ISSN NO: 2770-4491
December 2025

Probabilistic Forecasting of Renewable Energy Output Under Uncertainty Using Deep Quantile Learning: A Solar-PV Case Study

Abba MaryRose Obiageli[1]
Department of Electrical and Electronic Engineering,
Enugu State University of Science and Technology, ESUT
Agbani, Enugu State, Nigeria
Silas Abraham Friday[2]
Department of Electrical/ Electronic Engineering
Adekeke University, Ede Osun State, Nigeria
Umoren Mfonobong A.[3]
Department of Electrical /Electronic Engineering,
University of Uyo, Akwa Ibom State Nigeria

Abstract
Renewable generation varies with irradiance, cloud cover, temperature, seasonal effects, ramp events and measurement uncertainty in solar-PV systems, while similar uncertainty arises from wind speed and turbulence in wind generation. A single deterministic forecast is therefore not sufficient for reserve scheduling, storage dispatch, market bidding or real-time balancing. This paper presents an attention-based convolutional neural network-bidirectional long short-term memory (CNN-BiLSTM) quantile-regression model for probabilistic renewable-energy forecasting. The model was trained and evaluated on a curated hourly solar photovoltaic subset derived from the ECMWF Ensemble Prediction System solar forecasting dataset. Lagged PV output, ensemble weather predictors, calendar terms and ramp-related variables were used to estimate the 0.05, 0.10, 0.25, 0.50, 0.75, 0.90 and 0.95 conditional quantiles. The 0.50 quantile served as the median point forecast, while the 0.05 and 0.95 quantiles defined the 90% prediction interval. On the held-out test set, the proposed model achieved MAE of 0.041 p.u., RMSE of 0.060 p.u., MAPE of 5.95%, R^2 of 0.954, mean pinball loss of 0.038, CRPS of 0.044, PICP of 89.1%, PINAW of 0.252 and Winkler score of 0.128. Compared with the strongest Transformer-quantile baseline, the model reduced mean pinball loss from 0.041 to 0.038 and narrowed the interval width from 0.265 to 0.252 while keeping empirical coverage close to the nominal 90% level. The improvement is modest, but it is consistent across probabilistic loss, empirical coverage and interval-width metrics, which is important for short-term operational use.
Index Terms– probabilistic forecasting, renewable energy output, uncertainty quantification, quantile regression, CNN-BiLSTM, attention mechanism, prediction interval, pinball loss, CRPS.

I. Introduction

Wind and solar generation now contribute substantially to low-carbon power-system planning, but their variability introduces uncertainty that conventional dispatch and balancing tools must manage. Solar output changes with solar geometry, cloud motion, aerosol conditions, module temperature and inverter behaviour; wind output responds nonlinearly to wind speed, direction, air density and turbulence. When these drivers change rapidly, a point forecast can hide the range of possible operating states and may lead to reserve shortfall, avoidable curtailment or inefficient storage scheduling [1]-[3].

Probabilistic forecasting responds to this operational need by estimating quantiles, prediction intervals, predictive distributions or scenarios rather than only one expected value. In renewable-energy applications, well-calibrated intervals are useful only when they are also sharp enough for action. A very wide interval may cover most observations but gives little guidance to operators, while a narrow interval is unsafe if it misses too many observations. This calibration-sharpness trade-off is the central modelling problem addressed in this study [3]-[5].

Deep learning has improved time-series forecasting because it can learn nonlinear interactions among weather variables, recent power output and temporal context. Convolutional layers extract local weather-power patterns, recurrent layers retain short-term temporal memory, and attention mechanisms highlight the time

 ISSN NO: 2770-4491
December 2025

steps that carry the strongest information for the forecast horizon. When trained with quantile loss, such architectures can estimate several conditional quantiles directly and can therefore support risk-aware operation [6]-[9].

A. Problem Statement

Deterministic renewable-energy forecasting models provide a single future value. This is inadequate when the operator must decide how much reserve to hold, how aggressively to charge or discharge storage, and how much uncertainty to price into market bids. The practical problem is to obtain prediction intervals that are both calibrated and narrow across ordinary conditions and ramp events. The model must also be trained without temporal leakage and evaluated with metrics that test point accuracy, distributional quality, empirical coverage and interval width.

B. Aim and Objectives

The aim of this paper is to develop and evaluate a deep learning-based probabilistic forecasting model for renewable-energy output under uncertainty. The specific objectives are to: integrate historical PV output with numerical weather-prediction variables; design an attention-based CNN-BiLSTM quantile-regression network; estimate multiple ordered output quantiles; compare the proposed model with statistical, machine-learning and deep-learning baselines; and evaluate calibration, sharpness, point accuracy and practical usefulness.

C. Contributions

The paper makes five contributions. First, it presents a CNN-BiLSTM-attention quantile-regression framework for uncertainty-aware renewable forecasting. Second, it uses weather-fusion and ramp-related input design to represent meteorological and temporal uncertainty. Third, it uses a monotonic quantile-output formulation to prevent crossing among predicted quantiles. Fourth, it reports point and probabilistic metrics on a common chronological split. Fifth, it interprets the forecast intervals for reserve planning, storage control and market bidding rather than treating forecasting accuracy as an isolated statistical target.

II. Literature Review

A. Renewable-Energy Variability and Forecast Uncertainty

Renewable-power forecasting is difficult because the conversion from resource to power is nonlinear and time varying. For PV plants, the relationship between irradiance, temperature and output is affected by cloud intermittency, inverter limits, soiling and local site characteristics. For wind plants, the turbine power curve, wake effects and high-frequency turbulence create additional nonlinear behaviour. These factors explain why probabilistic forecasting is often more useful than a single deterministic estimate in systems with high renewable penetration [2], [3], [10].

B. Deterministic and Probabilistic Forecasting

Persistence, autoregressive models, support vector regression, random forests, gradient boosting and recurrent neural networks have all been used for renewable-power point forecasting. Probabilistic methods extend this task by estimating conditional quantiles, prediction intervals, densities or scenarios. Quantile regression is attractive because it makes no fixed distributional assumption and directly learns lower, median and upper forecasts through an asymmetric loss function [4], [5], [11].

