Bayesian-optimized LSTM networks for accurate day-ahead photovoltaic power prediction

Enas Ali Ahmed[1], Muna Hassan Hussein[1], Abdulkreem Mohammed Salih[2]
[1]Department of Medical Instrumentation Techniques Engineering, Northern Technical University (NTU), Mosul, Iraq
[2]Aldour Technical Institute, Northern Technical University (NTU), Mosul, Iraq

Article Info

Article history:

Received Feb 16, 2026
Revised Aug 8, 2026
Accepted Aug 24, 2026

Keywords:

Bayesian optimization
LSTM networks
Photovoltaic power forecasting
Solar energy prediction
Time-series forecasting

ABSTRACT

Accurate day-ahead photovoltaic (PV) power forecasting is essential for effective energy management and grid balancing. This study proposes a Bayesian-optimized long short-term memory (LSTM) network for day-ahead PV power prediction. The model was evaluated using a PV-meteorological time-series dataset collected from a 500-kWp grid-connected PV installation in Mosul, Iraq, between January 1, 2021 and December 31, 2023 at an hourly sampling interval. After data-quality screening, 25,842 valid synchronized observations were retained from 26,280 timestamps. Each forecasting sample used the previous 24 hours of PV power, solar irradiance, ambient temperature, relative humidity, wind speed, and temporal indicators to predict the subsequent 24-hour PV power profile. The data were divided chronologically into training, validation, and independent test subsets. Preprocessing included duplicate removal, missing-value treatment, IQR-based outlier handling, temporal alignment, and min-max normalization fitted only on the training subset. Bayesian optimization tuned the LSTM architecture and training hyperparameters, while a hybrid MSE-MAE loss balanced large deviations and overall error. The proposed model achieved an MAE of 15.2 kW, an RMSE of 20.5 kW, a MAPE of 8.3%, and an R^2 of 0.93, outperforming linear regression and autoregressive integrated moving average (ARIMA) under the evaluated conditions.

Corresponding Author:

Enas Ali Ahmed
Department of Medical Instrumentation Techniques Engineering, Northern Technical University (NTU)
Mosul 41002, Iraq
Email: enasali@ntu.edu.iq

1. INTRODUCTION

The energy sector is experiencing a fast-moving transformation, where the reduction of greenhouse gases (GHG) and the switch from carbon-intensive power-generation systems to cleaner and more sustainable alternatives drive the revolution. The energy sector is in a state of accelerated transformation, with the focus on reducing GHG emissions and replacing carbon-intensive power generation systems with clean and sustainable alternatives. Solar PV generation has been a hot topic, as the energy from the sun is omnipresent, the environmental impact of solar PV is relatively low, and the costs of deployment are still falling. But there are challenges that occur when PV panels are integrated on a large scale, as sunlight is intermittent and very weather-sensitive [1], [2].

PV output is influenced by solar irradiance, cloud cover, ambient temperature, relative humidity, wind speed, seasonal variation, and local geographical conditions [3]. Rapid changes in these variables can produce substantial fluctuations in generated power, increasing uncertainty in generation scheduling and grid balancing. Accurate day-ahead PV forecasting is therefore important for storage scheduling, demand–supply balancing, reserve allocation, and secure renewable-energy integration [4]. Reliable forecasts are also useful for smart-grid operation, demand-response programs, energy trading, and long-term planning.

The forecast of PV power is still challenging due to the nature of the underlying time series, which is nonlinear and non-stationary, weather-dependent, and shows daily and seasonal fluctuations. The traditional methods are persistence, autoregressive models, autoregressive integrated moving average (ARIMA), physical models, and numerical weather prediction (NWP) [5]. While persistence is inexpensive to compute, there are large errors when weather rapidly changes. Physical and NWP-based models can be used for longer time scales, yet they depend on accurate meteorological data, need careful configuration, and take up a lot of computational resources [6]. These limitations have encouraged data-driven forecasting approaches. Support vector machines, random forests, and artificial neural networks can learn nonlinear relationships between meteorological inputs and PV output [7], [8]. Figure 1 illustrates conceptual structures of support vector regression and a feedforward neural network used in conventional PV forecasting.

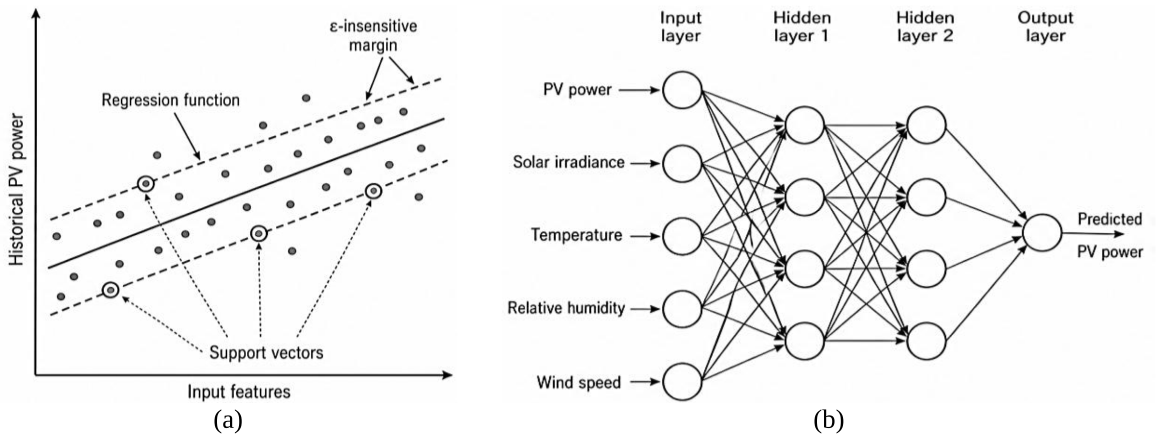

Figure 1. Conceptual structures of (a) support vector regression and (b) a feedforward artificial neural network used in conventional PV power forecasting

While traditional machine learning models can make accurate predictions, many of the process observations are taken as independent samples and do not explicitly account for the temporal dependence. Recurrent neural networks (RNNs) represent sequential information by means of recurrent hidden states [9]. Long short-term memory (LSTM) networks are an extension of recurrent models with the addition of memory cells and input, forget, and output gates, which makes it possible to store the relevant information for longer periods of time [10], [11]. In general, the LSTM sequence model is presented in Figure 2.

