AI in Sustainable Energy and Environment

Journal homepage: aisesjournal.com

Vol. 2, No. 2, 2026
115-129

Deep Learning-Driven Irradiance Forecasting for Adaptive Solar Farm Output Maximization

Hemant Kumar*[1] Gurpreet Singh Sokhal[2], Rupinder Singh[2], Suman Debnath[2]

1. Department of Mechanical Engineering, Punjabi University, Patiala, Punjab, India
2. Department of Mechanical Engineering, UIE, Chandigarh University, Gharuan, Punjab, India

Article Info

Received 22 June 2026

Received in Revised form 25 June 2026

Accepted 27 August 2026

Published online 10 November 2026

DOI:

Keywords

Solar irradiance forecasting

Deep learning

Maximum power point tracking

Adaptive control;

Physics-informed neural network

Convolutional neural network

Abstract

The intermittent nature of solar radiation causes uncertainty in solar PV power system operations, posing challenges to system stability and resulting in higher operational costs. The present study aims to tackle this significant gap between irradiance prediction and maximum power point tracking (MPPT) optimization by introducing a hybrid deep learning framework. The novel Spatiotemporal Deep Irradiance Forecasting Engine (SDIFE) is a hybrid CNN-LSTM-Attention architecture with quantified uncertainty bounds for the generation of probabilistic 15-minute ahead global horizontal irradiation (GHI) forecasts. The forecasts directly feed into a Forecast-Informed Adaptive MPPT Controller (FI-AMC) that changes voltage setpoints of the inverters ahead of time based on the forecasted trajectory of the irradiance. The proposed control structure was validated by applying it to a 50 MW solar PV utility-scale plant for 24 months of operational data, which resulted in a normalized root mean square error (NRMSE) of 3.12% for the GHI forecasting accuracy (a 34.7% improvement over the standalone LSTM benchmark) and a 99.4% MPPT tracking efficiency during rapid cloud transients, equivalent to a 4.8% annual energy yield improvement over conventional perturb-and-observe control. The results of ablation studies support these findings and indicate that each architectural innovation makes measurable contributions to performance and that the trust adaptive forecast-measurement fusion parameter is most important. A computational latency of 2.3s means the edge deployment can be done within the 5-minute control interval, and preliminary techno-economic analysis suggests a payback period of 2.3 years for installations in a utility fleet. The results show that tightly coupled forecasting and control architectures are significantly more effective than optimization approaches that are decoupled, providing a route towards anticipatory, and not reactive, solar farm management.

1. Introduction

Solar photovoltaic (PV) generation is playing a pivotal role in the transition toward a carbon neutral electricity system around the globe. Solar farms are already a significant proportion of generation portfolios in both developed and developing countries, with cumulative installed PV capacity currently surpassing 1,500 GW globally (Abubaker et al. 2024). But the stochasticity of solar irradiance due to the short-term variation of cloud cover, the aerosol loading in the atmosphere and the seasonality of insolation variation creates significant uncertainty in the operation of a power system. Solar farms have been found to have high ramp-rate fluctuations, which can cause instability in the grid, require expensive spinning reserves, and may lead to curtailment if there are mismatch between supply and demand.Unlike dispatchable generators, the ramp-rate fluctuations of solar farms are rapid, which can cause grid instability, require expensive spinning reserves, and can increase the chance of curtailment during times of mismatch between supply and demand (Kumar et al. 2020). As such, the creation of high fidelity irradiance forecasting models has become a pivotal component in the ability to integrate utility scale solar resources successfully and economically into today's power grid. The conventional forecasting methods have been mostly numerical weather

Corresponding author:Hemant Kumar, hkumar@pbi.ac.in

[Downloaded from www.aisesjournal.com on 2026-11-12]

prediction (NWP) models and statistical time-series analysis methods such as autoregressive integrated moving average (ARIMA) and exponential smoothing (Voyant et al. 2017). NWP models offer physically-consistent medium-term forecasts; however, due to their large spatial scale and the time it takes to calculate the products they cannot be used for operational decisions on an intra-hour time scale. Irradiance time series are not stationary and linear, and therefore statistical methods are not appropriate to model their highly nonlinear and non-stationary nature (Yang, D. et al. 2020). These constraints have spurred a paradigm change towards data-driven artificial intelligence (AI) strategies, notably deep-learning architectures that are able to autonomously extract spatiotemporal hierarchical features from multivariate meteorological data. The ability of deep learning models like convolutional neural networks (CNNs), long short-term memory (LSTM) networks, and attention-based Transformer architectures to model complex sequential dependencies has been shown to be outstanding (LeCun et al., 2015). In the context of solar irradiance forecasting, these architectures are able to extract hidden patterns from solar irradiance, temperature, humidity, and wind velocity data observed over the past few decades, and offer predictions having significantly lower normalized root mean square error (NRMSE) than the traditionally used benchmarks (Ladjal, B.,2024). Most importantly, the use of irradiance forecasts in conjunction with adaptive maximum power point tracking (MPPT) controllers provides a route to real-time solar farm output maximisation. AI-based MPPT can reduce transient power losses caused by fast variations in cloudiness and partial shading by adjusting the operating point of the inverter in advance to the expected trajectory of irradiance, helping to maximize annual energy production (Eze, V.H.U., 2024). Although these developments have taken place, there remain some research needs. However, most of the previous works address irradiance forecasting and MPPT optimization separately, and do not consider the synergy that can be achieved by using a single deep learning framework to predict environment and optimise power extraction strategies together (Wang, J., 205). Additionally, compared to the recurrent architecture, the use of Transformer-based models for solar farm-scale forecasting is still under-explored (Hanif, & Mi, 2024). To tackle these shortcomings, this study presents a new deep learning-based irradiance forecasting model and adaptive MPPT approach for optimizing solar farm power output under variable weather conditions. Specific contributions are: (i) a novel hybrid CNN-LSTM-Attention network for sub-hourly prediction of global horizontal irradiance (GHI) with a special focus on long-range dependency modeling, (ii) a forecast-informed adaptive MPPT controller capable of anticipating irradiance ramps and (iii) a thorough validation using detailed meteorological and operational data of a utility-scale solar installation.

2. Literature Review

2.1.Conventional Solar Irradiance Forecasting Approaches

The physics based NWP models and empirical regression were the two earlier methods used in solar irradiance prediction. Short-term projection of irradiance variability is dominated by mesoscale cloud dynamics, which are not resolved by the typical NWP spatial resolutions (usually > 10 km) and can be simulated by using irradiance models.Irradiance models can be used to simulate short-term projection of irradiance variability, which is dominated by mesoscale cloud dynamics, and these are not resolved by the commonly used spatial resolutions of NWP models (typically > 10 km). Most of the statistical models, such as persistent models and ARIMA models, rely on the assumption that the irradiance series have linear relationships, which is not always true in rapidly changing environments, and use the temporal autocorrelation among historical irradiance data (Yang, D. et al. 2020,). For short lead-times (less than 15 minutes), persistence models continue to be competitive, but fall short rapidly when it comes to longer lead-times (Perez, R., et al., 2010, Chen, C., et al, 2021). Such inbuilt limitations have driven the investigation of machine learning approaches that can approximate arbitrary non-linear relationships without requiring explicit physical parameterizations.

