Global Solar Irradiance Forcasting Using Deep Learning Algorithms

Ahmedalmamoun Awad, Hammad Hasan, Mohamed Atef Shahoob, Mohamed Montasir Babiker, Davor Svetinovic
*Department of Electrical and Computer Engineering
Khalifa University, Abu Dhabi
United Arab Emirates

***Abstract*—The worldwide expansion of solar PV installations is noteworthy, yet the inherent variability of solar power poses challenges for seamless grid integration. Solar irradiance, a pivotal factor influencing PV output, is thoroughly examined in this paper. This paper delves into the application of various deep learning models, including MLP, RNN, LSTM, LSTM-RNN, and Bidirectional LSTM, for day-ahead forecasting of Global Horizontal Irradiance (GHI). Using historical meteorological and solar irradiance data, the research evaluates models performance using metrics such as Mean Squared Error (MSE), Mean Absolute Error (MAE), and Root Mean Squared Error (RMSE). The findings underscore the accuracy of LSTM in predicting irradiance, offering valuable insights for effective day-ahead solar power forecasting.**

***Index Terms*—irradiance, forecasting, solar power, cloud days**

I. Introduction

In 2021, solar photovoltaic (PV) generation experienced significant growth, reaching over 1000TWh with a record increase of 179TWh (22%) [1]. Solar PV is increasingly becoming the most cost-effective choice for new electricity generation in many regions worldwide, leading to anticipated investments in the years to come.

The rapid expansion of solar power faces obstacles when it comes to widespread implementation and replacing fossil fuels. Although the average output of solar photovoltaic (PV) systems over longer periods can be anticipated by considering daily and seasonal solar variations, short-term predictions are challenging due to unpredictable weather conditions, particularly cloud coverage. In situations where there are intermittent clouds, the output of PV systems can plummet to nearly zero within minutes [2]. Presently, short-term forecasts for PV output suffer from inadequate precision, and minute-level predictions in partly cloudy conditions often exhibiting substantial errors.

Accurate short-term PV forecasts have various applications. Generation utilities can use them for real-time market bidding strategies, transmission operators and balancing authorities can determine the need for operating reserves or ancillary services, and PV users can modulate power demand to avoid peak penalties based on cloud cover.

Generally, the forecasting methods can be classified into three approaches: physical, statistical, and machine learning methods [3]. According to [4], PV power forecasting using machine learning models can be performed by two techniques. Either to predict PV power directly from power and irradiance data or to predict solar irradiance first and then predict the PV power based on the predicted solar irradiance.

II. Literature Review

Most of the work done in the literature focused on predicting the PV power directly from the previous PV power data and a few studies considered irradiance estimation methods [5]. The focus in this paper is on solar irradiance forecasting which can be used afterwards to predict the PV power. There is much work in the literature about solar irradiation forecasting using diverse families of machine learning algorithms therefore we will focus on comparing some of the widely used machine learning models in solar irradiance forecasting. [6] suggests a model utilizing LSTM algorithm designed to handle constrained input data. The inputs consist of the predicted weather for the following day and the daily solar irradiance taken in 1-hour resolution, employed to forecast the irradiance for the subsequent day. The model demonstrated commendable accuracy, yielding an annual Root Mean Square Error (RMSE) of 12%. However, a drawback of this model lies in its reliance on the assumption of 100% accuracy in the forecasted weather. In [4], LSTM model was built to forecast the solar irradiance to overcome the problem of gradient vanishing (exploding) that occur in the recurrent neural network (RNN) due to intense weather change. Also, the clearness index was considered as an input feature to improve the prediction accuracy. The data resolution is 1 hour, the data length is 5 years and the algorithm uses one day in advance. A hybrid model introduced in [7] merges a gated recurrent unit (GRU) neural network with an attention mechanism. Utilizing one year of solar irradiance data at a one-minute resolution, the model generates forecasts for irradiance at various time horizons ranging from 5 to 30 minutes ahead. The suggested model outperforms the accuracy of both LSTM and GRU models across all examined time horizons. Another hybrid model was proposed in [8]