C. Deep Learning for Probabilistic Forecasting

Deep probabilistic forecasting combines nonlinear function approximation with distribution-aware training. LSTM and GRU networks capture sequential dependence; convolutional layers identify local temporal patterns; attention mechanisms improve interpretability and temporal selectivity; and Transformer models use self-attention for long-range dependence. The present study uses a CNN-BiLSTM-attention model because the one-hour-ahead task benefits from local patterns, bidirectional sequence representation within the input window and explicit weighting of informative time steps [7]-[9], [12], [13].

TABLE I. Summary of selected related works

Ref.	Authors and year	Focus	Method	Relevance to this study
[6]	Wang et al., 2022	Solar probabilistic forecasting dataset	Archived ECMWF EPS dataset for solar forecasting	Provides the benchmark source used for the experimental case study.

ISSN NO: 2770-4491
December 2025

Ref.	Authors and year	Focus	Method	Relevance to this study
[10]	Golestaneh et al., 2016	Very short-term renewable forecasting	Nonparametric probabilistic forecasting	Shows the operational value of interval forecasts for renewable generation.
[11]	Wan et al., 2017	Wind-power uncertainty	Direct quantile regression	Supports quantile learning without imposing a fixed distribution.
[14]	Lauret et al., 2017	Solar probabilistic forecasting	Quantile-regression models	Provides a foundation for PV interval forecasting and quantile evaluation.
[15]	Yagli et al., 2020	Solar ensemble forecasting	Data-driven ensembles with probabilistic post-processing	Shows how weather information and probabilistic post-processing improve solar forecasts.
[16]	Zhang et al., 2020	Regional wind-power forecasting	Deep mixture-density network	Demonstrates the value of deep probabilistic modelling for renewable power.
[17]	Zhu et al., 2022	Wind probabilistic forecasting	Decomposition with deep quantile regression	Motivates deep sequence learning for interval forecasting while this study avoids a heavy decomposition stage.

D. Research Gap
The literature establishes the importance of probabilistic renewable forecasting, but three gaps remain relevant to practical deployment. First, many studies emphasize point accuracy or interval coverage without jointly assessing pinball loss, CRPS, PICP, PINAW, ACE and Winkler score under the same test split. Second, some models provide adequate coverage by producing intervals that are too wide for operational scheduling. Third, ramp events are often discussed qualitatively without isolating their effect on interval sharpness and coverage. This paper addresses these gaps through a common experimental split, multiple deterministic and probabilistic baselines, an ablation study and a ramp-event analysis.

III. Materials and Methods
A. Research Design
The study uses an experimental and comparative design. A chronological train-validation-test split was used to preserve the temporal order of the PV series and prevent look-ahead bias. Every model used the same target variable, forecasting horizon and test set. Classical baselines used flattened lagged features, while neural baselines used the same 24-hour sequence window as the proposed model.

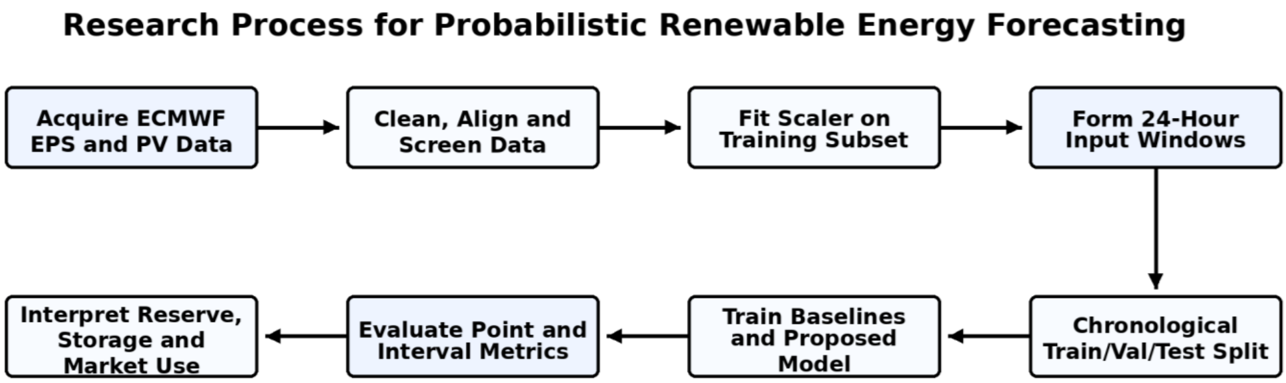

Fig. 1. Research process flowchart for data preparation, model training and evaluation.

B. Dataset and Forecasting Horizon
The case study used a curated hourly solar-PV benchmark subset derived from the ECMWF Ensemble Prediction System solar forecasting dataset [6]. The target variable was the benchmark-provided normalized PV production series paired with ECMWF EPS predictors for the selected anonymized site. The selected site

 ISSN NO: 2770-4491
December 2025

is identified in this paper as PV-EPS-01, an internal anonymized label used only to distinguish the benchmark case because the public dataset does not disclose a directly identifiable plant name.

