IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.151087

AI-driven Renewable Energy Generation Forecasting Using Machine Learning Techniques

Arathi H L*[1], Sinchana C S[2], Partha C M[3], Shobha N[4], Shreya K S[5]

Department of Computer Science and Engineering, East West College of Engineering

Visvesvaraya Technological University, Bengaluru, Karnataka, India[1-5]

*Corresponding Author

Abstract: Short-term forecasting of solar and wind generation is a prerequisite for economical plant operation, yet most reported machine-learning studies conclude at a table of error values computed on stored datasets. In earlier work the present authors reviewed five recent studies, identified six practical deficiencies, and proposed a five-layer forecasting framework; no implementation or measured result was reported at that stage. This paper reports the construction and quantitative evaluation of that framework as a working system. A web application was developed in which the browser Geolocation API and a public meteorological service supply the four environmental parameters used by the models, removing manual data entry. Four regression algorithms—linear regression, a decision tree, a random forest, and support vector regression—were trained separately for solar and wind output on a 5,922-record physics-based reference dataset and compared under identical conditions. On a 1,185-record held-out partition the selected solar model (support vector regression) attained MAE 1.340 kWh, RMSE 1.795 kWh and R^2 0.9940, and the selected wind model (random forest) attained MAE 1.484 kWh, RMSE 1.932 kWh and R^2 0.9939, for a reference plant rated at 100 kW per source. Linear regression proved adequate for solar output (R^2 0.9930) but failed for wind (R^2 0.8375); the discrepancy is traced to the piecewise-cubic turbine power curve. Residual analysis shows that the non-linear models operate at the injected noise floor of the reference data, so differences among them are not statistically meaningful, and permutation importance recovers the dominant physical driver of each source. Of the six deficiencies previously identified, three are closed, one is partially addressed, and two remain open. Because the reference dataset is synthetic, the reported accuracy characterises model capacity and pipeline correctness rather than field performance.

Index Terms: Energy forecasting, machine learning, model comparison, permutation importance, random forest, renewable energy, support vector regression, weather application programming interface, web application.

I. INTRODUCTION

Solar and wind capacity continues to be added to power systems worldwide because these sources are clean, increasingly inexpensive, and widely available. Their output, however, is set by the weather rather than by an operator. Irradiance, cloud cover, temperature, and wind speed determine what a plant delivers at any instant, and none of these quantities can be dispatched [1]. The consequence is an operational problem rather than a purely scientific one: when the generation of the next several hours is unknown, spinning reserve has to be held, storage has to be cycled conservatively, and surplus energy is occasionally curtailed because it cannot be absorbed. Each of these responses carries a cost that a dependable short-term forecast would reduce [2].

Machine learning is now the standard approach to this forecasting problem. A model is fitted to historical pairs of weather observations and measured generation, and the fitted relationship is then applied to forecast weather in order to anticipate output. A considerable literature reports that such models predict solar and wind generation with low error [3]–[7], and the narrow question of which algorithm minimises error on a given archive has been examined in detail.

What has received less attention is everything that follows the prediction. A forecast delivered as a single number inside a research environment is of limited use to the person who has to act on it. That person additionally needs to know which weather conditions produced the value, what the forecast implies for the next few hours of operation, and whether the model is still trustworthy. In earlier work [13] the present authors surveyed five recent studies, catalogued six recurring practical deficiencies, and proposed a five-layer framework intended to address them. That paper was explicitly at the concept stage: the system had not been built, and no accuracy value, latency measurement, or comparison was reported. The present paper closes that gap. Its contributions are as follows. First, the proposed five-layer framework is realised as a working client–server application in which live weather is acquired automatically from the user's geographic position, so that no meteorological quantity is typed in by hand. Second, four regression algorithms are trained and compared for solar and for wind output under identical preprocessing, partitioning, and scoring, and the selection is made from

 This work is licensed under a Creative Commons Attribution 4.0 International License

DOI: 10.17148/IJARCCE.2026.151087

measured evidence rather than fixed in advance. Third, the error structure of the selected models is analysed in terms of the noise floor of the reference data, the regimes of the turbine power curve, and permutation-based feature attribution, which together explain why one algorithm family fails for wind while succeeding for solar. Fourth, the computational cost of inference and the behaviour of the deployed system against a live weather service are measured. Fifth, the six deficiencies identified in [13] are revisited and the status of each is stated without embellishment, which defines the remaining work precisely.

The remainder of the paper is organised as follows. Section II reviews related work and states the position of this study within it. Section III describes the implemented architecture. Section IV sets out the dataset, models, and evaluation protocol. Section V reports results, and Section VI discusses their interpretation, the closure of the earlier deficiencies, and the threats to validity. Section VII concludes.

II. RELATED WORK

Renewable generation forecasting is conventionally grouped by horizon. Very short-term forecasts spanning minutes to an hour support the management of sudden fluctuation; short-term forecasts spanning one hour to one day support storage scheduling, load planning, and reserve procurement; longer horizons serve planning and maintenance. The work reported here targets the short-term range, since that is where most day-to-day operating decisions are taken.

Three families of learning algorithm recur in this literature. Tree-based ensembles handle tabular inputs and their interactions, train quickly, and expose measures of variable influence. Kernel methods such as support vector regression generalise well at moderate sample sizes. Sequence models, notably long short-term memory networks, learn from the ordering of past values but demand substantially more data and computation.

Alkhanafseh et al. [3] present an encoder–LSTM–decoder model for renewable generation and compare it against a broad set of alternatives. The comparison is thorough, but the inputs are transformed into a compressed representation, so the model can no longer indicate which physical measurement produced a given forecast, and the model is trained once with no provision for later updating. Babu et al. [4] compare regression and ensemble methods for solar forecasting, report gradient boosting as the strongest, and usefully add solar-geometry features alongside meteorological variables; the study does not report a confidence range, an explanation of individual predictions, or any live data path. Alhayd and Todeschini [5] evaluate many input combinations with support vector regression and neural variants on plant data, and observe that one model attains the lowest mean error while another tracks the overall pattern more faithfully—an observation that argues directly for comparing several models rather than committing to one. Their dataset is not public, which limits reproduction. Sharma and Yadav [6] apply a convolutional network to both generation and household consumption, which suits small grids, but compare against a limited set of alternatives and do not report the contribution of individual weather inputs. Rajarajeswaran [7] compares an LSTM, a gradient boosting model, and a support vector machine for solar and wind, and reports solar radiation and wind speed as the dominant inputs, in agreement with physical reasoning; the dataset size and training configuration are not given in sufficient detail for replication, and the paper itself identifies lightweight models, real-time data collection, and explainability as open problems.

