather conditions. Third, the hyperparameter configuration was not extensively optimized through large-scale search procedures because the objective of this study focused on evaluating temporal context rather than architecture tuning. Finally, the proposed LSTM-720 design was not benchmarked against other deep learning architectures such as GRU, transformer-based models, or hybrid CNN-RNN frameworks. These limitations highlight the need to expand the modelling scale, incorporate additional predictors, and conduct spatial transferability tests across different tropical regions.

VI. Conclusion

This study demonstrated the strong potential of deep recurrent neural networks, particularly the LSTM-720 architecture, for forecasting hourly global solar radiation in a tropical setting. Using data collected in Mompox, Colombia, from 2021 to 2024, the model accurately reproduced both short-term variations and long-term seasonal trends without requiring normalization or additional meteorological inputs. By employing a 720-hour input window, the network achieved consistently high predictive accuracy, with determination coefficients (R^2) exceeding 0.90 in all validation years. These results confirm the model's ability to generalize across distinct climatic conditions and emphasize the importance of selecting an appropriate temporal window to capture the intrinsic periodicity of solar radiation. Shorter input sequences (5–24 hours) were effective for representing diurnal behavior but fell short in modeling multi-day persistence and intra-monthly variability. In contrast, the 720-hour configuration provided a more complete view of atmospheric dynamics, allowing the LSTM to learn both daily and broader climatological dependencies relevant for solar energy forecasting. The model also demonstrated strong temporal stability when validated against independent datasets from 2023 and 2024, maintaining RMSE deviations within ± 8 Wh/m² of the internal benchmark.

Residual analysis revealed minimal bias, near-symmetric error distributions, and stable variance across all years, evidence of the robustness and reliability of the proposed LSTM-720 approach. Beyond accuracy improvements, the results reveal a structural implication: forecasting in tropical environments requires long-memory recurrent modelling rather than short-window designs. The 720-hour context enabled the model to capture intra-month atmospheric persistence and sub-seasonal variability, which short sequences cannot represent. This positions the proposed framework as a methodological contribution that advances solar forecasting research beyond parameter optimization, providing a foundation for long-context modelling in other tropical regions.

From a practical standpoint, these findings underscore the suitability of recurrent neural networks for operational solar forecasting in tropical regions with limited meteorological infrastructure. The framework developed here can support photovoltaic energy management, grid planning, and hybrid renewable energy system design, offering a scalable and data-efficient tool for real-world applications. Future research should explore the integration of additional predictors, such as air temperature, humidity, and cloud fraction, and test hybrid deep learning architectures to enhance both predictive accuracy and spatial transferability. These results also provide modelling guidance for future tropical forecasting studies by confirming that long-context horizons are structurally required to recover sub-seasonal atmospheric behavior, a principle that may transfer to other equatorial locations and to different renewable forecasting tasks.

Declarations

The datasets analyzed in this research were provided by the Instituto de Hidrología, Meteorología y Estudios Ambientales (IDEAM), Colombia. They include hourly measurements of global solar radiation for the period 2021–2024, corresponding to the Mompox meteorological station. Due to institutional data-sharing restrictions, the original datasets are not publicly available but can be obtained upon reasonable request from the corresponding author or directly from IDEAM at https://www.ideam.gov.co/. All model training scripts and supplementary materials used in this study are also available from the corresponding author upon reasonable request to ensure full reproducibility of the results.

Conflict of Interest

The authors declare no conflict of interest.

Author Contributions

M.O.: conceptualization, writing an original draft, investigation, data curation, reviewing and editing; G.C.: investigation, methodology, reviewing and editing; A.R.: research design, data analysis, funding acquisition and project administration. All authors have read and agreed to the published version of the manuscript.

Declaration of Generative AI and AI-Assisted Technologies in the Writing Process

During the preparation of this work, the authors used Grammarly and ChatGPT to improve their writing. After using this tool, the authors reviewed and edited the content as needed and took full responsibility for the publication's content.

References

[1] H. Shahbeik et al., "Biomass to biofuels using hydrothermal liquefaction: A comprehensive review," Renew. Sustain. Energy Rev., vol. 189, p. 113976, Jan. 2024, doi: 10.1016/j.rser.2023.113976.

[2] F. A. Ahmad, J. Liu, F. Hashim, and K. Samsudin, "Short-Term Load Forecasting Utilizing a Combination Model: A Brief Review," Int. J. Technol., vol. 15, no. 1, p. 121, Jan. 2024, doi: 10.14716/ijtech.v15i1.5543.

[3] B. Tian, N. Wang, Y. Lin, and S. Shao, "Enhancing solar irradiance prediction precision: A stacked ensemble learning-based correction paradigm," Next Energy, vol. 8, p. 100306, Jul. 2025, doi: 10.1016/j.nxener.2025.100306.

[4] W. Dou, K. Wang, S. Shan, K. Zhang, H. Wei, and V. Sreeram, "A hybrid correction framework using disentangled seasonal-trend representations and MoE for NWP solar irradiance forecast," Appl. Energy, vol. 397, p. 126295, Nov. 2025, doi: 10.1016/j.apenergy.2025.126295.