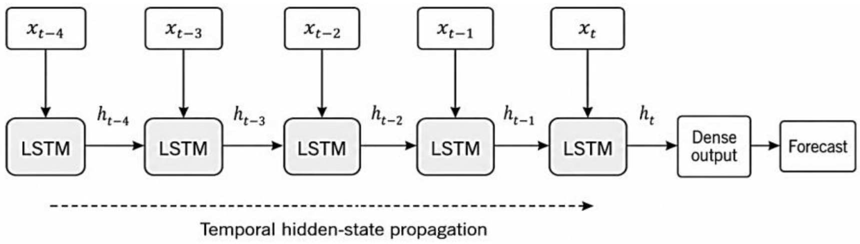

Figure 2. General LSTM sequence-model architecture

LSTM-based PV forecasting still faces two important challenges. First, PV and meteorological records may contain missing timestamps, duplicate records, sensor noise, communication failures, and physically invalid values. Without appropriate preprocessing, such defects can degrade forecasting performance [7]. Second, LSTM accuracy depends strongly on architectural and training hyperparameters, including the number of recurrent layers, units per layer, learning rate, dropout rate, batch size, and input-window length. Manual tuning is subjective, while exhaustive grid search can become computationally expensive. Bayesian optimization addresses this limitation by using a probabilistic surrogate model and an acquisition function to select promising configurations [12]. In this work, a hybrid loss function combining mean squared error (MSE) and mean absolute error (MAE) is used, along with a multivariate LSTM forecasting model and a Bayesian optimization method. The evaluation of the framework is performed with the PV–meteorological data recorded during 3 years at a 500-kWp grid-connected PV system deployed in Mosul city, Iraq, at the hourly level. The framework has been evaluated using PV–meteorological data recorded for 3 years at a 500-kWp grid-connected PV system installed in Mosul city, Iraq, at an hourly time scale. The historical PV power is combined with solar irradiance, ambient temperature, relative humidity, wind speed, and temporal indicators. It is not about introducing new individual algorithms such as LSTM or Bayesian optimization but about combining direct 24-step multi-variable forecasting, leakage-aware chronological preprocessing, systematic optimization of the most important architectural and training parameters, and hybrid MSE–MAE learning in one forecasting framework that is reproducible.

2. RELATED WORK

As solar power becomes a more prominent part of energy systems, forecasting PV generation has emerged as a critical part of grid operations, storage planning, reserve scheduling, and generation planning. Previous works range from statistical, physical, machine-learning, and deep-learning methods. It is often challenging to directly compare the results from published studies due to variations in geographic area, PV capacity, sampling time, forecast horizon, input variables, and normalization methods.

2.1. Early approaches to PV forecasting

Early forecasting studies relied mainly on historical irradiance and power-generation time series [13]. Persistence models make predictions for future generations based on the current value or a similar value from a prior point in time. They do not involve much computation, but their accuracy becomes poor if the meteorological conditions change quickly [5]. Temporal autocorrelation, trend, and periodic behavior can be represented by ARIMA and seasonal ARIMA models, respectively, and have been able to forecast short-term values under stable conditions with satisfactory results [14]-[16]. Their performance is, however, dependent on the selection of the order and proper treatment of stationarity, but they might not fully describe a nonlinear response and sudden changes in the irradiance. The linear and multiple-regression models have the advantages of interpretability and low computing cost, but lack the ability to establish a linear model to describe the complex interactions between PV and weather.

2.2. Physical and numerical weather prediction models

Physical forecasting methods involve estimating the PV output using equations from the atmosphere, solar geometry, module characteristics, and NWP information. NWP products can be provided to PV-performance models for medium- and long-range prediction [17]. These methods are physically based and rely on accurate initial conditions, boundary data, local weather data, and system specifications. They are also difficult to compute and cannot be configured easily [18]. Their accuracy will be reduced if the conditions in the regions for which the NPWs are computed are not representative of the local conditions of the cloud and irradiance, especially in the case of distributed PV.

2.3. Machine-learning approaches in PV forecasting

Non-linear relationships can be learnt from the historical PV and meteorological data using machine-learning methods [19], [20]. Solar radiation and PV generation forecasting have been attempted using artificial neural networks (ANN), support vector machines (SVM), random forests (RF), and gradient boosting (GB) [20], [21]. Many of the traditional approaches process observations independently and fail to explicitly model the temporal ordering of the process. Lagged variables can be manually entered, but may not sufficiently describe longer dependencies. Yet another difficulty arises when the hyperparameters of the network architecture and the training procedure are tuned by hand, grid search, or random search, which is inefficient when multiple parameters need to be optimized at once.

2.4. LSTM networks and hyperparameter optimization

LSTM networks are a type of recurrent network with memory cells and gating mechanisms that regulate the flow of information over time [11], [22]. Time series forecasting for hourly and day-ahead solar forecasting has been performed successfully using appropriately configured LSTMs [23]-[25]. Their accuracy nevertheless depends on the dataset, climatic conditions, feature set, forecast horizon, and evaluation protocol. LSTM performance is also sensitive to layer depth, hidden units, learning rate, dropout, batch size, and input-window length [25]. Bayesian optimization provides a systematic alternative by using previous trial outcomes to guide subsequent candidate selection [26]. The present study applies Bayesian optimization to the main LSTM architectural and training parameters and performs feature-level fusion of PV power, weather variables, and temporal indicators [23], [27]. Related work has also combined decomposition and LSTM-based PV forecasting [28], evaluated solar-power prediction without exogenous inputs [29], and reviewed artificial-intelligence techniques for photovoltaic applications [30]-[33].