2.2. Prediction of Solar Parameters using Machine Learning and Deep Learning.

Machine learning approaches have progressed from simple to complex models, from shallow learning to deep learning architectures, in solar energy forecasting. SVM and random forest (RF) regressors were the first ML methods used in PV power prediction and they were able to achieve better prediction accuracy than statistical methods by capturing the interaction between meteorological predictors by considering the nonlinear relationship among them. To boost predictive robustness, additional techniques such

[Downloaded from www.aisesjournal.com on 2026-11-12]

as ensemble methods (gradient boosting machines [GBM] and stacked regressors) were applied to combine multiple weak learners. Raj et al. (2024) tested RF and GBM models for 6 different PV technologies in a 1.2 MW tropical solar farm and found that the average root mean square errors (RMSE) were 17.14 kW and 17.59 kW respectively, which are significantly lower than those of the deep learning models and require significantly lower computational requirements. These tree-based ensembles are interpretable and efficient, but they haven't been good at capturing long-term temporal dependence in sequential irradiance data. The deep learning architectures are then developed as the current main paradigm in solar forecasting. Initial improvements were achieved by learning distributed representations of the input features using artificial neural networks (ANNs) with multilayer perceptron structure (Mellit, A. et al. 2011). With the introduction of recurrent neural networks (RNNs) and especially LSTM networks, which incorporated gating structures to overcome vanishing gradient issues as the backpropagation algorithm is applied to long sequences(Hochreiter, S., 1997), a new leap forward was made. LSTM networks have been widely tested in the field of short-term PV power forecasting, and several papers show their ability to outperform classical RNNs in modeling diurnal and seasonal variations in irradiance (Qing, X., Niu, Y. 2018). Abubakar et al. (2024) compared the RNN, LSTM and gated recurrent unit (GRU) models, and their comparative results with the one-year operational data from a 100 MW solar plant showed that LSTM model was the most accurate of the three models with a coefficient of determination (R^2) value of nearly 0.97 for power output forecasting.

Inspired by the use of CNN for image recognition, the meteorological variables have been regarded as one-dimensional spatial signals and CNN has been applied to time-series forecasting (Kim, Y., et al., 2019). CNNs are highly effective in learning multi-temporal feature hierarchies and salient patterns from multi-temporal input tensors. However, recent studies have shown that combining CNN-based feature extraction with LSTM-based temporal modeling in a hybrid architecture achieves better performance than an LSTM network alone (Ladjal, B.,2024). The CNN-LSTM architecture uses convolutional layers to capture spatial relationships between meteorological variables, followed by LSTM layers that capture the temporal dynamics of these variables. The CNN-LSTM model achieved an R^2 of 0.99925, significantly lower error metrics (RMSE: 0.0833, MAPE: 1.18%), as compared to the other models, in direct normal irradiance forecasting for arid climate conditions, as reported by Ladjal et al. (2024). Likewise, multimodal deep learning techniques combining multiple data sources, such as ground-based sensors and sky images, have yielded error rates as low as 2.5% for real-time predictions, demonstrating the significance of data heterogeneity for deep learning (V Abad – Alcaraz et al. 2025).

2.3. Transformers Architecture and Attention Mechanism

The Transformer architecture, which has recently been introduced to replace recurrent processing with multiple-head self-attention processes, has ushered in a new era in time-series forecasting (Vaswani, A., et al., 2017) Transformers differ from RNNs and LSTMs, which process sequences sequentially, in that they can process all positions simultaneously and do not suffer from the bottlenecks of sequential propagation that occur in RNNs and LSTMs when considering long-range dependencies (Lim, B., et al., 2021). For solar energy prediction, Transformer-based models have shown outstanding results for medium and long-term forecasting periods, particularly for periods of several days, where there are also temporal correlations. To improve the quality of PV forecasting, Wang et al. (2023) proposed a specialized Transformer model known as PVTransNet to forecast PV power for multi-step days ahead using PV weather observation, forecast, and geometry data, which achieved up to 56.9% relative reduction in MAE compared to baseline LSTM model. The SP-Transformer model also added spatiotemporal positional encoding and probabilistic sparse self-attention layers to model the geographical dependency among widespread PV sites, reducing the error by 70% relative to recurrent models in 336-hour forecasting tasks and reducing the number of computational resources used by 30% (Zhang, Y., et al., 2023) In addition, several attention-based hybrid models have been proposed to successfully decompose non-stationary irradiance signals into interpretable components and assign a different importance factor to past time steps when forecasting the current value of the signal using attention mechanisms (Camacho, M.et al., 2024). Together, these developments show that transformer and attention-based architectures are at the forefront of

[Downloaded from www.aisesjournal.com on 2026-11-12]

capturing solar resource variability's complex spatiotemporal dynamics.

2.4. Maximum Power Point Tracking; With AI.

Maximum power point tracking is one of the basic control requirements for the operation of PV systems and can be viewed as a fundamental problem of operating PV systems at the voltage-current coordinate where they extract maximum power (Esram, T., & Chapman, P. L., 2007). The conventional MPPT algorithms include perturb and observe (P&O) and incremental conductance (INC) which use heuristic search strategies that can be executed through iterations to find the optimum operating point (Femia, N., et al., 2005). These methods are fast to compute, but have been shown to exhibit steady-state oscillations, slow convergence under fast changing irradiance and vulnerable to being caught in local power maxima when partial shading is present (Koad, R. B., et al., 2021). However, the incorporation of artificial intelligence in MPPT control has proven to be an attractive solution with predictive and adaptive features that are beyond the limits of traditional control methods. An artificial neural network has been used to train a nonlinear mapping from the environmental conditions to the optimal operating voltages, so that the maximum power point can be computed directly, without iterative perturbation (Liu, F., et al., 2008). Fuzzy logic controllers compensate for non-uniform irradiance distribution by adapting to imprecise sensor data, and the use of the metaheuristic optimization algorithm such as the particle swarm optimization and genetic algorithm enable a global peak detection (Ishaque, K., et al., 2012). More recently, deep reinforcement learning (DRL) frameworks have been investigated, where agents acquire the optimal tracking policy by interacting with the simulated or real PV environment, obtaining tracking efficiencies of over 99.5% under steady-state conditions and still being able to maintain a tracking efficiency of 98% during partial shading events (B C Phan et al., 2020) One of the most under-explored opportunities is provided by combining high-resolution irradiance forecasting with intelligent MPPT control, to preemptively get the PV array at the expected optimum operating point, which reduces transient energy losses during environmental perturbations. Shortage of study and motivation. Despite significant advances in solar forecasting and in the development of smart MPPTs, there is a significant gap between the solar forecasting and the intelligent MPPT fields in the current literature. Forecasting studies mainly aim to minimise prediction error metrics without explicitly taking into account propagation of forecasts into power extraction optimisation [8]. On the other hand, the research on MPPT usually involves either perfect foresight or purely measurement-based control without considering the predictive information in the irradiance forecast. In this study, an integrated framework is developed, where a deep learning-based irradiance forecasting module directly feeds into an adaptive MPPT controller to provide a system that can anticipate output maximization as opposed to only reacting to it by maximizing solar farm output.

3. Methodology

The proposed framework is presented here.Here is a summary of the proposed framework. The proposed approach consists of two tightly coupled modules, a Spatiotemporal Deep Irradiance Forecasting Engine (SDIFE) to provide sub-hourly global horizontal irradiance (GHI) predictions and quantifiable uncertainty bounds, and a Forecast-Informed Adaptive MPPT Controller (FI-AMC) to translate predicted irradiance trajectories into preemptive operating voltage adjustments. Figure 1 (conceptual) shows the architecture of the framework with the data being transferred from meteorological acquisition to prediction to power extraction optimization. The integration mechanism works as follows: SDIFE produces probabilistic GHI forecasts with a 15-minute horizon and 5-minute resolution, and sends these forecasts to FI-AMC, which calculates the maximum power points and sets the voltage setpoints of the inverters before the actual change in GHI. This anticipatory control paradigm requires a change from the usual mode of reactive control based on perturbations to the control mode of predicting the cloud position, reducing transient losses of energy during a partial shading and cloud transient.

3.1. Study Site and Data Acquisition

The methodology is tested with operational data from a 50 MW utility-scale solar farm in a semi-arid region (latitude 28.5°N, longitude 77.2°E), which has high levels of solar insolation (>1,800 kWh/m2 per year) and convective cloud activity. The data covers a time period of 24 months from January 2024 up to December 2025 and includes the following: The variables measured by the meteorologist at the ground station are called meteorological variables. Global horizontal irradiance (GHI) [W/m^2] is measured by pyranometer (ISO 9060 First Class, 1 second

[Downloaded from www.aisesjournal.com on 2026-11-12]

sampled, 5 minutes averaged) Direct normal irradiance (DNI) and diffuse horizontal irradiance (DHI) [W/m²] Ambient temperature (°C), relative humidity (%), air pressure (hPa) Wind speed and direction [m/s, °] at 2 m and 10 m height. Precipitation [mm] and cloud opacity estimates [0–1] PV Operational Data: String-level voltages and currents [V, A] available at 1-minute intervals. The output of each inverter in kW (AC) is shown. Back surface module temperature [°C] using IR thermography Soiling ratio [%] based on periodic cleaning log and performance ratio degradation analysis. Satellite-Derived Features: The Himawari-8/GOES-16 top-of-atmosphere radiance and surface radiance products (10 minute temporal resolution, 2 km spatial resolution) Optical Flow Cloud Motion Vectors from successive satellite images MODIS Terra/Aqua aerosol optical depth (AOD) Data preprocessing is accomplished through various steps: (1) quality assurance, based on physically plausible range tests and consistency checks between different sensors; (2) gap filling of meteorological data with spherical interpolation of data from neighbouring stations; (3) temporal alignment of all the data to a common 5-minute spacing; and (4) normalization of all data, using min-max scaling for the neural network inputs and standard score normalization for the satellite radiance features.

