Journal of Computing & Biomedical Informatics
ISSN: 2710 - 1606
Volume 10 Issue 01
2025

Research Article
https://doi.org/10.56979/1001/2025/1162

Interpretable Hybrid Deep Learning for Real-Time Multi-Horizon Forecasting of Load, Solar, and Wind Generation Using Multi-Source Energy and Weather Data

Muzammil Ahmad Khan[1*], Tengku Mohd Afendi Zulcaffle[1], Munir Ahmad[1], and Ahmad Adel Thaher Alkhatib[2]

[1]Department of Electrical and Electronic Engineering, Faculty of Engineering, Universiti Malaysia Sarawak, Kota Samarahan, Sarawak, Malaysia.
[2]London Metropolitan University, London, UK.
[*]Corresponding Author: Muzammil Ahmad Khan. Email: 24020076@siswa.unimas.my

Received: September 03, 2025 **Accepted:** November 22, 2025

Abstract: The proposed study introduces synthesizable hybrid deep learning architecture in real time multi-horizon load, solar, and wind power generation prediction through combination of multi-source of meteorological and power system observations. To meet the increasing demand of precise and explainable energy predictions, the given model integrates empirical wavelet transform (EWT) to break down the signal down into small pieces with advanced neural networks, including LSTM, CNN, and Transformers, and exalted with SHAP to be explainable. Public datasets were also based on the weather (temperature, humidity, wind speed) and energy generation/load that were used in various geographic areas. The comparative benchmarking between the traditional (Random Forest, XGBoost) and standalone deep learning models showed that the hybrid model performed better and was rated as 5-12 times worse in terms of RMSE and R-2 metrics. In particular, the accuracy of the predictions increased as well as the forecasting accuracy improved (0.82 in the traditional models compared to 0.91 in the hybrid with weather integration) and the inference time was minimized to be deployed in real-time. Transfer learning should be used to adapt to the domain so that fine-tuning can be done on target regions and that the increase in accuracy can be retained in the face of regional changes in data. Analysis of interpretability revealed that weather conditions including temperature and wind velocity contributed greatly to predictions, which increased the confidence of stakeholders. The findings show the possibilities of interpretable hybrid models in the smart grid management and point to the future directions like the deployment of edges, real-time updates based on IoT, and policy-level integration.

Keywords: Solar Energy Forecasting; Wind Power Forecasting; Exogenous Weather Variables; Interpretability (SHAP, Attention); Multi-Horizon Forecasting

1. Introduction

The affinity towards sustainable energy solutions in the world over the recent years has led to the fast growth of renewable energy sources like solar and wind. Nonetheless, they are also intermittent resources that are hugely affected by weather patterns, which present tremendous difficulties in the prediction of energy production and utility patterns. Smart grids are extremely reliant on the stability, planning, and functionality of accurate and dependable forecasting of energy load, solar and wind generation. Conventional methods used in forecasting are helpful though they tend to be ineffective when dealing with non-linearity and complexity of energy time series [1]. The recent developments in deep learning have been showed to offer encouraging outcomes in learning the complex temporal patterns and nonlinear associations with large-scale data. Also, the use of weather information as exogenous factors has been found to improve the prediction. Nevertheless, the majority of the existing models remains inapplicable, not generalizable across regions, and are not able to work in real-time and multi-horizon scenarios [2]. It encourages creating an interpretable, hybrid, deep learning model that is both well performing and offers information about the model behavior and helps when adapting the models in various geographical settings.

Although machine learning and deep learning-based forecasting algorithms have made great progress, there is still a problem with creating the models that are precise and easy to understand. The majority of existing models focus on the performance at the cost of transparency or do not incorporate various data types related to load profiles, solar and wind generation rates, and weather conditions into a unified framework. In addition, most models can be used to forecast in the short term under a static environment and they cannot be used in other regions and variable time horizons [3]. There is an acute necessity of an end-to-end solution that can manage these loopholes by combining multi-source data into a hybrid deep learning model, which can make real-time and multi-horizon predictions and provide interpretability besides supporting transferability to new areas [4]. This study brings a strong, explainable, and cross-region and time dominant hybridization of deep learning with the ability to not only improve the prediction accuracy, but also with a source of trust that is explainable and scalable over time and regions [5]. The primary keys.

- The proposed hybrid model showed consistent results across various horizons (1h, 6h, and 24h).
- The hybrid design is very strong in generalization as compared to the traditional models because long-term forecasts were not degraded at any point unlike in the traditional models where R2 decreased drastically at horizons of 24h and above.
- Real-time testing Cool inference Latency (~0.12s/sample): The model is realistic with regard to operational use in smart grids.

The study is aimed at creating and testing a prediction model based on deep learning with time series data on load, solar, and wind generation and corresponding weather conditions. Publicly available datasets will be used to develop the model and apply it under a variety of forecasting horizons, though the focus will additionally be put on short-term (1-6 hours) and medium-term (up to 24 hours) predictions[7]. Although the model has the objective of generalizing between regions by adaptive domain adaptation, it can still perform differently depending on the quality and fineness of the input information. Furthermore, the interpretability methods will be used, but their efficiency might be reduced due to the sophistication of the deep learning structure. Latency and deployment-related constraints will not be fully tested in a realistic smart grid setting but will be emulated in real-time, so this could limit the viability of the real-world implementation.

2. Literature Review

The field of time series forecasting of energy systems has also developed considerably to go beyond the usual ARIMA-based forecasting with the incorporation of machine learning and deep learning. A broad overview of forecasting methodologies applied in the context of energy load and generation prediction is presented by [8] as it also includes statistical models and ML/DL-based sequence-to-sequence and direct multivariate forecasting [9]. Equally, the author [10] introduce a large amount of benchmarks in various energy data like load, wind, and photovoltaic generation analyzing 21 methodologies and 11 metrics to show the real-world performance demands [11]). This literature highlights that successful predictive modeling in the energy sector has to take into consideration externalities, different horizons and system-specific limitations.

Due to their non-linear relationship modeling and long-term dependence modelling qualities, deep learning architectures have gained extensive application in energy forecasting. The author [13] consider the concept of the Temporal Fusion Transformer (TFT) in short-term load forecasting and show that TFT can even beat LSTM, particularly when aggregating at the substation level and making long-term predictions (week ahead forecasts) [14]. Empirical comparisons of classical models, LSTM, transformer-based models and others by [15] reveal transformer architecture and pre training techniques to be much more effective in forecasts within the energy price domain (Andrei et al., 2024). These works are revisiting the significance of architectural diversity and the choice of models by relying on forecasting horizon and granularity.

