Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

A Photovoltaic Power Forecasting Method Integrating Physical Mechanisms and Deep Learning

Hongyu Gao
Northeast Petroleum University
Yuan Lu
Northeast Petroleum University
Lizhong Tang
Northeast Petroleum University
Jinrui Cai
Northeast Petroleum University
Fei Han
156683006250163.com

State Key Laboratory of Continental Shale Oil (SKL-CSO)

Research Article

Keywords: Photovoltaic power forecasting, Feature engineering, Adaptive fusion, Multi-scale feature extraction, Dynamic weight allocation

Posted Date: April 17th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-8576566/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

A Photovoltaic Power Forecasting Method Integrating Physical Mechanisms and Deep Learning

Hongyu Gao[1,2,3,4] Yuan Lu[1,2,3,4] Lizhong Tang[1,2,3,4] Jinrui Cai[1,2,3,4]
Fei Han[1,2,3,*]

[1]Artificial Intelligence Energy Research Institute, Northeast Petroleum University, Daqing 163318, China
[2]Heilongjiang Provincial Key Laboratory of Networking and Intelligent Control, Northeast Petroleum University, Daqing 163318, China
[3]State Key Laboratory of Continental Shale Oil (SKL-CSO), Daqing 163318, China
[4]School of Electrical and Information Engineering, Northeast Petroleum University, Daqing 163318, China
[*]Corresponding author. E-mail: tomcumt@126.com

Abstract

With the rapid expansion of distributed photovoltaic installations, the intermittent and fluctuating nature of their output power poses challenges to grid dispatch and operational stability. Physical models offer good interpretability but suffer from limited accuracy under complex meteorological conditions. Pure data-driven methods lack sufficient generalization under extreme operating conditions, and existing hybrid approaches often employ static weights that struggle to adapt to environmental changes. To address these issues, this paper proposes a photovoltaic power forecasting method that integrates physical mechanisms with multi-scale deep learning. The key innovations of this method include: (1) A non-uniform error compensation strategy based on irradiance interval segmentation to enhance the accuracy of the physical branch across different light intensity ranges; (2) Construction of a 37-dimensional physically enhanced feature system incorporating meteorological, temporal, physically derived, and statistical interaction features, utilizing a parallel Multi-scale CNN (kernel lengths 3/7/15) combined with BiLSTM architecture to extract multi-scale temporal characteristics; (3) A dynamic weighted fusion mechanism based on four-dimensional confidence levels (irradiance, temperature, time, and weather stability) to achieve complementary environmental perception between physical and data-driven models. Validation using 15-minute resolution annual measurement data from a prefecture-level city photovoltaic power station demonstrates significant improvements over pure physical and pure neural network approaches in metrics including MAE, RMSE, and R^2 (example results: MAE = 3.571 kW, RMSE = 6.384 kW, R^2= 0.9781). Ablation analysis further validated the effectiveness of each module. By balancing physical consistency with data-driven expressiveness, this approach provides practical value for enhancing the accuracy of photovoltaic power forecasting and improving engineering applicability.

Keywords: Photovoltaic power forecasting Feature engineering Adaptive fusion Multi-scale feature extraction Dynamic weight allocation

1 Introduction

In the context of the global energy transition toward decarbonization, photovoltaic (PV) power generation has emerged as a critical component of renewable energy systems due to its clean, low-carbon characteristics and sustainability (1; 2). According to the International Energy Agency (IEA), as illustrated in Figs. 1 and 2, global renewable energy capacity additions have maintained continuous growth for 22 consecutive years. In 2024, newly installed renewable capacity reached approximately 700 GW worldwide, with PV accounting for nearly 80% of this total. Annual electricity generation from solar PV and wind increased by a record 670

TWh, compared to increases of 170 TWh for natural gas and 90 TWh for coal. Notably, in the European Union, the combined generation share of solar PV and wind exceeded that of coal and natural gas for the first time. In the United States, solar PV and wind generation reached 16% of the total, surpassing coal-based generation. In China, solar PV and wind contributed approximately 20% of total electricity generation (3).

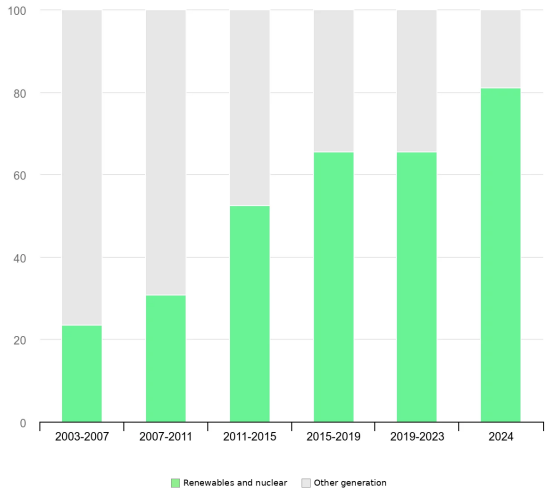


Figure 1: Global Electricity Generation Growth Share, 2003–2024

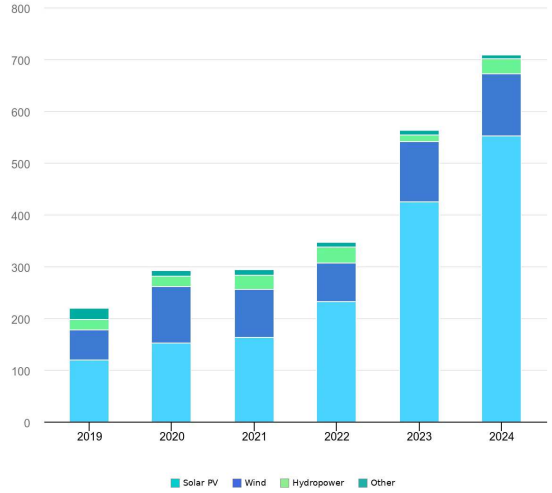


Figure 2: Annual Changes in Global Electricity Generation by Source, 2019–2024

High-precision photovoltaic power prediction serves as a crucial technical foundation for optimizing grid dispatching, enhancing the capacity to integrate renewable energy, and ensuring the safe and reliable operation of the power system. However, the large-scale integration of photovoltaic systems, while promoting clean energy, introduces new challenges for the safe dispatching and stable operation of the power grid. Various meteorological factors, including solar irradiance, environmental temperature, changes in cloud cover, and humidity, influence the output of photovoltaic power generation. These factors contribute to significant randomness, intermittency, and nonlinearity in power generation. This volatility complicates the stable prediction of power system output, thereby increasing the demand for reserve capacity, reducing the dispatching flexibility of the system, and leading to curtailment issues in certain regions (4; 5; 6; 7; 8). For example, in 2023, China experienced a waste of approximately 7.3 billion kilowatt-hours of photovoltaic power, resulting in a curtailment rate of 2.3% and direct economic losses exceeding 3 billion yuan. Consequently, the development of a high-precision and robust photovoltaic power prediction model is of considerable theoretical and engineering significance for improving the efficiency of power grid dispatching, facilitating the operation of the power market, and advancing the planning of energy storage systems.