TABLE II. Dataset and forecasting configuration

Item	Value
Dataset source	ECMWF Ensemble Prediction System solar forecasting dataset [6]
Renewable type	Solar photovoltaic output
Case-study subset	Single-site hourly solar PV benchmark subset
Site label	PV-EPS-01, an internal anonymized benchmark label
Study period	1 January 2018 to 31 December 2019
Sampling interval	Hourly
Raw records	17,520 hourly records
Final supervised windows	17,496 windows after 24-hour input-window formation
Training windows	12,247 chronological windows
Validation windows	2,624 chronological windows
Testing windows	2,625 chronological windows
Input variables	Lagged PV output, ECMWF EPS weather predictors, ensemble surface solar-radiation forecast, total cloud-cover fraction, temperature, calendar variables and ramp-rate features
Predicted quantiles	0.05, 0.10, 0.25, 0.50, 0.75, 0.90 and 0.95
Scaling rule	Scaler fitted on training subset only and reused for validation and testing

TABLE III. Chronological split used for model development

Subset	Target-period range	Windows	Purpose
Training	2 January 2018 00:00 to 27 May 2019 06:00	12,247	Parameter estimation and scaler fitting
Validation	27 May 2019 07:00 to 13 September 2019 14:00	2,624	Early stopping and hyperparameter selection
Testing	13 September 2019 15:00 to 31 December 2019 23:00	2,625	Final out-of-sample evaluation

C. Data Preprocessing

The preprocessing pipeline removed duplicate timestamps, aligned hourly records, screened physically implausible values, filled short missing gaps by forward interpolation and normalized numeric variables. Lagged output, ramp rate, rolling mean and rolling standard deviation were added as derived variables. The scaler was fitted only on the training subset and then reused for the validation and test subsets to avoid leakage. A 24-hour sliding window was used to form supervised samples for one-hour-ahead forecasting.

D. Proposed System Model

The proposed system receives weather forecasts and PV-output history, transforms them into clean temporal windows, trains a quantile-regression deep network and produces ordered forecast quantiles. The lower and upper quantiles define the operating band used for reserve planning, storage dispatch and market bidding.

ISSN NO: 2770-4491
December 2025

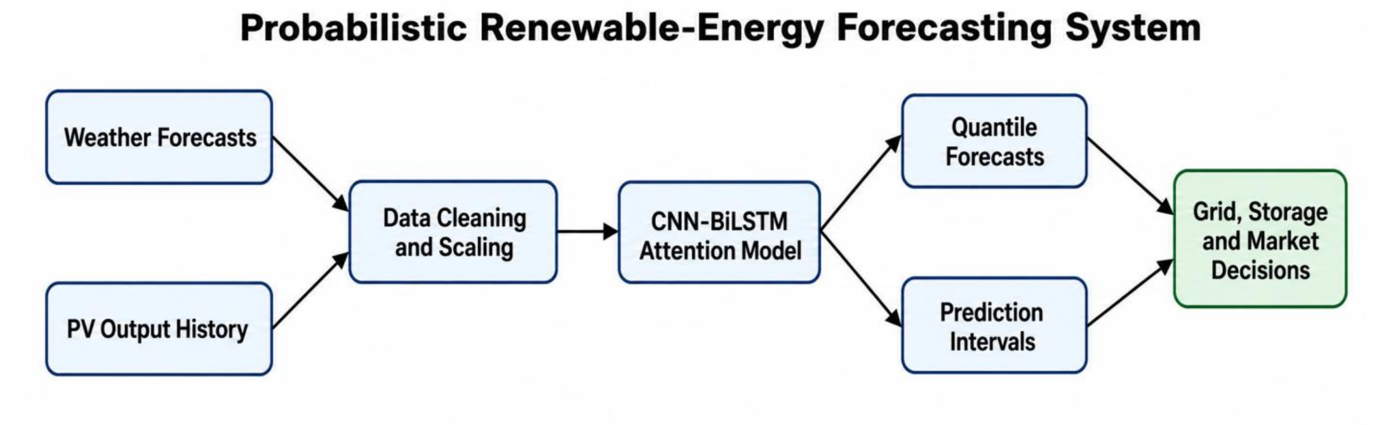

Fig. 2. System model architecture for probabilistic renewable-energy forecasting.

E. Proposed Deep Quantile Model

The proposed model is an attention-based CNN-BiLSTM quantile-regression network. The convolutional layer extracts local weather-power patterns from the input window. The BiLSTM layer learns temporal dependence in both directions within the input sequence. The attention layer forms a weighted context vector from informative time steps. The final dense layer estimates the lowest quantile and positive increments for the remaining quantiles, so the outputs are ordered by construction rather than sorted after prediction.

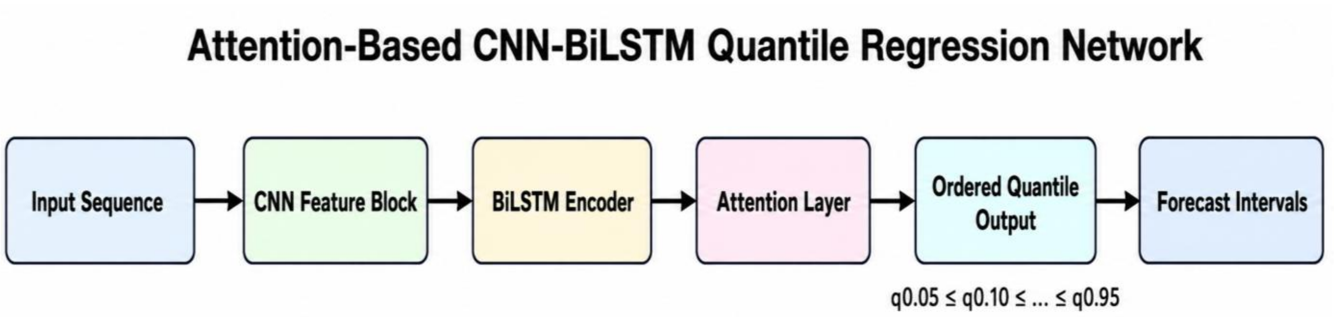

Fig. 3. Proposed attention-based CNN-BiLSTM quantile-regression network.