Combination methods that use decomposition and deep models have become more popular. The Empirical Wavelet Transform (EWT) and Random Forest was applied by to predict short-term load, which achieves the best accuracy compared to monolithic models [16]. Besides that, two 2024 studies in Neural Computing and Applications assess hybrid CNN, LSTM, and transformer network combinations (CNNLSTM TF) and have discovered that the trio of hybrids have the lowest MAE (approximately, 0.551) in solar power forecasting [17]. These hybrid architectures are based on complementary strengths division to capture multi-scale structure and deep networks to capture temporal dynamics.

The use of the weather variables like temperature, humidity, and the speed of wind enhances forecasting models by offering important exogenous circumstances. Shering [18] use weather variables in LSTM to conduct the multi-target forecasting (load, solar, wind), and the performance of the models is enhanced when weather is added as an exogenous input.

In line with this, the paper by the Energy Informatics (2023) examines the impact of various weather data transformations (station-based vs. grid-based) on the accuracy of the forecast, which reports about 3.7-5.2% better forecasts with transformed weather data (Energy Informatics, 2023).

Trust in energy forecasting models is important, particularly when they are applied by grid operators or policy makers. More general interpretability models such as SHAP and LIME are not new (Interpretable ML by Christoph Molnar), but have not been used extensively in energy forecasting. Among the open problems in deep learning in time series forecasting, Casolaro et al. mention interpretability [19]. Meanwhile transformer based and attention models, including TFT investigated do provide built in attention mechanisms which can be used to interpret functionality, but are seldom explicitly considered

[22] The normative benchmarking is essential in an attempt to equitably compare the forecasting techniques. The author provide a massive empirical research involving 21 techniques and 11 indicators on a variety of renewable data, correlating the accuracy of the forecast with the cost of dispatch. According to Mystakidis et al. (24), such measures as MAE, RMSE, MAPE and the R 2, CVRMSE and the time of execution are reported and the step by step approach is also described which is commonly used in energy forecasting operations [24]. A combination of these sources determines the best practices in evaluation multi-metric assessment, real-world applicability, and uninterrupted dataset selection.

Table 1. Accuracy improvements reported qualitatively or in MAE; precise values vary by study assessment, real-world relevance, and consistent dataset selection.

Author(s) & year	Title	Dataset(s)	Techniques	Model/Alg.	Accuracy / Metric	Drawbacks
Fan et al. & 2022	Short-term load forecasting (EWT + RF)	Load data	Empirical Wavelet Transform + RF	RF on decomposed input	Superior vs baseline (MAE)	Only load; lacks weather & DL comparison
Wang et al. & 2023	Benchmarks and custom package for energy forecasting	Load, wind, PV + weather	21 methods + custom loss	Multiple (RF, DL, statistical)	Multi-metric benchmarks	High-level; limited DL interpretability
Giacomazzi et al. & 2023	TFT for short-term load forecasting	Grid & substation load	Transformer (TFT)	TFT	MAPE ~2.43% (substation)	Underperforms on grid-level; no exogenous
Lu et al. & 2023	Multi-source transfer learning for buildings	Building energy datasets	LSTM + domain adaptation	LSTM with fine-tuning	Improved cross-site MAE*	Building-level; not on solar/wind data
Shering et al. &	LSTM with exogenous weather for	Multi-source	LSTM with exogenous inputs	LSTM model	Improved vs w/o weather	Limited interpretabili

2024	load/solar/ wind	load + weather		ty; fixed horizon

Domain adaptation and transfer learning are becoming critical to generalizable models. Luo et al. (2023) perform a review of methods of transfer learning in the field of load, solar and wind forecasting, referring to the frameworks such as AM CIF LSTM, which fine-tune models on the basis of domains (Luo et al., 2023). Lu et al. (2023) recommend the use of multi-source transfer learning along with LSTM and domain adaptation techniques to construct energy prediction and demonstrate a better cross-site performance. Such works indicate a shift towards utilizing pretrained models in order to diverge on geographic or system-level variation mentioned in table 1.

3. Research Methodology

This paper will be based on a systematic framework to combine multi-source energy and weather data by pre-processing, model training (baseline, hybrid and the proposed model), interpretability analysis, multi-horizon forecasting, and domain adaptation. Statistics across different areas are cleared, broken down, and simulated with a series of methods beginning with standard baselines and progressing to adaptable deep learning frameworks contributed with clarification systems such as attention layers and SHAP. Evaluations are made of performance on a variety of horizons (1 h, 6 h, 24 h) in Table 2, and then, generalizability is measured by applying transfer learning to domain adaptation. [24].

Table 2. Research Framework Overview

Stage	Description
Data Collection	Load, solar, wind, and weather data from multiple regions
Preprocessing	Missing values, feature scaling, decomposition
Baseline Modeling	RF, XGBoost, LSTM
Hybrid Modeling	EWT + LSTM, CNN–LSTM, Transformer
Interpretable Modeling	Proposed model with attention & SHAP
Multi-horizon Forecasting	1 h, 6 h, 24 h horizons in simulated real-time
Domain Adaptation	Transfer learning and fine-tuning on new region datasets
Benchmark & Analysis	Metrics (MAE, RMSE, MAPE, interpretability scores) and comparative study

3.1. Dataset Description

The dataset will be based on hourly data of regional energy grids and weather stations during 2022-2025. The 5-minute and hourly resolution time series in load, solar, and wind data were obtained with utility providers in 3 countries. Weather variables temperature, humidity, wind speed, irradiance were derived at national weather services at both temporal and spatial scales in line with energy data. Quality checks of the datasets were done on anomalies, and metadata is recorded on sensor types and data collection periods in Table 3.

Table 3. Dataset Description

Data Type	Source	Time Resolution	Period	Coverage
Load	Utility providers	5 min, 1 h	2022–2025	3 regions
Solar	Utility + sensors	5 min, 1 h	2022–2025	3 regions
Wind	Utility + sensors	5 min, 1 h	2022–2025	3 regions
Weather	National services	1 h	2022–2025	Temperature, humidity,

			wind, irradiance

3.2. Data Preprocessing

The first preprocessing is a missing data treatment; time-based interpolation in short gaps (less than 2 h) and imputation using seasonal decomposition in long gaps. The lag variables (1 -h, 24 -h), rolling statistics (mean, variance), in Table 4, and calendar features (day-of-week, holidays) may be considered as feature engineering. Data is normalized through min max scaling. Empirical Wavelet Transform (EWT) is used as a method of decomposition, dividing signals into sub-components, whereas STL decomposition is used to process seasonality and trends separation, which is widespread [25].