3. METHOD

The proposed methodology employs an LSTM network for multi-step day-ahead PV power forecasting. The framework includes data acquisition, data-quality screening, preprocessing, chronological sample construction, Bayesian hyperparameter optimization, final model training, and independent evaluation. Figure 3 illustrates the physical experimental setup and data flow, while Figure 4 summarizes the complete computational workflow.

The dataset was collected from a 500-kWp grid-connected PV installation in Mosul, Iraq, from January 1, 2021 and December 31, 2023. PV output power was obtained from the plant supervisory monitoring system, while solar irradiance, ambient temperature, relative humidity, and wind speed were obtained from a co-located meteorological station. Hour-of-day and day-of-year variables were derived from the timestamps to represent daily and seasonal patterns. The three-year observation period covered winter, spring, summer, and autumn and therefore represented seasonal variation in irradiance and local meteorological conditions.

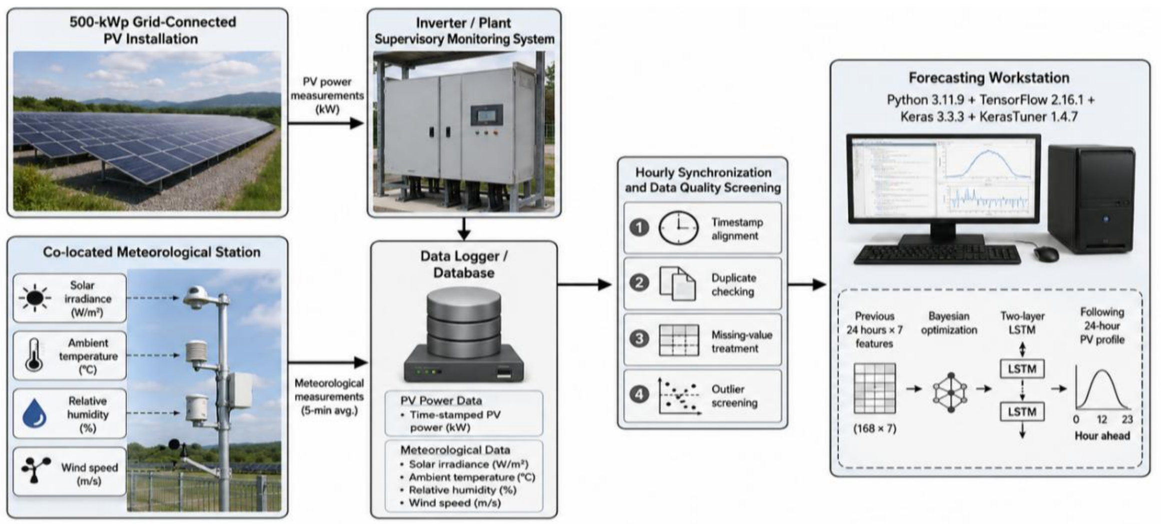

Figure 3. Physical experimental setup and data flow for day-ahead PV power forecasting

All measurements were synchronized at an hourly interval. The complete three-year timeline contained 26,280 timestamps. After duplicate removal, short-gap treatment, and exclusion of unresolved or physically invalid records, 25,842 synchronized hourly observations were retained; 438 timestamps (1.67%) were excluded. The hourly resolution provides 24 observations per day and supports a 24-step day-ahead forecast. This resolution supplies 24 measurements per day, aligns directly with the 24-output forecasting target, and represents the complete diurnal generation cycle at the temporal resolution adopted for this study.

The records were sorted chronologically and checked for duplicates, missing values, inconsistent units, and physically impossible measurements. Gaps of no more than three consecutive hours were reconstructed by linear interpolation within the corresponding data partition; longer unresolved gaps were excluded. Interpolation was not permitted across training, validation, or test boundaries.

Potential outliers were identified using physical operating limits and the interquartile-range (IQR) rule. The IQR and candidate interval were defined as (1) and (2).

$$\mathrm{IQR} = \mathrm{Q}_3 - \mathrm{Q}_1 \tag{1}$$

$$[\mathrm{Q}_1 - 1.5{\cdot}\mathrm{IQR}, \mathrm{Q}_3 + 1.5{\cdot}\mathrm{IQR}] \tag{2}$$

The IQR limits were estimated using the training subset only and then applied to the validation and test subsets. Values associated with clear measurement or communication errors were winsorized to the corresponding training-derived limits or removed when reliable correction was not possible. Physically meaningful nighttime zero-generation observations were retained.

Min–Max normalization transformed each feature to the interval [0,1]. For a feature x, the transformation was:

$$x' = (x - x_{min}) / (x_{max} - x_{min}) \tag{3}$$

The minimum and maximum values were estimated exclusively from the training subset. The same training-derived parameters were applied unchanged to the validation and test subsets to prevent information leakage. The final input vector contained historical PV power, solar irradiance, ambient temperature, relative humidity, wind speed, hour of day, and day of year.

The 25,842 retained observations were divided chronologically without random shuffling into 18,089 training observations (70%), 3,876 validation observations (15%), and 3,877 test observations (15%). Sliding windows were constructed separately within each subset. Each sample used the previous 24-hourly observations as input and the subsequent 24 PV power values as the target. This produced 18,042 training sequences, 3,829 validation sequences, and 3,830 independent test sequences, assuming continuous records after preprocessing.