Existing photovoltaic power generation forecasting methods primarily fall into three categories: physics-based models, data-driven models (statistical models and machine learning models), and hybrid models (combinations of the aforementioned approaches) (9; 10; 11; 12; 13). Physics-based models are typically constructed based on solar geometry, radiation transfer, and thermodynamic processes, offering good interpretability. However, as noted by several researchers (14; 15; 16; 17), their forecasting accuracy heavily relies on high-quality meteorological inputs and exhibits significant limitations under complex weather conditions.

Data-driven models encompass a range of statistical and machine learning approaches. Traditional statistical models, such as the Autoregressive Integrated Moving Average (ARIMA) and exponential smoothing methods, have been extensively investigated and demonstrate effectiveness for stationary or periodic sequences; however, their capacity to capture complex nonlinear relationships remains limited (18). To address this limitation, machine learning models have been introduced, which excel at handling nonlinear relationships by establishing specific mappings between input and output data through model training (19). Representative examples include Support Vector Regression (SVR), Random Forests (RF), and Artificial Neural Networks (ANN), all of which have been successfully applied to PV forecasting tasks (20; 21; 22).

Currently, artificial intelligence has emerged as the cornerstone of PV forecasting methodologies. In the era of big data, deep learning is progressively surpassing traditional neural networks and conventional machine

learning models. In recent years, deep learning methods have achieved significant advancements in PV power forecasting. Convolutional Neural Networks (CNNs) have been employed to extract local spatio-temporal features through multi-scale convolutional kernels. Long Short-Term Memory (LSTM) networks and their bidirectional variants (BiLSTM) effectively capture temporal dependencies; Wang et al. (23) proposed an accurate solar PV power prediction interval method based on frequency-domain decomposition and LSTM model. Meanwhile, Gated Recurrent Units (GRU) offer comparable predictive performance while reducing computational complexity. The Transformer architecture and its variants—including Informer, Autoformer, and Temporal Fusion Transformer—have demonstrated exceptional capabilities in modeling long sequences through multi-head self-attention mechanisms. Piantadosi et al. (24) developed a Transformer-based framework specifically for photovoltaic power forecasting. Furthermore, Yuan et al. (25) proposed a multistep photovoltaic power forecasting method based on multi-timescale fluctuation aggregation attention mechanism and contrastive learning, enabling adaptive learning of dependencies across different temporal scales. Given the inherent complexity of time series data and the limitations of single-model approaches, hybrid architectures combining Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs)—such as CNN-LSTM, LSTM-CNN, and CNN-GRU—have been widely adopted and have demonstrated strong capabilities in photovoltaic power generation forecasting (26). Nevertheless, purely data-driven approaches suffer from three core deficiencies, as identified in recent studies (27; 28): (1) Physical inconsistency—models may generate predictions that violate fundamental physical laws (e.g., non-zero power output during nighttime hours, output exceeding rated capacity); (2) Insufficient generalization—model performance significantly degrades under extreme weather conditions that fall outside the training data distribution; (3) Lack of interpretability—the inherent "black-box" nature of these models renders predictions difficult to trace and limits the potential for human intervention and domain knowledge integration.

To address these limitations, several hybrid prediction frameworks that integrate physical prior knowledge with data-driven models have been developed in recent years (29; 30; 31). These approaches aim to balance model interpretability with data learning capabilities by embedding physical constraints within neural networks or by incorporating outputs from physical models. However, current fusion strategies primarily utilize static weighting or simple concatenation methods, which do not adaptively adjust the credibility of submodels in response to varying environmental conditions. This limitation constrains the models' adaptability and robustness in non-stationary time-series scenarios.

Addressing the limitations of traditional data-driven models lacking physical constraints and the susceptibility of physical models to environmental influences, this study introduces a novel photovoltaic power prediction framework termed PhysFusionNet. This framework employs a dual approach combining data and models to achieve comprehensive deep fusion learning, seamlessly integrating the photovoltaic physical model with a deep neural network at both the feature and decision-making layers. Notably, the predictions generated by the physical model not only contribute to the final fusion output but also serve as crucial input features for the deep network, thereby harmonizing mechanism constraints and data-driven learning within the network. Comprising five essential modules—data preprocessing and quality control, enhanced feature engineering, physical model branch, deep learning branch, and environmental perception fusion layer—this framework represents a significant advancement in the field.

1. Non-uniform error compensation based on irradiance intervals and physical model improvement: Addressing the issue of uneven error distribution under varying light conditions in traditional physical models, this paper proposes a non-uniform error compensation strategy based on irradiance interval segmentation. Irradiance intensity is categorized into three intervals: low light ($\leq$200 W/m^2), medium light (200–500 W/m^2), and high light ($>$500 W/m^2). A robust statistical measure (median ratio) is employed to determine calibration factors, effectively reducing systematic biases in the physical model across multiple operating conditions. Furthermore, a wind speed–temperature coupling correction mechanism was incorporated into the traditional NOCT model, constructing an improved physical model that accounts for wind cooling effects. This significantly enhances the benchmark prediction capability and environmental adaptability of the physical branch.

2. Feature engineering for physical knowledge enhancement: Construct a 37-dimensional physical knowledge enhancement feature space, integrate basic meteorological features, cycle time coding, physical derived features, interactive features, and temporal statistical features to achieve deep integration of multi-source information. This feature system significantly improves the representation ability of the

model while maintaining interpretability. The prediction results of the physical model not only participate in the final fusion but also enter the deep network as important input features to achieve end-to-end fusion of mechanism information and data learning, so that the model naturally follows physical constraints during the training process.

3. Multi-scale Feature Extraction and Physical Constraint Embedding: A parallel multi-scale convolutional–bidirectional LSTM (BiLSTM) architecture is designed with convolutional kernel lengths of 3, 7, and 15 to capture short-term fluctuations, medium-term trends, and long-term periodic features, respectively. BiLSTM further models temporal dependencies to achieve global spatio-temporal feature extraction. Simultaneously, we incorporate physical consistency constraint layers (non-negativity, capacity limits, nighttime zero power, etc.) into the network to ensure predictions align with fundamental PV power physics, enhancing the model's engineering credibility.

4. Environment-Aware Dynamic Fusion Mechanism: We propose a four-dimensional dynamic weighting strategy based on irradiance, temperature, time, and weather stability confidence to adaptively adjust the fusion ratio between physical and data-driven models. This mechanism maintains prediction robustness and generalization capability under variable meteorological conditions, overcoming the limitations of traditional static weighting methods.

The proposed dual-driven framework, PhysFusionNet, effectively balances physical consistency with data learning capabilities, thereby significantly enhancing prediction accuracy and robustness. The model demonstrates strong generalization performance across diverse meteorological conditions, offering a viable new approach for intelligent and engineering-oriented PV power forecasting applications.

The remainder of this paper is organized as follows: Section 2 details the proposed methodological framework; Section 3 introduces the experimental setup and datasets; Section 4 presents experimental results with in-depth analysis; Section 5 discusses the method's strengths, limitations, and future directions; Section 6 concludes the paper.

2 Methods

The proposed model-data dual-driven photovoltaic power forecasting framework comprises five core modules: data preprocessing and quality control, enhanced feature engineering, physical model branch, deep learning branch, and intelligent fusion decision layer. Adopting an end-to-end learning paradigm, the framework achieves a complete workflow from raw data to final predictions.