*All authors are with the Department of Electrical and Computer Engineering, Khalifa University, Abu Dhabi, United Arab Emirates. Email: {100062546, 100062594, 100062625, 100062618, davor.svetinovic}@ku.ac.ae

which combines LSTM and residual modeling to achieve a day-ahead probabilistic forecasting. The used data contains hourly solar irradiance and meteorological data over 10 years. Furthermore, the author of [9] hybridized a forecasting model integrated three methodologies: wavelet decomposition (WD), the Convolutional Neural Network (CNN), and Long Short-Term Memory (LSTM). Through a comparison with six alternative models for irradiance, it is evident that the DWT-CNN-LSTM model exhibits significant superiority in solar irradiance forecasting, particularly in challenging weather conditions. In [3] the author developed a structured output LSTM prediction model. The predicted output is the hourly day-ahead solar irradiance using the hourly weather forecasts of the same day. The results showed performance improvements over backpropagation algorithm (BPNN) even by short data length (2 years) or long data length (10 years). Some work was done on a micro grid level as in [10]. Also, in [11] a hybrid LSTM-RNN is proposed to forecast the solar irradiance especially in microgrids, where historical data of solar irradiance is not easily available. So, the model uses other features rather than solar irradiance as dry-bulb temperature, dew-point temperature, and relative humidity. The model was compared with the feedforward neural network (FFNN) because this method showed accurate results in solar irradiance forecasting. This comparison showed that the LSTM-CNN even performs better than FFNN. In [12] another hybrid model is used which is LSTM-CNN which predicts half-hourly, daily and monthly solar radiations. The results showed that the proposed model outperforms the standalone models.

III. Data and Models

A. Dataset Collection and Description

The National Renewable Energy Laboratory (NREL) has compiled a solar and meteorological dataset spanning the last 5 years, with readings sampled at a 5-minute temporal resolution. Obtained from solar monitoring stations, the dataset includes measurements of global horizontal irradiance (GHI), direct normal irradiance (DNI), and diffuse horizontal irradiance (DHI), along with meteorological parameters like temperature, dew point, wind speed, and direction. Time stamps, denoting year, month, day, hour, and minute, identify each observation. NREL employs various instruments, such as pyranometers and pyrheliometers, to collect these measurements. The dataset undergoes preprocessing steps, including handling missing values and normalization. The overall data collection and storage adhere to quality control standards, enabling researchers to access this resource for solar energy analysis.

The dataset used for this project spans the last 5 years, with a temporal resolution of 5 minutes. The data, obtained from the National Renewable Energy Laboratory (NREL), includes information about solar irradiance (GHI, DNI, DHI), weather conditions (Temperature, Dew Point, Relative Humidity, Wind Direction, Wind Speed, Pressure, Precipitable Water), and time stamps (Year, Month, Day, Hour, Minute). Preprocessing involved handling missing values, removing outliers, and normalizing numerical features. GHI (Global Horizontal Irradiance), DNI (Direct Normal Irradiance), and DHI (Diffuse Horizontal Irradiance) along with other features were utilized throughout the analysis to streamline the representation of solar irradiance components.

According to the correlation matrix, the features shown in Table I were shown. Surface Albedo and Pressure didn't have a direct correlation with the choosen output parameter (GHI). So, it was more computationally easier to remove them from the dataset. Cloudy conditions cause a lot of variations in direct normal irradiation that diminish electricity generation. An example can be seen in Figure I. The left image shows a clear sky day whilst the other one shows a cloudy day.