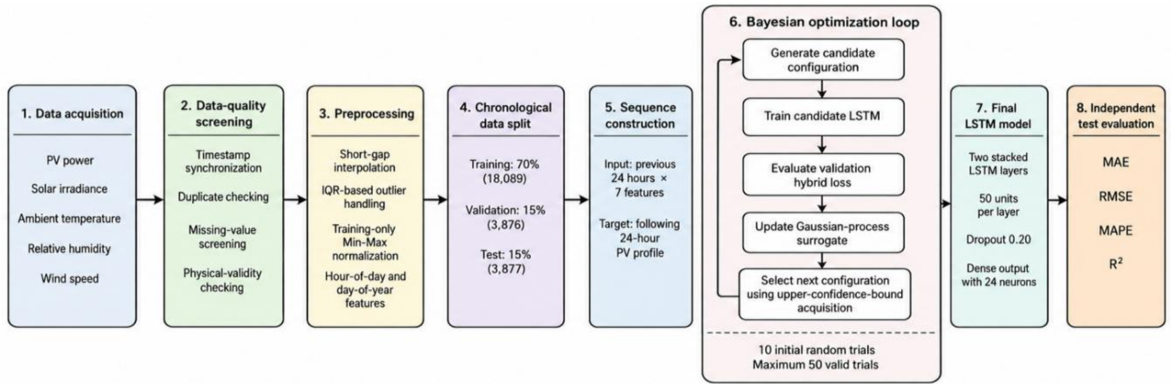

Figure 4. Complete methodology of the proposed Bayesian-optimized LSTM forecasting framework

The selected network contained two stacked LSTM layers with 50 units each, a dropout layer with a rate of 0.20, and a dense output layer with 24 neurons, one for each hour of the forecast horizon. Figure 5 shows the implemented architecture. The model was trained using a hybrid MSE–MAE loss:

$$L = \alpha \cdot MSE + (1 - \alpha) \cdot MAE \tag{4}$$

The loss weight was $\alpha = 0.50$. Adam optimization was used with a learning rate of 0.001 and a batch size of 64. Training was permitted for a maximum of 150 epochs, with early stopping patience of 15 epochs and restoration of the weights associated with the lowest validation loss.

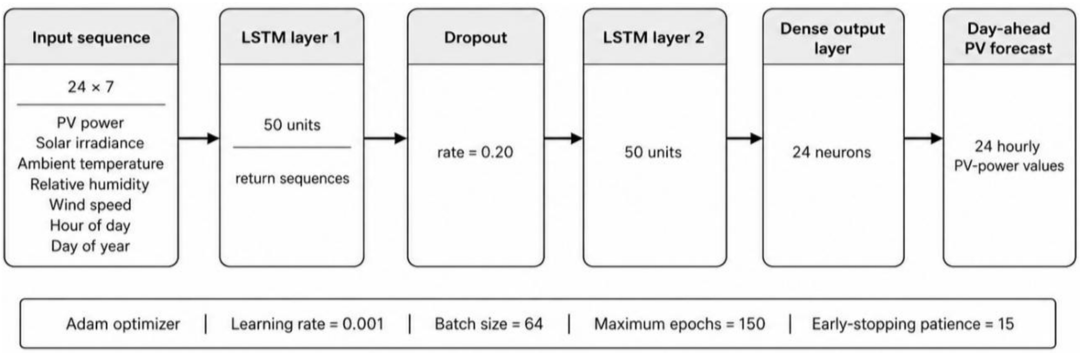

Figure 5. Implemented a two-layer LSTM architecture for direct 24-hour PV power forecasting

Bayesian optimization was implemented with KerasTuner. A Gaussian-process surrogate approximated the validation objective, and an upper-confidence-bound acquisition rule selected candidate configuration. For minimization, the acquisition score was expressed as (5):

$$a(\theta) = \mu(\theta) - \beta \cdot \sigma(\theta) \quad (5)$$

where $\mu(\theta)$ and $\sigma(\theta)$ are the surrogate mean and uncertainty for configuration θ, respectively, and $\beta = 2.6$ controls exploration versus exploitation. The Gaussian-process noise parameter was set to 10^{-4}. The search began with 10 random configurations and evaluated a maximum of 50 valid trials using a fixed random seed of 42. The optimization objective was the validation hybrid loss, and the independent test subset was not accessed during tuning.

4. EXPERIMENTAL WORK AND RESULTS

The model was implemented using the software environment summarized in Table 1. All computations used 32-bit floating-point precision, chronological data loading, and random seed 42. Linear regression and ARIMA (2,1,2) were used as baseline methods. The linear regression model used the same input variables and chronological partitions, while ARIMA was fitted to the PV power series using the training period and selected using validation performance. Linear regression was implemented as a multi-output regression model. Each 24-hour input window, containing seven variables, was flattened into a 168-dimensional feature vector, and the model simultaneously generated the 24-hourly PV-power values of the subsequent day. Thus, the linear regression baseline used the same input information, chronological partitions, and 24-hour prediction horizon as the proposed LSTM. ARIMA was implemented as a univariate ARIMA (2,1,2) model using the historical PV-power series. The model generated the subsequent 24-hour forecast recursively: each predicted hourly value was incorporated into the forecasting process to obtain the next prediction until the complete 24-hour horizon was produced. The order was selected from candidate configurations using the lowest validation RMSE, with the Akaike information criterion used as a secondary criterion when configurations produced similar validation performance. The Augmented Dickey–Fuller test was applied to assess stationarity. The original PV power series was found to be non-stationary at the adopted significance level, while first-order differencing produced a stationary series; therefore, d = 1 was used. The ARIMA model was fitted only to the training period; its order was selected using the validation period, and the independent test period remained unused until final evaluation $(2{,}1{,}2)(p{,}d{,}q)d = 1$.