As shown in Fig. 3, raw data enters through the data acquisition module, which includes the collection and integration of multi-source meteorological parameters such as solar irradiance, temperature, humidity, and wind speed. This data proceeds through the data quality control module for cleaning, missing value imputation, and outlier handling. Subsequently, the signal separation module employs a multi-stage IIR filter to decompose the data into trend, seasonal, and residual components. These components are then fed into the enhanced feature engineering module, followed by processing through the physical model and data-driven model branches. The outputs of these branches ultimately converge in the environmental perception fusion layer for dynamic weight fusion. Figure 3 illustrates the overall system architecture and detailed data flow.

2.1 Data Preprocessing

The quality of raw data directly impacts model performance. The data cleaning process includes:

2.1.1 Missing Value Handling

A segmented strategy is employed to handle missing values, accommodating different lengths of missing data:

Short-term missing data: When missing values occur consecutively for no more than four time steps, linear interpolation is used to fill in the gaps. Based on the linear relationship between known data points, this method reasonably estimates missing values using the formula:

$$x_i = x_{i-1} + \frac{x_{i+n} - x_{i-1}}{n+1} \times k, \quad k \in \{1, 2, \ldots, n\} \tag{1}$$

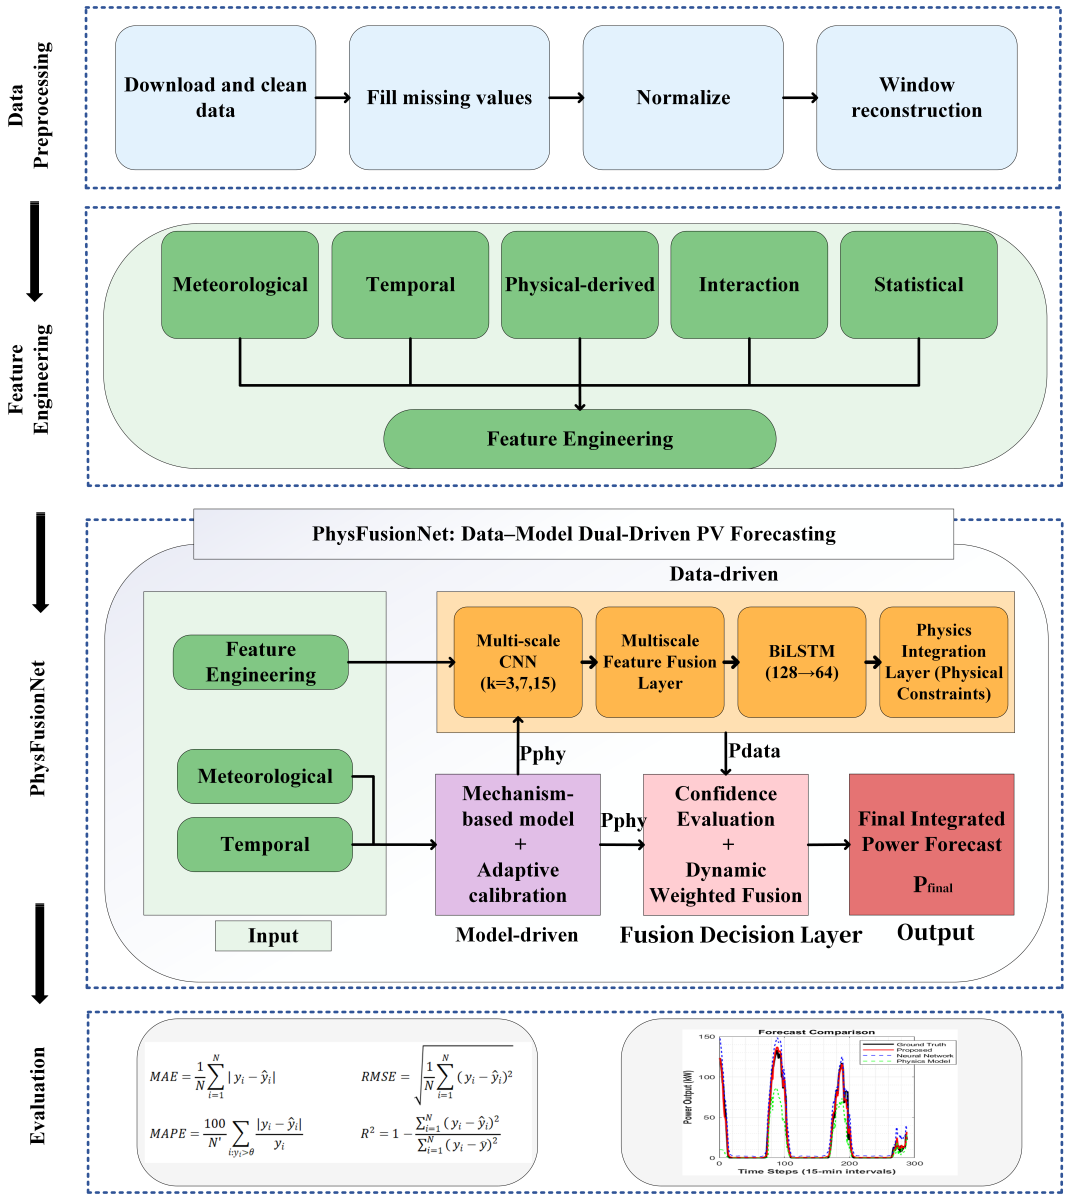


Figure 3: Overall Architecture of the Hybrid PV Power Forecasting Framework

where x_i denotes the estimated value of the ith missing value, x_{i-1} denotes the known value preceding the missing value, x_{i+n} denotes the nth known value following the missing value, and k denotes the step size from x_{i-1} to the current missing value.

Long-term missing values: For missing values spanning more than 4 consecutive time steps, first use nearest-neighbor interpolation to obtain initial fill values. Then apply smoothing to the filled data to reduce abruptness caused by interpolation, ensuring the data better aligns with the overall trend.

$$x(t) = \begin{cases} x_{t_{\text{next}}}, & \text{if } t < t_{\min} \\ x_{t_{\text{prev}}}, & \text{if } t > t_{\max} \\ \text{interpolated by Eq.(1)}, & \text{otherwise} \end{cases} \tag{2}$$

where t_{next} denotes the most recent subsequent non-missing value time point, and t_{prev} denotes the most recent preceding non-missing value time point. Specifically, when time t is less than the minimum time $t_{\min}$, the most recent subsequent non-missing value is used to fill the gap; when t exceeds the maximum time $t_{\max}$, the most recent preceding non-missing value is used to fill the gap.

