International Journal of Energy
ISSN: 2957-9473 | Vol. 7, No. 2, 2025

Research on Photovoltaic Power Generation Energy Forecasting Model Driven by Deep Learning

Kai Yuan [1], Le'an Yi [2, *], Shaozhou Rao [1], Lei Tang [1], Ying Yu [1], Jianguang Jin [1]

[1] China Energy Jiangxi New Energy Industry Co. Ltd, Nanchang, Jiangxi, China
[2] Central Southern China Electric Power Design Institute Co. Ltd, Of China Power Engineering Consulting Group, Wuhan, Hubei, China
*** Corresponding author:** Le'an Yi (Email: yla5995@csepdi. Com)

Abstract: To improve the accuracy of photovoltaic power generation forecasting, this paper constructs two deep learning forecasting models: LSTM and CNN-LSTM. Historical power generation and meteorological data from photovoltaic power plants are collected and preprocessed for model training and optimization. Comparative experiments are conducted with models such as ARIMA, CNN, and GRU. Results show that the CNN-LSTM hybrid model achieves the best performance, with a MAE of 2.35 kW·h, an RMSE of 3.26 kW·h, a MAPE as low as 3.65%, and an R^2 of 0.958, significantly outperforming both ARIMA (MAE 6.85 kW·h, MAPE 10.56%) and a single LSTM (MAE 3.56 kW·h, MAPE 5.92%). Simulation curves demonstrate that the model accurately captures power generation fluctuations, particularly during peak and valley periods. Research has shown that a hybrid model that integrates the spatial characteristics of meteorological data with the temporal characteristics of power generation can effectively improve forecasting accuracy and provide a reliable basis for power system scheduling.

Keywords: Deep Learning; Photovoltaic Power Generation Forecasting; LSTM; CNN-LSTM; Model Comparison.

1. Introduction

Against the backdrop of global energy crisis and environmental challenges, developing and utilizing renewable energy has become a key strategic option for countries to realize sustainable development. As a clean, pollution-free, abundant renewable energy source, PV generation has seen explosive growth over the last few decades. Based on data, global average PV installed capacity increased by more than 20% every year and its share of the energy mix continues to grow [1]. However, the output of PV generation is greatly influenced by various natural factors such as solar radiation, temperature, humidity, wind speed, etc. These characteristics pose major challenges for energy systems' safe and stable operation, efficient energy planning and efficient energy market transactions [2]. For example, an abrupt drop in PV power generation could lead to fluctuations in the frequency of the grid or even power outages without timely replenishment. In addition, where PV generation is much higher than the demand for power generation, there is energy wastage [3]. As a result, accurate prediction of PV power generation is very important to improve the flexibility of power system, optimize resource allocation, and promote efficient PV consumption.

This study aims to develop a high-precision photovoltaic power generation prediction model based on deep learning [4]. First, photovoltaic power generation data, including historical power generation data and meteorological data, is collected and preprocessed. Next, data characteristics are analyzed to explore the key factors influencing photovoltaic power generation. Next, an LSTM model and a CNN-LSTM hybrid model are constructed, trained, and optimized [5]. Finally, through system simulation experiments, the constructed model is compared with traditional models and a single deep learning model to verify its superiority and provide more effective support for photovoltaic power generation prediction.

2. Deep Learning Prediction Model Construction

2.1. Model Design Principle

Variations in photovoltaic power generation not only have significant temporal dependencies with historical data but also exhibit complex nonlinear correlations with meteorological factors [6]. Historical power generation data exhibits distinct diurnal, weekly, and seasonal cycles, making it a typical form of time series data. Meteorological data, however, includes factors such as solar radiation intensity and temperature, which exhibit distinct characteristics in both spatial distribution and temporal variation. For example, solar radiation intensity exhibits a unimodal trend from sunrise to sunset, and temperature distribution varies across regions. This study designed two deep learning prediction models: an LSTM model and a CNN-LSTM hybrid model. The LSTM model primarily targets the temporal characteristics of historical power generation data, capturing long-term dependencies within the data for prediction. The CNN-LSTM hybrid model combines the spatial feature extraction capabilities of CNN with the temporal feature processing capabilities of LSTM, taking into account both meteorological and historical power generation data to fully utilize multi-source information and improve prediction accuracy [7]. The model inputs relevant data from the previous 24 hours (including historical power generation and selected meteorological factors), and outputs a forecast for photovoltaic power generation for the next hour.