TABLE I
SELECTED FEATURES

Feature	Unit
Temperature	Degrees celsius
Clearsky DHI	(W/m^2)
Clearsky DNI	(W/m^2)
Clearsky GHI	(W/m^2)
Cloud Type	Type of Cloud (0-12)
Dew Point	Temperature (C)
DHI	(W/m^2)
DNI	(W/m^2)
GHI	(W/m^2)
Relative Humidity	-
Solar Zenith Angle	Degrees
Precipitable Water	cm
Wind Direction	Direction
Wind Speed	m/s

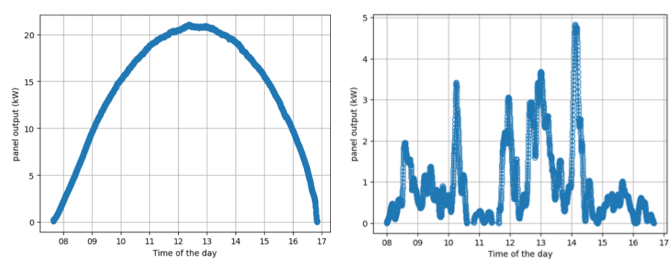

Fig. 1. Sample Irradiance for a sunny day and cloudy day respectively.

B. Data Preprocessing

It is necessary to apply a filtering process to eliminate inaccurate, missing data and ensure the reliability of the trained model. The data is divided into training, testing, and validation using 60% for the training and 20% each for validation and testing.

TABLE II
DATASET

Total Dataset	Training	Testing	Validation
525600	315360	105120	1025120

C. Timeseries Data

The first step implemented for handling the collected data is conversion of pertinent timestamp columns, namely 'Year,' 'Month,' 'Day,' 'Hour,' and 'Minute,' into a unified format "YYYY-MM-DD HH:MM:SS". This transformation generates a new column labeled 'Datetime' which is then used as an index in pandas dataframe, facilitating chronological sorting and efficient time-based queries. This methodology offers a systematic approach to preprocess time-series data, ensuring that it is appropriately formatted for subsequent analysis. After that, relevant features are selected for the training purpose and then the tabular data is converted into a sequential structure that is suitable for timeseries forecasting. Each input sequence is paired with the corresponding target. A sequence of 14 days is used as a historical window to capture temporal dependencies and predict next day's GHI.

D. Models Used for Prediction

1) Multilayer Perceptron Neural Networks: The first part of the study introduces an MLP neural network (NN) model for day-ahead GHI prediction, characterized by a structured architecture. The NN comprises one input layer and two hidden layers, with 300 and 200 neurons, respectively. This design emphasizes the network's capacity to capture intricate patterns within historical irradiance and meteorological data. The proposed 1-300-200 NN architecture strikes a balance between model expressiveness and computational efficiency, offering a robust framework for accurate GHI forecasting.

2) RNN: A Recurrent Neural Network is a type of artificial neural network designed for sequence modeling and processing. Unlike traditional feedforward neural networks, which process input data in a single pass without maintaining any internal state, RNNs have connections that form directed cycles, allowing them to maintain a hidden state that captures information about previous inputs in the sequence. The RNN model structure that is used here consists of four layers: an input layer, two hidden layers, and an output layer. The number of neurons are as follow: 64 neurons in the input layer, 64 neurons in the first hidden layer, and 16 neurons in the second hidden layer.

3) LSTM: LSTM networks have emerged as a powerful tool for solar photovoltaic (PV) forecasting, leveraging their capacity to capture sequential dependencies in time-series data. The benefits of using LSTMs in this context include their aptitude for modeling non-linear relationships, accommodating variable input lengths influenced by changing daylight cycles, and adapting to dynamic environmental conditions affecting solar power generation. However, challenges such as substantial data requirements, potential overfitting complexities, interpretability issues, computational resource demands, and the need for meticulous hyperparameter tuning should be acknowledged. Despite these shortcomings, the adaptability and temporal awareness of LSTMs make them a valuable asset for accurate and dynamic solar PV predictions.

4) LSTM-RNN: Employing a hybrid LSTM (Long Short-Term Memory) and RNN (Recurrent Neural Network) model for time series forecasting offers a potent blend of capabilities, leveraging the strengths of both architectures. LSTMs are adept at capturing long-term dependencies in sequential data, making them well-suited for understanding intricate patterns in time series. The inclusion of traditional RNN components complements the LSTM's capabilities by providing a continuous recurrence mechanism, enabling the model to capture short-term dependencies efficiently. This hybrid approach strikes a balance, harnessing the memory-retention capabilities of LSTMs while maintaining the simplicity and interpretability of RNNs. The result is a robust forecasting model that excels in capturing both short and long-term temporal patterns, making it particularly effective in scenarios where the time series exhibits a mix of complex dynamics.