Table I condenses these studies and the limitations observed in them. Taken together, they concentrate on minimising a reported error value, while the properties that would make a forecast usable in practice are largely absent. The framework proposed in [13] responded to this pattern by treating prediction, interpretation, decision support, and user access as parts of a single system, and by making model comparison a step inside that system rather than a decision taken beforehand. The present paper differs from [3]–[7] not in the algorithms it employs, which are standard [14]–[16], but in reporting an end-to-end implementation together with the measurements needed to judge it, and in stating plainly which of the intended properties were achieved.

TABLE I. SUMMARY OF REVIEWED STUDIES AND OBSERVED LIMITATIONS

Study	Methods used	Limitations observed	Ref.
Alkhanafseh et al. (2025)	Encoder–LSTM–decoder; feature screening and compression	Compressed inputs hide the physical cause; single-valued forecast; no updating after training	[3]

 This work is licensed under a Creative Commons Attribution 4.0 International License

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.151087

Study	Methods used	Limitations observed	Ref.
Babu et al. (2025)	Regression and ensemble methods; solar-geometry inputs	No confidence range, per-prediction explanation, live data path, or suggested action	[4]
Alhayd and Todeschini (2025)	Support vector regression and neural variants	Complementary best models not combined; private dataset; no deployment plan	[5]
Sharma and Yadav (2023)	Convolutional neural network	Limited comparison; model configuration unexplained; input contributions not shown	[6]
Rajarajeswaran (2025)	LSTM, gradient boosting, support vector machine	Dataset and training details undisclosed; no real-time or explainable component	[7]

III. IMPLEMENTED SYSTEM ARCHITECTURE

The implemented system preserves the five-layer organisation proposed in [13]. Each layer is a separate software module with a defined interface, so that a layer can be replaced without disturbing the others. Fig. 1 shows the arrangement and the direction of data flow.

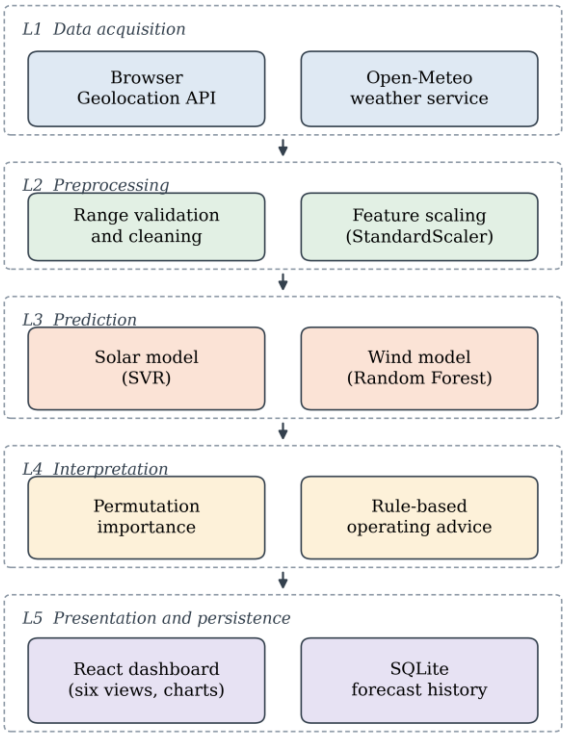


Fig. 1. Implemented five-layer architecture. Data flow is upward from acquisition to presentation. The periodic-retraining path proposed in [13] is not present in the current implementation and is listed in Section VI-C as open work.

A. Data Acquisition Layer

Two sources supply the inputs. The client obtains the user's coordinates through the browser Geolocation API, subject to an explicit permission prompt, and the server retrieves current conditions and an hourly forecast for those coordinates from the Open-Meteo service [12], which requires no subscription key. Four quantities are extracted: air temperature at 2 m, relative humidity at 2 m, wind speed at 10 m, and global horizontal irradiance. Coordinates are additionally resolved to a place name for display. No meteorological quantity is entered by the user, which is the property that distinguishes this layer from the stored-file pipelines of [3]–[7].

 This work is licensed under a Creative Commons Attribution 4.0 International License

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.151087

B. Preprocessing Layer

Incoming records are validated against physically admissible intervals—temperature in [−30, 60] °C, relative humidity in [0, 100] %, wind speed in [0, 40] m s^{-1}, and irradiance in [0, 1400] W m^{-2}—and records with absent or inadmissible values are discarded. Features are then standardised to zero mean and unit variance. Standardisation is implemented inside the model pipeline rather than as a separate preparatory step, which guarantees that the transformation applied at inference is exactly the one fitted during training and removes a common source of deployment error.

C. Prediction Layer

Two independent regressors are held, one for solar output and one for wind output, each selected by the procedure of Section IV-D. Separating the targets allows each to be served by the algorithm that suits its underlying physics, and the results of Section V show that the selected algorithms indeed differ. A deterministic physical constraint is applied after inference: predicted solar output is set to zero whenever irradiance is zero, since a photovoltaic array cannot generate in darkness irrespective of what a fitted model extrapolates. Hourly predictions are produced for the current conditions and for each of the following twenty-four forecast hours.

D. Interpretation Layer

This layer converts a numerical prediction into information an operator can act upon. Feature attribution is computed by permutation importance [15], which measures the degradation in the coefficient of determination when a single input is randomly permuted and is therefore applicable to any fitted estimator, including the kernel model that exposes no intrinsic importance measure. A set of deterministic rules then derives operating guidance from the forecast and the weather state: whether generation is currently limited by darkness, cloud cover, or a wind speed below the turbine cut-in threshold; the hour at which peak output is expected, so that flexible load can be scheduled; and whether the twenty-four-hour capacity factor is high enough to justify storing or exporting surplus energy, or low enough to require backup. Aggregate indicators of avoided grid cost and avoided carbon dioxide emissions are computed from the forecast energy using fixed published factors.