2.2. LSTM Prediction Model

2.2.1. Model Structure

The LSTM model consists of input layer, LSTM layer and output layer. There is 1 neuron in the input layer corresponding to previous 24 hours power generation data input in sequence with 1 time step. The first LSTM hidden layer consists of 64 neurons, which are used to extract

temporal features in power generation. To avoid overfitting, it uses the tanh activation function with a dropout rate of 0.2. A second hidden layer of LSTM comprises 32 neurons that are further processed and compressed by the first LSTM layer. It uses a tanh activation function with a dropout rate of 0.2. Both LSTM layers return the last output of a sequence as input [8]. The output layer consists of one neuron and linear activation function, which outputs predicted PV power generation values in the following hour. A Keras framework is used to build a model structure with TensorFlow used to transfer data between layers to ensure efficient information flow.

2.2.2. Parameter Settings

All models were trained with the adam optimizer. Initially, the learning rate was set at 0.001, and then adjusted adaptively over time. The learning rate is halved if the validation set loss does not decrease for five consecutive epochs. A loss function of Mean Square Error (MSE) is used to measure the squared error between predicted and true values. Batch size 32 and iteration number 100. In addition, an early stopping strategy has been implemented that stops training 10 consecutive epochs without improving validation set loss to avoid overtraining. In order to improve the initial learning capability of model, the bias of forget gate, input gate and output gate of LSTM layer was initialized at 0.1. A grid search method was used to optimize parameters such as neuron count and dropout rate in LSTM layer. Finally, the parameters were determined so as to ensure both the performance of the model and the training efficiency [9].

2.3. CNN-LSTM Hybrid Prediction Model

2.3.1. Model Structure

The CNN-LSTM hybrid model is composed of CNN feature extraction module, LSTM processing module and output layer. The CNN feature extraction module is composed of two convolutional layers and a pooling layer. A ReLU activation function [10] uses 32 3 × 3 convolutional kernels to convolve the input weather data matrix (24 hours × 5 weather factors, where 5 weather factors are solar radiation intensity, humidity, wind velocity, and air pressure). The pooling layer uses 2 × 2 max pooling in order to reduce the dimensionality of each convolutional layer, while preserving important features. The output dimension is 12 × 2 × 32. In the second convolutional layer, the paper uses 16 2 × 2 convolution kernels to extract deep spatial features. Use ReLU activation to produce 11 × 1 × 16 output dimensions that are then flattened to 176-dimensional vector. The LSTM processing module is composed of LSTM and full connection layer. The LSTM layer receives CNN module's weather characteristic vector as well as previous 24 hours' power generation data (1 × 24 dimensions). It uses the tanh activation function, with a dropout rate of 0.2. A fully connected layer consists of 16 neurons, uses ReLU activation function, and integrates LSTM layer output. The output layer consists of one neuron and one linear activation function.

2.3.2. Parameter Settings

The CNN-LSTM model is based on the Adam optimizer, and the initial learning rate is 0.001. Using the same learning rate adjustment and early termination strategy as LSTM model. Loss function is MSE, batch size 32, iteration number 100. Each convolutional layer is initialized using the He normal distribution, which adapts to the ReLU activation function. The LSTM layer's parameter settings are similar to those in LSTM models, such as forget gate initialization 0.1. These parameters were determined by comparing the effects of different convolution kernel size and number of LSTM neurons on model performance. For example, a 3 × 3 convolution kernel is used to extract the best spatial features of meteorological data. The LSTM layer of 48 neurons achieves the balance between capturing temporal features and preserving the complexity of model.