2.1.2 Outlier Detection and Handling

Identifying Outliers Using Physical Constraints and Statistical Methods. Physical constraints are limitations on data ranges based on actual physical significance. Define the physical constraint function Γ:

$$\Gamma(x, l, u) = \begin{cases} l & \text{if } x < l \\ u & \text{if } x > u \\ x & \text{otherwise} \end{cases} \tag{3}$$

The constraint boundaries for each variable are:

$$\begin{aligned} G^*_{\text{GHI}} &= \Gamma(G_{\text{GHI}}, 0, 1500) \\ G^*_{\text{DNI}} &= \Gamma(G_{\text{DNI}}, 0, 1000) \\ G^*_{\text{DHI}} &= \Gamma(G_{\text{DHI}}, 0, 800) \\ G^*_{\text{tilt}} &= \Gamma(G_{\text{tilt}}, 0, 1500) \\ T^*_{\text{amb}} &= \Gamma(T_{\text{amb}}, -20, 60) \\ v^*_{\text{wind}} &= \Gamma(v_{\text{wind}}, 0, 30) \\ H^*_{\text{rel}} &= \Gamma(H_{\text{rel}}, 0, 100) \end{aligned} \tag{4}$$

The function Γ constrains each variable (such as global horizontal irradiance G_{GHI}, direct normal irradiance G_{DNI}, ambient temperature T_{amb}, etc.) within reasonable physical limits (taking the lower bound l if below the lower limit, taking the upper bound u if above the upper limit, and otherwise retaining the original value).

For power output P, a fractional-based outlier detection method is employed, with the formula defined as:

$$P^* = \begin{cases} 0 & \text{if } P < 0 \\ Q_{99} & \text{if } P > 2 \cdot Q_{99} \\ P & \text{otherwise} \end{cases} \tag{5}$$

where $Q_{99} = \text{quantile}(P, 0.99)$ represents the 99th percentile of the power. Specifically, when the power is less than 0, it is adjusted to 0; when it exceeds twice the 99th percentile, it is adjusted to Q_{99}; otherwise, it retains its original value.

These methods ensure data remains within reasonable bounds through physical principles and statistical properties, providing high-quality inputs for subsequent analysis or modeling.

2.1.3 Data Standardization

Z-score standardization is employed:

$$x_{\text{norm}} = \frac{x - \mu_{\text{train}}}{\sigma_{\text{train}}} \tag{6}$$

Anti-standardization:

$$x = x_{\text{norm}} \cdot \sigma_{\text{train}} + \mu_{\text{train}} \tag{7}$$

where μ_{train} and σ_{train} represent the mean and standard deviation of the training set, respectively.

2.2 Physical-Model Enhancement and Nonuniform Error Compensation Strategy

2.2.1 Photovoltaic Physical Model

Developing a physical model based on the mechanism of photovoltaic power generation:

$$P_{\text{physical}}(t) = \frac{G_{\text{tilt}}(t)}{G_{\text{ref}}} \cdot P_{\text{nom}} \cdot \eta_{\text{irrad}}(t) \cdot \eta_{\text{temp}}(t) \cdot \eta_{\text{system}} \tag{8}$$

where $G_{\text{tilt}}(t)$ is the irradiance on the tilted surface (W/m^2), P_{nom} is the nominal power capacity (kW), $\eta_{\text{irrad}}(t)$ is the irradiance-related efficiency factor, $\eta_{\text{temp}}(t)$ is the temperature correction factor, and η_{system} represents system-level losses (including inverter efficiency, line losses, etc.). The reference irradiance G_{ref} is typically set to 1000 W/m^2 in the photovoltaic field.

The temperature efficiency factor is calculated as follows:

$$\eta_{\text{temp}}(t) = 1 + \gamma \cdot [T_{\text{module}}(t) - T_{\text{ref}}] \tag{9}$$

where γ is the temperature coefficient (typically -0.004/°C for crystalline silicon modules), and T_{ref} is the reference temperature (25°C).

Estimation of thermal models for module temperature:

$$T_{\text{module}}(t) = T_{\text{amb}}(t) + k_{\text{NOCT}} \cdot G_{\text{tilt}}(t) + \Delta T_{\text{wind}}(t) \tag{10}$$

where $T_{\text{module}}(t)$ represents the module temperature; $k_{\text{NOCT}} = 0.03\,^\circ\text{C} \cdot \text{m}^2/\text{W}$ denotes the coefficient related to the module's Nominal Operating Cell Temperature (NOCT), which represents the cell temperature when the module operates under open-circuit conditions at an irradiance of 800 W/m^2, ambient temperature of 20°C, and wind speed of 1 m/s.

Wind Speed Correction Factor (physically meaning: higher wind speeds result in more pronounced module heat dissipation and lower temperatures):

$$\Delta T_{\text{wind}}(t) = -2.0 \times \min\left(\frac{v_{\text{wind}}(t)}{10}, 1.0\right) \tag{11}$$

Irradiance efficiency factor accounts for the low-light effect:

$$\eta_{\text{irrad}}(t) = \begin{cases} 0.6 \cdot \dfrac{G_{\text{GHI}}(t)}{100}, & \text{if } G_{\text{GHI}}(t) \leq 100 \text{ W/m}^2 \\ 0.6 + 0.4 \times \min\left(1, \dfrac{G_{\text{GHI}}(t) - 100}{900}\right), & \text{if } G_{\text{GHI}}(t) > 100 \text{ W/m}^2 \end{cases} \tag{12}$$

This piecewise function captures the nonlinear response characteristics of photovoltaic modules under different irradiance levels.

2.2.2 Non-uniform Error Compensation Strategy Based on Irradiation Interval Partitioning

To address system bias and site-specific characteristics, we implemented a non-uniform error compensation strategy based on irradiation interval segmentation:

$$P_{\text{calib}}(t) = P_{\text{physical}}(t) \cdot k_{\text{calib}}(G_{\text{GHI}}(t)) \tag{13}$$

The calibration factor k_{calib} is divided into three intervals based on global horizontal irradiance (GHI):

$$k_{\text{calib}}(G) = \begin{cases} k_{\text{low}}, & \text{if } G \leq 200 \text{ W/m}^2 \\ k_{\text{mid}}, & \text{if } 200 < G \leq 500 \text{ W/m}^2 \\ k_{\text{high}}, & \text{if } G > 500 \text{ W/m}^2 \end{cases} \tag{14}$$

The calibration factor is determined using the median ratio between observed and predicted power during training, employing robust statistics to minimize the impact of outliers:

$$k_{\text{segment}} = \text{median} \left\{ \frac{P^{(i)}_{\text{observed}}}{P^{(i)}_{\text{physical}}} \mid i \in \text{segment} \right\} \tag{15}$$

This segmentation method recognizes the systematic variation in model error under different operating conditions, providing more precise calibration. Combining Equations (8) through (15), the overall calculation formula for the physical model is:

$$\widehat{P}_{\text{physics}}(t) = P_{\text{nom}} \cdot \eta_{\text{ref}} \cdot [1 + \gamma\,(T_{\text{module}}(t) - 25)] \cdot \frac{G_{\text{tilt}}(t)}{1000} \cdot k_{\text{calib}}(G_{\text{GHI}}(t)) \tag{16}$$

where η_{ref} denotes the reference conversion efficiency of the PV module under standard test conditions (STC: irradiance 1000 W/m^2, cell temperature 25°C, AM1.5 spectrum), which is typically provided by the module manufacturer and ranges from 0.15 to 0.22 for crystalline silicon modules. In this study, η_{ref} is set to 0.18 based on the specifications of the installed modules. Compared to Equation (8), the system loss factor η_{system} is implicitly incorporated into η_{ref} as an average derating factor, while the irradiance-dependent efficiency $\eta_{\text{irrad}}(t)$ is captured by the piecewise calibration factor $k_{\text{calib}}(G_{\text{GHI}})$, which inherently accounts for low-light efficiency degradation through segment-specific calibration.