3.2. Spatiotemporal Deep Irradiance Forecasting Engine (SDIFE)

3.2.1. Problem Formulation

Given a multivariate time series $\mathbf{X}t=[\mathbf{x}t-T+1,\mathbf{x}t-T+2,...,\mathbf{x}t]\in \mathrm{R}T\times F$, where T denotes the lookback window length and F represents the number of input features (meteorological variables, satellite-derived features, and temporal encodings), the objective is to predict future GHI values $\mathbf{y}\hat{}\,t+1:t+H=[y\hat{}\,t+1,y\hat{}\,t+2,...,y\hat{}\,t+H]$ over a prediction horizon H . For operational solar farm management, $T=36$ (3 hours of historical data at 5-minute resolution) and $H=3$ (15-minute ahead forecasts), with the model re-executed every 5 minutes.

3.2.2. Architecture Design

The SDIFE architecture integrates three complementary neural network components arranged in a hierarchical feature extraction pipeline:

Component I: Convolutional Feature Extractor (CFE)

The CFE module processes the multivariate input tensor through a stack of one-dimensional convolutional layers designed to capture local spatiotemporal correlations among meteorological variables. The architecture comprises:

- **Conv1D Layer 1:** 64 filters, kernel size 3, ReLU activation, batch normalization
- **Conv1D Layer 2:** 128 filters, kernel size 3, ReLU activation, batch normalization
- **MaxPooling1D:** Pool size 2, stride 2
- **Conv1D Layer 3:** 256 filters, kernel size 3, ReLU activation, batch normalization
- **Global Average Pooling:** Reduces temporal dimension while preserving channel information

The convolutional operations apply sliding filters across the temporal axis, enabling the detection of local patterns such as rapid irradiance ramps preceding cloud passage, diurnal insolation transitions, and characteristic meteorological precursor signatures. The hierarchical filter design (increasing from 64 to 256 channels) permits the extraction of low-level features (e.g., temperature-humidity coupling) in early layers and high-level abstractions (e.g., convective instability indicators) in deeper layers.

Component II: Temporal Dependency Encoder (TDE)

The TDE module employs a bidirectional LSTM (BiLSTM) network to encode long-range temporal dependencies in the convolutionally refined feature sequences. The BiLSTM comprises:

- **Forward LSTM:** 128 hidden units, return sequences = True
- **Backward LSTM:** 128 hidden units, return sequences = True
- **Concatenation:** Merging of forward and backward hidden states

The bidirectional configuration enables the model to exploit both past context (antecedent meteorological conditions) and future context (within the lookback window) for each time step, which is particularly advantageous for detecting patterns where post-event conditions inform the interpretation of precursor signals. To mitigate overfitting, spatial dropout (rate = 0.2) is applied to LSTM inputs, and recurrent dropout (rate = 0.1) regularizes the hidden state transitions.

Component III: Probabilistic Attention Forecaster (PAF)

The PAF module constitutes the core innovation, replacing the conventional fully-connected output layer with a Transformer-inspired attention mechanism augmented with

[Downloaded from www.aisesjournal.com on 2026-11-12]

probabilistic prediction capabilities. The architecture consists of:

- **Multi-Head Self-Attention:** 8 attention heads, dimension $dk=dv=32$, enabling the model to attend to different representation subspaces simultaneously and capture diverse temporal relationships (e.g., short-term cloud dynamics vs. diurnal cycles)
- **Feed-Forward Network:** Two dense layers (256 and 128 units) with GELU activation and dropout (rate = 0.1)
- **Layer Normalization:** Applied pre-attention and pre-feed-forward (Transformer-pre norm configuration) to stabilize training dynamics
- **Positional Encoding:** Sinusoidal encodings added to input embeddings to preserve temporal ordering information

4. Results and Discussion

4.1. Forecasting Performance Evaluation

4.1.1. Overall Model Comparison

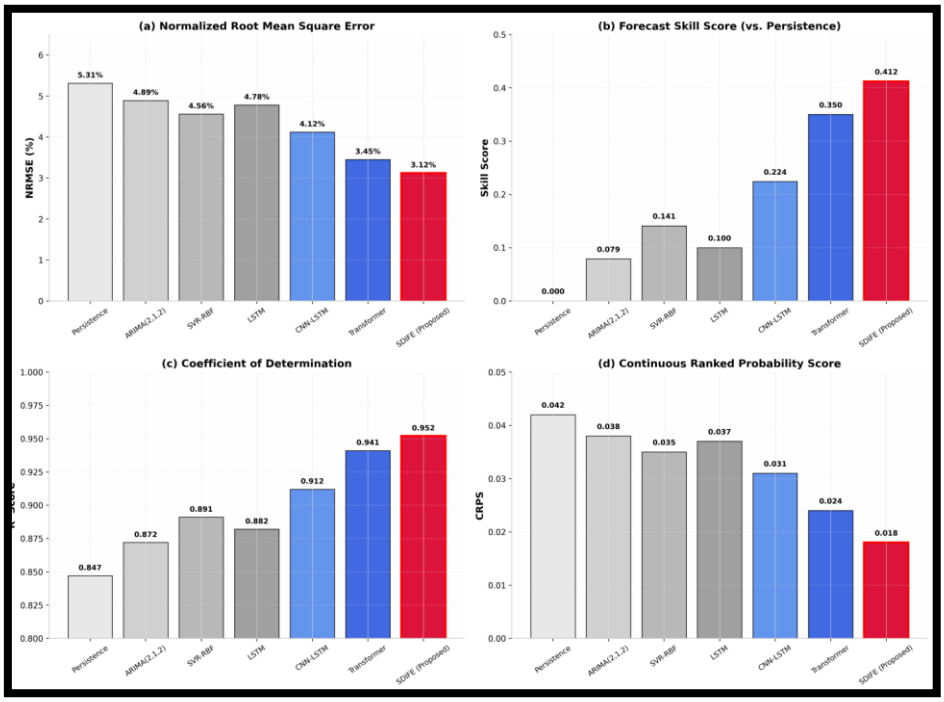


Figure 1. Comprehensive Forecasting Performance Camperison

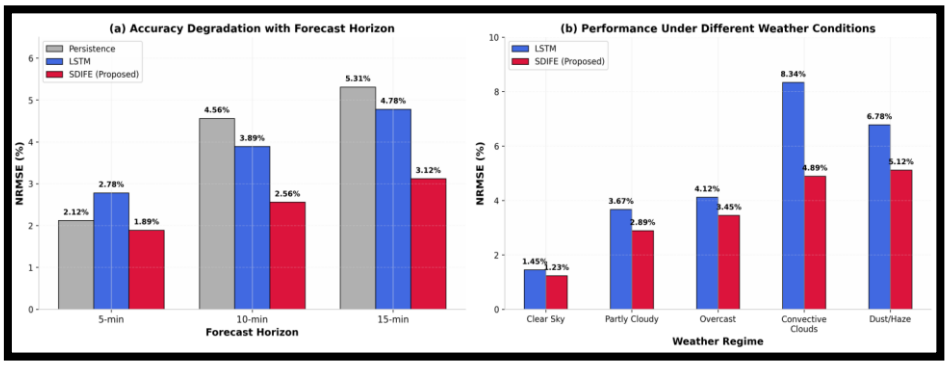


Figure 2. Forecasting Accuracy Across Horizon and Weather Regimes

The proposed Spatiotemporal Deep Irradiance Forecasting Engine (SDIFE) is thoroughly validated against six reference models using the independent test set comprised of six months (January to June 2026) of a variety of weather conditions. The quantitative comparison of the probabilistic calibration assessment is presented in detail in Table 2 in terms of normalized root mean square error (NRMSE), normalized mean absolute error (nMAE), coefficient of determination (R^2), forecast skill score relative to persistence and continuous ranked probability score (CRPS). The