3. System Simulation Experiment and Result Analysis

3.1. Experimental Environment

This experiment uses hardware configuration: Intel Core i9-13900K, 24 cores, 32 threads, 3.0 GHz base frequency, 5.8 GHz maximum turbo frequency, providing powerful computational support for model training [11]. NVIDIA GeForce RTX 4080 has 16GB video memory support and CUDA 12.1 acceleration technology to improve the training speed of deep learning. It has a storage capacity of 128 GB for DDR5 4800 MHz. It can satisfy large scale data processing and model training. It is a 2 TB NVMe solid state drive that ensures efficient data read and write. Regarding the software environment, the operating system used was Windows 11 Professional 64-bit; the programming language was Python 3.10, paired with the deep learning frameworks PyTorch 2.0 and TensorFlow 2.12. Data processing and visualization tools included Pandas 2.0.3, NumPy 1.23.4, and Matplotlib 3.7.2. In this system simulation platform, MATLAB R2023b is used to build a simulation model for PV power generation system. The GPU's parallel computing ability is fully utilized during the experiment, enabling CUDA acceleration to effectively reduce model training time.

3.2. Evaluation Metrics

This experiment used mean absolute error (MAE), root mean square error (RMSE), mean absolute percentage error (MAPE), and coefficient of determination (R^2) as metrics to evaluate model prediction performance.

3.3. Experimental Design

3.3.1. Comparison Model Selection

To comprehensively evaluate the performance of the constructed model, the paper selected traditional time series models and other deep learning models for comparison. These models include:

- The traditional ARIMA model, a classic model in the field of time series forecasting, achieves prediction by modeling the autoregressive and moving average properties of data.
- A single LSTM model, which uses only historical power generation data for forecasting, is used to compare the forecasting performance of a single time series feature with that of a model that integrates multiple features.
- The GRU model, a simplified version of the LSTM model, is also used to process time series data and allows for comparison of the performance of different gated recurrent neural networks.
- A CNN model, which is used solely to extract spatial features from meteorological data and perform forecasting, highlights the importance of processing time series features in hybrid models.

3.3.2. Experimental Procedure

First, the preprocessed training set was fed into each model for training. The change in loss for each model was recorded during training. The batch size for training was set to 64, the number of iterations was set to 150, and an early stopping

strategy was employed. Model training was stopped when the loss on the validation set did not decrease for 15 consecutive iterations to prevent overfitting.

Next, the trained models were optimized using the validation set. Key parameters for each model, such as the number of hidden layer neurons and the learning rate in the LSTM model, were adjusted using a grid search method to obtain the optimal parameter configuration.

Finally, the test set was fed into each optimized model for prediction. Evaluation metrics such as MAE, RMSE, MAPE, and R^2 were calculated for each model. To reduce random error in the experiment, each model was independently run 10 times, and the average of the evaluation metrics was taken as the final result. The prediction results were visualized, with plots comparing the predicted values to the true values and error distribution histograms, to more intuitively analyze the prediction performance of each model.