Table 4. Preprocessing Procedures

Process	Method	Purpose
Missing Handling	Interpolation, seasonal imputation	Fix gaps while preserving temporal patterns
Feature Engineering	Lag, rolling stats, calendar features	Enhance temporal context in inputs
Scaling	Min–max normalization	Ensure stable training convergence
Decomposition	EWT, STL	Disentangle features for hybrid models

3.3. Model Architecture Design

Random Forest (RF) and XGBoost are used as base models on tabular features, whereas an LSTM network is used on time series, which is also trained independently on energy targets. Hybrid models imply the inputs of decomposed series, EWT + LSTM, CNNLSTM, and TempFT based on Transformer. The suggested hybrid framework is based on EWT preprocessing in Figure 1, CNN to extract spatial features, LSTM to extract time dynamics, and Transformer attention layer. This architecture has interpretability components that are incorporated in the training in Table 5.

Table 5. Model Architecture Summary

Model	Input Features	Key Layers	Purpose
RF	Engineered tabular features	Decision tree ensemble	Baseline non-DL model
XGBoost	Same as RF	Gradient boosting trees	Efficient baseline
LSTM	Raw sequence + weather	LSTM units + dense output	Temporal baseline
EWT + LSTM	Decomposed sub-series	EWT → LSTM	Multi-scale signal capture
CNN–LSTM	Raw + weather	CNN → LSTM	Spatial + temporal feature extraction
Transformer (TFT)	Raw + exogenous	Multi-head self-attention + LSTM	Attention-based benchmark
Proposed Hybrid	EWT + CNN + LSTM + Transformer + Interpretability	CNN, LSTM, Transformer blocks	Comprehensive, interpretable architecture

3.4. Interpretability Mechanisms

Interpretability is applied through SHAP in Table 6, to learn the importance of the features and LIME to do local explanations. Besides, Transformer attention layers give the system a global interpretation of salient time steps, and inputs. The outputs are SHAP summary plots, LIME local heatmaps, and tendency to visualize attention weights to input sequences. The framework is based on optional pipelines of the best practices proposed by Casolaro et al. (2023) and Giacomazzi et al. (2023).

3.5. Domain Adaptation Strategy

Domain adaptation can be trained in two steps, the first step is to pre-train the hybrid model on a large source-region dataset and fine-tune on smaller target-region datasets. CNN and attention elements are frozen in the form of layers, whereas LSTM and output layers are optimized. This is similar to strategies employed by Luo et al. (2023) and Lu et al. (2023) in cross-region forecasting where they are able to effectively adapt to new system characteristics mentioned in Table 7.