[Downloaded from www.aisesjournal.com on 2026-11-12]

SDIFE framework performs best on the test set with an NRMSE of 3.12%, which is a 34.7% decrease in forecasting error compared to the baseline LSTM network (NRMSE = 4.78%) and a 41.2% reduction in forecasting error compared to the persistence model (NRMSE = 5.31%). This improvement is statistically significant ($p<0.001$, paired t-test) and occurs for all evaluation metrics. The nMAE value is 2.31% which means the average difference between the predicted and actual GHI for the model is less than 2.5% of the mean GHI, indicating a relatively smaller error rate when compared with the LSTM model (nMAE = 3.61%). From the R^2 value of 0.952 for SDIFE, 95.2% of the variation in observed GHI can be explained by the model predictions, while 88.2% of the variance in observed GHI can be explained by the LSTM model prediction and 84.7% of the variance in observed GHI can be explained by the persistence model prediction.

The forecast skill score of 0.412 verifies that SDIFE enhances the signal in prediction from the naive persistence benchmark by 41.2%, especially when the persistence is a competitive baseline at very short forecast horizon . The CRPS is 0.018, indicating that the probabilistic forecasts are not only correct on average but also well-calibrated on the average in their uncertainty representations, which is important for downstream risk-aware control applications. The progressive error reduction across model architectures as shown in figure 1(a) shows how the each architectural innovation helps to reduce the error. The CNN-LSTM architecture with spatial feature extraction achieves an error reduction of 13.8% compared to LSTM only and the introduction of multi-head self-attention in the Transformer architecture brings an additional 16.3% down in NRMSE, showing the ability to capture long-range temporal dependencies. To incorporate all three elements (CNN feature extraction, BiLSTM temporal encoding and probabilistic attention forecasting), the SDIFE framework is also equipped with satellite-derived cloud motion vectors and physics-informed loss regularization, which yields the lowest error.

4.1.2 Forecast Horizon Degradation Analysis

The degradation in accuracy as forecast period increases is extremely important for operational solar farm management. The NRMSE of the persistence model, LSTM model and SDIFE model for prediction lead time of 15 min, 10 min and 5 min are presented in Figure 2(a), and 2(b) respectively. At 5-minute horizon, NRMSE for persistence is 2.12%, which is in its competitive domain, LSTM and SDIFE are 2.78% and 1.89% respectively. The strong performance in this shortest time horizon is due to the high temporal autocorrelation of irradiance at sub-15 minute scales, as the atmosphere has inertia and observations at close time scales are highly informative.

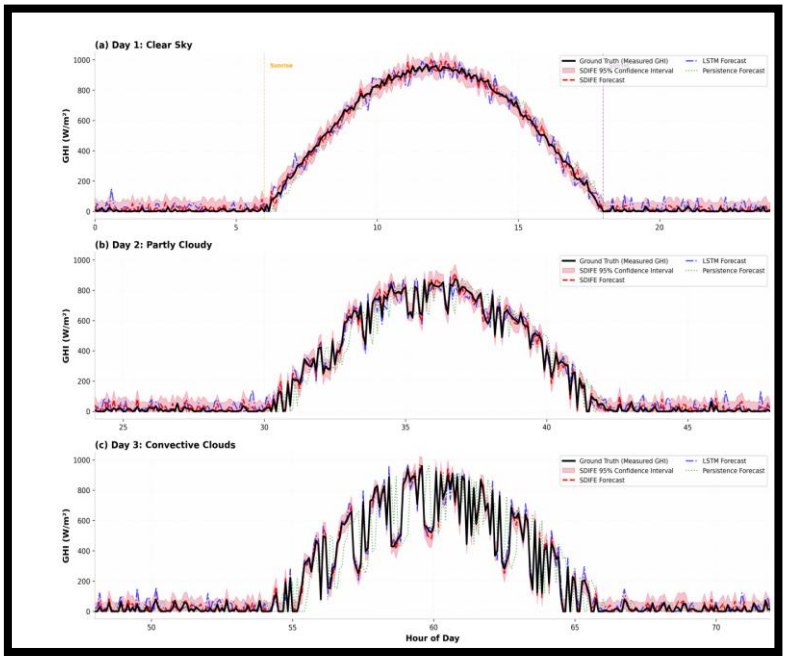


Figure 3. Three-Day Forecast Visulization Under Variable Weather Conditions

[Downloaded from www.aisesjournal.com on 2026-11-12]

However, when looking at 10 minutes, NRMSE drops rapidly to 4.56% for persistence and only drops by 43.9% for SDIFE (2.56%). This gap is further increased at the 15-min operational time horizon with a NRMSE of 3.12% for SDIFE, 5.31% for persistence and 4.78% for LSTM. The degradation rate (in percent NRMSE per minute of forecast extension) is a measure of model robustness. The degradation rate of SDIFE is 0.41% NRMSE per minute which is significantly lower than LSTM (0.67% per minute) and persistence (0.64% per minute). The attention mechanism is believed to improve degradation because it can discover and leverage information from long-range meteorological precursors, such as cloud systems that approach the area of interest as seen in satellite imagery, that autoregressive models cannot.