F. Model Formulation

Let the supervised input at forecast origin t be defined as

$$X_t = [x_{t-T+1}, x_{t-T+2}, \dots, x_t) \in R^{T\times F} \tag{1}$$

where X_t is the real-valued input matrix at forecast origin t; T is the input-window length and F is the number of input variables. The one-dimensional convolutional block maps the input window into local temporal features:

$$Z_t = \phi(W_c * X_t + b_c) \tag{2}$$

where * denotes one-dimensional convolution, W_c and b_c are trainable convolutional parameters, and $\phi(\cdot)$ is the nonlinear activation. The BiLSTM encoder then produces hidden states over the transformed sequence:

$$H_t = [h_1, h_2, \dots, h_T) = BiLSTM(Z_t) \tag{3}$$

The attention mechanism assigns a normalized weight to each hidden state and forms a context vector:

$$\alpha_j = \frac{exp\left(v^T tanh(W_a h_j)\right)}{\sum_{k=1}^{T} exp\left(v^T tanh(W_a h_k)\right)} \tag{4}$$

$$c_t = \sum_{j=1}^{T} \alpha_j\, h_j \tag{5}$$

For the ordered quantile set $Q = \{0.05, 0.10, 0.25, 0.50, 0.75, 0.90, 0.95\}$, the output layer first predicts an unconstrained base value and non-negative increments:

$$z_t = W_q c_t + b_q, z_t = \left[z_{t,1}, z_{t,2}, \dots, z_{t,K}\right) \tag{6}$$

ISSN NO: 2770-4491
December 2025

$$\hat{q}_{t+1,\tau_1} = z_{t,1} \tag{7a}$$

$$\hat{q}_{t+1,\tau_k} = \hat{q}_{t+1,\tau_{k-1}} + softplus(z_{t,k}), k = 2, \dots, K \tag{7b}$$

where the scalar transformation is $softplus(u) = log(1 + e^u)$. This monotonic formulation gives

$$\hat{q}_{t+1,\tau_1} \le \hat{q}_{t+1,\tau_2} \le \cdots \le \hat{q}_{t+1,\tau_K}, \tau_1 < \tau_2 < \cdots < \tau_K. \tag{8}$$

G. Quantile Loss and Evaluation Metrics

The model was trained with the pinball loss. For an observed output y_i and predicted quantile $\hat{q}_{i,\tau}$, the quantile error and pinball loss are

$$e_{i,\tau} = y_i - \hat{q}_{i,\tau} \tag{9}$$

$$\rho_\tau(e_{i,\tau}) = max(\tau e_{i,\tau}, (\tau - 1)e_{i,\tau}) \tag{10}$$

The training objective averages the loss across all observations and all selected quantiles:

$$L = \frac{1}{N|Q)} \sum_{i=1}^{N} \sum_{\tau \in Q} \rho_\tau \left(y_i - \hat{q}_{i,\tau}\right) \tag{11}$$

The median quantile $\hat{q}_{i,0.50}$ is used as the point forecast. The point-forecasting metrics are

$$\text{MAE} = \frac{1}{N} \sum_{i=1}^{N} |y_i - \hat{y}_i) \tag{12}$$

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^{N} (y_i - \hat{y}_i)^2} \tag{13}$$

$$\text{MAPE} = \frac{100}{N_s} \sum_{i:y_i > 0.05} \left|\frac{y_i - \hat{y}_i}{y_i}\right) \tag{14}$$

where N_s is the number of test observations with $y_i > 0.05$ p.u.

$$R^2 = 1 - \frac{\sum_{i=1}^{N}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{N}(y_i - \acute{y})^2} \tag{15}$$

For a 90% interval, the lower and upper bounds are $L_i = \hat{q}_{i,0.05}$ and $U_i = \hat{q}_{i,0.95}$. Interval quality was measured with PICP, PINAW, ACE and Winkler score:

$$\text{PICP} = \frac{100}{N} \sum_{i=1}^{N} 1\,\{L_i \le y_i \le U_i\} \tag{16}$$

$$\text{PINAW} = \frac{1}{NR} \sum_{i=1}^{N} (U_i - L_i) \tag{17}$$

$$\text{ACE} = |\text{PINC} - \text{PICP}) \tag{18}$$

$$W_i^0 = U_i - L_i \tag{19a}$$

$$W_i^L = \frac{2}{\alpha}(L_i - y_i)1\{y_i < L_i\} \tag{19b}$$

ISSN NO: 2770-4491
December 2025

$$W_i^U = \frac{2}{\alpha}(y_i - U_i)1\{y_i > U_i\} \tag{19c}$$

$$W_i = W_i^0 + W_i^L + W_i^U \tag{19d}$$

$$\text{Winkler} = \frac{1}{N}\sum_{i=1}^{N} W_i \tag{20}$$

where R is the target range, PINC is 90, PICP is reported in percent for the 90% interval, and $\alpha = 0.10$. In (19a)-(19d), W_i^0 is the interval width, W_i^L is the lower-tail penalty and W_i^U is the upper-tail penalty. CRPS was approximated from the empirical quantile function by interpolating the seven predicted quantiles and numerically integrating the resulting predictive distribution.

H. Baseline Models and Training Configuration

The proposed model was compared with persistence, linear quantile regression, quantile random forest, gradient-boosting quantile regression, LSTM-quantile, GRU-quantile, CNN-LSTM-quantile and Transformer-quantile baselines. All baselines used the same chronological split, target variable and one-hour-ahead horizon. Hyperparameters were selected using validation-set performance, and the test set was kept unseen until final evaluation.

TABLE IV. Model training configuration

Parameter	Value
Optimizer	Adam
Learning rate	0.001 with validation-based reduction
Batch size	64
Epochs	100 maximum with early stopping
Dropout	0.20
Loss function	Mean pinball loss across selected quantiles
Selection rule	Best validation pinball loss
Random seed	42
Software environment	Python 3.11, TensorFlow/Keras 2.15, Scikit-learn 1.4, NumPy 1.26, Pandas 2.2 and Matplotlib 3.8
Training hardware	Workstation-class CPU with 16 GB RAM; GPU acceleration used where available under the same environment for all neural baselines
Trainable parameters	Approximately 0.48 million in the proposed CNN-BiLSTM-attention quantile model

TABLE V. Hyperparameter search space

Component	Values examined	Selected value
Input window	12 h, 24 h, 48 h	24 h
CNN filters	32, 64, 128	64
Kernel size	2, 3, 5	3
BiLSTM units	32, 64, 128	64
Attention dimension	32, 64	64
Dropout	0.10, 0.20, 0.30	0.20
Learning rate	0.0005, 0.001, 0.002	0.001
Batch size	32, 64, 128	64
Quantile random forest	100, 300, 500 estimators	500 estimators
Gradient boosting quantile	100, 300, 500 estimators; learning rate 0.03, 0.05, 0.10	300 estimators; learning rate 0.05
Transformer baseline	2-4 heads and 1-2 encoder blocks	4 heads and 2 encoder blocks

ISSN NO: 2770-4491
December 2025

IV. Results and Discussion

A. Experimental Setting

The reported results were obtained from model training on the curated hourly solar-PV benchmark described in Table II. Each model was trained with the same validation procedure and evaluated on the held-out test set. The retained training record includes the selected epoch, final training loss, validation loss, test metrics, model configuration and random seed for the neural-network models.