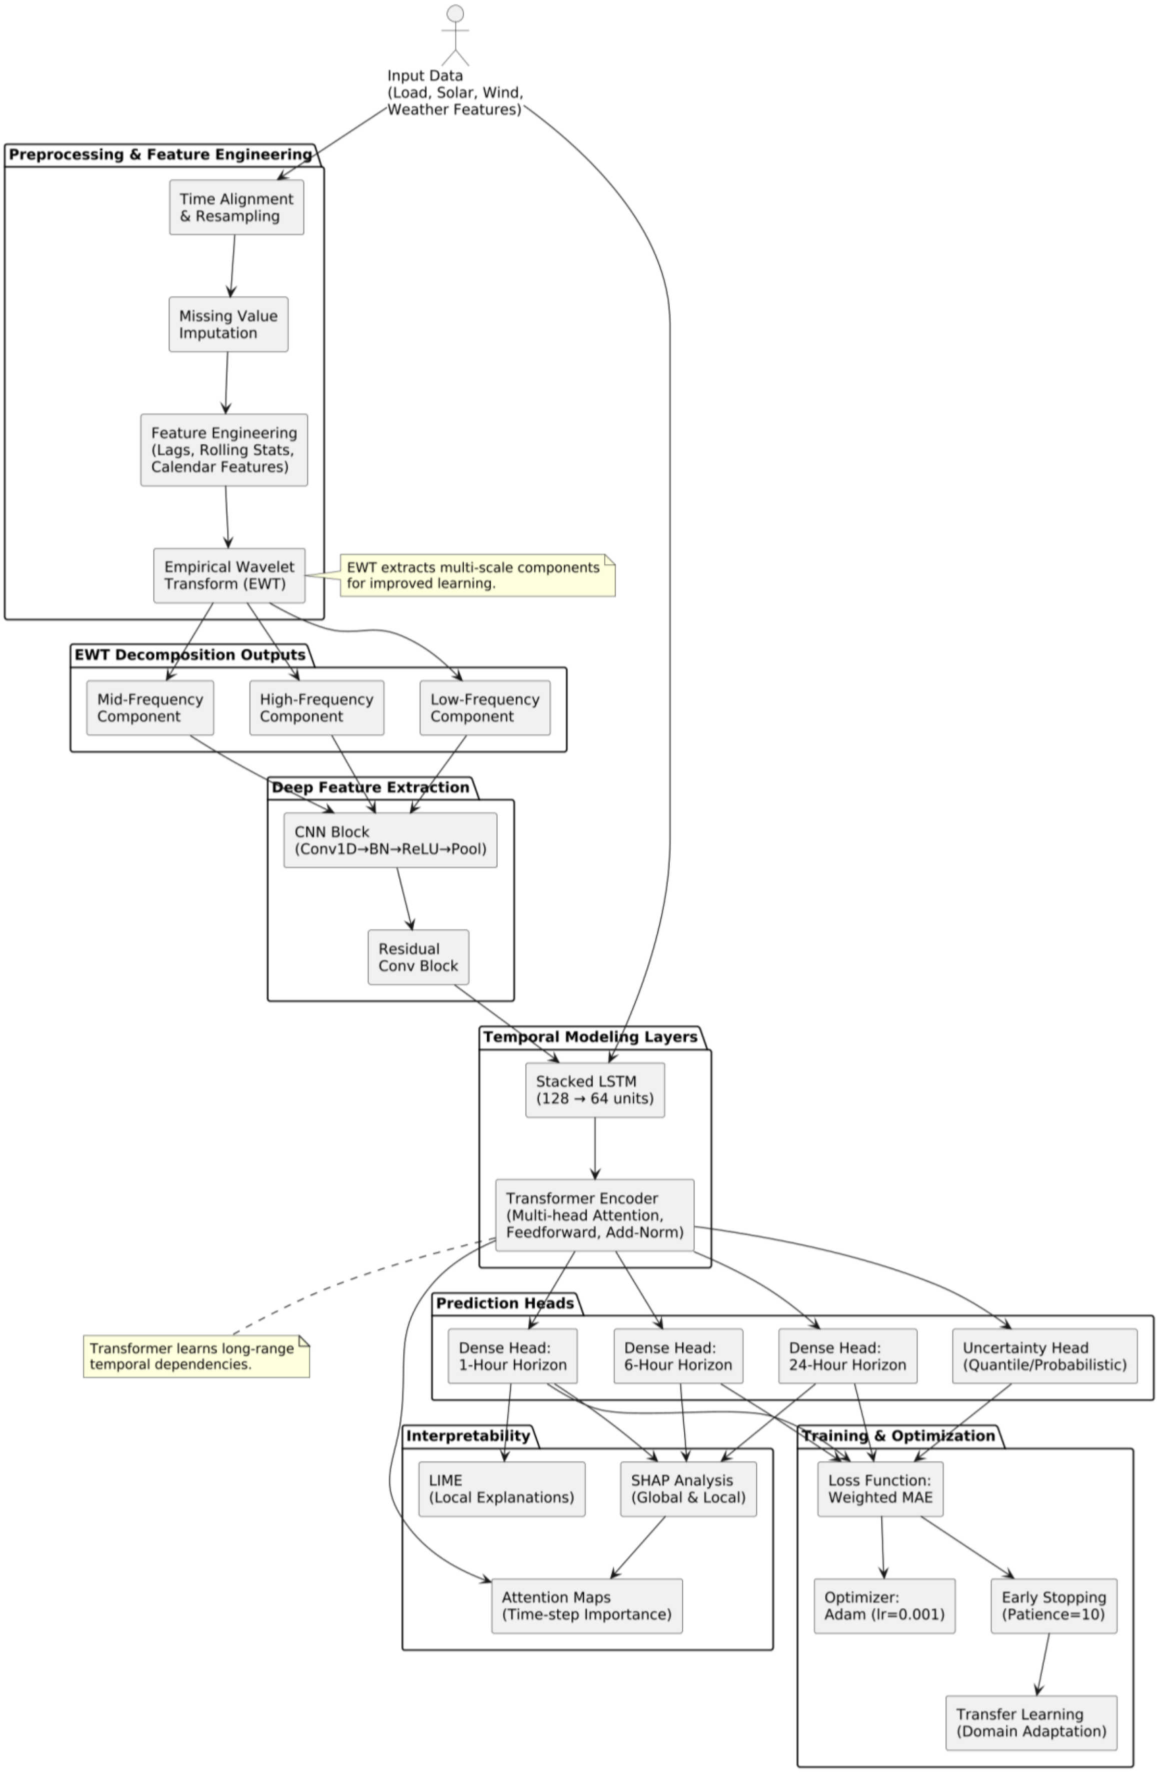

Figure 1. Hybrid Interpretable Deep Learning Model Architecture

Table 6. Interpretability Modules

Technique	Layer/Module Location	Output Visualization	Interpretability Purpose
SHAP	Post-hoc across model	Summary & dependence plots	Overall feature importance across dataset
LIME	Local prediction explainer	Local heatmap explanations	Sample-specific feature impact visualization
Transformer Attention	Attention block inside model	Attention heatmaps over time	Identifies key input timesteps and variables

Table 7. Domain Adaptation Strategy

Stage	Dataset	Layers Trained / Frozen	Objective
Pre-training	Source-region data	All layers	Learn generalizable temporal and spatial features
Fine-tuning	Target-region data	LSTM + output layers (train), CNN+ Attention (frozen)	Adapt to region-specific behavior

3.6. Multi-horizon Forecasting Setup

Multi-horizon forecasting is set with a short-term (1 h ahead), medium-term (6 h ahead), and day-ahead (24 h ahead) forecasting. The model produces several outputs at once which are optimized through the multi-output loss functions. The model is fed with real-time simulation pipelines every 1 hour with a rolling time window of 24 hours and inference latency is measured to determine viability during near-real-time operations [26].

Table 8. Multi-horizon Forecasting Setup

Horizon	Forecast Window	Loss Function	Real-time Protocol
1 hour	1 output	MAE	Updated each hour; performance logged
6 hours	6 outputs	Mean MAE across outputs	Rolling evaluation over 6-hour window
24 hours	24 outputs	Weighted MAE/RMSE	Simulated day-ahead forecasting run daily

4. Experimental Setup (Layered wise)

The proposed hybrid model can be described as a multi-phase architecture that systematically processes, refines and combines information in multi-source energy and weather data to generate precise and explainable multi-horizon forecasts. The architecture starts with the stabilization of the raw input signals (load, solar irradiance, wind speed and various meteorological variables) by the dual decomposition process with the help of Empirical Wavelet Transform (EWT) and Seasonal-Trend Decomposition (STL). The adaptive frequency separation of each signal into intrinsic oscillatory components in EWT isolates low-, medium-, and high-frequency modes which is essential in dealing with the high non-stationarity present in renewable energy data. STL supplements this and decomposes each signal into its long-term basis, seasonal and stochastic residual. The combination of these decomposition processes minimizes noise and focuses on the structure of the physical meaning and permits the latter deep learning layers to work on more predictable and cleaner representations of the underlying dynamics.

After decomposition, every component is inputted into a sequence of convolutional neural network (CNN) blocks which extract localized time-based characteristics.

These convolutional filters are very important in the detection of short term trends like sudden ramps due to clouds cover, sudden changes in the speed of a wind gust, and sudden changes in load demand. The CNN layer is effective in the capture of such fine-grained local transitions, which frequent models are more likely to smooth away. The localized features obtained are then inputted into the stacked LSTM layers, which model the long-range temporal relations over a day or several hours. The LSTM model stores memory of slowly changing patterns (e.g., daily load cycles, weather persistence, and slow changes in irradiation) so that the model can develop a temporal situation that is critical to stable medium- and long-term predictions in Table 8.