[5] A. Kumar and S. Kerdsuwan, "Efficient Energy Conversion of Wastes and Fuels in Power Systems," KMUTNB Int. J. Appl. Sci. Technol., vol. 7, no. 2, pp. 1–26, Jun. 2014, doi: 10.14416/j.ijast.2014.02.001.

[6] F. J. L. de Lima et al., "Evaluation of the short and medium-term forecast quality of global solar irradiance from GFS-MOS and WRF-Solar models for the northeast region of Brazil," Energy Reports, vol. 13, pp. 2187–2203, Jun. 2025, doi: 10.1016/j.egyr.2025.01.073.

[7] M. Wild, D. Folini, C. Schär, N. Loeb, E. G. Dutton, and G. König-Langlo, "The global energy balance from a surface perspective," Clim. Dyn., vol. 40, no. 11–12, pp. 3107–3134, Jun. 2013, doi: 10.1007/s00382-012-1569-8.

[8] H. Mokhtarzadeh, S. Gorjian, and S. Minaei, "Design, development, and evaluation of a low-cost smart solar-powered weather station for use in agricultural environments," Results Eng., vol. 26, p. 104848, Jun. 2025, doi: 10.1016/j.rineng.2025.104848.

[9] U. H. Ramadhani, L. N. Ramdhania, H. Iskandar, A. Fudholi, and H. Abimanyu, "Review of solar thermal technologies in sustainable animal agriculture farms: Current and potential uses," Sol. Energy, vol. 291, p. 113374, May 2025, doi: 10.1016/j.solener.2025.113374.

[10] A. O. Adelakun and F. M. Adelakun, "Mathematical modeling and seasonal solar radiation variability in Nigeria's geo-political zones: A recurrence and multifractal analysis," J. Atmos. Solar-Terrestrial Phys., vol. 261, p. 106290, Aug. 2024, doi: 10.1016/j.jastp.2024.106290.

[11] X. Wang, P. Xie, Y. Xie, and H. Jiang, "Localized solar radiation zoning by combining spatially continuous estimates and Gaussian mixture models," J. Atmos. Solar-Terrestrial Phys., vol. 268, p. 106432, Mar. 2025, doi: 10.1016/j.jastp.2025.106432.

[12] J. Polo et al., "Preliminary survey on site-adaptation techniques for satellite-derived and reanalysis solar radiation datasets," Sol. Energy, vol. 132, pp. 25–37, Jul. 2016, doi: 10.1016/j.solener.2016.03.001.

[13] E. Piña Henríquez and C. Millán Aguilar, "Propuesta para incrementar el uso de energía solar fotovoltaica en Ámsterdam. = Strategic guidelines to increase the use of photovoltaic solar energy in Amsterdam," Cuad. Investig. Urbanística, no. 153, pp. 01–66, Apr. 2024, doi: 10.20868/ciur.2024.153.5251.

[14] O. F. Al-Rawi, Y. Bicer, and S. G. Al-Ghamdi, "Sustainable solutions for healthcare facilities: examining the viability of solar energy systems," Front. Energy Res., vol. 11, Jul. 2023, doi: 10.3389/fenrg.2023.1220293.

[15] I. Akhtar, S. Kirmani, M. Jameel, and F. Alam, "Feasibility Analysis of Solar Technology Implementation in Restructured Power Sector With Reduced Carbon Footprints," IEEE Access, vol. 9, pp. 30306–30320, 2021, doi: 10.1109/ACCESS.2021.3059297.

[16] S. Keeratianant, S. Chirarattananon, A. Nathakaranakule, and P. Rakkwamsuk, "Cost-effectiveness of Solar Cooling for Office and Hypermarket," Appl. Sci. Eng. Prog., Sep. 2019, doi: 10.14416/j.asep.2019.09.004.

[17] M. M. de Souza Grilo, A. F. C. Fortes, R. P. G. de Souza, J. A. M. Silva, and M. Carvalho, "Carbon footprints for the supply of electricity to a heat pump: Solar energy vs. electric grid," J. Renew. Sustain. Energy, vol. 10, no. 2, Mar. 2018, doi: 10.1063/1.4997306.

[18] W. Dou et al., "A multi-modal deep clustering method for day-ahead solar irradiance forecasting using ground-based cloud imagery and time series data," Energy, vol. 321, p. 135285, Apr. 2025, doi: 10.1016/j.energy.2025.135285.

[19] N. T. Anh et al., "Evaluation Of Solar Radiation Forecast Models: Lightgbm, Lstm And Gru," J. Sci. Technol. - HaUI, vol. 60, no. 8, pp. 3–10, Aug. 2024, doi: 10.57001/huih5804.2024.256.

[20] E. Gulay, M. Sen, and O. B. Akgun, "Forecasting electricity production from various energy sources in Türkiye: A predictive analysis of time series, deep learning, and hybrid models," Energy, vol. 286, p. 129566, Jan. 2024, doi: 10.1016/j.energy.2023.129566.

[21] A. Verdone, M. Panella, E. De Santis, and A. Rizzi, "A review of solar and wind energy forecasting: From single-site to multi-site paradigm," Appl. Energy, vol. 392, p. 126016, Aug. 2025, doi: 10.1016/j.apenergy.2025.126016.