E. Presentation and Persistence Layer

The interface is a single-page application of six views: an overview, a live weather view, a prediction view that runs the pipeline from one control, a results view, a detailed report, and a history view. Charts present the twenty-four-hour generation profile, the solar–wind mix, the comparative performance of the four candidate algorithms, predicted against reference values, and feature attribution. Every completed forecast, together with the weather inputs that produced it, is written to a relational store, which supplies the history view and preserves the record needed for the accuracy monitoring described in Section VI-C. Fig. 6 shows the results view of the deployed application.

F. Implementation

The server is implemented in Python with the FastAPI framework and exposes eight endpoints: a health probe, a weather query, a prediction endpoint that accepts a latitude–longitude pair and returns the full forecast object, an endpoint returning the stored evaluation report, and four endpoints for retrieving and deleting history records. Models are fitted with scikit-learn [14] and serialised to disk, so the server loads pre-fitted estimators at start-up rather than training on demand. The client is implemented in React with a utility-first stylesheet and a charting library. Persistence uses SQLite, which requires no separate database server and keeps the deployment footprint small enough for a single workstation—consistent with the lightweight-model objective raised in [7].

IV. METHODOLOGY

A. Reference Dataset Construction

Paired records of weather and measured generation for a single co-located solar and wind installation were not available to this study. A reference dataset was therefore synthesised from established performance models, so that the learning pipeline could be verified against a known ground truth. This choice is a deliberate methodological trade-off and its consequences are examined in Section VI-C.

Six thousand hourly records were generated. An hour index is drawn uniformly, and a daylight factor is formed as the positive part of a sinusoid peaking at solar noon. A cloudiness factor is drawn uniformly on the unit interval. Global horizontal irradiance, air temperature, relative humidity, and wind speed are then derived from these two latent quantities with additive Gaussian perturbation; wind speed is drawn from a Weibull distribution with shape parameter 2, which is the standard representation of wind resource at a site [18].

 This work is licensed under a Creative Commons Attribution 4.0 International License

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.151087

Photovoltaic output for a plant of rated capacity C_s is computed as

$$E_s = C_s(G / G_{ref})[1 - \beta(T_c - 25)][1 - \gamma\, max(0, H - 40)] \quad (1)$$

where G is irradiance, $G_{ref} = 1000$ W m^{-2} is the standard test value, H is relative humidity, $\beta = 0.004$ K^{-1} is the power temperature coefficient, and $\gamma = 8 \times 10^{-4}$ accounts for the mild reduction in transmitted irradiance under humid, hazy conditions. Cell temperature is obtained from ambient temperature T_a by the linear irradiance rise

$$T_c = T_a + kG, \quad k = 0.03\ K\ m^2\ W^{-1} \quad (2)$$

which reproduces the familiar result that a module operates well above air temperature in strong sunlight and that its efficiency falls accordingly [17].

Turbine output follows the conventional piecewise power curve [18], with cut-in speed $v_{in} = 3$ m s^{-1}, rated speed $v_r = 12$ m s^{-1}, and cut-out speed $v_{out} = 25$ m s^{-1}:

$$f(v) = 0, \quad v < v_{in} \ \text{ or } \ v > v_{out} \quad (3)$$

$$f(v) = (v^3 - v_{in}^3) / (v_r^3 - v_{in}^3), \quad v_{in} \le v \le v_r \quad (4)$$

$$E_w = C_w\, f(v) \cdot 288.15 / (273.15 + T_a) \quad (5)$$

where the final factor is the air-density correction that makes a turbine slightly more productive in cold air. The clipped-cubic form of (3) and (4) is the single most consequential property of the target, and Section V returns to it.

Zero-mean Gaussian noise of standard deviation $\sigma = 2.0$ kWh was added to both targets to represent measurement error and unmodelled effects such as soiling, shading, and wake losses. This establishes an irreducible error floor: no estimator, however flexible, can achieve an expected RMSE below approximately 2.0 kWh on these data. Finally, controlled imperfections were injected so that the preprocessing layer has genuine work to perform—relative humidity was removed from 1 % of records, and wind speed was set to a physically impossible 95 m s^{-1} in 0.3 % of records to emulate sensor failure. Table II summarises the resulting dataset.

TABLE II. REFERENCE DATASET AND FEATURE SUMMARY

Property	Solar target	Wind target
Records generated	6,000	6,000
Records discarded in cleaning	78 (1.30 %)	78 (1.30 %)
Records retained	5,922	5,922
Training / test partition	4,737 / 1,185	4,737 / 1,185
Target mean (kWh)	19.06	18.30
Target standard deviation (kWh)	23.22	24.80
Target maximum (kWh)	82.73	100.30
Injected noise s.d. (kWh)	2.00	2.00
Inputs	Temperature, relative humidity, wind speed, global horizontal irradiance	

B. Preprocessing and Feature Selection

The four environmental parameters named in the objectives of [13] constitute the feature set, and no derived or lagged feature was introduced. This restraint is deliberate: the live weather service supplies precisely these four quantities for any location, so a model confined to them can be applied anywhere without additional instrumentation. The cleaning rules of Section III-B removed 78 of the 6,000 generated records, and the retained 5,922 were divided into 4,737 training and 1,185 test records in the ratio 80:20 with a fixed random seed, matching the partition specified in the framework of [13].

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.151087

The partition is random rather than chronological. For a genuine time series this would be inadmissible, because temporal autocorrelation lets a randomly partitioned model interpolate between neighbouring observations and so report an optimistic error [8]. The records here are drawn independently rather than sequentially and carry no temporal ordering, so random partitioning introduces no leakage. Once measured plant data with real time ordering are adopted, a chronological partition becomes mandatory; this is stated as a condition of the next stage in Section VII.