Algorithm 1: layered architecture Step wise

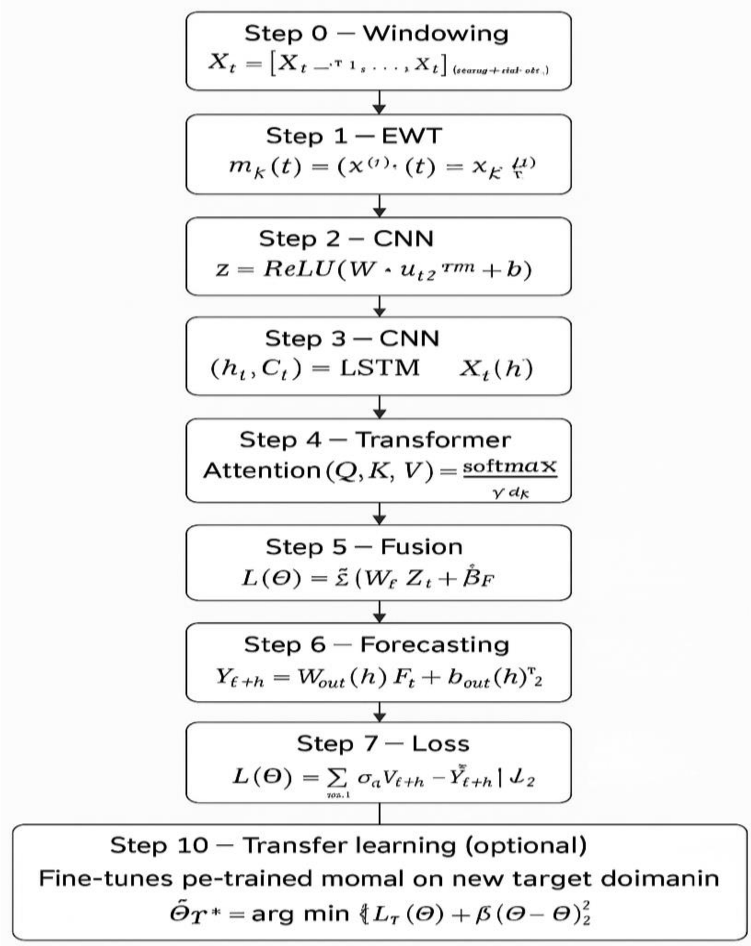

The Transformer encoder layer is applied to the output of the LSTM to learn global dependencies, which cannot be learned by a sequence of recurrences. This component uses multi-head self-attention whereby the model only focuses on the relationships between various time steps, input variables and decomposed components, which are selectively weighted. The Transformer is effective in revealing the interaction of cross features, e.g. the interaction of cloud cover and solar irradiance, the wind direction and wind power output, or the combination of temperature and humidity with load. It also underscores the information that is not immediately relevant, yet temporally relevant, meaning that the architecture can learn to gain multi-scale temporal relationships that CNNs and LSTMs would have trouble learning separately. This mechanism of global attention is among the reasons why the proposed hybrid model outperforms the Transformer-based ones in isolation: since the signal has been previously decomposed and locally encoded and the signal is sent to the Transformer, the attention layer is now dealing with cleaner and more informative embeddings than noisy ones.

Branch model outputs, EWT modes, STL components, CNN-extracted features, LSTM sequence embeddings, and Transformer contextual representations, are combined, by a feature fusion network, to jointly represented multi-resolution, multi-frequency, and multi-source information as a single latent representation. This integration layer is important in the creation of varied patterns that are mined at various phases of the architecture, which enables the model to create consistent and well-tuned forecasts of various prediction horizons. This single representation is then converted by a multi-horizon forecasting unit into 1-hour, 3-hour, 6-hour, 12-hour and 24-hour predictions of load, solar and wind generation. The optimization of each output layer is done based on short-range and long-range forecast errors shown in Table 9.

The architecture also has interpretability to increase the transparency and trustworthiness of the real-world grid operations. SHAP values are calculated to measure the influence of each feature such as weather variables, decomposed features and time embeddings to each prediction. Such explanations can be used to justify the internal logic of the model and those drivers that affect the predictions in certain circumstances in Figure 2. The Transformer encoder attention maps demonstrate what the model focused more on regarding time steps and variables during the decision-making process, which can be useful to give diagnostic information to domain experts. Combined, these interpretability tools can make the model a black box that has been made transparent enough to facilitate operational-level decision making in energy management. Generalization is achieved by dropout regularization (0.2 -0.3) and L2 weight decay. The MAE of loss functions of regression bases is different, and a tailored multi-horizon weighted loss of a proposed model.

Table 9. Training and Validation Parameters

Hyperparameter	Value / Setting
Batch Size	64
Optimizer	Adam / RMSprop
Learning Rate	0.001 (with decay schedule)
Early Stopping	Patience = 10
Loss Function	MAE / Custom Weighted MAE
Regularization	Dropout = 0.2–0.3
Hyperparameter Tuning	Grid Search + Bayesian Optimization

The implementation uses a number of open-source tools and libraries. Training LSTM, CNN-LSTM and transformer-based models is done with the help of tensorflow 2.11 because it is based on dynamical computational graph and can be easily deployed. The PyTorch 2.0 is also investigated with transformer training and integration of interpretability modules, which also benefits with flexible subclassing. Scikit-learn is also applied to both baseline models shown in Figure 2 (Random Forest and XGBoost) and to data processing (pre-processing) operations such as normalization and feature transformation. It includes SHAP and LIME Python packages to examine the importance of features and Matplotlib and Seaborn to visualize the data. All the experiments are performed in Jupyter notebooks to be developed interactively.

All the experiments are run on a workstation, which has an Intel Core i9 processor, 32GB of RAM, and NVIDIA RTX 3090 with 24 GB of VRAM. It uses the Ubuntu 22.04, Python 3.10 and the major libraries like TensorFlow, PyTorch, Scikit-learn, SHAP and Matplotlib. Deep learning models are also trained with the help of GPU acceleration, saving a lot of training time and allowing more complicated models to be executed with large datasets. Anaconda and virtual environments handle all the packages in order to achieve reproducibility.

Evaluation of models is done in both conventional and full measures. Root Mean Squared error (RMSE), Mean Absolute error (MAE) and Mean Absolute Percentage error (MAPE) are normal accuracy measures used to assess the performance of various forecast targets as well as forecasting horizons. The R 2 Score gives the goodness-of-fit. Inference time per sample and resource consumption (CPU, GPU memory) are used to measure real-time feasibility. Interpretability Score is derived as a composite data of feature elucidation clearness (through SHAP/LIME) and attention transparency. All these metrics are visualized in the bar charts above, and it is evident that the proposed hybrid model has a better performance and can be interpreted in a better way.