Forecasting performance was evaluated using MAE, RMSE, MAPE, and the coefficient of determination (R^2). MAE represents the average absolute deviation, whereas RMSE assigns greater weight to large errors. MAPE was computed only for observations whose measured PV power exceeded 5 kW to avoid instability from nighttime or near-zero generation. MAPE was calculated only for independent-test target values whose measured PV power exceeded 5 kW. This cutoff corresponds to 1% of the 500-kWp rated plant capacity and was selected to prevent nighttime and near-zero generation values from producing disproportionately large and numerically unstable percentage errors. The independent test set contained 3,830 forecasting sequences, each comprising 24 target hours, yielding 91,920 evaluated target values. Of these, approximately 47,800 target values, corresponding to approximately 52.0%, had measured PV power less than or equal to 5 kW and were excluded from the MAPE calculation. The remaining approximately 44,120 target values were used to calculate the reported MAPE of 8.3%. The 5-kW threshold was applied only to MAPE; all 91,920 target values were retained when calculating MAE, RMSE, and R^2.R^2

On the independent test set, the proposed Bayesian-optimized LSTM achieved an MAE of 15.2 kW, an RMSE of 20.5 kW, a MAPE of 8.3%, and an R^2 of 0.93 based on the selected configuration in Table 2. Linear regression produced an MAE of 25.4 kW, an RMSE of 30.1 kW, and a MAPE of 12.5%. ARIMA (2, 1, 2) achieved an MAE of 22.8 kW, an RMSE of 27.6 kW, and a MAPE of 11.0%. Table 3 summarizes these results.

Relative to linear regression, the proposed model reduced MAE by 40.16%, RMSE by 31.89%, and MAPE by 33.60%, as shown in Table 4. Compared with ARIMA, the corresponding reductions were 33.33%, 25.72%, and 24.55%. These are controlled comparisons based on the same dataset, chronological test interval, sampling resolution, and 24-hour forecast horizon.

Figure 6 presents measured and predicted PV power over a representative 96-hour interval from the independent test subset. The model follows the principal daily generation pattern and reproduces the main rising and falling trends. The largest deviations occur near peak-generation periods, where rapid irradiance changes produce sharper variations. The one-hour sampling interval provides 24 observations per day and is aligned with the 24-step day-ahead prediction objective. It captures the diurnal PV generation profile at the temporal resolution used for the scheduling-oriented evaluation in this study.

Table 1. Experimental setup and reproducibility settings

Item	Setting
PV installation	500-kWp grid-connected system, Mosul, Iraq
Observation period	1 January 2021–31 December 2023
Sampling interval	1 hour
Valid hourly observations	25,842
Training/validation/test observations	18,089/3,876/3,877
Training/validation/test sequences	18,042/3,829/3,830
Input/forecast horizon	24 hours/24 hours
Input features	PV power, irradiance, temperature, humidity, wind speed, hour, day of year
Software	Python 3.11.9; TensorFlow 2.16.1; Keras 3.3.3; KerasTuner 1.4.7
Optimizer/learning rate/batch size	Adam/0.001/64
Maximum epochs/patience	150/15
Initial/maximum Bayesian trials	10/50
Random seed	42
Chronological partition sizes	18,089/3,876/3,877 observations (70%/15%/15%)
Data-acquisition hardware	Plant supervisory monitoring system and co-located meteorological station; exact manufacturers and models unavailable in the retained study records
Excluded-record breakdown	438 timestamps excluded in total; category-level counts unavailable; interpolated and winsorized values were retained

Table 2. Bayesian-optimization search space and selected configuration

Hyperparameter	Search space	Selected value
LSTM layers	1–3	2
Units per layer	32–128, step 2	50, 50
Dropout rate	0.00–0.50	0.20
Learning rate	10^{-4}–10^{-2}, logarithmic	0.001
Batch size	16, 32, 64, 128	64
Input window	24, 48, 72 hours	24 hours
Hybrid-loss weight α	0.20–0.80	0.50

Table 3. Test-set performance of the proposed model and baseline methods.

Method	MAE (kW) ↓	RMSE (kW) ↓	MAPE (%) ↓	R^2 ↑
Linear regression	25.4	30.1	12.5	0.85
ARIMA (2,1,2)	22.8	27.6	11.0	0.87
Proposed BO-LSTM	15.2	20.5	8.3	0.93

Table 4. Relative error reduction achieved by the proposed model

Baseline	MAE reduction (%)	RMSE reduction (%)	MAPE reduction (%)
Linear regression	40.16	31.89	33.60
ARIMA (2,1,2)	33.33	25.72	24.55

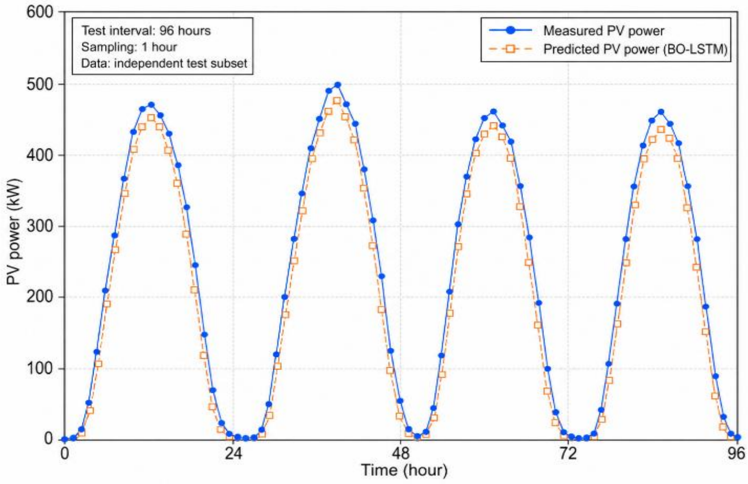

Figure 6. Hourly measured and predicted PV generation over a representative 96-hour interval from the independent test subset

The prediction-error distribution is shown in Figure 7. Most residuals are concentrated near zero, indicating no large systematic positive or negative bias. A smaller number of larger residuals occur during high-variability and peak-generation periods. Figure 8 compares the test-set MAE, RMSE, and MAPE values of the three methods. The proposed model records the lowest error for all metrics. The improvement over linear regression demonstrates the benefit of nonlinear temporal modeling, while the improvement over ARIMA indicates that the multivariate LSTM captures relationships not fully represented by a conventional univariate time-series model.