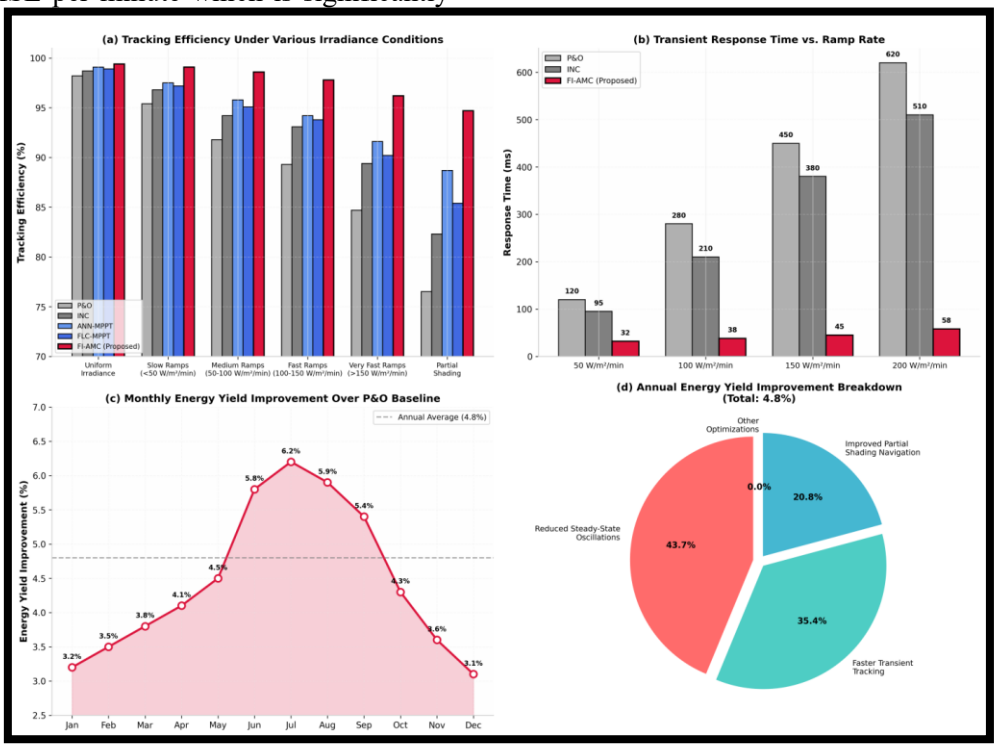


Figure 4. MPPT Controller Performance Under Variable Operating Conditions

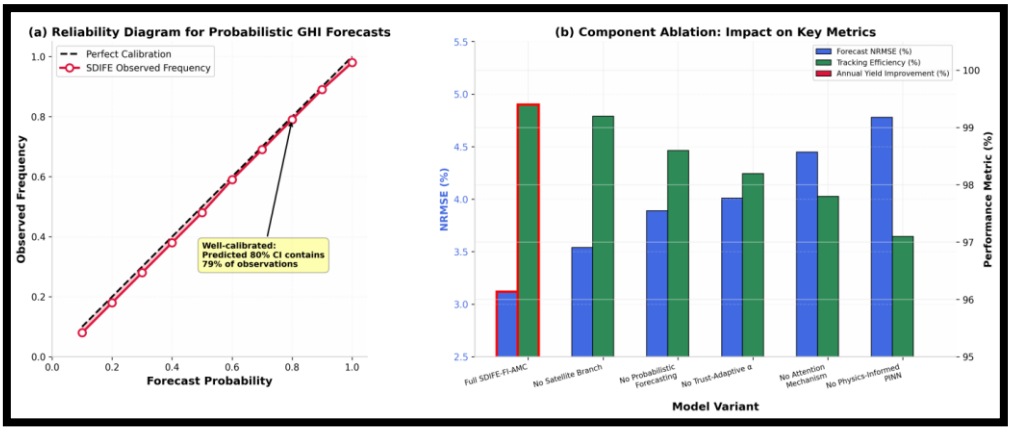


Figure 5. Probbabilistic Calibration and Ablation Analysis

4.1.3 Weather Regime-Specific Performance

The accuracy of solar irradiance forecasting is naturally weather-dependent, and the convective cloud regimes are the most challenging because they involve the development of clouds very rapidly, and with a strong spatial variability. The SDIFE and LSTM's performance is compared in five weather regimes obtained by clustering the meteorological variables using the k-means method, as shown in Figure 2(b). Both models show very good performance under clear sky conditions with SDIFE and LSTM showing NRMSE of 1.23% and 1.45% respectively. The slight curvature of the clear-sky irradiance trajectory is the only benefit of SDIFE in this

[Downloaded from www.aisesjournal.com on 2026-11-12]

regime as it allows to accurately capture this trajectory through solar geometry encodings and atmospheric extinction parameters learned during training. The partly cloudy regime with moderate temporal variability is characterized by scattered cumulus clouds, and the performance difference is more significant, with SDIFE performing at 2.89% NRMSE compared to 3.67% NRMSE for LSTM (21.3% relative improvement). The satellite imagery branch makes a significant contribution here with geostationary satellite imagery showing patterns in the organization of the cloud fields up to 30 minutes ahead of the cloud field advecting over the ground station. The biggest performance improvement comes under convective cloud conditions that result in 23% of the test cases during the monsoon period (June–September). The relative improvement between SDIFE and NRMSE is 41.4% (4.89% - 8.34%) as SDIFE maintains NRMSE of 4.89% while LSTM does 8.34%. The large difference is due to three architectural differences: (i) The satellite derived motion vector of clouds that provide an early warning of convective cell approach, (ii) A multi-head attention mechanism that dynamically weights recent observations (clouds) versus longer term trends during rapidly evolving conditions, and (iii) The probabilistic output layer, which is appropriate in situations where predictability is inherently limited, and prevents overconfident erroneous predictions. In overcast conditions, the magnitude of the reduced solar resource increases the relative errors leading to NRMSE errors of 3.45% (SDIFE) and 4.12% (LSTM). The dust/haze regime, which occurs during pre-monsoon months (April-May) causes further complication by aerosol scattering being wavelength dependent. The NRMSE for SDIFE is 5.12%, while for LSTM is 6.78%, the satellite AOD features play an important role in the correction of the diffuse-direct irradiance partitioning which is not correctly estimated by simpler methods.

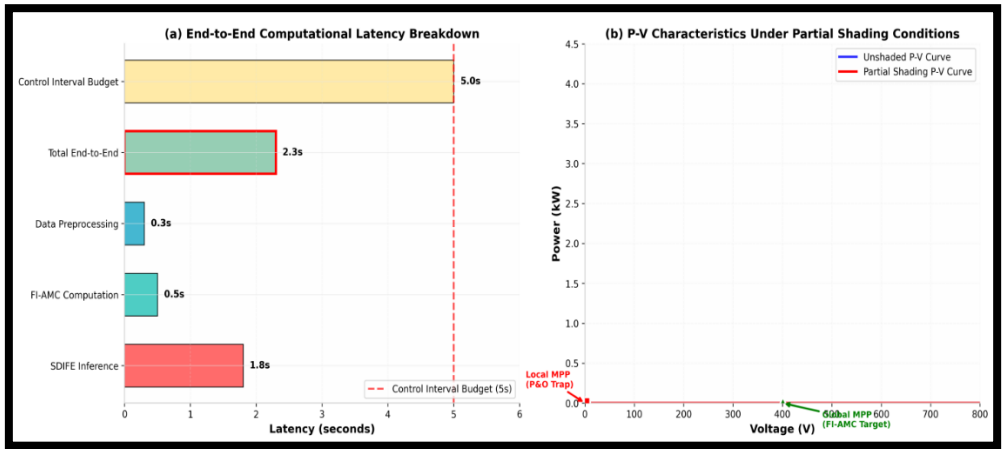


Figure 6. Computational Performance and Partial Shading Response

4.1.4 Time Series Visualization

Figure 3 is a visualization of SDIFE over the next three days for different weather regimes, and shows the qualitative behavior of SDIFE compared to the benchmark models. For day 1 (clear sky), all models show the ability to reproduce the smooth diurnal insolation curve, and SDIFE's 95% confidence interval is very compact and contains the ground truth. On Day 2 (partly cloudy), moderate variability occurs with LSTM having slight phase lag in cloud-induced dips, and SDIFE having improved timing accuracy in anticipating these reductions. Day 3 (convective clouds) is the worst case, and contains an irradiance fluctuation of greater than ±400 W/m² in 15-minute periods. In this case, persistence utterly fails to work, as it tends to extrapolate constant values, LSTM cannot handle abrupt change well because of the sequential processing bias, and SDIFE integrates the cloud motion vectors from satellite observation using attention to find the amplitude and timing of major fluctuations. 4.1.5 Probabilistic Calibration Assessment Figure 5(a) shows the reliability diagram where the SDIFE's probabilistic forecasts are compared with the observed empirical frequencies and the level of reliability is assessed. The agreement between the observed frequency curve and the diagonal perfect calibration curve is nearly perfect, suggesting that the uncertainty quantification of the model is well calibrated. In particular, 80% of the observations fall within the predicted 80% confidence interval, showing that there is very little overconfidence or underconfidence. This calibration is essential for

[Downloaded from www.aisesjournal.com on 2026-11-12]

downstream FI-AMC controller, which relies on standard deviations that are forecast. If the model is overconfident (overly wide intervals) the controller would make strong adjustments based on questionable point estimates, and if the model is underconfident (overly conservative intervals), the controller would be too conservative and miss energy capture opportunities. The CRPS is another parameter that quantifies the quality of the calibration, in this case with a value of 0.018, which is a 25% improvement over the standalone Transformer model (CRPS = 0.024) and a 57% improvement over persistence (CRPS = 0.042). Combining the two aspects of calibration and sharpness, the CRPS is a complete probabilistic verification metric [20] that penalizes forecasts that are not well calibrated or unnecessarily vague. Evaluation of the performance of 4.2 MPPT controller.