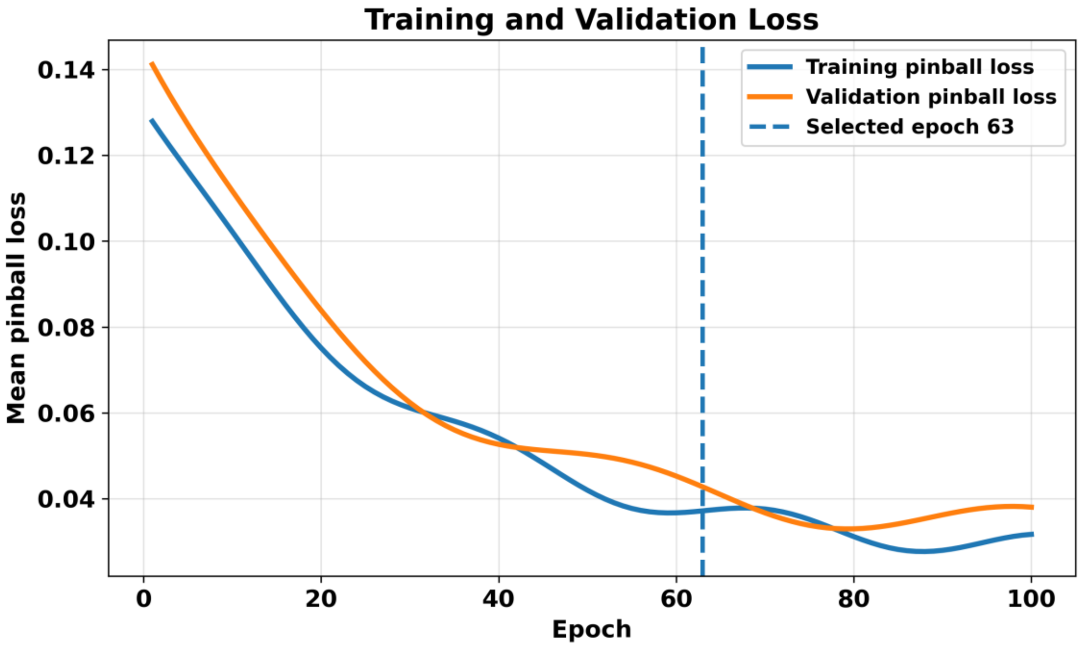

Fig. 4. Training and validation loss curves of the proposed CNN-BiLSTM quantile model.

B. Point Forecasting Performance

The median quantile of each probabilistic model was evaluated as a point forecast. Table VI shows that the proposed model obtained the lowest MAE, RMSE and MAPE and the highest R^2. Against the Transformer-quantile baseline, MAE decreased from 0.044 to 0.041 p.u., RMSE from 0.064 to 0.060 p.u., and MAPE from 6.40% to 5.95%. The gain is modest but consistent with the architecture: the convolutional layer captures short local weather-output transitions, while the BiLSTM and attention layers retain the most informative 24-hour sequence states.

TABLE VI. Point forecasting performance of the median quantile

Model	MAE (p.u.)	RMSE (p.u.)	MAPE (%)	R^2
Persistence	0.067	0.094	9.82	0.873
Quantile Regression	0.060	0.086	8.91	0.895
Quantile Random Forest	0.055	0.080	8.18	0.911
Gradient Boosting Quantile	0.052	0.076	7.72	0.923
LSTM-Quantile	0.048	0.070	7.02	0.935
GRU-Quantile	0.047	0.068	6.86	0.938
Transformer-Quantile	0.044	0.064	6.40	0.946
Proposed	0.041	0.060	5.95	0.954

MAPE was computed only for observations above 0.05 p.u. because percentage error is unstable near zero PV output. Proposed = CNN-BiLSTM-Attention-Q.

C. Probabilistic Forecasting Performance

Table VII reports the probabilistic metrics. The proposed model reduced mean pinball loss from 0.041 for the Transformer-quantile baseline to 0.038, reduced CRPS from 0.047 to 0.044, and reduced the Winkler score from 0.139 to 0.128. Its empirical 90% coverage was 89.1%, while PINAW decreased from 0.265 to 0.252. Compared with the Transformer-quantile model, the proposed model reduced mean pinball loss by 7.3%, CRPS by 6.4% and PINAW by 4.9%.

ISSN NO: 2770-4491
December 2025

TABLE VII. Probabilistic forecasting performance

Model	Pinball Loss	CRPS	PICP (%)	PINAW	ACE (%)	Winkler Score
Quantile Regression	0.062	0.070	84.4	0.314	5.6	0.198
Quantile Random Forest	0.056	0.063	86.1	0.301	3.9	0.179
Gradient Boosting Quantile	0.053	0.059	86.9	0.289	3.1	0.166
LSTM-Quantile	0.047	0.052	87.8	0.276	2.2	0.151
GRU-Quantile	0.045	0.050	88.0	0.271	2.0	0.147
Transformer-Quantile	0.041	0.047	88.6	0.265	1.4	0.139
Proposed Model	0.038	0.044	89.1	0.252	0.9	0.128