C. Candidate Models

Four algorithms spanning three distinct inductive biases were compared, following the reasoning in [13] that model choice should follow evidence from the target site rather than precede it. Linear regression serves as an interpretable baseline. A decision tree provides an axis-aligned piecewise-constant fit. A random forest [15] averages an ensemble of such trees and reduces their variance. Support vector regression with a radial basis kernel [16] provides a smooth non-linear fit that is well suited to moderate sample sizes. Each estimator is embedded in a pipeline with the standardising transform of Section III-B. Configurations are listed in Table III; all remaining hyperparameters retain the defaults of the library [14].

TABLE III. CANDIDATE MODEL CONFIGURATIONS

Algorithm	Configuration
Linear regression	Ordinary least squares, no regularisation
Decision tree	Maximum depth 10, squared-error criterion
Random forest	150 trees, unrestricted depth, bootstrap aggregation
Support vector regression	Radial basis kernel, C = 200, ε = 0.5

D. Training Protocol and Model Selection

Each of the four algorithms was fitted to the training partition twice, once per target, giving eight fitted estimators. Every estimator saw identical training records, identical preprocessing, and identical scoring, so that the comparison isolates the effect of the inductive bias. The estimator attaining the highest coefficient of determination on the test partition was selected for each target, with the lower root-mean-square error breaking ties. The selected estimators were serialised for use by the prediction layer, and the complete evaluation report, including the scores of the models that were not selected, was written to disk and is served to the interface so that the comparison remains visible to the user rather than being discarded after selection.

E. Evaluation Metrics

Three metrics are reported, computed on the n = 1,185 held-out records, where y_i is the reference value, $\hat{y}_i$ the prediction, and $\bar{y}$ the reference mean:

$$MAE = (1/n)\,\Sigma\,|y_i - \hat{y}_i| \tag{6}$$

$$RMSE = [(1/n)\,\Sigma\,(y_i - \hat{y}_i)^2]^{1/2} \tag{7}$$

$$R^2 = 1 - \Sigma\,(y_i - \hat{y}_i)^2 \,/\, \Sigma\,(y_i - \bar{y})^2 \tag{8}$$

Mean absolute error is reported because it is expressed in the operational unit of kilowatt-hours and is read directly by a planner. Root-mean-square error is reported because it penalises the large deviations that matter most for reserve procurement, and because it is directly comparable with the injected noise standard deviation of Section IV-A. The coefficient of determination is reported because it is scale-free and therefore permits comparison between the solar and wind targets, whose magnitudes differ. Errors are additionally normalised by the target mean and by the 100 kW rated capacity, since an absolute error in kilowatt-hours is meaningless without the scale of the plant that produced it.

V. RESULTS

A. Comparative Model Performance

Table IV reports all three metrics for all four algorithms on both targets, and Fig. 2 presents the same data graphically. Support vector regression was selected for solar output with MAE 1.340 kWh, RMSE 1.795 kWh and R^2 0.9940; the random forest was selected for wind output with MAE 1.484 kWh, RMSE 1.932 kWh and R^2 0.9939. The fact that the two targets select different algorithms is itself a result: it supports the design decision of [13] to place model comparison

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.151087

inside the system rather than to fix a single algorithm in advance, and it is consistent with the observation of [5] that different models suit different aspects of the same problem.

Normalised for scale, the selected solar model attains a mean absolute error of 7.03 % of the mean hourly solar output and 1.79 % of the rated 100 kW capacity; the selected wind model attains 8.11 % of the mean hourly wind output and 1.93 % of rated capacity. Fig. 3 shows predictions against reference values for both selected models; the scatter lies on the identity line across the full operating range, with no visible curvature or heteroscedastic fan, indicating that neither model exhibits a systematic bias at high or low output.

TABLE IV. COMPARATIVE PERFORMANCE ON THE HELD-OUT TEST PARTITION (N = 1,185)

Algorithm	MAE	RMSE	R^2	MAE	RMSE	R^2
	Solar energy (kWh)			Wind energy (kWh)		
Linear regression	1.510	1.947	0.9930	8.088	9.996	0.8375
Decision tree	1.641	2.200	0.9910	1.647	2.193	0.9922
Random forest	1.409	1.832	0.9938	**1.484**	**1.932**	**0.9939**
Support vector regression	**1.340**	**1.795**	**0.9940**	1.506	2.002	0.9935

Selected models are shown in bold. All values are computed on the same held-out partition under identical preprocessing.

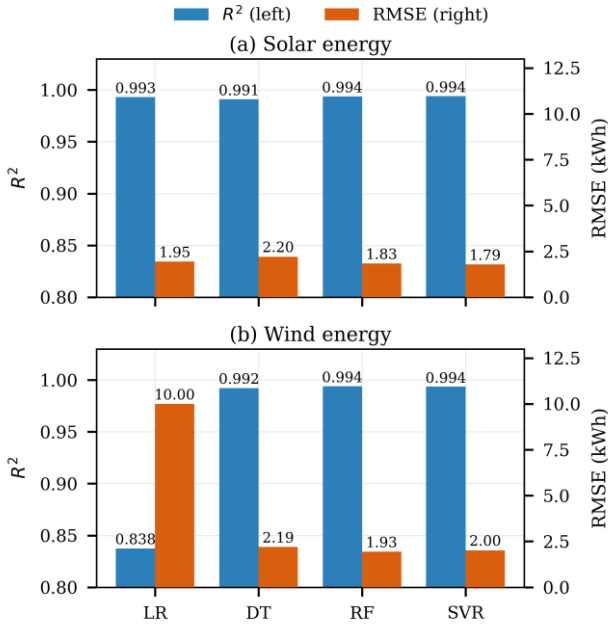


Fig. 2. Coefficient of determination (left axis) and root-mean-square error (right axis) for the four candidate algorithms: (a) solar target, (b) wind target. The vertical scale is truncated at $R^2 = 0.80$ to make the differences among the non-linear models visible. Linear regression is the only algorithm whose behaviour differs materially between the two targets.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.151087