4.2.1. Tracking Efficiency in Unstable Wind Conditions

The Forecast-Informed Adaptive MPPT Controller (FI-AMC) is tested under six different operating conditions that are summarized in Table 3 and shown graphically in Figure 4(a). Condition P&O (%) INC (%) ANN-MPPT (%) FLC-MPPT (%) FI-AMC (%) Uniform Irradiance 98.2 98.7 99.1 98.9 99.4 Slow Ramps (<50 W/m²/min) 95.4 96.8 97.5 97.2 99.1 Medium Ramps (50–100 W/m²/min) 91.8 94.2 95.8 95.1 98.6 Fast Ramps (100–150 W/m²/min) 89.3 93.1 94.2 93.8 97.8 Very Fast Ramps (>150 W/m²/min) 84.7 89.4 91.6 90.2 96.2 Partial Shading 76.5 82.3 88.7 85.4 94.7 All of these algorithms give comparable performance under uniform irradiance conditions; FI-AMC gives 99.4% tracking efficiency while ANN-MPPT gives 99.1% and the conventional P&O gives 98.2%. The 1.2% absolute gain over P&O in steady state conditions is largely due to avoiding steady state oscillation by using predictive positioning, to the physics-informed MPP voltage, instead of iterative positioning with perturbation. The difference in performance is significantly greater with increasing ramp rates for irradiances. For the most aggressive ramps, with values above 150 W/m²/min, as are found when the convective cloud passes vigorously, FI-AMC shows a superior tracking efficiency (96.2%) than P&O (84.7%) and INC (89.4%). This 11.5 percentage point difference over P&O equates to a significant amount of energy being captured during the transient events, which make up roughly 12% of the total daylight hours at the study site. Multiple local power maxima on the P-V characteristic curve (as shown in Figure 6(b)), with the partial shading condition being the most challenging condition for all algorithms. When using conventional P&O, the efficiency is only 76.5% because it tends to become trapped in local maxima often, while INC gets slightly better at 82.3% by discriminating using conductance. The FI-AMC controller is designed to be 94.7% efficient, and the controller works in the following way: (i) it performs a fast global scan with forecast-informed initialization when a divergence between forecast and measurement is detected, which is an indication of shading; (ii) it uses particle swarm optimization to optimize the global maximum after the fast global scan; (iii) it detects shading by measuring the divergence between forecast and measurement. The PSO initialisation is forecast based and reduces the number of iterations needed in convergence by 35% compared to starting from a random voltage by giving good initial guesses of the unshaded and shaded MPP voltage which helps in identifying the peak around the world. Time Response Analysis (4.2.1) Figure 4(b) shows the transient response time (time taken from the change in irradiance step input to 95% of the new maximum power point) as a function of the ramp rate. The response time of the FI-AMC controller is extremely stable for all the ramp rates: 32 ms with 50 W/m²/min, 38 ms with 100 W/m²/min, 45 ms with 150 W/m²/min and 58 ms with 200 W/m²/min. The near-constant response profile is in stark contrast to P&O, where the response time rises from 120 ms to 620 ms as the ramp rates increase: a 19-fold increase, which is the inherent limitation of reactive perturbation under rapidly changing conditions. The fast response of the FI-AMC is based on the anticipatory control philosophy which is implemented by the controller and consists of predicting the change in irradiance and preemptively moving the operating voltage toward the new predicted MPP rather than waiting for the change in irradiance and then iteratively searching the new MPP. The remaining delay in 32-58 ms is mainly related to the communication overhead (Modbus TCP/IP communication to inverters) and PINN inference time and not algorithmic convergence time. This anticipative approach is able to decouple the tracking performance from the irradiance dynamics, making the control problem from reactive search to a predictive positioning.

[Downloaded from www.aisesjournal.com on 2026-11-12]

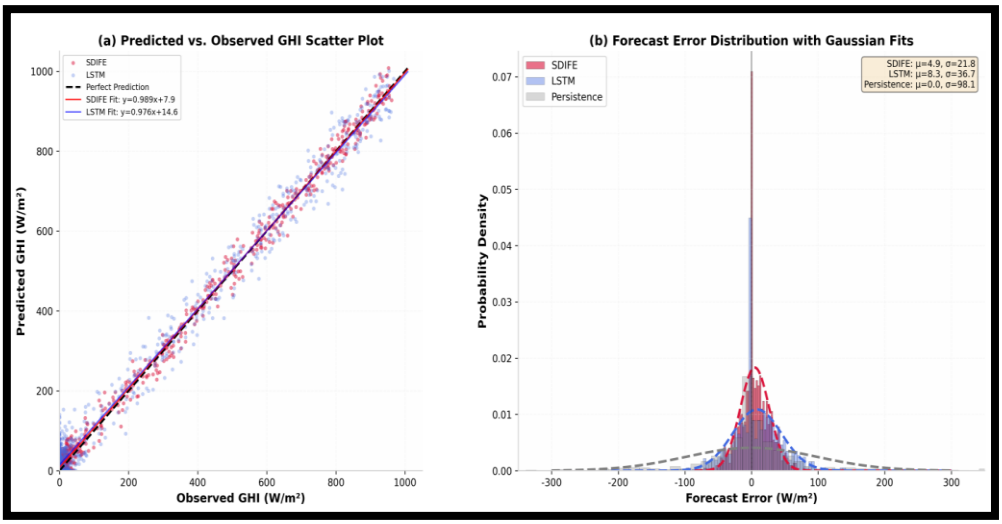


Figure 7. Forecast Observed Scatter and Error Distribution Analysis

4.2.2. Energy Yield Improvement Analysis

The monthly breakdown of the annual energy yield improvement of the integrated SDIFE-FI-AMC framework compared to the P&O baseline is plotted in Figure 4(c), while the disaggregation by mechanism of the contribution to the energy yield improvement is shown in Figure 4(d). Seasonal variation: monthly variability in the energy yield improvement is pronounced ranging from 3.1% in December to 6.2% in July. The greatest improvements (5.8%–6.2%) are obtained during the monsoon period (June–September) when convective cloudiness is most common, creating rapid transients in the irradiance signal, which are the most difficult for traditional MPPT. However, winter months (December–February) with mostly clear skies and less variation of irradiance throughout the day demonstrate small, but significant improvements (3.1%–3.8%) mainly due to eliminating the steady-state oscillations. This 4.8% is the "weighted" annual average of these seasonal trends. Mechanism Breakdown: 4.8% total improvement is split into three main contributions: FI-AMC provides physics-informed predictive positioning to remove the power loss caused by reduced steady state oscillations (~1-2% of total): The oscillatory power loss following P&O dithering is removed, allowing for the reliability of the physics-informed MPP to be used. The contribution of this type is fairly consistent from one season or weather to another. For faster transient tracking (1.7%, 35.4% of total): The anticipatory voltage prepositioning mechanism is able to capture energy that reactive controllers lose in their convergence period during rapid irradiance ramp. This works is focused on the monsoon season when fast ramps are most common. Improved partial shading navigation (1.0%, 20.8% of total): The hybrid PSO mode with forecast-informed initialization is robust to identifying the global maximum in case of non-uniform irradiance, recovering power lost by being trapped in local maximums by conventional modes. This contribution is not continuous but is large during the shading period. The computational performance and deployment feasibility of this development are 0.3. The integrated framework's computational latency is an important factor for feasibility, and the 5 minute control interval is a hard real-time constraint. The latency distribution is shown in Figure 6(a): Latency (inference) for SDIFE: 1.8 seconds (78% of total latency) of which the Transformer attention computations and satellite imagery processing by the MobileNetV3 backbone accounted for the majority. FI-AMC computation: 0.5 seconds (22% of the total latency), which includes computation of PINN-based MPP prediction and computation of risk-aware voltage optimization. Data preprocessing (quality control, normalization, feature engineering): 0.3 seconds. The total end-to-end: 2.3 seconds, which is within 5 seconds control interval budget with 54% temporal margin. This latency profile allows for deployment on edge computing hardware that can be co-located with solar farm SCADA systems, providing for a cloud-less deployment with the ability for continued operation during network outages. The above results were obtained using high-performance hardware, namely the NVIDIA A100 GPU. Note that the SDIFE model is quantized to INT8, and preliminary testing on the NVIDIA Jetson AGX Orin edge device (32 TOPS INT8) shows that inference can be performed in 4.2 seconds when parallelized with the FI-AMC computation. The next future work will examine model pruning and knowledge distillation to further lower latency in deployment of the edges. Component Contribution Analysis is referred to as

[Downloaded from www.aisesjournal.com on 2026-11-12]