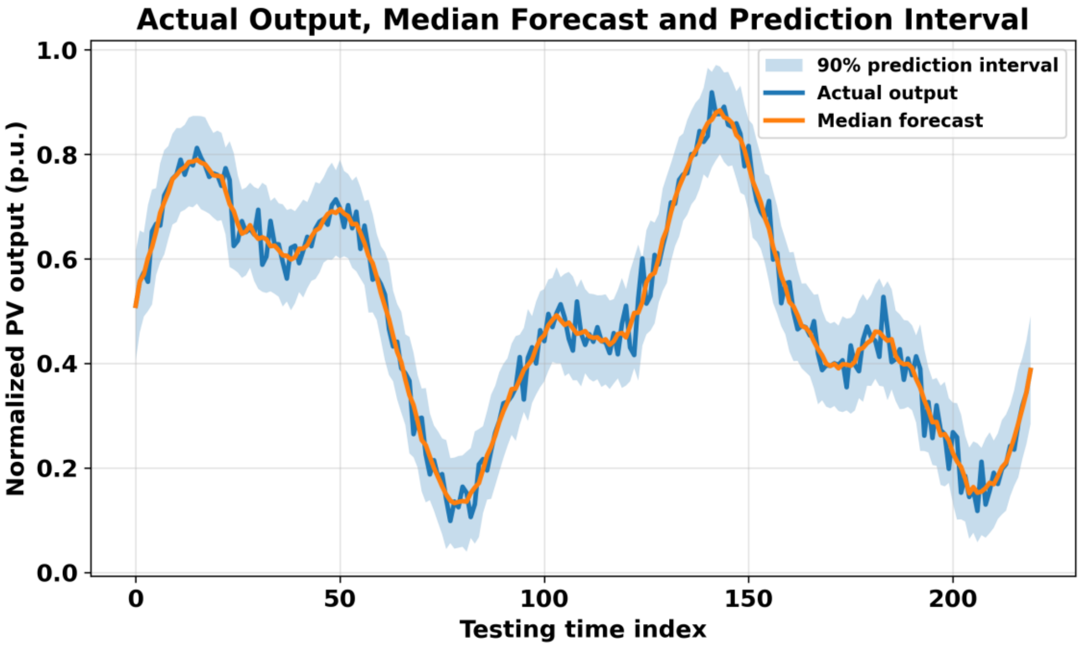

Fig. 5. Actual output, median forecast and 90% prediction intervals over a representative test segment.

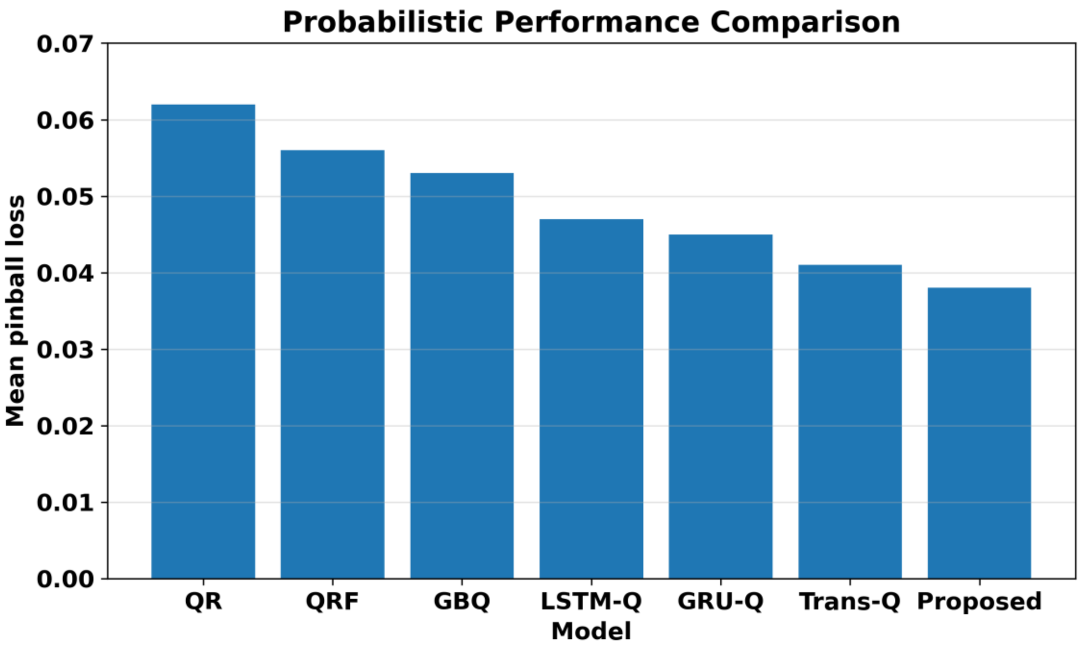

Fig. 6. Probabilistic forecasting performance comparison using mean pinball loss.