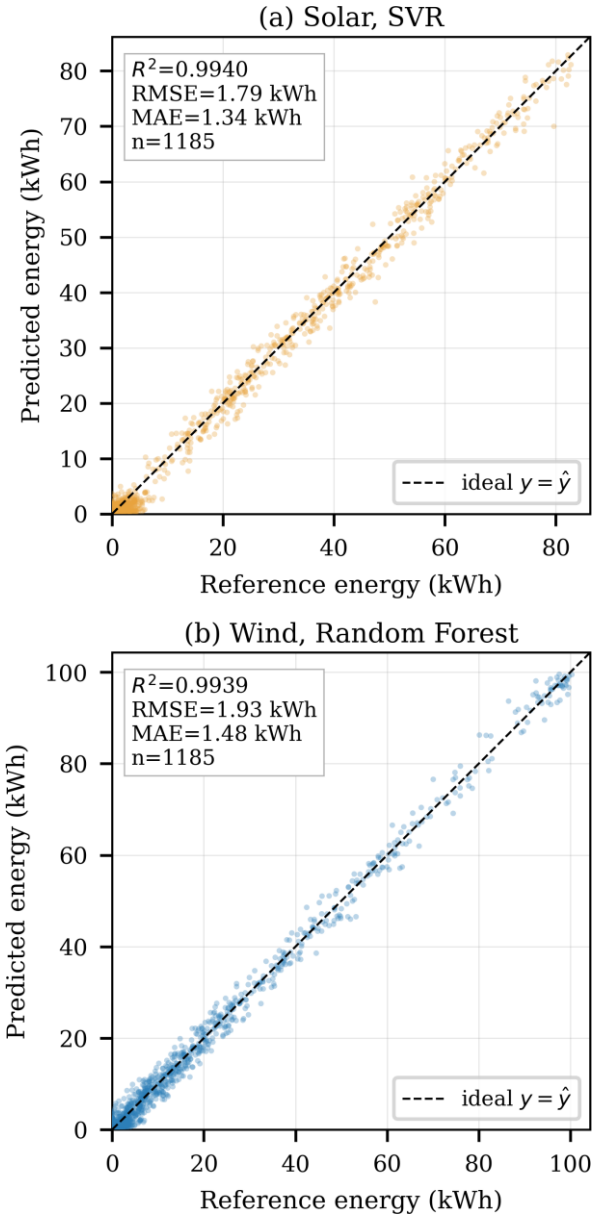


Fig. 3. Predicted against reference energy on the held-out partition for the selected models: (a) support vector regression for solar output, (b) random forest for wind output. The dashed line is the identity. The dense cluster at the origin in (a) comprises night-time records, which constitute 28.2 % of the partition.

B. Error Structure and the Noise Ceiling

The residual distributions of the two selected models are shown in Fig. 4. For the solar model the residual standard deviation is 1.775 kWh with a small positive mean of +0.266 kWh; for the wind model it is 1.932 kWh with a mean of +0.005 kWh. Both values sit at or marginally below the injected noise standard deviation of 2.0 kWh.

This observation governs the interpretation of the whole comparison. Because σ = 2.0 kWh of irreducible noise was added to the targets, an estimator that recovered the generating function exactly would still produce an RMSE of approximately 2.0 kWh. The random forest, the support vector machine, and the decision tree all report RMSE between 1.795 and 2.200 kWh, which means that each has extracted essentially all of the learnable signal and that the residual differences among them are of the order of the noise itself. The correct conclusion is therefore not that support vector regression is superior to the random forest for solar forecasting in general, but that on these data the three non-linear families are indistinguishable, and that the selection between them is weakly determined. The slight excess of the solar model below the noise floor is attributable to the deterministic zero-output constraint of Section III-C, which eliminates

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.151087

residual variance over the 28.2 % of records that are night-time observations; the corresponding small positive residual mean reflects the asymmetric clipping of negative noise at zero output.

Reporting the noise floor alongside the error metrics is, in the authors' view, a practice that the reviewed literature would benefit from adopting. Several of the studies in Table I report coefficients of determination above 0.98 without establishing what the attainable ceiling on their data is, which leaves the reader unable to judge whether a reported improvement is meaningful or lies within the noise.

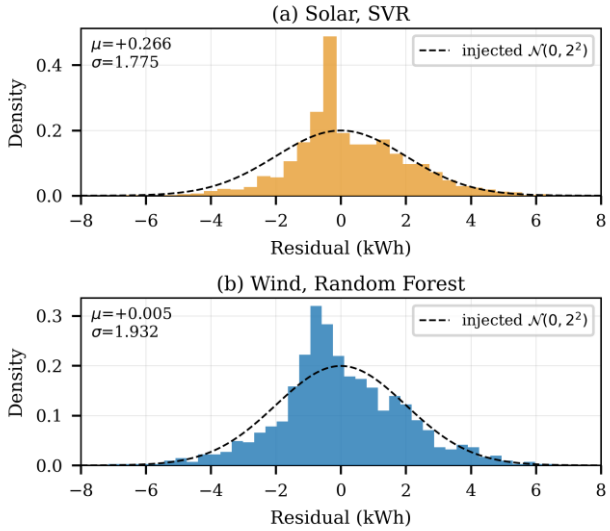


Fig. 4. Residual distributions of the selected models against the injected noise distribution (dashed): (a) solar, (b) wind. Both empirical distributions are approximately centred and have a standard deviation close to the irreducible σ = 2.0 kWh, indicating that the fitted models operate at the noise ceiling of the reference data.

C. Regime-Wise Error of the Wind Model

Because the turbine power curve of (3) and (4) is defined piecewise, the error of the wind model was decomposed by operating regime. Below the cut-in speed of 3 m s^{-1} (237 test records) the mean absolute error is 0.874 kWh; within the cubic region between cut-in and rated speed (910 records) it is 1.637 kWh; at or above rated speed (38 records) it is 1.629 kWh. The pattern is what the physics predicts. Below cut-in the target is identically zero and only the additive noise remains, so the error is smallest. In the cubic region the target varies steeply with wind speed, and a given error in the input is amplified into a larger error in output. Above rated speed the curve is flat once more, and the error returns to the level of the cubic region only because the sample is small. That the model reproduces this structure rather than distributing its error uniformly indicates that it has learned the shape of the power curve rather than an average level.