5) Bidirectional-LSTM: Bidirectional Long Short-Term Memory (BiLSTM) networks present a powerful solution for sequential data processing, including applications such as time-series forecasting, owing to their unique ability to capture information from both past and future time steps. The major advantage of BiLSTMs lies in their enhanced contextual understanding, allowing them to discern intricate patterns and dependencies within sequential data. This bidirectional processing capability is particularly beneficial in scenarios where both preceding and succeeding context contribute to the predictive task, as seen in solar photovoltaic forecasting. However, it's important to note the associated challenges. The increased computational complexity of processing sequences in both directions can result in longer training times and higher resource requirements. Additionally, potential overfitting may occur, especially in situations with limited data. Effective hyperparameter tuning becomes crucial to balance model complexity and prevent overfitting. Despite these challenges, the bidirectional nature of BiLSTMs offers a valuable tool for tasks demanding a comprehensive understanding of sequential data dependencies.

E. Training

A stochastic gradient descent optimizer, Adam, is used to adjust the learning rate. Mean Aboslute Error (MAE) of the GHI prediction is used as the loss function to update the new weights after each iteration.

$$MAE = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i| \tag{1}$$

where n is the number of samples, y_i is the actual output and $\hat{y}_i$ is the true output for sample i. The loss is measured in KW^2.

IV. Results

The Root Mean Squared Error (RMSE) and Mean Squared Error (MSE) metrics are used to evaluate the performances of the different models.

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

$$\text{RMSE} = \sqrt{\sum_{i=1}^{n} \frac{(\hat{y}_i - y_i)^2}{n}} \qquad (3)$$

where $\hat{y}_i$ is the predicted value, y_i is the ground truth, and n is the number of samples.

A. Results for MLP

Figure 3 shows the predicted GHI versus the actual GHI using MLP model for three days from the testing data. It can be seen that there is a mismatch between the actual and predicted GHI since MLP is not designed for sequential data.

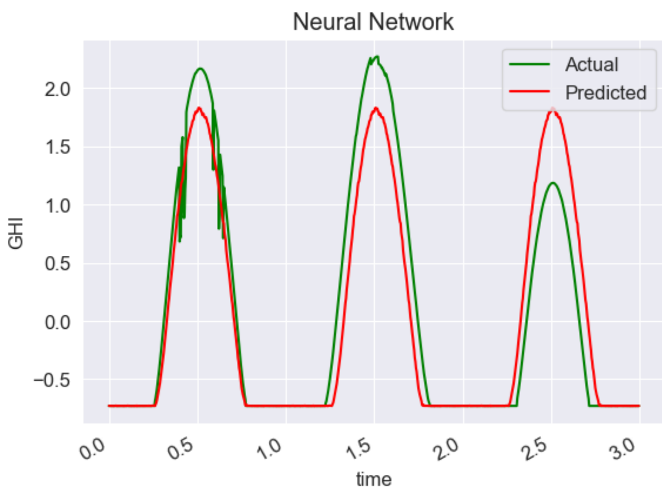

Fig. 2. Results for MLP

B. Results for RNN

Figure 4 shows the predicted GHI versus the actual GHI using RNN model. Besides its short memory RNN also suffers from vanishing and exploding Gradients affecting its ability to generalize and caputre dependencies. These are some reasons why the prediction couldn't match the actual GHI very well.