2.3 Deep Learning Architectures for Physically Informed Enhancement

2.3.1 Feature Engineering

In photovoltaic power forecasting, relying solely on raw meteorological and electrical inputs often fails to fully capture the nonlinear effects of radiation, temperature, and environmental conditions on PV output. To fully leverage the multi-source information embedded in PV power forecasting, this study constructs a five-tier feature engineering system. This system encompasses raw meteorological features, temporal feature engineering, physically derived features, interaction features, and time-series statistical features, collectively forming a 37-dimensional input feature vector that aligns with the input layer structure of neural network models. This feature system design not only ensures comprehensive utilization of raw measurement data but also incorporates domain-knowledge-driven physical information and historical statistical patterns, thereby enabling a deeper characterization of the photovoltaic power evolution mechanism.

Raw Meteorological Features (10 dimensions): Directly collected via power plant monitoring system sensors, including global horizontal irradiance (GHI), direct normal irradiance (DNI), diffuse horizontal irradiance (DHI), tilt-angle irradiance, ambient temperature, atmospheric pressure, relative humidity, wind speed, wind direction, and real-time active power.

Time Feature Engineering (7 Dimensions): Time features exhibit distinct daily and annual cyclical patterns. To avoid boundary discontinuities (e.g., 23:00 → 0:00), sine-cosine cyclic encoding is employed:

$$h_{\sin}(t) = \sin\left(\frac{2\pi \cdot h(t)}{24}\right), \quad h_{\cos}(t) = \cos\left(\frac{2\pi \cdot h(t)}{24}\right) \tag{17}$$

$$d_{\sin}(t) = \sin\left(\frac{2\pi \cdot d(t)}{365}\right), \quad d_{\cos}(t) = \cos\left(\frac{2\pi \cdot d(t)}{365}\right) \tag{18}$$

$$m_{\sin}(t) = \sin\left(\frac{2\pi \cdot m(t)}{1440}\right), \quad m_{\cos}(t) = \cos\left(\frac{2\pi \cdot m(t)}{1440}\right) \tag{19}$$

Additionally, monthly integer encoding (1–12) is introduced to capture seasonal effects.

Physical derivative features (10 dimensions) are combined with mechanistic modeling results of PV power forecasting to introduce the following features. The calculation of the component temperature is the same as that of Equation (10).

Theoretical power:

$$P_{\text{theory}}(t) = \frac{G_{\text{tilt}}(t)}{1000} \cdot P_{\text{nom}} \cdot \eta_{\text{ref}} \tag{20}$$

Radiation Quality Indicators:

$$Q_{\text{rad}}(t) = \begin{cases} \dfrac{G_{\text{DNI}}(t)}{G_{\text{GHI}}(t)}, & \text{if } G_{\text{GHI}}(t) > 1 \\ 0, & \text{otherwise} \end{cases} \tag{21}$$

Scattering ratio:

$$R_{\text{diffuse}}(t) = \begin{cases} \dfrac{G_{\text{DHI}}(t)}{G_{\text{GHI}}(t)}, & \text{if } G_{\text{GHI}}(t) > 1 \\ 0, & \text{otherwise} \end{cases} \tag{22}$$

Inclined Surface Advantage:

$$A_{\text{tilt}}(t) = \max\left(0.8, \min\left(1.5, \frac{G_{\text{tilt}}(t)}{G_{\text{GHI}}(t) + \epsilon}\right)\right) \tag{23}$$

Sunny Sky Index:

$$K_c(t) = \min\left(1, \frac{G_{\text{GHI}}(t)}{E_0(t) \cdot \cos(\theta_z(t)) + \epsilon}\right) \tag{24}$$

where $E_0(t)$ is the extraterrestrial irradiance and $\theta_{\text{zt}}(t)$ is the solar zenith angle.

Air density factor:

$$\rho_{\text{air}}(t) = \frac{P_{\text{atm}}(t)}{P_0} \times \frac{T_0}{T_0 + T_{\text{amb}}(t)} \tag{25}$$

where $P_0 = 1013.25$ hPa is the standard atmospheric pressure, $T_0 = 273.15$ K is the reference temperature, and $T_{\text{amb}}(t)$ is the ambient temperature in °C.

Shadow Indicator:

$$I_{\text{shadow}}(t) = \mathbf{1}_{[G_{\text{DNI}}(t)<50 \text{ and } G_{\text{GHI}}(t)>100]} \tag{26}$$

The calculation of the power predicted by the physical model is the same as that of formula (16).

Interaction Features (3 dimensions): To capture nonlinear coupling effects between variables, this study designed three types of interaction features: Radiation-temperature interaction features describe efficiency degradation under high radiation and high temperature conditions; Wind-cooling-temperature interaction features reflect the regulatory effect of wind cooling on module temperature, the baseline temperature of 15°C is the empirically optimal value; Air quality interaction features reveal the combined influence of atmospheric transmittance and cloud cover conditions. These interaction features are generated through element-wise multiplication, offering both simplicity and clear physical interpretation.

Radiation-temperature interaction features:

$$f_{\text{int}}^{(1)}(t) = G_{\text{GHI}}(t) \times \eta_{\text{temp}}(t) \tag{27}$$

Wind-cooling-temperature interaction features:

$$f_{\text{int}}^{(2)}(t) = v_{\text{wind}}(t) \times [T_{\text{amb}}(t) - 15] \tag{28}$$

Air quality interaction features:

$$f_{\text{int}}^{(3)}(t) = Q_{\text{rad}}(t) \times K_c(t) \tag{29}$$

Temporal Statistical Features (7 dimensions): Photovoltaic power sequences exhibit significant autocorrelation characteristics. Based on a sliding window approach (96 time steps, equivalent to 24 hours), we extract temporal statistical features, including moving average of power output, moving standard deviation of power output, moving maximum of power output, moving average of GHI, and lagged power values at 1-step, 2-step, and 4-step intervals. The moving average and standard deviation are employed to characterize short-term stationarity, while the lagged terms incorporate historical information to enable the model to capture temporal dependencies.

In summary, the final input feature vector comprises five components, as shown in Table 1.

Table 1: Feature Fusion and Input Construction

Feature Type	**Count**	**Example Features**	**Effect**
Raw Meteorological	10	GHI, DNI, temperature	Meteorological environment
Time Feature	7	Hour/Day Cyclic Encoding	Cyclical Patterns
Physical Derivative	10	Module temperature	Physical constraints
Interaction Features	3	Radiation-temperature	Nonlinear coupling
Time-series Statistical	7	Moving average, lag term	Historical dependency

To validate the effectiveness of the proposed 37-dimensional feature system and elucidate the intrinsic correlations between individual features and photovoltaic power output, we employed a comprehensive evaluation framework comprising Pearson correlation coefficient, Spearman rank correlation coefficient, and partial correlation analysis. Specifically, the Pearson correlation coefficient quantifies the degree of linear association between variables, the Spearman coefficient assesses monotonic relationships through rank transformation, and partial correlation analysis reveals the net correlation between two variables while controlling for confounding effects. Based on 34,835 measured data points, the correlation coefficients between raw meteorological variables and power output were calculated, with results presented in Table 2.