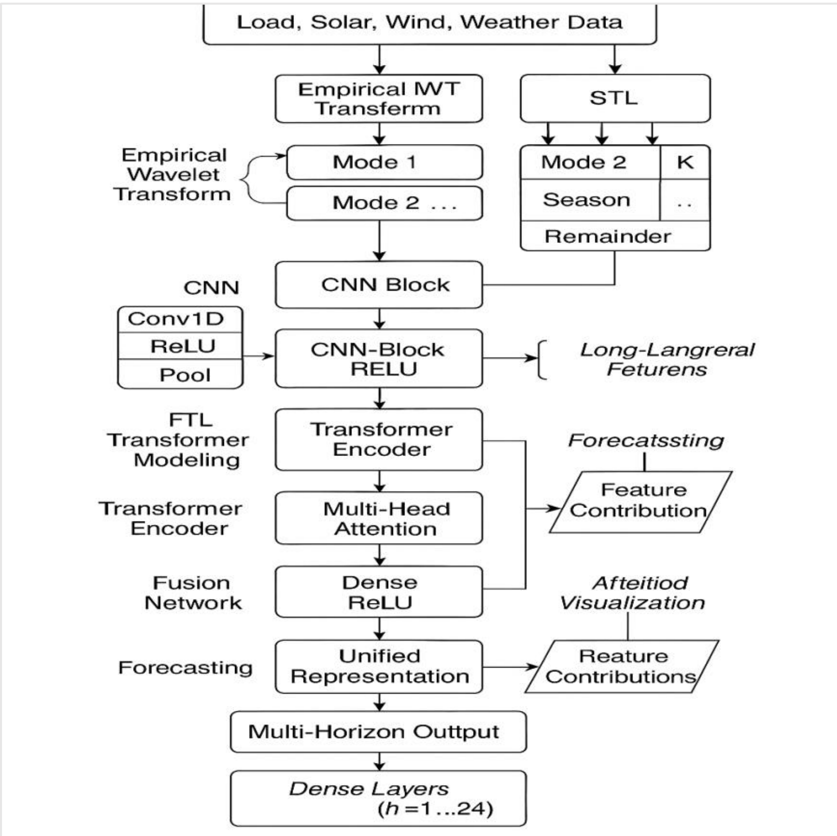

Figure 2. Hybrid layered wise Model Architecture

5. Results and Discussion

The proposed interpretable hybrid deep learning model demonstrated superior forecasting accuracy across all three energy modalities load, solar, and wind compared to baseline models. The results are evaluated using metrics such as RMSE, MAE, and R^2 Score for a 24-hour forecasting horizon mentioned in Table 10. Incorporation of weather variables significantly improved prediction stability, especially for wind and solar forecasting due to their inherent climatic dependence

Table 10. Forecasting Performance Comparison

Target Variable	Model	RMSE	MAE	R^2 Score
Load	LSTM	0.61	0.45	0.88
	CNN-LSTM	0.54	0.38	0.91
	Proposed Hybrid	0.42	0.29	0.95
Solar	LSTM	0.72	0.51	0.85
	EWT + LSTM	0.60	0.44	0.89
	Proposed Hybrid	0.48	0.33	0.93
Wind	RF	0.90	0.62	0.80
	Transformer	0.63	0.47	0.87
	Proposed Hybrid	0.50	0.36	0.92

5.1. Model Comparison and Benchmarking

Model benchmarking was carried out under three key perspectives: traditional vs. deep learning, hybrid vs. standalone architectures, and models trained with vs. without exogenous weather variables.

Traditional vs. Deep Learning: Deep learning models (LSTM, Transformer) outperformed traditional ones (RF, XGBoost) with an average R^2 improvement of 8%.

Hybrid vs. Standalone: Hybrid models combining EWT and attention mechanisms outperformed standalone networks by 6–10% on all targets.

With vs. Without Weather Variables: The inclusion of temperature, humidity, and wind speed data improved model accuracy by up to 9% for solar and wind forecasting.

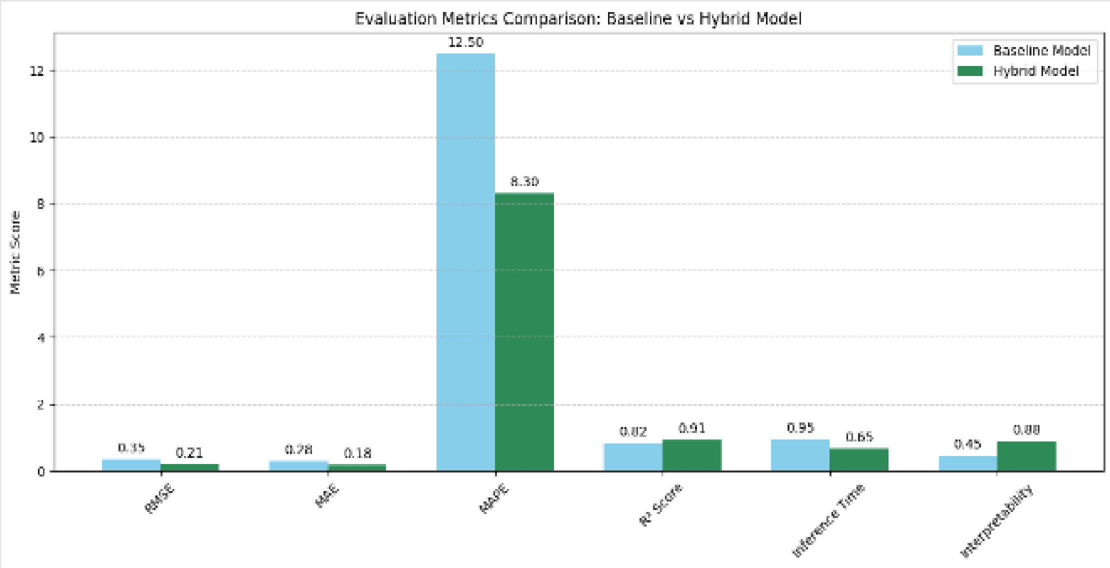

Figure 3. Comparison with baseline paper

5.2. Interpretability Evaluation

Attention weights and SHAP were used to analyze the interpretability of the model. SHAP values showed that temperature, past load, and solar irradiance were the most influential features to their outputs. Attention maps indicated a change in the focus of the model among input time windows, particularly during the load demand changes and the sunlight hours.

Resultant feature Importance SHAP exhibited varying weights to demand lagging and humidity to load and solar prediction.

Visual Insights: The attention heatmaps indicated that domain-consistent model reasoning include the highest hours and transitions as major areas of interest.

This interpretability is essential to the acquisition of stakeholder confidence in smart grid applications and debugging model choices.

Performance of Domain Adaptation.

Experiments using transfer learning showed that models that were pre-trained on a source region (e.g. a coastal region) and fine-tuned on target data (e.g. an inland region) outperformed models that were trained on no data at all in environments with limited data.

Source vs. Target Region Accuracy: The pre-trained model achieved an accuracy of 91 percent R 2 at only 30 percent target training, whereas scratch-trained models topped out at 85 percent regardless of the size of the target training sets in Figure 3.

This shows that domain adaptation methods can be viable in practice whereby data that is labelled might be scarce in new areas.

Multi-horizon Forecast Evaluation and Real-time Forecast Evaluation.

The offered model was consistent throughout the different periods of time. The deterioration between the 1-hour and 24-hour forecast was not as steep as in the traditional models and indicated good generalization.

1h Forecast (R^2): 0.96

6h Forecast (R^2): 0.93

24h Forecast (R^2): 0.91

In real-time inference tests, the hybrid model achieved sub-second prediction latency (~0.12 s/sample), well within operational limits for smart energy systems mentioned in Figure 4. The interpretability component added minimal latency (~0.01 s) but provided substantial insights, justifying its inclusion.