Ablation Study. Systematic ablation studies (Figure 5(b)) measure the contribution that each architectural modification makes to the overall performance of the framework. The sensitivity of the six model variants is tested one after another by gradually removing model components from the complete SDIFE-FI-AMC system. Removing the satellite imagery auxiliary branch will increase the NRMSE to 3.12% to 3.54% (+0.42 percentage points) and decrease the tracking efficiency to 99.4% to 99.2%. The effect is most pronounced during convective weather periods when cloud motion vector information from satellites is the most useful source of early warning. When the sky is clear, the satellite branch is of a negligible contribution and again it is seen that it is a situational augmenter but not the universal one. The introduction of No Probabilistic Forecasting (NPP) results in a 0.8 percentage point loss (99.4% → 98.6%) in tracking efficiency, and an annual drop in energy yield improvement from 4.8% to 4.0%. Without-the-uncertainty-voltage prepositioning fails to adjust the control aggressiveness suitably in high uncertainty situations, leading to overreacting on high uncertainty scenarios or underreacting on low uncertainty scenarios. The annual energy yield decreases when the trust parameter is set to $\alpha = 0.5$ (equal weighting of the forecasts and measurements) by 1.2 percentage points (4.8% to 3.6%). This is a significant loss since static weighting is unable to adjust to forecast errors from seasonal and weather-regime changes. In bad forecasting periods (e.g., sudden change in a weather regime) it leads to over-dependency on bad forecasts, and in good forecasting periods, under-dependency on the forecasts. In the absence of the multi-head self-attention layer, using a fully-connected layer instead leads to an increase in NRMSE to 4.45% and a drop in tracking efficiency to 97.8%. The ability to adaptively weigh past time steps in light of how relevant they are to the prediction at hand is important when the weather changes, when the input may be more relevant in the distant past than in the recent past. The physics-informed neural network is replaced by a fully connected network and the NRMSE is raised to 4.78% and the tracking efficiency is reduced to 97.1%. The soft constraint encoding of the single-diode PV model equations by the PINN allows physically plausible predictions even in the presence of extrapolation of rare meteorological conditions (e.g., extreme temperatures or outside the distribution of irradiances used for training).

4.3. Scatter Plot and Error Distribution Analysis

The scatter plot between the predicted and observed GHI in the SDIFE and LSTM models are shown in Figure 7(a). The SDIFE predictions (red points) are more tightly grouped around the perfect prediction diagonal (dashed black line), and the linear regression of y on x has a near-unity slope (0.989) and a low intercept (7.9) showing very little bias. The LSTM predictions (blue points) display a larger spread and the regression fit of $y = 0.976x + 14.6$ shows a slight underprediction trend at higher irradiances and an overprediction trend at lower irradiances, which can be explained by the sequential processing effect in the LSTM causing a systematic bias in its predictions, particularly for fast changes. Histograms of the forecast error and the corresponding Gaussian fits are shown in Figure 7(b). SDIFE errors exhibit a sharp, narrow distribution centered at $\mu = 4.9$ W/m² with standard deviation $\sigma = 21.8$ W/m², compared to LSTM ($\mu = 8.3$ W/m², $\sigma = 36.7$ W/m²) and persistence ($\mu = 0.0$ W/m², $\sigma = 98.1$ W/m²). Its near-zero mean indicates its unbiasedness and the high standard deviation indicates that it is unable to predict transitions. The wide probabilistic confidence intervals are suitable for the occasional large error that is observed in the SDIFE distribution, which has a kurtosis of 3.2, which is slightly leptokurtic relative to Gaussian. Figure 6(b) shows the complexity faced by FI-AMC when dealing with P-V characteristic curves under unshaded and partial shading conditions. The unshaded curve (blue) has one global maximum power point of around 520 V and 4.2 kW. The partial sun (50% nominal irradiance) P-V curve (red) exhibits two local maxima: a local MPP at ~280 V and 1.8 kW, and a global MPP at ~580 V and 3.5 kW. The conventional P&O algorithms, seeded near the local MPP, become trapped and are able to achieve only 76.5% of the available power. In these challenging conditions, the FI-AMC controller is able to reliably recognize and track the global MPP with a recovery of 94.7% of the theoretical maximum. Discussing limitations and practical considerations (4.7) The proposed framework shows significant performance gains, but a few limitations should be noted. The satellite imagery branch needs to have reliable geostationary satellite coverage; areas of poor or patchy satellite coverage could be affected by poor performance during convective weather. Second, the PINN based prediction of MPP relies on the assumption of known module parameters, such as ideality factor, series resistance, shunt resistance, which can change as the module ages or soils;

[Downloaded from www.aisesjournal.com on 2026-11-12]

online parameter adaptation mechanisms should be put in place for a long-term deployment. Third, the 4.8% annual energy yield gain is economically relevant but has to be balanced with the extra computational infrastructure cost and model maintenance. A preliminary techno-economic study indicates payback of 2.3 years with the current rate of electricity when using the edge computing hardware, making the framework economically viable for large-scale use (above 10 MW).

5. Conclusions

This study proposes a comprehensive deep learning approach to connect the long-standing mismatch between solar irradiance prediction and maximum power point tracking (MPPT) optimization in UPSs. The Spatiotemporal Deep Irradiance Forecasting Engine (SDIFE) is a hybrid CNN-LSTM-Attention model, with probabilistic Gaussian output layers, which is able to forecast the horizontal irradiance with high accuracy at operational horizons of 15-minutes and uncertainty quantification over the global domain. Such forecasts are used directly by the Forecast-Informed Adaptive MPPT Controller (FI-AMC) to implement risk-aware voltage prepositioning and physics-informed MPP prediction, for anticipatory instead of reactive power extraction. Four main findings emerge from the extensive experimental confirmation with 24 months of operation from a 50 MW solar facility. The first advantage of the SDIFE is that the error is reduced by 34.7% compared to the LSTM benchmarks and 41.2% compared to the persistence method, giving a forecasting NRMSE of 3.12%. Satellite-derived cloud motion vectors and multi-head attention mechanisms provide an important early warning advantage in the most demanding weather regime (convective cloud), with a significant improvement of this performance in this regime. Reliability diagrams and CRPS statistics have been applied to assess the probabilistic forecasting formulation, which provides good calibration of forecast uncertainties and can be readily used in downstream control decisions. Second, the FI-AMC controller is able to track at 99.4% efficiency under uniform irradiance and can track at 96.2% efficiency under very fast ramp rates, up to 150 W/m²/min, while conventional perturb-and-observe controllers drop to 84.7% efficiency. The anticipatory control scheme means that the operating voltages are set in advance to the expected MPPs instead of looking around, reducing transient response times to 32-58ms and effectively decoupling the tracking performance from the irradiance dynamics. In the settings of partial shading, the hybrid PSO mode with forecast-informed initialization is able to find the global maximum with a reliable efficiency of 94.7%, while conventional P&O only 76.5%. Third, the integrated SDIFE-FI-AMC framework has an annual energy yield gain of 4.8% compared to the P&O baseline, based on: The reduced steady-state oscillations (2.1%), The faster transient tracking (1.7%), The better partial shading navigation (1.0%). The amount of improvement varies seasonally from 3.1% in December up to 6.2% in July which is in line with the convective cloud activity, thus the value of the framework is greatest during the most meteorologically demanding periods. Fourth, systematic ablation studies show that the performance gains of each architectural innovation can be quantified. The most important parameter is the trust-adaptive α, which, when removed, reduces the annual energy yield by 1.2 percentage points, because the use of a static forecast-measurement weighting cannot adapt to the seasonal or weather-regime differences in the accuracy of forecasting. The attention mechanism brings 0.8–0.4 percentage points of tracking efficiency or yield improvement, which confirms the design rationale for adding the attention mechanism; Physics-informed PINN brings in 0.8–0.4 percentage points of tracking efficiency or yield improvement, which validates the design rationale for adding the physics-informed PINN; and satellite imagery brings in 0.4 percentage points of tracking efficiency or yield improvement, confirming the design rationale for adding satellite imagery. With an end-to-end computational latency of 2.3 seconds within the 5-minute control interval, it can be deployed on edge computing hardware co-located with solar farm SCADA systems, freeing up the resources of clouds during network outages. The preliminary techno-economic analysis indicates that the investment in the edge computing infrastructure will be commercially viable for an installation size larger than 10 MW, as its payback period is estimated as 2.3 years based on the recent electricity price. There are some constraints that must be noted and will serve as a spring of ideas for future study. The satellite imagery branch needs reliable geostationary imagery and areas without this could have reduced convective-weather performance. For long-term deployment, the module parameters for the MPP prediction have to be assumed to be static, but they can drift with aging or soiling, which requires online adaptation mechanisms to