3.4. Experimental Results Analysis

3.4.1. Comparison of Prediction Results of Different Models

Table 1 shows the prediction performance evaluation metrics of each model on the test set. As can be seen from the data in this table, the traditional ARIMA model has the lowest prediction accuracy, with a MAE of 6.85 kW·h, RMSE of 9.23 kW·h, MAPE of 10.56%, and R^2 of 0.786. This is because the ARIMA model can only capture linear relationships in the data and has difficulty handling the complex nonlinear relationships and long-term temporal dependencies between PV power generation and meteorological factors. The CNN model outperformed the ARIMA model, with a MAE of 5.21 kW·h, RMSE of 7.32 kW·h, MAPE of 8.15%, and R^2 of 0.843. This demonstrates that the CNN model can effectively extract the spatial characteristics of meteorological data and is somewhat helpful for predicting PV power generation. However, due to its inability to handle the temporal characteristics of the data, prediction accuracy still has significant room for improvement. The GRU model achieved better prediction performance than the CNN model, with a MAE of 3.12 kW·h, an RMSE of 5.87 kW·h, a MAPE of 6.78%, and an R^2 of 0.892. As a gated recurrent neural network, the GRU model can capture temporal dependencies in data, making it superior when processing time series data such as photovoltaic power generation. However, compared to the LSTM model, its ability to capture long-term temporal dependencies is slightly weaker. The prediction accuracy of the single LSTM model was further improved, with a MAE of 3.56 kW·h, an RMSE of 3.98 kW·h, a MAPE of 5.92%, and an R^2 of 0.915. The LSTM model, through its unique gating mechanism, is better able to handle long-term temporal dependencies, resulting in superior performance compared to the GRU model when using only historical power generation data for prediction. The CNN-LSTM hybrid model constructed in this paper achieved the highest prediction accuracy, with a MAE of only 2.35 kW·h, an RMSE of 3.26 kW·h, a MAPE of 3.65%, and an R^2 of 0.958. This model combines the CNN model's ability to extract spatial features with the LSTM model's advantage in processing temporal features. It can simultaneously utilize the spatial characteristics of meteorological data and the temporal characteristics of power generation data, resulting in significantly better prediction performance than other models.

Table 1. Comparison of Prediction Performance Evaluation Metrics of Different Models

model	MAE(kW·h)	RMSE(kW·h)	MAPE(%)	R^2	Training time (min)	Single prediction time (ms)
ARIMA	6.85	9.23	10.56	0.786	2.1	8.5
CNN	5.21	7.32	8.15	0.843	15.6	12.3
GRU	4.12	5.87	6.78	0.892	28.3	15.7
LSTM	3.56	4.98	5.92	0.915	32.5	18.2
CNN-LSTM	2.35	3.26	3.65	0.958	45.8	22.6

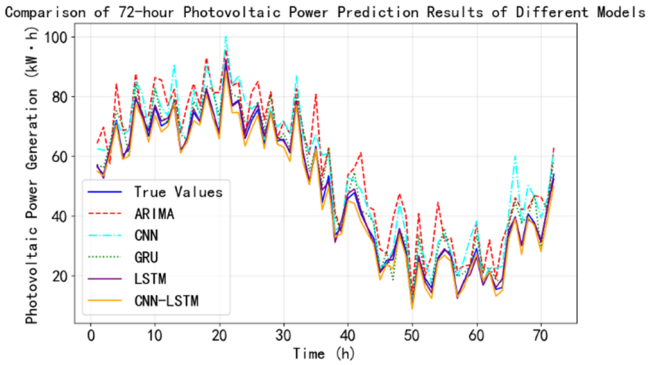


Fig 1. Comparison of predicted and true values for each model

In order to demonstrate each model's prediction performance better, the paper selected 72 consecutive hours of PV power generation data from a test set and plotted a comparison curve between each model's predicted values and actual values, as illustrated in Figure 1. The horizontal axis represents time (hours) ranging from 1 to 72, while the vertical axis indicates PV power generation (kW · h). The figure clearly shows a large deviation between the ARIMA model and the real value curve, especially in periods of large power generation. The prediction curve of CNN model is relatively smooth, however, there is large error when it comes to peak or trough. The prediction curves of GRU model and LSTM model closely match true value curves. LSTM model performs a little better than GRU model. The prediction curve of CNN-LSTM hybrid model is nearly identical to that of real value, which accurately captures changes in PV power generation and demonstrates extremely high prediction precision during periods of stable power generation as well as periods of large fluctuations.