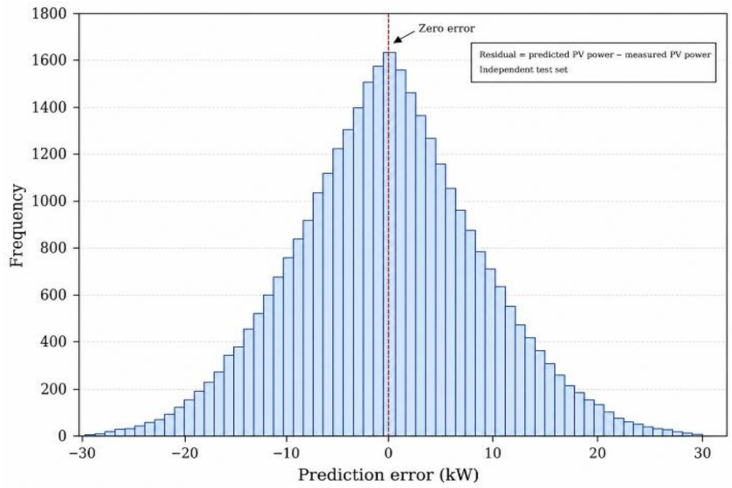

Figure 7. Distribution of prediction residuals for the proposed Bayesian-optimized LSTM on the independent test set

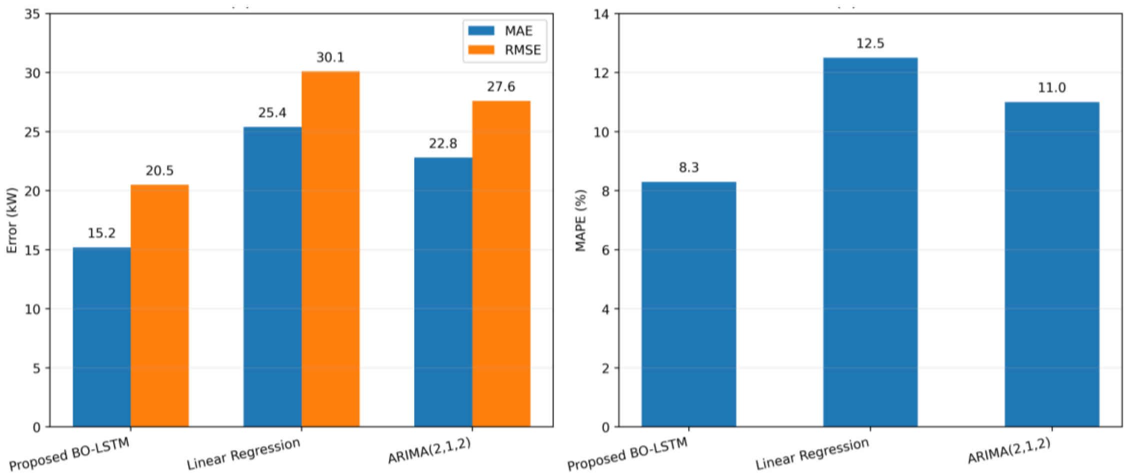

(a) (b)

Figure 8. Test-set performance comparison of the proposed BO-LSTM, linear regression, and ARIMA (2,1,2): (a) MAE and RMSE and (b) MAPE

The lower errors of the proposed model are consistent with its ability to represent nonlinear temporal dependencies while jointly processing PV power, meteorological variables, and temporal indicators. Linear regression is restricted to a fixed linear mapping, whereas the ARIMA (2,1,2) baseline is a univariate time-series model and therefore does not directly exploit the complete multivariable input representation used by the LSTM. The time-series and residual analyses also identify the remaining weakness of the proposed model. Forecasts follow the principal daily profile over most of the displayed interval, but larger deviations occur around peak-generation and high-variability periods, where rapid irradiance changes create sharper transitions. Consequently, the results demonstrate improvement only relative to the evaluated baselines under the same single-site test interval and should not be interpreted as universal superiority over all modern forecasting architectures. Overall, the results indicate that the proposed Bayesian-optimized LSTM improved day-ahead PV forecasting relative to the evaluated baselines under the studied single-site conditions. The findings are limited to the evaluated installation, observation period, and climatic conditions in Mosul; broader multi-site validation is required before generalization to other PV systems can be established.

5. CONCLUSION

This study presented a Bayesian-optimized LSTM framework for 24-hour-ahead PV power forecasting using synchronized PV and meteorological measurements from a 500-kWp grid-connected installation in Mosul, Iraq. The evaluation used 25,842 valid hourly observations collected from 1 January 2021 to 31 December 2023. Historical PV power, solar irradiance, ambient temperature, relative humidity, wind speed, hour of day, and day of year were integrated within a multivariate sequential-learning framework. Chronological training, validation, and test partitions were employed, and preprocessing parameters were estimated from the training subset only to reduce temporal leakage. On the independent test set, the model achieved an MAE of 15.2 kW, an RMSE of 20.5 kW, a MAPE of 8.3%, and an R^2 of 0.93. Relative to linear regression, MAE, RMSE, and MAPE were reduced by 40.16%, 31.89%, and 33.60%, respectively; relative to ARIMA (2, 1, 2), the reductions were 33.33%, 25.72%, and 24.55%. Bayesian optimization provided a systematic procedure for selecting the LSTM architecture and training parameters, while the hybrid MSE–MAE loss balanced large deviations and average error. The findings support the effectiveness of the proposed framework under the evaluated single-site conditions. Future work should investigate multi-site and cross-climate validation, probabilistic prediction intervals, repeated-run statistical analysis, additional deep-learning baselines, and real-time deployment in PV energy-management systems.

FUNDING INFORMATION

Authors state no funding involved.