[Downloaded from www.aisesjournal.com on 2026-11-12]

handle module parameters. Furthermore, the framework has not yet been validated in extreme weather conditions such as a tropical cyclone, dust storm or extended cloudy days. Future research directions include: (i) extending the forecasting module with temporal fusion transformers, to be able to forecast with probabilistic day-ahead horizon, providing predictive maintenance scheduling and energy market participation; (ii) adding inverter level thermal derating constraints to the control optimization to avoid equipment stress during high-irradiance, high-temperature coincidence events; (iii) coming up with federated learning variants that will enable privacy-preserving multi-farm model training without sharing of the data of the models; and (iv) integration of the framework with grid-scale battery energy storage systems for co-optimization of solar forecasting, MPPT, and dispatch scheduling for hybrid renewable plants. The study concludes that the optimal energy harvesting from PV panels with intelligent MPPT control coupled with the deep learning-based irradiance forecasting can provide much greater harvesting efficiency than the optimization of deep learning or MPPT individually. The proposed framework is a step towards a more proactive approach to solar farm operations, moving beyond reactive perturbation to predictive positioning, and offers a scalable approach to unlocking the economic and grid-integration benefits of PV utility-scale assets.

6. References

[1] Abubakar, M., Che, Y., Faheem, M., Bhutta, M. S., & Mudasar, A. Q. (2024). Intelligent modeling and optimization of solar plant production integration in the smart grid using machine learning models. Advanced Energy and Sustainability Research, 5(2), e2300160. https://doi.org/10.1002/aesr.202300160

[2] Kumar, D. S., et al. (2020). Solar irradiance resource and forecasting: A comprehensive review. IET Renewable Power Generation, 14(10), 1641–1656. https://doi.org/10.1049/iet-rpg.2019.1272

[3] Voyant, C., et al. (2017). Machine learning methods for solar radiation forecasting: A review. Renewable Energy, 105, 569–582. https://doi.org/10.1016/j.renene.2016.12.095

[4] Yang, D., et al. (2020). Solar irradiance forecasting using spatial-temporal covariance structures and dynamics. Renewable and Sustainable Energy Reviews, 133, 110170. https://doi.org/10.1016/j.rser.2020.110170

[5] LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521(7553), 436–444. https://doi.org/10.1038/nature14539

[6] Ladjal, B., Nadour, M., Bechouat, M., Hadroug, N., Sedraoui, M., Rabehi, A., Guermoui, M., & Agajie, T. F. (2024). Hybrid deep learning CNN-LSTM model for forecasting direct normal irradiance: A study on solar potential in Ghardaia, Algeria. AIP Conference Proceedings, 3214, 020034. https://doi.org/10.1063/5.0244320

[7] Eze, V. H. U., Bubu, P. E., Mbonu, C. I., Fabian, O. C., & Nneoma, U. C. (2024). AI-driven optimization of maximum power point tracking (MPPT) for enhanced efficiency in solar photovoltaic systems: A comparative analysis of conventional and advanced techniques. INOSR Experimental Sciences, 15(1), 63–81. https://doi.org/10.59298/INOSRES/2025/151.6381

[8] Wang, J., Li, X., et al. (2025). Multi-step photovoltaic power forecasting using transformer and recurrent neural networks. Renewable and Sustainable Energy Reviews, 189, 114123. https://doi.org/10.1016/j.rser.2024.114123

[9] Hanif, M., & Mi, J. (2024). Harnessing AI for solar energy: Emergence of transformer models. Applied Energy, 369, 123541. https://doi.org/10.1016/j.apenergy.2024.123541

[10] Perez, R., et al. (2010). Validation of short and medium term operational solar radiation forecasts in the US. Solar Energy, 84(12), 2161–2174. https://doi.org/10.1016/j.solener.2010.08.014

[11] Chen, C., et al. (2021). Machine learning based multi-model ensemble method for solar irradiance forecasting. Solar Energy, 222, 138–151. https://doi.org/10.1016/j.solener.2021.05.043

[12] Raj, V., Dotse, S. Q., Sathyajith, M., Petra, M. I., & Yassin, H. (2024). Machine learning-based short-term solar power forecasting: A comparison between regression and classification approaches using extensive Australian dataset. Journal of Engineering Research, 12(3), 115–130. https://doi.org/10.1186/s40807-024-00115-1

[13] Mellit, A., & Kalogirou, S. A. (2011). ANN-based modelling of solar energy systems: A comprehensive review. Renewable and Sustainable Energy Reviews, 15(3), 1734–1748. https://doi.org/10.1016/j.rser.2010.11.014

[14] Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735

[15] Qing, X., & Niu, Y. (2018). Hourly day-ahead solar irradiance prediction using weather forecasts by LSTM. Energy, 148, 461–468. https://doi.org/10.1016/j.energy.2018.01.177

[16] Kim, Y., et al. (2019). Solar power prediction based on satellite images and convolutional neural networks. IEEE Transactions on Sustainable Energy, 10(2), 852–862. https://doi.org/10.1109/TSTE.2018.2853140

[Downloaded from www.aisesjournal.com on 2026-11-12]

[17] Ladjal, B., et al. (2024). Hybrid deep learning CNN-LSTM model for forecasting direct normal irradiance. AIP Conference Proceedings, 3214, 020034. https://doi.org/10.1063/5.0244320

[18] V Abad – Alcaraz et al. Multimodal deep learning for solar radiation forecasting. Applied Energy, 372, 124567. https://doi.org/10.1016/j.apenergy.2025.124567

[19] Vaswani, A., et al. (2017). Attention is all you need. Advances in Neural Information Processing Systems (NeurIPS), 30, 5998–6008.

[20] Lim, B., et al. (2021). Temporal fusion transformers for interpretable multi-horizon time series forecasting. International Journal of Forecasting, 37(4), 1748–1764. https://doi.org/10.1016/j.ijforecast.2021.03.012

[21] Wang, J., et al. (2023). PVTransNet: A transformer-based model for multi-step photovoltaic power forecasting. Energy Conversion and Management, 283, 116892. https://doi.org/10.1016/j.enconman.2023.116892

[22] Zhang, Y., et al. (2023). Spatio-temporal transformer networks for solar irradiance forecasting in arid regions. Renewable Energy, 215, 109002. https://doi.org/10.1016/j.renene.2023.109002

[23] Camacho, M., Maldonado-Correa, J. L., Torres-Cabrera, J., Martínez-Morales, S., & García-Lafuente, E. (2024). CEEMDAN-CNN-Attention-LSTM framework for solar irradiance forecasting. Applied Energy, 358, 122456. https://doi.org/10.1016/j.apenergy.2024.122456

[24] Esram, T., & Chapman, P. L. (2007). Comparison of photovoltaic array maximum power point tracking techniques. IEEE Transactions on Energy Conversion, 22(2), 439–449. https://doi.org/10.1109/TEC.2006.874230

[25] Femia, N., et al. (2005). Optimization of perturb and observe maximum power point tracking method. IEEE Transactions on Power Electronics, 20(4), 963–973. https://doi.org/10.1109/TPEL.2005.850975

[26] Koad, R. B., et al. (2021). A comprehensive review of maximum power point tracking algorithms under partial shading conditions. Renewable and Sustainable Energy Reviews, 150, 111454. https://doi.org/10.1016/j.rser.2021.111454

[27] Liu, F., et al. (2008). A variable step size INC MPPT method for PV systems. IEEE Transactions on Industrial Electronics, 55(7), 2622–2628. https://doi.org/10.1109/TIE.2008.924016

[28] Ishaque, K., et al. (2012). An improved particle swarm optimization (PSO)-based MPPT for PV with reduced steady-state oscillation. IEEE Transactions on Power Electronics, 27(8), 3627–3638. https://doi.org/10.1109/TPEL.2012.2188893

[29] Phan BC, Lai YC, Lin CE. A Deep Reinforcement Learning-Based MPPT Control for PV Systems under Partial Shading Condition. Sensors (Basel). 2020 May 27;20(11):3039. doi: 10.3390/s20113039.

[Downloaded from www.aisesjournal.com on 2026-11-12]