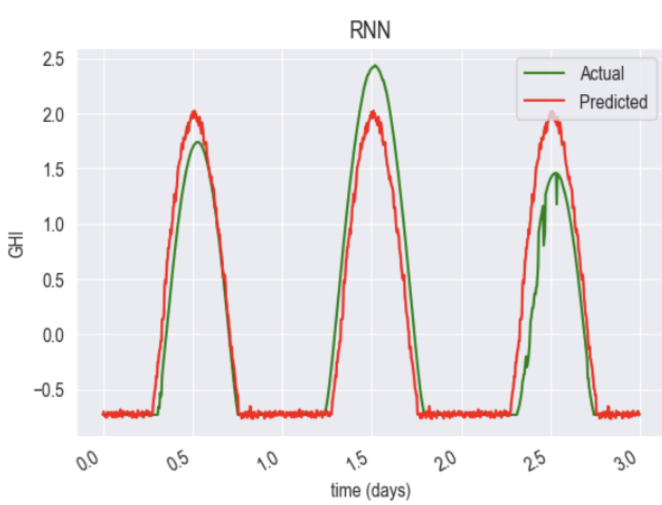

Fig. 3. Results for RNN

C. Results for LSTM

LSTM network test prediction for three testing days can be seen in Figure 5. Since this neural network achieved the lowest error, the predicted GHI is overlapping with the actual GHI throughout the entire period. However, rapid irradiance fluctuation occurs within one minute timeframe and can not be captured by the model since the data is collected every 5 mins.

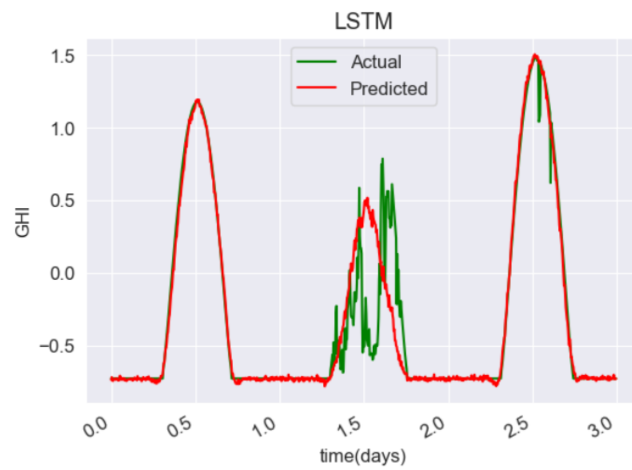

Fig. 4. Results for RNN

D. Results for LSTM-RNN

Figure 6 shows the prediction of three days using the hybrid model LSTM-RNN. LSTM-RNN is good in short and medium term time series forecasting [13]. In this case, the other models are performing better. This model when trained with different number of layers, learning rate, and nodes couldn't get a high accuracy.

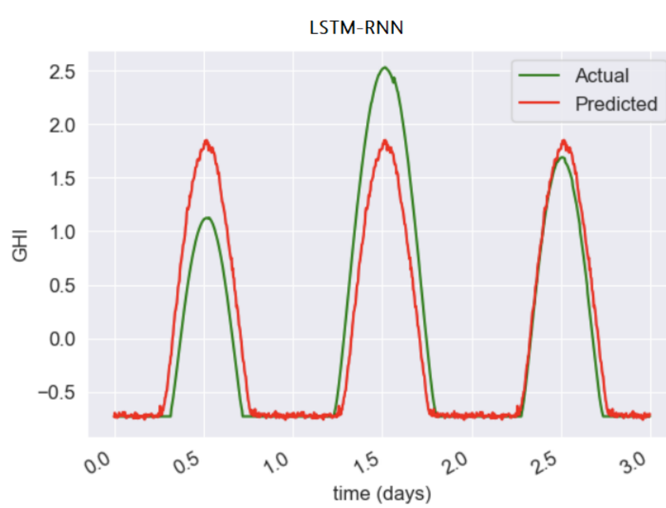

Fig. 5. Results for LSTM-RNN

E. Results for Bidirectional LSTM

The prediction of three days using Bidirectional LSTM can be seen in Figure 7. The bidirectional nature of BiLSTM introduces additional complexity which is the main reason for the large fluctuations in the prediction. The bidirectional architecture, adding parameters and complexity, increases the risk of memorizing training examples, hindering effective generalization to unseen data while also being more computationally expensive than LSTM.