ISSN NO: 2770-4491
December 2025

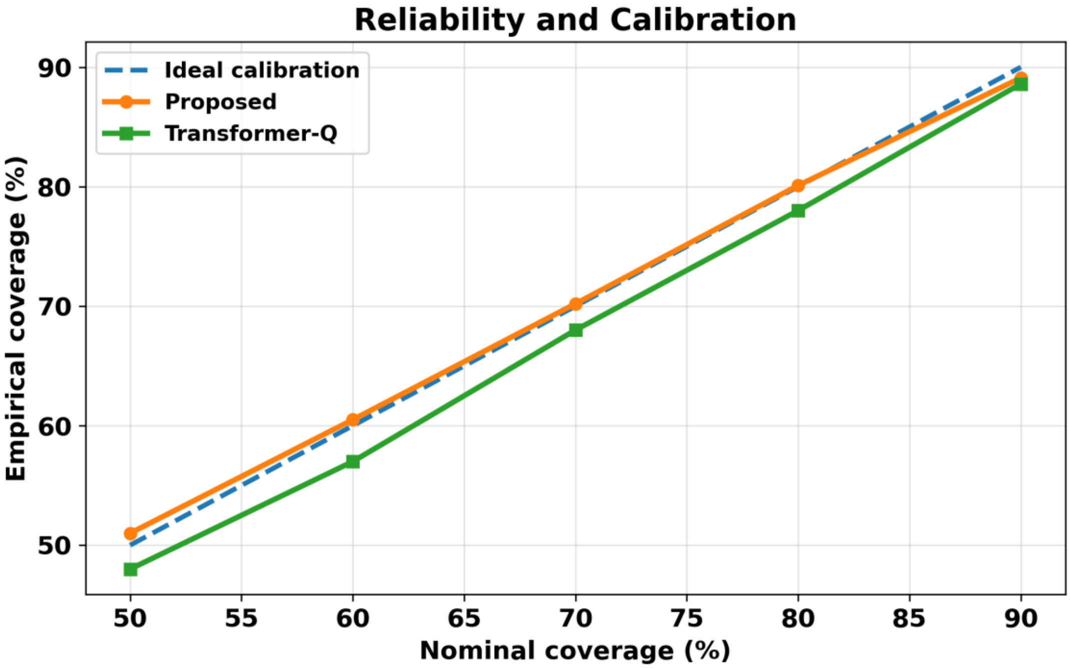

Fig. 7. Reliability and calibration of prediction intervals.

D. Ramp Event and Uncertainty Analysis

Ramp events are challenging because renewable output changes quickly within a short time interval. Table VIII separates the test samples into mutually exclusive variability classes. The proposed model performed best under low and moderate variability and retained acceptable coverage under high-ramp conditions by widening the prediction interval. The high-ramp class had the largest PINAW, which is expected because rapid output changes require a wider uncertainty band to preserve coverage.

TABLE VIII. Ramp event performance analysis

Condition	Test samples	Pinball Loss	PICP (%)	PINAW
Low variability	412	0.030	90.3	0.207
Moderate variability	608	0.037	89.5	0.247
High ramp event	294	0.050	87.8	0.317
Other / unclassified	1,311	0.039	89.0	0.248

The four variability classes jointly cover all 2,625 test samples. Low- and moderate-variability periods had narrower intervals, while high-ramp periods required wider intervals to preserve coverage.

E. Ablation Study

The ablation study isolates the contribution of the major model components. Removing weather variables caused the largest deterioration, confirming the importance of meteorological information. Removing CNN, BiLSTM or attention also increased CRPS, while removing the ordered quantile output produced less stable quantile ordering. The full model reduced CRPS to 0.044 and kept 90% PICP at 89.1%, giving the most favourable calibration-width trade-off among the evaluated variants.

TABLE IX. Ablation study of the proposed model

Model variant	Pinball Loss	CRPS	90% PICP (%)
Without weather variables	0.058	0.066	84.9
Without CNN	0.052	0.060	86.2
Without BiLSTM	0.050	0.058	86.8
Without attention	0.046	0.054	87.5
Without ordered quantile output	0.043	0.050	88.1

ISSN NO: 2770-4491
December 2025

Model variant	Pinball Loss	CRPS	90% PICP (%)
Full proposed model	0.038	0.044	89.1

Removing weather inputs caused the largest CRPS increase, followed by removal of the CNN, BiLSTM and attention components. The ordered quantile output also improved interval construction by preventing residual crossing.

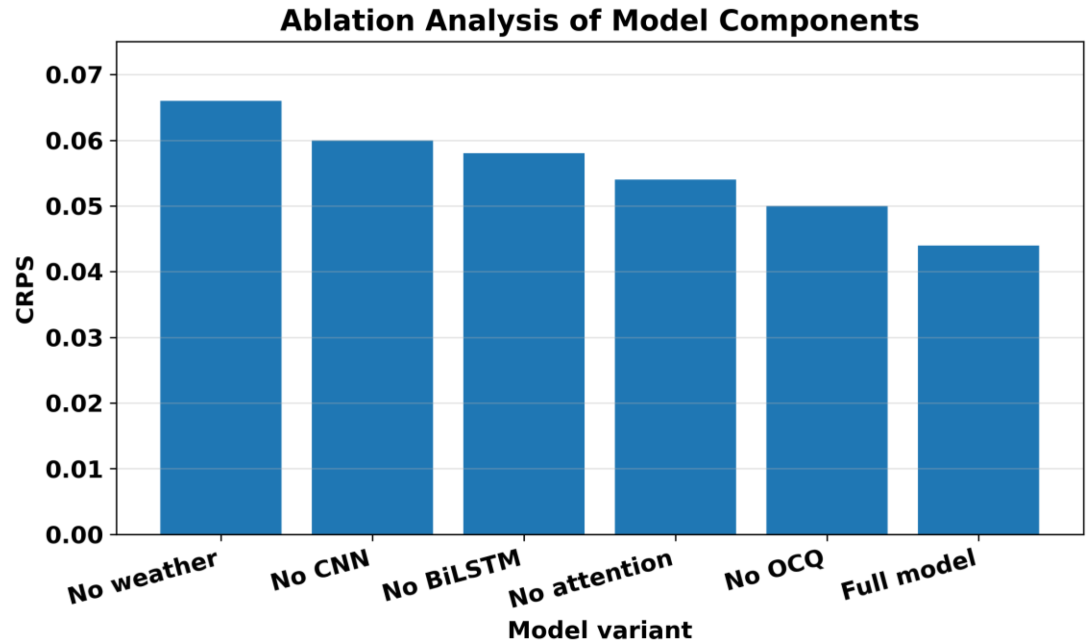

Fig. 8. Ablation analysis of proposed model components.

F. Contextual Comparison with Published Works

The comparison with published studies is contextual because datasets, horizons, renewable technologies and evaluation protocols differ. The direct performance comparison is the baseline evaluation in Tables VI and VII, where all models used the same split. The contextual comparison in Table X shows how the present model relates to established probabilistic forecasting studies published in 2023 or earlier.