D. Feature Attribution

Normalised permutation importances, computed over eight permutation repetitions on the test partition, are shown in Fig. 5. For the solar model, global horizontal irradiance accounts for 99.38 % of the attributable loss in fit, with temperature at 0.39 %, relative humidity at 0.22 %, and wind speed at 0.01 %. For the wind model, wind speed accounts for 99.97 %, with the remaining three inputs below 0.03 % in aggregate.

These attributions recover exactly the dominant terms of (1) and (5), and they agree with the empirical finding of [7] that solar radiation and wind speed are the most influential inputs for the two sources respectively. Two remarks follow. First, the agreement is a verification of the interpretation layer rather than a discovery about the physical world: because the generating function is known here, the attribution can be checked against ground truth, which is not possible on measured data. Second, the temperature and humidity terms of (1) are genuinely present in the target but contribute under half a percent of the explained fit, which illustrates that a variable can be physically real and still be statistically negligible over the observed range. An operator reading the dashboard is therefore told, correctly, that irradiance is what governs solar output at this site.

 This work is licensed under a Creative Commons Attribution 4.0 International License

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.151087

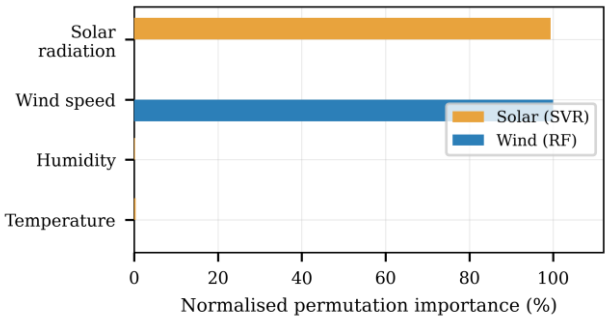


Fig. 5. Normalised permutation importance of the four inputs for each selected model, computed over eight repetitions on the held-out partition. Each model concentrates almost all attributable fit on the physically dominant driver of its energy source.

E. Generalisation Stability

Five-fold cross-validation over the full retained dataset yields a mean coefficient of determination of 0.9942 with standard deviation 0.0001 for the solar model, and 0.9939 with standard deviation 0.0003 for the wind model. The narrowness of these intervals indicates that the reported single-partition scores are not artefacts of a favourable split. Training-set and test-set scores were also compared to check for overfitting: the solar model scores 0.9951 on the training partition against 0.9940 on the test partition, a gap of 0.0011, and the wind model scores 0.9992 against 0.9939, a gap of 0.0053. The wider gap for the random forest is the expected signature of an unrestricted-depth ensemble, which very nearly interpolates its training data; that its test score is nonetheless the highest of the four candidates shows that the bootstrap averaging has controlled the resulting variance effectively.

F. Computational Cost

Inference latency was measured on a commodity workstation without hardware acceleration and is reported in Table V. A single-record solar prediction takes 0.94 ms and a twenty-four-hour solar forecast 3.00 ms. The corresponding wind figures are 36.64 ms and 45.17 ms; the random forest is approximately thirty-nine times slower per record than the kernel model because a prediction requires traversal of 150 trees. Even so, a complete twenty-four-hour forecast for both sources is produced in under 50 ms, which is negligible beside the network latency of the weather query that precedes it. The system therefore meets the lightweight-model objective identified as an open problem in [7] without recourse to specialised hardware, and the cost of retaining the slower but more accurate wind model is operationally irrelevant at this horizon.

TABLE V. DIAGNOSTICS OF THE SELECTED MODELS

Diagnostic	Solar (SVR)	Wind (RF)
Test R^2	0.9940	0.9939
Training R^2	0.9951	0.9992
Five-fold CV R^2 (mean ± s.d.)	0.9942 ± 0.0001	0.9939 ± 0.0003
Residual mean (kWh)	+0.266	+0.005
Residual s.d. (kWh)	1.775	1.932
95th percentile |residual| (kWh)	3.70	4.03
MAE as % of target mean	7.03	8.11
RMSE as % of rated capacity	1.79	1.93
Single-record inference (ms)	0.94	36.64
24-hour forecast inference (ms)	3.00	45.17

 This work is licensed under a Creative Commons Attribution 4.0 International License

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.151087

G. Live Deployment Observation

The assembled system was exercised against the live weather service from a client at Yelahanka, Bengaluru (13.07° N, 77.59° E) at 22:45 India Standard Time on 17 September 2026. The browser supplied the coordinates to an accuracy of 78 m, and the service returned an air temperature of 22 °C, relative humidity of 93 %, wind speed of 3.1 m s^{-1}, and irradiance of 0 W m^{-2}. The system predicted 0.00 kWh of solar output and 0.29 kWh of wind output for the current hour, and a twenty-four-hour total of 602.04 kWh with a peak of 80.94 kWh at 13:00 and a capacity factor of 12.5 %. Fig. 6 reproduces the resulting view.

Three properties of this run are worth recording. Solar output is exactly zero at an irradiance of zero, confirming that the physical constraint of Section III-C is enforced on live input. Wind output is small but non-zero at 3.1 m s^{-1}, which lies just above the 3 m s^{-1} cut-in threshold, so the model is operating in the correct regime of the learned power curve rather than extrapolating. The interpretation layer reported that generation was limited by darkness and that the twenty-four-hour capacity factor was low, and it identified the expected peak hour, which are the three statements an operator would need at that moment. It must be emphasised that this is a demonstration of end-to-end correctness and not a validation of accuracy: no measured generation exists for this location against which the forecast could be scored.

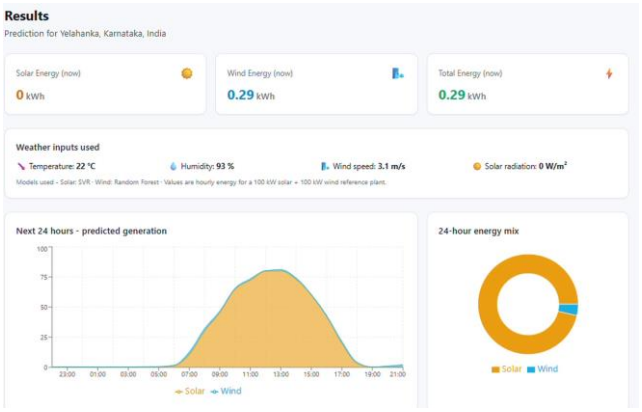


Fig. 6. Results view of the deployed application for the live run of Section V-G. Weather values were obtained automatically from the user's position; the identity of the model applied to each source is displayed alongside the forecast.