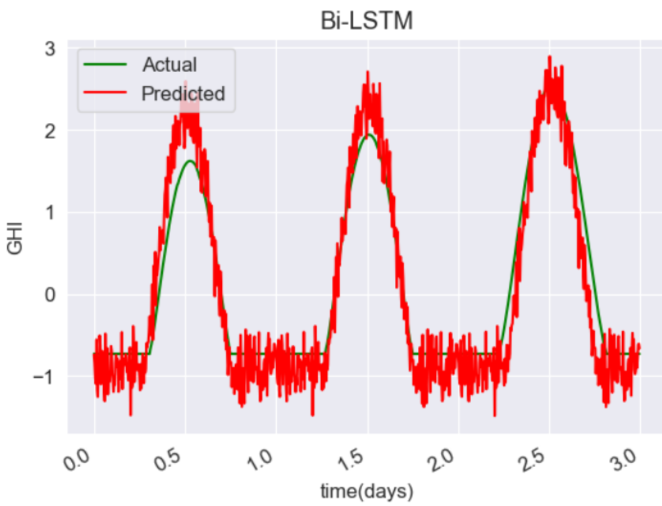

Fig. 6. Results for Bi-LSTM

F. Comparison

In the training and testing comparison tables, the LSTM model emerges as the most effective among the considered architectures. It achieves the lowest Mean Squared Error (MSE), Mean Absolute Error (MAE), and Root Mean Squared Error (RMSE) values in both datasets. This superior performance indicates that the Long Short-Term Memory (LSTM) architecture is well-suited for the specific task at hand when compared to Multi-Layer Perceptron Neural Network (MLP NN), Recurrent Neural Network (RNN), LSTM-RNN, and Bidirectional LSTM (Bi-LSTM) models. The consistently lower error metrics across multiple evaluation criteria position the LSTM model as the most reliable and accurate choice for the given predictive task.

TABLE III
TRAINING COMPARISON

Model	MSE	MAE	RMSE
MLP NN	0.2095	0.2621	0.4563
RNN	0.2709	0.4559	0.2130
LSTM	**0.0813**	**0.1142**	**0.2836**
LSTM-RNN	0.2188	0.2721	0.4677
Bi-LSTM	0.1149	0.1578	0.3328

V. FUTURE WORK

The following models can be implemented are known to work for time series data in competition with the LSTM.

TABLE IV
TESTING COMPARISON

Model	MSE	MAE	RMSE
MLP NN	0.2120	0.2658	0.4594
RNN	0.2666	0.4687	0.2228
LSTM	**0.0768**	**0.1247**	**0.2780**
LSTM-RNN	0.2029	0.2629	0.4516
Bi-LSTM	0.2562	0.3530	0.5116

A. Transformer

The Transformer architecture, introduced in the groundbreaking paper "Attention is All You Need" by Vaswani et al. (2017), has significantly influenced the field of time series analysis. Originally devised for natural language processing tasks, the Transformer's attention mechanism has proven highly adaptable to sequential data, including time series. Traditionally, recurrent neural networks (RNNs) were widely used for modeling temporal dependencies, but the Transformer introduces a parallelized and attention-based approach that addresses the limitations of sequential processing. In the context of time series data, the Transformer excels at capturing long-range dependencies, enabling more effective modeling of intricate temporal patterns. The self-attention mechanism allows the model to selectively focus on relevant timestamps, facilitating robust feature extraction and forecasting. This transformative architecture has not only redefined the landscape of sequence modeling but has also emerged as a powerful tool for advancing the state-of-the-art in time series analysis, offering improved efficiency and performance across a variety of applications [14].

B. RetNet

RetNet is proposed as a successor to Transformer for large language models, but it is not specifically compared to Transformer for time series in the provided sources. However, the paper highlights the advantages of RetNet, such as training parallelism, low-cost inference, and good performance in language modeling. RetNet also achieves favorable scaling results, parallel training, low-cost deployment, and efficient inference. Additionally, RetNet is shown to be consistently competitive with Transformer in terms of scaling curves and in-context learning in language modeling experiments. The mentioned advantages and competitive performance of RetNet suggest that it could be a promising alternative for time series tasks as well [15].