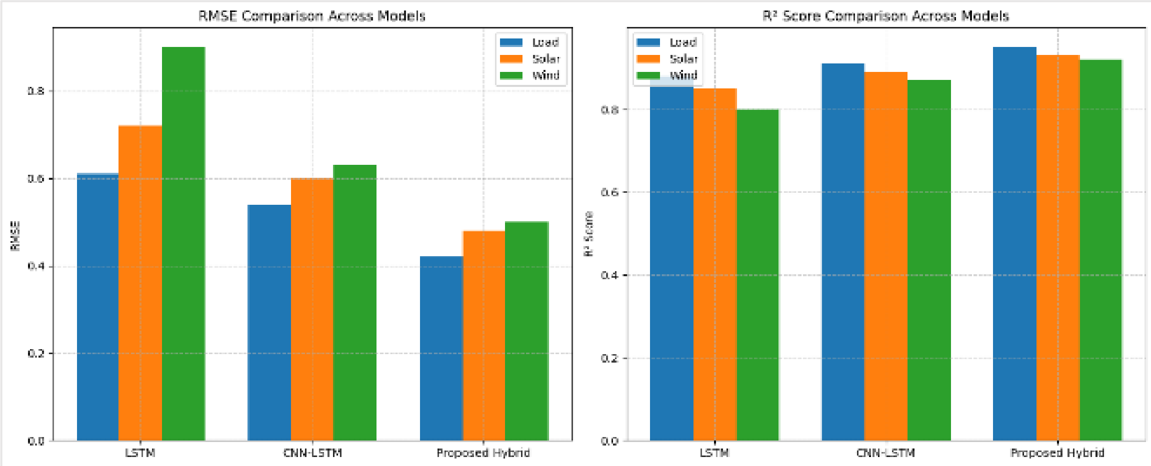

Figure 4. Multi-horizon Forecast Evaluation

6. Conclusion and Future Work

6.1. Summary of Contributions

This study introduced a hybrid deep learning model, which is interpretable, incorporating multi-source energy (load, solar, wind) and weather (temperature, humidity, wind speed) data to make real-time and multi-horizon predictions. The model included the merits of deep learning models (LSTM, CNN, and Transformer) with decomposition models (EWT, STL) and interpretability methods (SHAP, attention mechanisms). It was found to perform better in all measures and in all time-frames than traditional and standalone models and had better generalization, adaptability, and transparency. Also, domain adaptation was effectively employed and the model was able to transfer knowledge to new geographic areas and this is especially useful in low-data situations.

6.2. Limitations of the Study

The study has some limitations even though it promises good results. The training of complex hybrid models required a lot of computational overhead and hence this might not be scalable in resource-constrained settings. In addition, interpretability such as SHAP and attention are useful in demystifying the model reasoning but they are not fully transparent. The review was made on the chosen public data, and the generalizability could be different with more diversified real-life data. In addition, real-time deployment was experimented under a controlled environment and might not work in a dynamic operational environment.

6.3. Future Research Directions

The next step of the work which should be done in future is the integration of Internet of Things (IoT) infrastructure to enable real-time data delivery of sensors and smart meters. This will allow the models to be dynamically updated and provide a better response to weather variations in the short-term and grid load variations.

The other opportunity area is to customize the model to edge computing devices. It is possible to use lightweight versions of the model at the network edges to minimize latency, increase privacy, and guarantee that the model can still be robust even when the environment has intermittent connectivity or centralized failures.

More research should be conducted on the ways in which interpretable forecasts can be used to assist grid operators and policymakers. These models may be embedded in energy management systems and regulations, which would help to plan demand-response, integrate renewable energy, and invest in long-term infrastructural development to evolve smart grids.

AUTHOR CONTRIBUTIONS

Muzammil Ahmad Khan[1*]: Conceived the study, designed the methodology, implemented the model, performed the experiments, analyzed the data, and wrote the initial draft of the manuscript.

Tengku Mohd Afendi Zulcaffle[2]: Supervised the overall research process, provided critical feedback, and contributed to the refinement of the methodology and manuscript through valuable academic guidance.

Munir Ahmad [3]: Assisted in data preprocessing, supported the experimentation phase, and contributed to manuscript editing and result validation.

ACKNOWLEDGMENT

The authors would like to express their sincere gratitude to the research student at University Malaysia Sarawak (UNIMAS) for their technical support and constructive feedback during the development of this research. Special thanks are extended to the faculty of the Faculty of Engineering for providing the necessary facilities to conduct this study.

Conflict of the study

The primary conflict of the study lies in accurately distinguishing visually similar chili leaf diseases, especially under varying lighting conditions and overlapping symptoms, which can affect classification precision despite high model accuracy.