AUTHOR CONTRIBUTIONS STATEMENT

This journal uses the Contributor Roles Taxonomy (CRediT) to recognize individual author contributions, reduce authorship disputes, and facilitate collaboration.

Name of Author	C	M	So	Va	Fo	I	R	D	O	E	Vi	Su	P	Fu
Enas Ali Ahmed	✓	✓	✓	✓	✓	✓		✓	✓	✓	✓		✓	
Muna Hassan Hussein	✓	✓		✓	✓	✓	✓	✓		✓				
Abdulkreem Mohammed Salih	✓	✓		✓	✓	✓	✓			✓			✓	✓

C : **C**onceptualization
M : **M**ethodology
So : **So**ftware
Va : **Va**lidation
Fo : **Fo**rmal analysis
I : **I**nvestigation
R : **R**esources
D : **D**ata Curation
O : Writing - **O**riginal Draft
E : Writing - Review & **E**diting
Vi : **Vi**sualization
Su : **Su**pervision
P : **P**roject administration
Fu : **Fu**nding acquisition

CONFLICT OF INTEREST STATEMENT

Authors state no conflict of interest.

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author, Enas Ali Ahmed, upon reasonable request. The dataset consists of synchronized photovoltaic power and meteorological measurements collected from the 500-kWp grid-connected photovoltaic installation and the co-located meteorological station in Mosul, Iraq, for the period from 1 January 2021 to 31 December 2023. The data are not currently deposited in a public repository.

REFERENCES

[1] Q. Hassan *et al.*, "The renewable energy role in the global energy transformations," *Renewable Energy Focus*, vol. 48, p. 100545, Mar. 2024, doi: 10.1016/j.ref.2024.100545.

[2] F. Rozon, C. McGregor, and M. Owen, "Long-term forecasting framework for renewable energy technologies' installed capacity and costs for 2050," *Energies*, vol. 16, no. 19, p. 6874, Sep. 2023, doi: 10.3390/en16196874.

[3] M. Bošnjaković, M. Stojkov, M. Katinić, and I. Lacković, "Effects of extreme weather conditions on PV systems," *Sustainability*, vol. 15, no. 22, p. 16044, Nov. 2023, doi: 10.3390/su152216044.

[4] E. Ejuh Che, K. Roland Abeng, C. D. Iweh, G. J. Tsekouras, and A. Fopah-Lele, "The impact of integrating variable renewable energy sources into grid-connected power systems: challenges, mitigation strategies, and prospects," *Energies*, vol. 18, no. 3, p. 689, Feb. 2025, doi: 10.3390/en18030689.

[5] J. Antonanzas, N. Osorio, R. Escobar, R. Urraca, F. J. Martinez-de-Pison, and F. Antonanzas-Torres, "Review of photovoltaic power forecasting," *Solar Energy*, vol. 136, pp. 78–111, Oct. 2016, doi: 10.1016/j.solener.2016.06.069.

[6] S. E. Haupt, P. A. Jiménez, J. A. Lee, and B. Kosović, "Principles of meteorology and numerical weather prediction," in *Renewable Energy Forecasting*, Elsevier, 2017, pp. 3–28. doi: 10.1016/B978-0-08-100504-0.00001-9.

[7] S. Al-Dahidi *et al.*, "Forecasting solar photovoltaic power production: A comprehensive review and innovative data-driven modeling framework," *Energies*, vol. 17, no. 16, p. 4145, Aug. 2024, doi: 10.3390/en17164145.

[8] M. Abdelsattar, A. AbdelMoety, and A. Emad-Eldeen, "Advanced machine learning techniques for predicting power generation and fault detection in solar photovoltaic systems," *Neural Computing and Applications*, vol. 37, no. 15, pp. 8825–8844, May 2025, doi: 10.1007/s00521-025-11035-6.

[9] I. D. Mienye, T. G. Swart, and G. Obaido, "Recurrent neural networks: a comprehensive review of architectures, variants, and applications," *Information*, vol. 15, no. 9, p. 517, Aug. 2024, doi: 10.3390/info15090517.

[10] E. Sarmas, E. Spiliotis, E. Stamatopoulos, V. Marinakis, and H. Doukas, "Short-term photovoltaic power forecasting using meta-learning and numerical weather prediction independent Long Short-Term Memory models," *Renewable Energy*, vol. 216, p. 118997, Nov. 2023, doi: 10.1016/j.renene.2023.118997.

[11] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997, doi: 10.1162/neco.1997.9.8.1735.

[12] J. Wang, J. Xu, and X. Wang, "Combination of hyperband and Bayesian optimization for hyperparameter optimization in deep learning," *arXiv*, Jan. 2018, [Online]. Available: http://arxiv.org/abs/1801.01596

[13] D. Yang, J. Kleissl, C. A. Gueymard, H. T. C. Pedro, and C. F. M. Coimbra, "History and trends in solar irradiance and PV power forecasting: A preliminary assessment and review using text mining," *Solar Energy*, vol. 168, 2018, doi: 10.1016/j.solener.2017.11.023.

[14] J. Zhang, R. Verschae, S. Nobuhara, and J.-F. Lalonde, "Deep photovoltaic nowcasting," *Solar Energy*, vol. 176, pp. 267–276, Dec. 2018, doi: 10.1016/j.solener.2018.10.024.

[15] M. P. Almeida, M. Muñoz, I. de la Parra, and O. Perpiñán, "Comparative study of PV power forecast using parametric and nonparametric PV models," *Solar Energy*, vol. 155, pp. 854–866, Oct. 2017, doi: 10.1016/j.solener.2017.07.032.

[16] V. Arumugam and V. Natarajan, "Time series modeling and forecasting using autoregressive integrated moving average and seasonal autoregressive integrated moving average models," *Instrumentation Mesure Métrologie*, vol. 22, no. 4, pp. 161–168, Aug. 2023, doi: 10.18280/i2m.220404.