VI. CONCLUSION

This paper predicts the solar irradiance (GHI) using modern deep learning models that are popular in time series forecasting. These models include MLP, RNN, LSTM, LSTM-RNN, and Bi-directional LSTM. The models are trained and then the results are compared. LSTM was computationally less expensive but not as expensive as some other methods like the Bi-LSTM. It performed the best in both the testing and the training scenario. The accuracy of the models could be further improved by creating time-series data for a lower resolution than 5 minutes or choosing a larger window size which is currently 14 days. Modern and recent methods like the transformers, retnet, or some hybrid combination can be used to get a better test accuracy.

REFERENCES

[1] International Energy Agency (IEA), "Solar PV." https://www.iea.org/reports/solar-pv, 2022. License: CC BY 4.0.

[2] A. Zaher, S. Thil, J. Nou, A. Traoré, and S. Grieu, "Comparative study of algorithms for cloud motion estimation using sky-imaging data," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 5934–5939, 2017. 20th IFAC World Congress.

[3] X. Qing and Y. Niu, "Hourly day-ahead solar irradiance prediction using weather forecasts by lstm," *Energy*, vol. 148, pp. 461–468, 2018.

[4] Y. Yu, J. Cao, and J. Zhu, "An lstm short-term solar irradiance forecasting under complicated weather conditions," *IEEE Access*, vol. 7, pp. 145651–145666, 2019.

[5] A. Alcañiz, D. Grzebyk, H. Ziar, and O. Isabella, "Trends and gaps in photovoltaic power forecasting with machine learning," *Energy Reports*, vol. 9, pp. 447–471, 2023.

[6] B.-k. Jeon and E.-J. Kim, "Next-day prediction of hourly solar irradiance using local weather forecasts and lstm trained with non-local data," *Energies*, vol. 13, no. 20, 2020.

[7] K. Yan, H. Shen, L. Wang, H. Zhou, M. Xu, and Y. Mo, "Short-term solar irradiance forecasting based on a hybrid deep learning methodology," *Information*, vol. 11, no. 1, p. 32, 2020.

[8] H. He, N. Lu, Y. Jie, B. Chen, and R. Jiao, "Probabilistic solar irradiance forecasting via a deep learning-based hybrid approach," *IEEJ Transactions on Electrical and Electronic Engineering*, vol. 15, no. 11, pp. 1604–1612, 2020.

[9] F. Wang, Y. Yu, Z. Zhang, J. Li, Z. Zhen, and K. Li, "Wavelet decomposition and convolutional lstm networks based improved deep learning model for solar irradiance forecasting," *applied sciences*, vol. 8, no. 8, p. 1286, 2018.

[10] M. Aslam, J.-M. Lee, H.-S. Kim, S.-J. Lee, and S. Hong, "Deep learning models for long-term solar radiation forecasting considering microgrid installation: A comparative study," *Energies*, vol. 13, no. 1, p. 147, 2020.

[11] M. Husein and I.-Y. Chung, "Day-ahead solar irradiance forecasting for microgrids using a long short-term memory recurrent neural network: A deep learning approach," *Energies*, vol. 12, no. 10, p. 1856, 2019.

[12] S. Ghimire, R. C. Deo, N. Raj, and J. Mi, "Deep solar radiation forecasting with convolutional neural network and long short-term memory network algorithms," *Applied Energy*, vol. 253, p. 113541, 2019.

[13] L. Lin, M. Li, L. Ma, A. Baziar, and Z. M. Ali, "Hybrid rnn-lstm deep learning model applied to a fuzzy based wind turbine data uncertainty quantization method," *Ad Hoc Networks*, vol. 123, p. 102658, 2021.

[14] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," 2023.

[15] Y. Sun, L. Dong, S. Huang, S. Ma, Y. Xia, J. Xue, J. Wang, and F. Wei, "Retentive network: A successor to transformer for large language models," 2023.