3.4.2. Model Stability Analysis

Figure 2 shows the MAE boxplots for 10 runs of each model, with the horizontal axis representing the different models and the vertical axis representing MAE (kW·h). As can be seen from the figure, the ARIMA model's boxplot has wider boxes, a larger interquartile range, and a high number of outliers, further indicating its poor stability. The CNN model's boxplot is the next narrowest, while the GRU and LSTM models have narrower boxes and fewer outliers. The CNN-LSTM hybrid model has the narrowest boxplot, the smallest interquartile range, and no outliers, demonstrating its excellent stability.

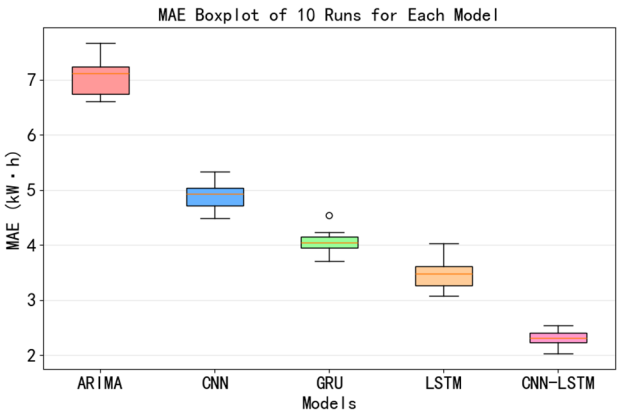


Fig 2. MAE boxplots for 10 runs of each model

3.4.3. Analysis of the Impact of Different Input Variables on Forecast Results

To explore the impact of different input variables on the prediction results of the CNN-LSTM hybrid model, four comparative experiments were designed: inputting only historical power generation data (Experiment 1); inputting historical power generation data + solar radiation intensity (Experiment 2); inputting historical power generation data + solar radiation intensity + temperature (Experiment 3); and inputting historical power generation data + solar radiation intensity + temperature + humidity + wind speed (Experiment 4). The prediction performance evaluation metrics for each experiment are shown in Table 2. Experiment 1 achieved the lowest prediction accuracy, with a MAE of 4.87 kW·h, RMSE of 6.52 kW·h, MAPE of 7.93%, and R^2 of 0.867. This indicates that relying solely on historical power generation data for predictions fails to fully capture the factors affecting PV power generation, resulting in limited prediction accuracy. The prediction accuracy of Experiment 2 was significantly higher than that of Experiment 1, with a MAE of 3.21 kW·h, an RMSE of 4.35 kW·h, a MAPE of 5.68%, and an R^2 of 0.908. This indicates that solar radiation intensity is a key factor affecting photovoltaic power generation, and incorporating it into the input variables significantly improves the model's prediction accuracy. The prediction accuracy of Experiment 3 was further improved, with a MAE of 2.76 kW·h, an RMSE of 3.82 kW·h, a MAPE of 4.52%, and an R^2 of 0.932. This indicates that temperature is also an important factor affecting photovoltaic power generation, and when used as an input variable along with solar radiation intensity, it further improves the model's prediction performance. Experiment 4 achieved the highest prediction accuracy, with a MAE of 2.35 kW·h, an RMSE of 3.26 kW·h, a MAPE of 3.65%, and an R^2 of 0.958. This indicates that meteorological factors such as humidity and wind speed also have a certain impact on photovoltaic power generation. By incorporating all of these factors into the input variables, the model can more comprehensively capture the factors affecting photovoltaic power generation, thereby achieving the highest prediction accuracy.

Table 2. Prediction Performance Evaluation Metrics of the CNN-LSTM Model for Different Input Variable Combinations

experiment number	Input variable combination	MAE(kW·h)	RMSE(kW·h)	MAPE(%)	R^2
1	Historical power generation data	4.87	6.52	7.93	0.867
2	Historical power generation data + solar radiation intensity	3.21	4.35	5.68	0.908
3	Historical power generation data + solar radiation intensity + temperature	2.76	3.82	4.52	0.932
4	Historical power generation data + solar radiation intensity + temperature + humidity + wind speed	2.35	3.26	3.65	0.958