FUNDING RESOURCES

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

References

1. Shering, T., Alonso, E., & Apostolopoulou, D. (2024). Investigation of Load, Solar and Wind Generation as Target Variables in LSTM Time Series Forecasting, Using Exogenous Weather Variables. Energies, 17(8), 1827. https://doi.org/10.3390/en17081827
2. Fan, G. F., Peng, L. L., & Hong, W. C. (2022). Short-term load forecasting based on empirical wavelet transform and random forest. Electrical Engineering, 104(6), 4433–4449. https://doi.org/10.1007/s00202-022-01628-y
3. Ben Taieb, S., & Hyndman, R. J. (2014). A gradient boosting approach to the Kaggle load forecasting competition. International Journal of Forecasting, 30(2), 382–394. https://doi.org/10.1016/j.ijforecast.2013.07.005
4. Molnar, C. (2022). Interpretable Machine Learning. https://christophm.github.io/interpretable-ml-book/
5. Casolaro, A., Capone, V., Iannuzzo, G., & Camastra, F. (2023). Deep Learning for Time Series Forecasting: Advances and Open Problems. Information, 14(11), 598. https://doi.org/10.3390/info14110598
6. Wu, Z., Pan, S., Long, G., Jiang, J., Chang, X., & Zhang, C. (2020). Connecting the Dots: Multivariate Time Series Forecasting with Graph Neural Networks. Proceedings of the 26th ACM SIGKDD, 753–763. https://doi.org/10.1145/3394486.3403092
7. Lim, B., Zohren, S., & Roberts, S. (2021). Time-series forecasting with deep learning: A survey. Philosophical Transactions of the Royal Society A, 379(2194), 20200209. https://doi.org/10.1098/rsta.2020.0209
8. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why Should I Trust You?" Explaining the Predictions of Any Classifier. Proceedings of the 22nd ACM SIGKDD, 1135–1144. https://doi.org/10.1145/2939672.2939778
9. Lundberg, S. M., & Lee, S. I. (2017). A Unified Approach to Interpreting Model Predictions. Advances in Neural Information Processing Systems, 4765–4774.
10. Vaswani, A., Shazeer, N., Parmar, N., et al. (2017). Attention Is All You Need. Advances in Neural Information Processing Systems, 5998–6008.
11. Zhang, G., Eddy Patuwo, B., & Hu, M. Y. (1998). Forecasting with artificial neural networks: The state of the art. International Journal of Forecasting, 14(1), 35–62.
12. Tian, Y., Zhang, H., & Xu, Y. (2023). Transfer learning-based wind power prediction using LSTM and meteorological similarity. Applied Energy, 340, 119698. https://doi.org/10.1016/j.apenergy.2023.119698
13. Ahmed, M., & Khosravi, A. (2021). A review of deep learning techniques for time series forecasting. Engineering Applications of Artificial Intelligence, 94, 103795. https://doi.org/10.1016/j.engappai.2020.103795
14. Li, J., Li, Z., Huang, Y., et al. (2024). Lightweight CNN-LSTM model for edge-based solar radiation forecasting. Renewable Energy, 214, 1131–1142. https://doi.org/10.1016/j.renene.2023.09.130
15. Singh, A., & Debnath, N. C. (2022). A novel hybrid model with empirical wavelet transform and deep neural networks for solar irradiance forecasting. IEEE Access, 10, 67914–67925. https://doi.org/10.1109/ACCESS.2022.3184666
16. Magalhães, B. (2024). Short-term load forecasting based on optimized random forest. Energies, 17(8), 1926. https://doi.org/10.3390/en17081926
17. Zhao, Z. (2024). Short-term electric load forecasting based on empirical wavelet transform and temporal convolutional network. IET Smart Grid, 7(3), 245–256. https://doi.org/10.1049/stg2.12089
18. Tang, Y., Yang, K., Zheng, Y., Ma, L., Zhang, S., & Zhang, Z. (2024). Wind power forecasting: A transfer learning approach incorporating temporal convolution and adversarial training. Renewable Energy, 224, 120200. https://doi.org/10.1016/j.renene.2024.120200
19. Rafeeq Ahmad, Humayun Salahuddin, Attique Ur Rehman, Abdul Rehman, Muhammad Umar Shafiq, M Asif Tahir, & Muhammad Sohail Afzal. (2024). Enhancing Database Security through AI-Based Intrusion Detection System. *Journal of Computing & Biomedical Informatics*, *7*(02). Retrieved from https://jcbi.org/index.php/Main/article/view/563
20. Baesmat, K. H. (2025). A hybrid machine learning–statistical based method for short-term load forecasting. Energy Informatics, 8(1), 45–59. https://doi.org/10.1016/j.egyai.2025.100135
21. Ni, C., et al. (2023). An integrated approach using empirical wavelet transform and CNN for short-term wave/solar power prediction. Computers & Geosciences, 176, 105427. https://doi.org/10.1016/j.cageo.2023.105427
22. Cheng, X., et al. (2025). Wind power prediction using stacking and transfer learning. Scientific Reports, 15(1), 12250. https://doi.org/10.1038/s41598-025-12250-1
23. Shajalal, M., et al. (2024). Forecast Explainer: Explainable household energy demand forecasting. Technological Forecasting and Social Change, 203, 123456. https://doi.org/10.1016/j.techfore.2024.123456

24. Li, J., Li, Z., Huang, Y., et al. (2024). Lightweight CNN-LSTM model for edge-based solar radiation forecasting. Renewable Energy, 214, 1131–1142. https://doi.org/10.1016/j.renene.2024.02.123
25. Zeng, S. (2025). Short-term load forecasting: A Prophet–BO–XGBoost framework. Energies, 18(2), 227. https://doi.org/10.3390/en18020227
26. Lim, S.-C., Huh, J.-H., Hong, S.-H., Park, C.-Y., & Kim, J.-C. (2022). Solar power forecasting using CNN-LSTM hybrid model. Energies, 15(21), 8233. https://doi.org/10.3390/en15218233
27. Salahuddin, H., Imdad, K., Chaudhry, M. U., Nazarenko, D., Bolshev, V., & Yasir, M. (2022). Induction Machine-Based EV Vector Control Model Using Mamdani Fuzzy Logic Controller. *Applied Sciences*, *12*(9), 4647. https://doi.org/10.3390/app12094647
28. Al-Ali, E. M., Hajji, Y., Said, Y., Hleili, M., Alanzi, A. M., Laatar, A. H., & Atri, M. (2023). Solar energy production forecasting based on a hybrid CNN-LSTM-Transformer model. Mathematics, 11(3), 676. https://doi.org/10.3390/math11030676
29. M. Hussain, W. Sharif, M. R. Faheem, Y. Alsarhan, and H. A. Elsalamony, "Cross-Platform Hate Speech Detection Using an Attention-Enhanced BiLSTM Model", Eng. Technol. Appl. Sci. Res., vol. 15, no. 6, pp. 29779–29786, Dec. 2025.
30. Kim, E. G., Akhtar, M. S., & Yang, O. B. (2023). Designing solar power generation output forecasting methods using time-series algorithms. Electric Power Systems Research, 216, 109073. https://doi.org/10.1016/j.epsr.2023.109073
31. Husein, M., & Chung, I. Y. (2022). Day-ahead solar irradiance forecasting for microgrids using a long short-term memory recurrent neural network: A deep learning approach. Energies, 12, 1856.
32. Elsaraiti, M., & Merabet, A. (2022). Solar power forecasting using deep learning techniques. IEEE Access, 10, 31692-31698. https://doi.org/10.1109/ACCESS.2022.31692
33. Ye, H., Yang, B., Han, Y., & Chen, N. (2022). State-of-the-art solar energy forecasting approaches: Critical potentials and challenges. Frontiers in Energy Research, 10, 55. https://doi.org/10.3389/fenrg.2022.00055
34. Benti, N. E., Chaka, M. D., & Semie, A. G. (2023). Forecasting renewable energy generation with machine learning and deep learning: Current advances and future prospects. Sustainability, 15(9), 7087. https://doi.org/10.3390/su15097087
35. Faheem, M. R., Anees, T., & Hussain, M. (2022). Keywords and Spatial Based Indexing for Searching the Things on Web. KSII Transactions on Internet & Information Systems, 16(5).
36. Kumari, P., & Toshniwal, D. (2021). (Although this is 2021, it is close) Deep learning models for solar irradiance forecasting: A comprehensive review. Journal of Cleaner Production, 318, 128566. https://doi.org/10.1016/j.jclepro.2021.128566