[17] D. Markovics and M. J. Mayer, "Comparison of machine learning methods for photovoltaic power forecasting based on numerical weather prediction," *Renewable and Sustainable Energy Reviews*, vol. 161, p. 112364, Jun. 2022, doi: 10.1016/j.rser.2022.112364.

[18] M. J. Mayer, "Benefits of physical and machine learning hybridization for photovoltaic power forecasting," *Renewable and Sustainable Energy Reviews*, vol. 168, p. 112772, Oct. 2022, doi: 10.1016/j.rser.2022.112772.

[19] J. Gaboitaolelwe, A. M. Zungeru, A. Yahya, C. K. Lebekwe, D. N. Vinod, and A. O. Salau, "Machine learning based solar photovoltaic power forecasting: a review and comparison," *IEEE Access*, vol. 11, 2023, doi: 10.1109/ACCESS.2023.3270041.

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

[21] J. Polo, N. Martín-Chivelet, M. Alonso-Abella, C. Sanz-Saiz, J. Cuenca, and M. de la Cruz, "Exploring the PV power forecasting at building façades using gradient boosting methods," *Energies*, vol. 16, no. 3, p. 1495, Feb. 2023, doi: 10.3390/en16031495.

[22] S. M. Al-Selwi *et al.*, "RNN-LSTM: From applications to modeling techniques and beyond—Systematic review," *Journal of King Saud University - Computer and Information Sciences*, vol. 36, no. 5, p. 102068, Jun. 2024, doi: 10.1016/j.jksuci.2024.102068.

[23] M. Aslam, S.-J. Lee, S.-H. Khang, and S. Hong, "Two-stage attention over LSTM with Bayesian optimization for day-ahead solar power forecasting," *IEEE Access*, vol. 9, pp. 107387–107398, 2021, doi: 10.1109/ACCESS.2021.3100105.

[24] S. Malakar *et al.*, "Deep-learning-based adaptive model for solar forecasting using clustering," *Energies*, vol. 15, no. 10, p. 3568, May 2022, doi: 10.3390/en15103568.

[25] Y. Yu, J. Cao, and J. Zhu, "An LSTM short-term solar irradiance forecasting under complicated weather conditions," *IEEE Access*, vol. 7, pp. 145651–145666, 2019, doi: 10.1109/ACCESS.2019.2946057.

[26] R. Turner *et al.*, "Bayesian optimization is superior to random search for machine learning hyperparameter tuning: Analysis of the Black-Box Optimization Challenge 2020," *Proceedings of Machine Learning Research*, vol. 133, pp. 3–26, 2020.

[27] J. Sumithra, J. C. Vinitha, M. J. Suganya, M. Anuradha, P. Sivakumar, and R. Balaji, "Machine learning techniques for solar energy generation prediction in photovoltaic systems," *International Journal of Power Electronics and Drive Systems (IJPEDS)*, vol. 16, no. 3, p. 2055, Sep. 2025, doi: 10.11591/ijpeds.v16.i3.pp2055-2062.

[28] Z. Hou and Y. Zhang, "A novel photovoltaic power forecasting approach using RF, LSTM with VMD decomposition," 2024. doi: 10.2139/ssrn.4831133.

[29] H. T. C. Pedro and C. F. M. Coimbra, "Assessment of forecasting techniques for solar power production with no exogenous inputs," *Solar Energy*, vol. 86, no. 7, pp. 2017–2028, Jul. 2012, doi: 10.1016/j.solener.2012.04.004.

[30] A. Mellit and S. A. Kalogirou, "Artificial intelligence techniques for photovoltaic applications: A review," *Progress in Energy and Combustion Science*, vol. 34, no. 5, pp. 574–632, Oct. 2008, doi: 10.1016/j.pecs.2008.01.001.

[31] M. F. Alsarraj, Y. E. Ahmed, S. B. Ezzat, and J. Pasupuleti, "Modeling and optimization of solar panel cooling using machine learning and deep learning techniques," *NTU Journal of Renewable Energy*, vol. 10, no. 1, pp. 1–10, 2026, doi: 10.56286/zczemr04.

[32] N. Waleed, R. K. Antar, and A. A. A. Al-Karakchi, "Solar power load energy management system – A review," in *1st International Conference on Sustainable Development Techniques (ICSDT2022) - AIP Conference Proceedings*, 2023, p. 020017. doi: 10.1063/5.0171732.

[33] N. Waleed, A. A. A. Al-Karakchi, and R. K. Antar, "Design and power load management of PV-battery-grid system," in *2nd International Conference on Renewable Energy (ICRE2022) - AIP Conference Proceedings*, 2024, p. 020022. doi: 10.1063/5.0181714.

BIOGRAPHIES OF AUTHORS

Enas Ali Ahmed is a lecturer in the Department of Medical Instrumentation Techniques Engineering, Northern Technical University, Mosul, Iraq. She received B.Eng. and M.Eng. degrees in computer engineering from the University of Mosul and Northern Technical University in 2014 and 2019, respectively. Her research interests include digital design, machine learning, power electronics, renewable energy, FPGA applications, embedded systems, artificial intelligence, and intelligent control. She can be contacted at email: enasali@ntu.edu.iq.

Muna Hassan Hussein is a lecturer in the Department of Medical Instrumentation Techniques Engineering, Northern Technical University, Mosul, Iraq. She received B.Eng. and M.Eng. degrees in computer engineering from the University of Mosul in 2013 and 2022, respectively. Her research interests include power electronics, renewable-energy systems, power-system protection, machine learning, and control. She can be contacted at email: mti.lec306.Muna@ntu.edu.iq.

Abdulkreem Mohammed Salih is an assistant professor at Al-Dour Technical Institute, Northern Technical University, Iraq. He received B.Sc. and M.Sc. degrees in computer engineering from the University of Mosul in 2008 and 2010, respectively. His research interests include image processing, digital watermarking, signal processing, and intelligent systems. He can be contacted at email: abdulkreem86@ntu.edu.iq.