Table 2: Correlation coefficients between meteorological variables and power output

Variable	**Pearson**	**Spearman**	**Kendall**	**Sig.**
GHI	0.9273	0.9426	0.8852	***
Tilted irradiance	0.8387	0.9124	0.8061	***
DNI	0.7549	0.8668	0.7483	***
DHI	0.6375	0.7714	0.6611	***
Humidity	-0.4363	-0.3929	-0.2821	***
Wind speed	0.2794	0.2816	0.2056	***
Temperature	0.1148	0.2106	0.1537	***
Wind direction	0.1058	0.0745	0.0530	***
Pressure	0.0009	-0.0969	-0.0692	ns

Note: *** $p < 0.001$; ns = not significant

As shown in Table 2, irradiance-related variables exhibited strong positive correlations with power output. Global horizontal irradiance (GHI) demonstrated the highest Pearson correlation coefficient of 0.9273, indicating its predominant influence on PV power generation. Tilted irradiance ($r = 0.8387$) and direct normal irradiance (DNI, $r = 0.7549$) also exhibited substantial correlations. Relative humidity showed a moderate negative correlation with power output ($r = -0.4363$), consistent with the physical mechanism in which increased humidity reduces atmospheric transmittance. Notably, atmospheric pressure showed a near-zero correlation coefficient, which was not statistically significant, suggesting a negligible direct impact on power generation.

To visualize the interrelationships among variables more intuitively, Fig. 4 presents heatmap matrices of both Pearson and Spearman correlation coefficients.

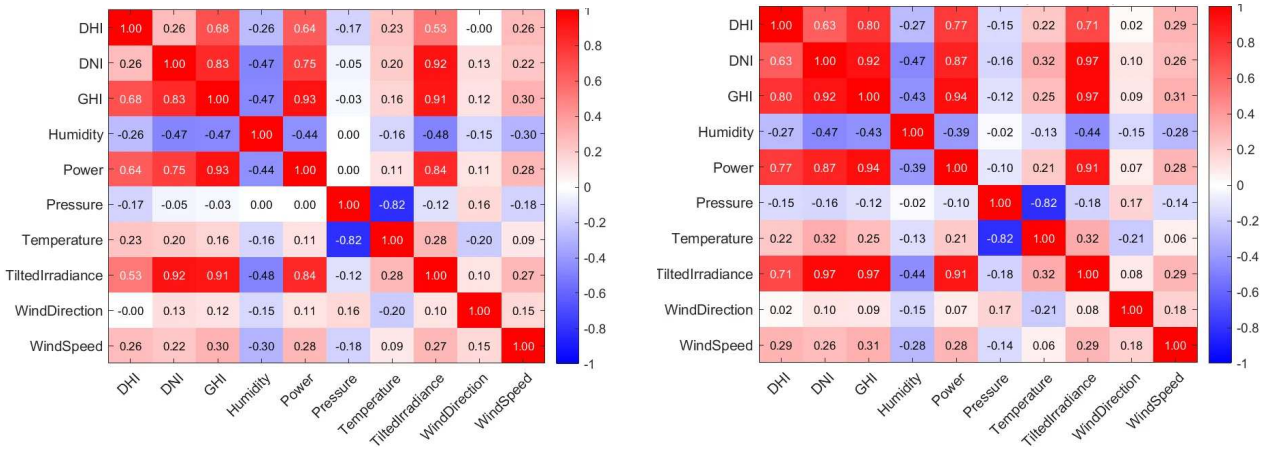


Figure 4: Correlation heatmaps: (a) Pearson correlation coefficient matrix; (b) Spearman correlation coefficient matrix

Several key observations can be drawn from Fig. 4: (1) Irradiance-related variables (GHI, DNI, DHI, and Tilted Irradiance) exhibit strong positive correlations, forming a distinct red cluster region, indicating that these variables carry similar solar radiation information. (2) Power demonstrates the strongest correlation with GHI (Pearson $r = 0.93$, Spearman $\rho = 0.94$), followed by TiltedIrradiance ($r = 0.84$, $\rho = 0.91$), confirming that irradiance is the dominant factor influencing photovoltaic power generation. (3) Humidity shows negative correlations with irradiance-related variables (blue regions), reflecting the physical phenomenon of reduced atmospheric transmittance under high-humidity conditions. (4) Pressure and Temperature exhibit a strong negative correlation ($r = -0.82$), consistent with atmospheric thermodynamic principles. (5) Wind direction displays generally weak correlations with other variables, suggesting its limited direct contribution to power prediction. The Spearman correlation coefficients are slightly higher than their Pearson counterparts overall, indicating the presence of certain nonlinear monotonic relationships among variables.

Further correlation analysis was conducted on the constructed 37-dimensional enhanced feature set. The average correlations between each feature category and power output are summarized in Table 3.

Table 3: Average correlations between feature categories and power output

Feature Category	**Dimensionality**	**Mean $\lvert r \rvert$**
Interaction features	3	0.568
Time-series Statistical Features	7	0.539
Raw Meteorological Features	10	0.510
Physical Derivative Features	10	0.430
Time Features	7	0.278

The analysis reveals that interaction features and temporal statistical features yield the highest average correlations of 0.568 and 0.539, respectively, validating the effectiveness of feature engineering in enhancing model representational capability. These correlation analysis results provide statistical justification for the 37-dimensional physics-enhanced feature framework adopted in this study, while simultaneously demonstrating the rationality of integrating physics-based models with data-driven approaches.

The model incorporates a total of 37 dimensional features. In the model implementation, these features undergo Z-score normalization and are then expanded into a matrix based on a 96-step (24-hour) sliding window:

$$\mathbf{X} = \begin{bmatrix} \mathbf{x}_1^\top \\ \mathbf{x}_2^\top \\ \vdots \\ \mathbf{x}_T^\top \end{bmatrix} \in \mathbb{R}^{T \times D} \tag{30}$$

where $T = 96$ represents the time steps and $D = 37$ denotes the fused feature dimensionality. This input tensor is directly connected to the neural network's input layer, serving as the foundation for downstream multi-scale convolution, temporal recurrence, and attention-based modeling.