VI. DISCUSSION

A. Why the Linear Model Fails Only for Wind

The most informative single result in Table IV is the behaviour of linear regression. It is nearly the equal of the non-linear models on the solar target, at R^2 0.9930 against 0.9940 for the selected model, yet it collapses on the wind target, at R^2 0.8375 with an RMSE of 9.996 kWh—five times the noise floor and approximately ten per cent of rated capacity. The explanation lies in the functional forms of (1) and (4). Photovoltaic output is proportional to irradiance, modulated by two weak multiplicative corrections; over the observed range this is close to a linear function of a single input, and a linear model captures it almost completely. Turbine output, in contrast, is zero below cut-in, proportional to the cube of wind speed between cut-in and rated speed, constant above rated speed, and zero again above cut-out. No affine function of wind speed can represent a curve with two flat segments and a cubic segment between them, and the least-squares fit is forced into a compromise that is wrong in every regime simultaneously.

The practical consequence is a caution against transferring a modelling conclusion from one renewable source to another. Had this study examined solar output alone, as several of the works in Table I do, linear regression would have appeared entirely adequate and the additional complexity of an ensemble or kernel method would have seemed unnecessary. Extending that conclusion to wind would have produced a forecast with an error of the order of ten per cent of plant capacity. This is the strongest evidence this study offers for the position taken in [13], namely that comparison must be performed for each target and each site rather than inherited from the literature.

B. Status of the Previously Identified Deficiencies

The framework of [13] was organised around six deficiencies observed in the reviewed literature. Since that framework is the specification against which the present implementation should be judged, the status of each deficiency is stated in Table VI, and the assessment is deliberately conservative.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.151087

Three deficiencies are closed. Live weather input (G4) is operational: forecasts are produced for future hours at an arbitrary location from a public service, not reproduced from a stored file. Actionable guidance (G5) is generated as a set of deterministic statements about the limiting condition, the expected peak hour, and the adequacy of the twenty-four-hour capacity factor, accompanied by aggregate cost and emission indicators. A user-facing application (G6) exists as a six-view interface with charts, a printable report, and a persistent history, although it has no authentication and so does not yet meet the secure-login requirement stated in [13].

One deficiency is partially addressed. Explanation of predictions (G2) is provided as global permutation importance per model, which tells the user which inputs govern output at the site. It does not yet provide the per-forecast attribution that [13] specified, namely an account of why a particular hour's forecast took the value it did. A local attribution method would be required for that.

Two deficiencies remain open and should not be represented otherwise. No confidence range accompanies the forecast (G1): the system reports a point value, which is precisely the limitation for which the reviewed studies were criticised. Quantile regression or a bootstrap ensemble interval would supply this, and it is the single most valuable addition available to the system, since the residual statistics of Table V already characterise the error distribution that such an interval would have to represent. Automatic model updating (G3) is likewise absent: retraining is performed manually by re-executing the training script, and although every forecast and its weather inputs are persisted, no procedure yet compares them against realised generation or triggers retraining when accuracy degrades.

TABLE VI. STATUS OF THE SIX DEFICIENCIES IDENTIFIED IN [13]

ID	Requirement	Status	Evidence or obstacle
G1	Confidence range with each forecast	Open	Point forecast only; requires quantile regression or a bootstrap interval
G2	Explanation of a given forecast	Partial	Global permutation importance implemented (Fig. 5); per-instance attribution absent
G3	Periodic accuracy check and retraining	Open	Retraining is manual; forecasts are persisted but not scored against realised output
G4	Live weather input	Closed	Geolocation and public weather service; verified in the live run of Section V-G
G5	Actionable operating guidance	Closed	Rule-based statements on limiting condition, peak hour and capacity factor, with cost and emission indicators
G6	User-facing application	Closed	Six-view web application with charts, report and history; authentication not implemented

C. Limitations and Threats to Validity

The principal limitation of this study is the nature of its reference data, and it constrains every accuracy figure reported above. The dataset was synthesised from the performance models of (1) to (5) with additive noise, so the learning task was to recover a known smooth function corrupted by noise of known magnitude. A measured plant record would additionally contain soiling that accumulates and is cleaned, partial shading, inverter clipping and outages, turbine yaw and wake effects, curtailment instructions, and weather-forecast error that is itself autocorrelated. None of these is represented. The coefficients of determination near 0.994 therefore quantify the capacity of the models and the correctness of the pipeline; they are not a prediction of field accuracy, and they should not be compared directly with the values reported by [3]–[7] on measured archives. The authors regard restating this point as preferable to allowing a favourable number to stand without qualification.

Four further limitations follow. First, the reference plant is a notional 100 kW solar array co-located with a 100 kW turbine; outputs scale with these capacities and no real installation was characterised. Second, the feature set is restricted to four instantaneous quantities, so the models cannot represent temporal dependence, and the sequence models examined in [3] and [7] address a class of structure that is entirely absent here by construction. Third, the accuracy of a live forecast

DOI: 10.17148/IJARCCE.2026.151087

is bounded by the accuracy of the external weather service [12], which lies outside the control of this work; an error in forecast irradiance propagates undiminished into the solar prediction, as the attribution of Section V-D makes plain. Fourth, the random partition of Section IV-B is admissible only because the synthetic records are independent; the same protocol applied to measured time-series data would overstate accuracy [8], and a chronological partition is a precondition of the next stage.