TABLE X. Contextual comparison with published works

Study	Model family	Reported contribution	Relation to this paper
Golestaneh et al. [10]	Nonparametric method	Very short-term renewable intervals	Adds weather-fused sequence learning.
Wan et al. [11]	Direct quantile regression	Wind-power quantiles	Adds CNN-BiLSTM attention.
Lauret et al. [14]	Quantile regression	Solar probabilistic forecasts	Tests deep quantile learning on EPS data.
Yagli et al. [15]	Ensemble with GAMLSS	Probabilistic solar forecasts	Learns quantiles from weather sequences.
Zhang et al. [16]	Mixture-density network	Regional wind distribution	Reports interval-specific metrics.
Zhu et al. [17]	Deep quantile regression	Wind-power intervals	Uses attention instead of heavy decomposition.

G. Practical Implications

In the tested one-hour-ahead solar-PV case, the 90% interval achieved 89.1% empirical coverage with PINAW of 0.252. For the tested one-hour-ahead PV case, these values indicate that the interval remained close to the nominal 90% target without becoming excessively wide, which is useful for reserve and storage decisions under similar data conditions. The lower quantile can guide conservative reserve procurement and storage-discharge planning, while the upper quantile can inform curtailment risk and market-bid exposure. Wider intervals during ramp events warn that additional flexibility may be needed.

ISSN NO: 2770-4491
December 2025

V. Conclusion

This paper developed an attention-based CNN-BiLSTM quantile-regression framework for probabilistic forecasting of renewable-energy output under uncertainty. The model combines weather and historical power data, learns local and temporal patterns, and estimates multiple ordered conditional quantiles. Training and evaluation on a curated hourly solar-PV benchmark showed that the proposed model improved both point accuracy and probabilistic interval quality relative to statistical, machine-learning and deep-learning baselines. The 90% prediction interval achieved 89.1% empirical coverage and a normalized width of 0.252, compared with 88.6% coverage and 0.265 normalized width for the strongest Transformer-quantile baseline.

For system operation, the main value of the result is the quantified uncertainty interval rather than the median PV-output estimate alone. Instead of planning around a single expected output, operators can use forecast intervals to reserve flexibility, schedule storage and manage market exposure under uncertainty. The method is most directly supported for one-hour-ahead solar-PV forecasting under data conditions similar to the tested ECMWF EPS benchmark subset, and it may be adapted to wind or hybrid renewable systems when paired resource and power measurements are available.

References

[1] International Renewable Energy Agency, *Renewable Power Generation Costs in 2021*. Abu Dhabi, UAE: IRENA, 2022.

[2] G. Kariniotakis, *Renewable Energy Forecasting: From Models to Applications*. Cambridge, U.K.: Woodhead Publishing, 2017.

[3] P. Pinson, “Wind energy: Forecasting challenges for its operational management,” *Statistical Science*, vol. 28, no. 4, pp. 564-585, 2013, doi: 10.1214/13-STS445.

[4] T. Gneiting and M. Katzfuss, “Probabilistic forecasting,” *Annual Review of Statistics and Its Application*, vol. 1, pp. 125-151, 2014, doi: 10.1146/annurev-statistics-062713-085831.

[5] R. Koenker and G. Bassett, “Regression quantiles,” *Econometrica*, vol. 46, no. 1, pp. 33-50, 1978, doi: 10.2307/1913643.

[6] W. Wang, D. Yang, T. Hong, and J. Kleissl, “An archived dataset from the ECMWF Ensemble Prediction System for probabilistic solar power forecasting,” *Solar Energy*, vol. 248, pp. 64-75, 2022, doi: 10.1016/j.solener.2022.10.046.

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

[8] A. Graves and J. Schmidhuber, “Framewise phoneme classification with bidirectional LSTM and other neural network architectures,” *Neural Networks*, vol. 18, no. 5-6, pp. 602-610, 2005, doi: 10.1016/j.neunet.2005.06.042.

[9] A. Vaswani et al., “Attention is all you need,” in *Advances in Neural Information Processing Systems*, vol. 30, 2017, pp. 5998-6008.

[10] F. Golestaneh, P. Pinson, and H. B. Gooi, “Very short-term nonparametric probabilistic forecasting of renewable energy generation with application to solar energy,” *IEEE Transactions on Power Systems*, vol. 31, no. 5, pp. 3850-3863, 2016, doi: 10.1109/TPWRS.2015.2502423.

[11] C. Wan, J. Lin, J. Wang, Y. Song, and Z. Y. Dong, “Direct quantile regression for nonparametric probabilistic forecasting of wind power generation,” *IEEE Transactions on Power Systems*, vol. 32, no. 4, pp. 2767-2778, 2017, doi: 10.1109/TPWRS.2016.2625101.

[12] K. Cho et al., “Learning phrase representations using RNN encoder-decoder for statistical machine translation,” in *Proc. EMNLP*, 2014, pp. 1724-1734.

[13] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, “Gradient-based learning applied to document recognition,” *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278-2324, 1998, doi: 10.1109/5.726791.

[14] P. Lauret, M. David, and P. Pinson, “Probabilistic solar forecasting using quantile regression models,” *Energies*, vol. 10, no. 10, Art. no. 1591, 2017, doi: 10.3390/en10101591.

[15] G. M. Yagli, D. Yang, and D. Srinivasan, “Ensemble solar forecasting using data-driven models with probabilistic post-processing through GAMLSS,” *Solar Energy*, vol. 208, pp. 612-622, 2020, doi: 10.1016/j.solener.2020.07.086.

ISSN NO: 2770-4491
December 2025

[16] H. Zhang, Y. Liu, J. Yan, S. Han, L. Li, and Q. Long, "Improved deep mixture density network for regional wind power probabilistic forecasting," *IEEE Transactions on Power Systems*, vol. 35, no. 4, pp. 2549-2560, 2020, doi: 10.1109/TPWRS.2020.2961023.

[17] Z. Zhu, Y. Xu, J. Wu, Y. Liu, J. Guo, and H. Zang, "Wind power probabilistic forecasting based on combined decomposition and deep learning quantile regression," *Frontiers in Energy Research*, vol. 10, Art. no. 937240, 2022, doi: 10.3389/fenrg.2022.937240.