The feature vector comprises:

$$\mathbf{x}_t = [\mathbf{f}_{\text{raw}}(t); \mathbf{f}_{\text{time}}(t); \mathbf{f}_{\text{physics}}(t); \mathbf{f}_{\text{interact}}(t); \mathbf{f}_{\text{stat}}(t)] \tag{31}$$

2.3.2 Deep Neural Network

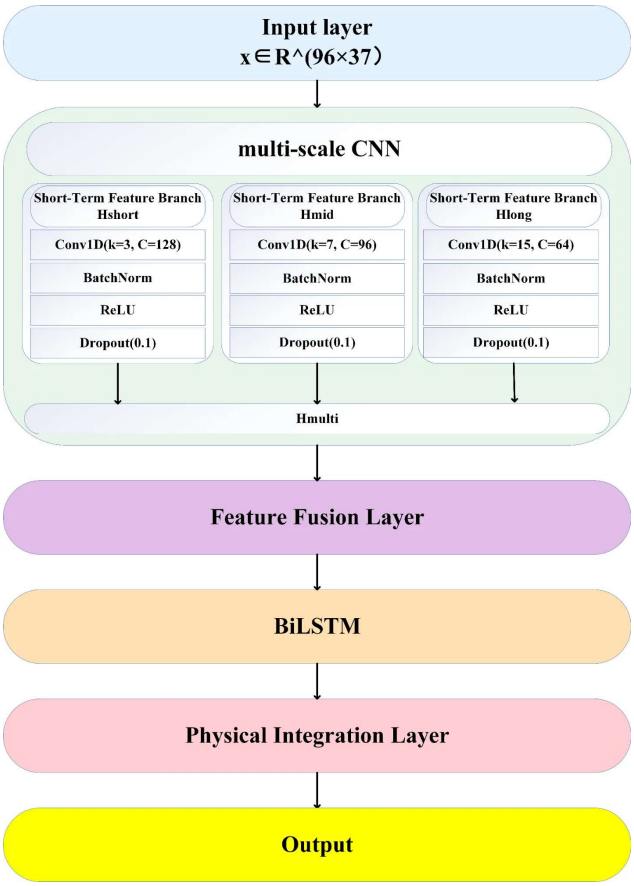


Figure 5: Deep Neural Network Architecture for the End-to-End Learning Paradigm

We propose a deep neural network architecture based on an end-to-end learning paradigm, as illustrated in Fig. 5, for efficient feature extraction and predictive modeling of complex time-series data. The overall architecture comprises an input layer, multi-scale convolutional feature extraction layers, feature fusion layers, BiLSTM temporal modeling layers, a physical integration layer, and an output layer. These components work synergistically to capture multi-scale temporal dependencies while achieving deep integration of physical prior knowledge with data-driven features. Specifically, multi-scale convolutions extract feature patterns across varying time spans; BiLSTM modules model bidirectional dependencies to effectively represent temporal contextual information; the physical integration layer fuses prior knowledge with learned features to enhance model interpretability and prediction stability; and the regression output layer generates final continuous-value predictions.

The deep neural network architecture for end-to-end learning comprises the following components:

Input Layer: The input feature matrix $\mathbf{X} \in \mathbb{R}^{96\times37}$ directly receives preprocessed normalized data without additional standardization, thereby preserving proportional relationships among features and original distribution characteristics. This provides high-fidelity input for subsequent multi-scale parallel convolutions.

Multi-scale Convolutional Feature Extraction Layers: The core function of Convolutional Neural Networks (CNNs) is local feature extraction. By sliding convolutional kernels across the input, CNNs effectively capture local patterns while reducing dimensionality, yielding concise yet information-rich feature representations for subsequent processing. Since photovoltaic power fluctuations are influenced by multi-timescale factors (e.g., instantaneous cloud cover, diurnal irradiance variations, and seasonal patterns), a single convolutional kernel cannot fully capture cross-scale feature patterns. To capture these cross-scale patterns, we employ a three-branch parallel Conv1D architecture with varying kernel sizes, as detailed in Table 4. Each branch follows an identical processing pipeline:

$$\mathbf{H}_i = \text{Dropout}\Big(\text{ReLU}\big(\text{BN}(\text{Conv1D}_{k_i}^{C_i}(\mathbf{X}))\big), p = 0.1\Big), \quad i \in \{\text{short}, \text{mid}, \text{long}\} \tag{32}$$

Table 4: Multi-scale convolutional branch configurations

Branch	Kernel Size	Filters	Target Scale
Short-term	3	128	Rapid fluctuations
Mid-term	7	96	Hourly variations
Long-term	15	64	Diurnal patterns

The outputs from the three branches are concatenated along the channel dimension to form a multi-scale temporal feature representation.

$$\mathbf{H}_{\text{multi}} = \text{Concat}([\mathbf{H}_{\text{short}}, \mathbf{H}_{\text{mid}}, \mathbf{H}_{\text{long}}]) \in \mathbb{R}^{T \times 288} \tag{33}$$

Feature Fusion Layer: To further integrate multi-scale feature information and map it to a unified high-dimensional feature space, a fusion convolutional layer is designed. This layer employs a Conv1D structure with a kernel size of 1 and 256 channels, achieving feature compression and mapping without altering the temporal dimension. Subsequently, batch normalization, ReLU activation, and Dropout (with a probability of 0.2) are applied sequentially to enhance training stability and prevent overfitting:

$$\mathbf{H}_{\text{fused}} = \text{Dropout}(\text{ReLU}(\text{BN}(\text{Conv1D}_{k=1}^{C=256}(\mathbf{H}_{\text{multi}}))), p = 0.2) \tag{34}$$

This layer effectively fuses features across multiple temporal scales, providing a unified and comprehensive feature representation for subsequent temporal modeling and prediction.

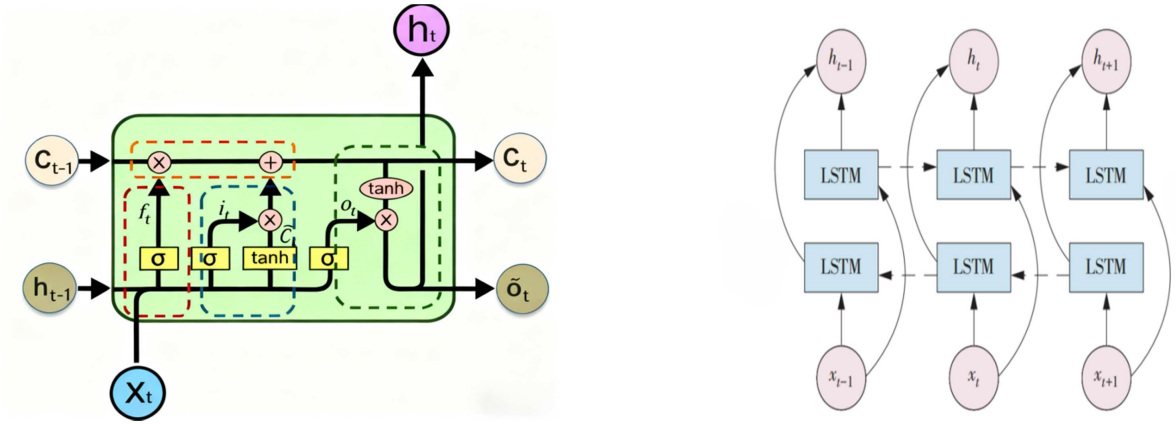


Figure 6: Neural Network Architectures: (a) LSTM; (b) BiLSTM

To further model the global temporal dependencies of photovoltaic power sequences, this study introduces a bidirectional long short-term memory (BiLSTM) architecture based on the $\mathbf{H}_{\text{fused}}$ fusion feature, as shown in Fig. 6. By simultaneously processing both forward and backward information in the time series, the

BiLSTM captures bidirectional contextual features, enabling a more comprehensive modeling of dynamic trends.