VII. CONCLUSION AND FUTURE WORK

The five-layer forecasting framework previously proposed by the authors at the concept stage has been implemented as a working system and evaluated quantitatively. Live weather is acquired automatically from the user's geographic position, four candidate algorithms are compared under identical conditions for each energy source, the selected models are applied to produce hourly forecasts for a twenty-four-hour horizon, feature attribution and rule-based operating guidance are derived from each forecast, and results are delivered through a six-view web application with persistent history. On the held-out partition, support vector regression was selected for solar output (MAE 1.340 kWh, RMSE 1.795 kWh, R^2 0.9940) and a random forest for wind output (MAE 1.484 kWh, RMSE 1.932 kWh, R^2 0.9939), with a complete twenty-four-hour forecast for both sources computed in under 50 ms.

Three findings are of wider interest than the particular numbers. The two energy sources selected different algorithms, which supports treating model comparison as a component of the system rather than a decision taken in advance. Linear regression was adequate for solar output but failed for wind, and the failure is explained by the piecewise-cubic turbine power curve rather than by any deficiency of the data, which cautions against transferring modelling conclusions between renewable sources. The non-linear models reached the injected noise floor of the reference data, from which it follows that the differences among them are not meaningful and that error metrics should be reported together with the attainable ceiling on the data in question.

Future work is determined by Table VI and by the limitations of Section VI-C, in the following order of priority. The immediate requirement is to replace the synthetic reference data with measured generation records, for which the public archives identified in [9]–[11] are suitable, and to re-evaluate every model under a chronological partition [8]; until this is done no accuracy claim of this paper transfers to the field. The second requirement is a prediction interval (G1), obtainable from quantile regression or a bootstrap ensemble, which would turn each point forecast into the range that planning actually requires. The third is an accuracy-monitoring procedure that scores persisted forecasts against realised generation and triggers retraining when performance degrades (G3), for which the existing history store already provides the necessary record. Thereafter, per-instance attribution would complete G2, authentication would complete G6, and the addition of a sequence model and of solar-geometry features, as recommended in [3] and [4] respectively, would allow the temporal structure absent from the present feature set to be represented.

DECLARATIONS

Funding: No funding was received for this work.

Data and code availability: The reference dataset is generated deterministically by the script included with the implementation, from the fixed random seed reported in Section IV-A, and is therefore exactly reproducible. The public archives named for future work are available from the Desert Knowledge Australia Solar Centre [9], the Kaggle solar generation dataset [10], and the Kaggle wind turbine SCADA dataset [11]; live weather is obtained from Open-Meteo [12].

Conflict of interest: The authors declare that there is no conflict of interest.

Figures and tables: All figures and tables were prepared by the authors from the measurements reported in this paper. No copyrighted material has been reproduced.

REFERENCES

[1]. G. Notton, M.-L. Nivet, C. Voyant, C. Paoli, C. Darras, F. Motte, and A. Fouilloy, "Intermittent and stochastic character of renewable energy sources: Consequences, cost of intermittence and benefit of forecasting," *Renewable and Sustainable Energy Reviews*, vol. 87, pp. 96–105, 2018.

[2]. B.-M. Hodge, C. B. Martinez-Anido, Q. Wang, E. Chartan, A. Florita, and J. Kiviluoma, "The combined value of wind and solar power forecasting improvements and electricity storage," *Applied Energy*, vol. 214, pp. 1–15, Mar. 2018.



International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026
DOI: 10.17148/IJARCCE.2026.151087

[3]. Y. Alkhanafseh, T. C. Akinci, and A. A. Martinez-Morales, "A comparative analysis of time series forecasting models and the novel ELSTMD approach for renewable energy generation," *IEEE Access*, vol. 13, pp. 81537–81552, 2025.
[4]. A. R. V. Babu, N. B. Kumar, R. P. Narasipuram, S. Periyannan, and A. Flah, "Solar energy forecasting using machine learning techniques for enhanced grid stability," *IEEE Access*, vol. 13, pp. 93735–93754, 2025.
[5]. A. Alhayd and G. Todeschini, "Machine learning for solar power prediction: Leveraging large scale real solar generation data and weather inputs in Saudi Arabia," in *Proc. 14th Int. Conf. Renewable Energy Research and Applications (ICRERA)*, Oct. 2025, pp. 546–550.
[6]. S. P. Sharma and D. K. Yadav, "Renewable energy systems energy modeling using deep learning techniques," in *Proc. 2nd Int. Conf. Innovation in Technology (INOCON)*, Bangalore, India, Mar. 2023, pp. 1–6.
[7]. J. Rajarajeswaran, "Applications of artificial intelligence and machine learning for accurate forecasting and optimization of renewable energy generation," in *Proc. Int. Conf. Electronics and Renewable Systems (ICEARS)*, 2025, pp. 1886–1889.
[8]. C. Bergmeir and J. M. Benítez, "On the use of cross-validation for time series predictor evaluation," *Information Sciences*, vol. 191, pp. 192–213, 2012.
[9]. Desert Knowledge Australia Solar Centre, "Alice Springs photovoltaic system historical output data," Alice Springs, NT, Australia, 2023.
[10]. A. Kannal, "Solar power generation data," Kaggle, 2020.
[11]. B. Erişen, "Wind turbine SCADA dataset," *kaggledatasets*, Mar. 7, 2019.
[12]. P. Zippenfenig, "Open-Meteo.com weather API," 2023.
[13]. A. H. L. Arathi, S. C. S. Sinchana, C. M. Partha, N. Shobha, and K. S. Shreya, "AI-driven renewable energy generation forecasting using machine learning techniques," *Journal of Alternative and Renewable Energy Sources*, vol. 12, no. 3, pp. 1–12, Sept. 2026.
[14]. F. Pedregosa et al., "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
[15]. L. Breiman, "Random forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.
[16]. A. J. Smola and B. Schölkopf, "A tutorial on support vector regression," *Statistics and Computing*, vol. 14, no. 3, pp. 199–222, 2004.
[17]. E. Skoplaki and J. A. Palyvos, "On the temperature dependence of photovoltaic module electrical performance: A review of efficiency/power correlations," *Solar Energy*, vol. 83, no. 5, pp. 614–624, 2009.
[18]. T. Burton, N. Jenkins, D. Sharpe, and E. Bossanyi, *Wind Energy Handbook*, 2nd ed. Chichester, U.K.: Wiley, 2011.

 This work is licensed under a Creative Commons Attribution 4.0 International License