4. Conclusion

This study designed two deep learning models, LSTM and CNN-LSTM, for the photovoltaic power generation forecasting problem. System simulation experiments validated the model performance. Results show that the CNN-LSTM hybrid model performs optimally across all metrics, achieving a MAE of 2.35 kW·h, an RMSE of 3.26 kW·h, a MAPE of 3.65%, and an R^2 of 0.958. This model achieves a MAE reduction of 65.7% and a MAPE reduction of 65.5% compared to the ARIMA model, and a MAE reduction of 34% and a MAPE reduction of 38.2% compared to the single LSTM model. Boxplots demonstrate the model's excellent stability, with a standard deviation of the MAE across 10 runs of only 0.18 kW·h. Input variable experiments demonstrate that integrating meteorological data such as solar radiation and temperature can reduce the MAPE from 7.93% to 3.65%. In summary, the CNN-LSTM model effectively integrates features from multiple data sources, achieving both accurate and stable forecasts. In the future, the paper can further expand the data scope and optimize the model structure to reduce computing costs and enhance generalization capabilities across different scenarios, providing technical support for the efficient use of renewable energy.

References

[1] Chen, M., Zhu, Y., Sun, Y., Xie, Z. & Wu, P. Comprehensive energy system forecasting and control considering high-penetration photovoltaic consumption and deep reinforcement learning. Transactions of the Chinese Society of Electrical Engineering, Vol. 39(2024) No. 19, p. 6056071.

[2] Luo, X., Hu, Z., Ma, Z. & Lü, Z. Photovoltaic cluster power generation forecasting based on hybrid deep learning. Journal of Southwest University (Natural Science Edition), Vol. 47(2025) No. 3, p. 200-210.

[3] Tan, Y., Huang, Y., Liu, J. & Yang, R. Long-term and short-term time series forecasting of multivariate loads in power systems based on MMoE-CNN-Informer model. Transactions of the Chinese Society of Electrical Engineering, Vol. 20(2025) No. 2, p. 253-263.

[4] Han, J., Yan, L., Zhang, F., Wang, Y. & Mao, H. Research on the technical path of high-quality development of photovoltaic industry driven by digital-physical integration. Laser Journal, Vol. 45(2024) No. 11, p. 193-196.

[5] Wei, W., Lü, Y., Qi, X., Liu, J. & Fang, F. State prediction method of wind turbine in power station based on CNN-LSTM-

AM dynamic integration model. Chinese Journal of Scientific Instrument, Vol. 44(2023) No. 4, p. 19-27.

[6] Zhang, S., Wang, J., Wang, C., Zhao, C. & Zhang, Y. Energy prediction of edge photovoltaic power generation based on hierarchical fuzzy neural network. Modern Electric Power, Vol. 41(2024) No. 3, p. 490-499.

[7] Fan, L., Zhang, Q., Li, G. & Wu, J. Short-term photovoltaic power generation prediction based on long short-term memory network digital twin. Modern Electric Power, Vol. 40(2023) No. 6, p. 899-905.

[8] Zhu, J., Dong, H., Li, S., Chen, Z. & Luo, T. A review of data-driven integrated energy system load forecasting. Proceedings of the CSEE, Vol. 41(2021) No. 23, p. 7905-7923.

[9] Zhao, Y., Wang, H., Kang, L. & Zhang, Z. Short-term power load forecasting based on temporal convolutional network. Transactions of the Chinese Society of Electrotechnical Engineering, Vol. 37(2022) No. 5, p. 1242-1251.

[10] Feng, P., Jiang, G., Xu, J., Ye, J. & Li, S. Short-term photovoltaic power forecasting based on CEEMD-BiLSTM-RFR. Science, Technology and Engineering, Vol. 24(2024) No. 5, p. 1955-1962.

[11] Zheng, X., Yang, X. & Lai, W. Classification and research progress of renewable energy power generation forecasting methods. Science, Technology and Engineering, Vol. 24(2024) No. 34, p. 14503-14521.