[22] S. Ghimire, R. C. Deo, D. Casillas-Pérez, and S. Salcedo-Sanz, "Boosting solar radiation predictions with global climate models, observational predictors and hybrid deep-machine learning algorithms," Appl. Energy, vol. 316, p. 119063, Jun. 2022, doi: 10.1016/j.apenergy.2022.119063.

[23] L. S. Hoyos-Gómez, J. F. Ruiz-Muñoz, and B. J. Ruiz-Mendoza, "Short-term forecasting of global solar irradiance in tropical environments with incomplete data," Appl. Energy, vol. 307, p. 118192, Feb. 2022, doi: 10.1016/j.apenergy.2021.118192.

[24] Z. He, X. Zhang, M. Li, S. Wang, and G. Xiao, "A novel solar radiation forecasting model based on time series imaging and bidirectional long short-term memory network," Energy Sci. Eng., vol. 12, no. 11, pp. 4876–4893, Nov. 2024, doi: 10.1002/ese3.1875.

[25] N. E. Michael, R. C. Bansal, A. A. A. Ismail, A. Elnady, and S. Hasan, "A cohesive structure of Bi-directional long-short-term memory (BiLSTM) -GRU for predicting hourly solar radiation," Renew. Energy, vol. 222, p. 119943, Feb. 2024, doi: 10.1016/j.renene.2024.119943.

[26] A. Sebastianelli, F. Serva, A. Ceschini, Q. Paletta, M. Panella, and B. Le Saux, "Machine learning forecast of surface solar irradiance from meteo satellite data," Remote Sens. Environ., vol. 315, p. 114431, Dec. 2024, doi: 10.1016/j.rse.2024.114431.

[27] M. Mohammadi, S. Jamshidi, A. Rezvanian, M. Gheisari, and A. Kumar, "Advanced fusion of MTM-LSTM and MLP models for time series forecasting: An application for forecasting the solar radiation," Meas. Sensors, vol. 33, p. 101179, Jun. 2024, doi: 10.1016/j.measen.2024.101179.

[28] S. C. Nwokolo et al., "Machine learning and physics-based model hybridization to assess the impact of climate change on single- and dual-axis tracking solar-concentrated photovoltaic systems," Phys. Chem. Earth, Parts A/B/C, vol. 138, p. 103881, Jun. 2025, doi: 10.1016/j.pce.2025.103881.

[29] L. Barancsuk, V. Groma, and B. Kocziha, "Hybrid ultra-short term solar irradiation forecasting using resource-efficient multi-step long-short term memory," Renew. Energy, vol. 247, p. 122962, Jul. 2025, doi: 10.1016/j.renene.2025.122962.

[30] S. Al-Dahidi et al., "Techno-economic implications and cost of forecasting errors in solar PV power production using optimized deep learning models," Energy, vol. 323, p. 135877, May 2025, doi: 10.1016/j.energy.2025.135877.

[31] G. Jerse and A. Marcucci, "Deep Learning LSTM-based approaches for 10.7 cm solar radio flux forecasting up to 45-days," Astron. Comput., vol. 46, p. 100786, Jan. 2024, doi: 10.1016/j.ascom.2024.100786.

[32] J. Kapica et al., "The potential impact of climate change on European renewable energy droughts," Renew. Sustain. Energy Rev., vol. 189, p. 114011, Jan. 2024, doi: 10.1016/j.rser.2023.114011.

[33] Y. Fan et al., "Assessing the potential and complementary characteristics of China's solar and wind energy under climate change," Renew. Energy, vol. 249, p. 123213, Aug. 2025, doi: 10.1016/j.renene.2025.123213.

[34] A. Rathore, P. Gupta, R. Sharma, and R. Singh, "Day ahead solar forecast using long short term memory network augmented with Fast Fourier transform-assisted decomposition technique," Renew. Energy, vol. 247, p. 123021, Jul. 2025, doi: 10.1016/j.renene.2025.123021.

[35] M. Madhiarasan and M. Louzazni, "Combined Long Short-Term Memory Network-Based Short-Term Prediction of Solar Irradiance," Int. J. Photoenergy, vol. 2022, pp. 1–19, Aug. 2022, doi: 10.1155/2022/1004051.

[36] A. M. Assaf, H. Haron, H. N. Abdull Hamed, F. A. Ghaleb, S. N. Qasem, and A. M. Albarrak, "A Review on Neural Network Based Models for Short Term Solar Irradiance Forecasting," Appl. Sci., vol. 13, no. 14, p. 8332, Jul. 2023, doi: 10.3390/app13148332.

[37] R. A. Rajagukguk, R. A. A. Ramadhan, and H.-J. Lee, "A Review on Deep Learning Models for Forecasting Time Series Data of Solar Irradiance and Photovoltaic Power," Energies, vol. 13, no. 24, p. 6623, Dec. 2020, doi: 10.3390/en13246623.

[38] IDEAM, "Banco de Datos Meteorológicos," 2024. https://www.ideam.gov.co