Forward LSTM:

$$\begin{cases} i_t = \sigma\left(W_i \cdot [h_{t-1}, x_t] + b_i\right) \\ f_t = \sigma\left(W_f \cdot [h_{t-1}, x_t] + b_f\right) \\ o_t = \sigma\left(W_o \cdot [h_{t-1}, x_t] + b_o\right) \\ \tilde{C}_t = \tanh\left(W_C \cdot [h_{t-1}, x_t] + b_C\right) \\ C_t = f_t \cdot C_{t-1} + i_t \cdot \tilde{C}_t \\ h_t = o_t \odot \tanh(C_t) \\ \overrightarrow{\mathbf{h}}_t = \mathrm{LSTM}_{\mathrm{forward}}(\mathbf{x}_t, \overrightarrow{\mathbf{h}}_{t-1}, \overrightarrow{\mathbf{C}}_{t-1}) \end{cases} \tag{35}$$

The backward LSTM uses similar equations but processes the sequence in reverse order. The final hidden state is the concatenation of the forward and backward states:

$$\overleftarrow{\mathbf{h}}_t = \mathrm{LSTM}_{\mathrm{backward}}(\mathbf{x}_t, \overleftarrow{\mathbf{h}}_{t+1}, \overleftarrow{\mathbf{C}}_{t+1}) \tag{36}$$

Bidirectional State Concatenation:

$$\mathbf{h}_t^{\mathrm{bi}} = [\overrightarrow{\mathbf{h}}_t; \overleftarrow{\mathbf{h}}_t] \tag{37}$$

This bidirectional structure simultaneously leverages historical and future contextual information (during training), significantly enhancing the model's ability to represent complex temporal dependencies.

The first layer BiLSTM is configured with 128 hidden units, and sets return_sequences=True to preserve the complete sequence of temporal features:

$$\{\mathbf{h}_t^{(1)}\}_{t=1}^{T} = \mathrm{BiLSTM}_{n=128}^{(1)}(\mathbf{H}_{\mathrm{fused}}) \tag{38}$$

Subsequently, a Dropout layer (0.25) is applied to enhance regularization. The second BiLSTM layer is configured with 64 hidden units and sets return_sequences=False, outputting only the final state of the sequence as a compressed representation of the temporal features. This is followed by another Dropout layer (0.25):

$$\mathbf{h}_{\mathrm{final}} = \mathrm{BiLSTM}_{n=64}^{(2)}(\{\mathbf{h}_t^{(1)}\}_{t=1}^{T}) \tag{39}$$

Following the BiLSTM layer, we introduced a dedicated physics integration layer to explicitly incorporate physical model predictions:

$$\mathbf{F}_{\mathrm{physics}} = \mathrm{Dropout}(\mathrm{ReLU}(\mathrm{BN}(\mathbf{W}_{\mathrm{fusion}} \cdot [\mathbf{h}_{\mathrm{final}}; \widehat{P}_{\mathrm{physics}}; \mathbf{f}_{\mathrm{physics}}] + \mathbf{b}_{\mathrm{fusion}})), p = 0.3) \tag{40}$$

This layer contains 128 neurons and employs batch normalization with ReLU activation functions to maintain stable feature distributions and preserve nonlinear numerical expression capabilities. The introduction of Dropout (0.3) further reduces the risk of model overfitting, enabling the network to achieve stronger generalization performance under physical consistency constraints.

Next is the final output layer, which consists of two fully connected mapping stages. First, an intermediate representation is obtained through nonlinear transformation of the hidden layer:

$$\mathbf{z} = \mathrm{ReLU}(\mathbf{W}_{\mathrm{hidden}} \cdot \mathbf{F}_{\mathrm{physics}} + \mathbf{b}_{\mathrm{hidden}}) \tag{41}$$

The final prediction result is then obtained through a linear transformation:

$$\widehat{f}_{\mathrm{data}} = \mathbf{W}_{\mathrm{output}}^{\top} \cdot \mathbf{z} + b_{\mathrm{output}} \tag{42}$$

This output layer achieves efficient regression mapping of fused features while maintaining temporal consistency, providing the final output for precise photovoltaic power prediction.

2.3.3 Loss function

The loss function employs a composite loss function, where the composite loss function $\mathcal{L}$ is composed of three weighted components. Its mathematical expression is:

$$\mathcal{L} = \alpha\mathcal{L}_{\text{MAE}} + \beta\mathcal{L}_{\text{MSE}} + \gamma\mathcal{L}_{\text{physics}} \tag{43}$$

The weights α, β, and γ represent the respective loss terms, balancing their contributions to the overall loss. To balance the effects of data fitting and physical constraints, sensitivity experiments were conducted within the γ range $\{0.01, 0.05, 0.1, 0.2, 0.5\}$. Results indicate that when $\gamma < 0.05$, insufficient physical constraints lead to non-zero outputs at night; When $\gamma > 0.2$, the physical constraints become overly stringent, causing underfitting and increasing the RMSE by approximately 12%. After optimization, $\gamma = 0.1$ was selected as the optimal weight. This value demonstrates an effective balance between physical consistency and high prediction accuracy under various meteorological conditions, exhibiting good stability and physical plausibility. We set $\alpha = 0.5$, $\beta = 0.4$, and $\gamma = 0.1$ based on grid search optimization over the validation set, where $\alpha \in \{0.3, 0.4, 0.5, 0.6\}$ and $\beta \in \{0.2, 0.3, 0.4, 0.5\}$ with γ fixed at 0.1. The selected combination achieved the minimum validation error.

$$\mathcal{L}_{\text{MAE}} = \frac{1}{N}\sum_{i=1}^{N}|y_i - \widehat{y}_i| \tag{44}$$

where N is the sample size, y_i is the true value of the i-th sample, $\hat{y}_i$ is the predicted value of the i-th sample, and $|y_i - \hat{y}_i|$ represents the absolute error between the true value and the predicted value. This metric is relatively insensitive to outliers and reflects the overall bias in predictions.

The mean squared error loss measures the average of the squared errors between predicted values and true values, expressed mathematically as:

$$\mathcal{L}_{\text{MSE}} = \frac{1}{N}\sum_{i=1}^{N}(y_i - \widehat{y}_i)^2 \tag{45}$$

By squaring the error, this approach amplifies the impact of larger errors, placing greater emphasis on the deviation between predicted and actual values. This allows the model to learn more thoroughly from samples with significant discrepancies.

The physical constraint loss is used to incorporate prior knowledge from the physical domain, imposing bounds on the range of predicted values. Its mathematical expression is:

$$L_{\text{physics}} = \frac{1}{N}\sum_{i=1}^{N}\max(0, -\hat{y}_i) + \frac{\lambda}{N}\sum_{i=1}^{N}\max(0, \hat{y}_i - 1.1P_{\text{nom}}) \tag{46}$$

where the first part, $\frac{1}{N}\sum_{i=1}^{N}\max(0, -\hat{y}_i)$, constrains the predicted value $\hat{y}_i$ from being negative. When $\hat{y}_i < 0$, $\max(0, -\hat{y}_i) = -\hat{y}_i$, resulting in loss; when $\hat{y}_i \geq 0$, this component's loss is zero.

The second part, $\frac{\lambda}{N}\sum_{i=1}^{N}\max(0, \hat{y}_i - 1.1P_{\text{nom}})$, constrains the predicted value $\hat{y}_i$ from exceeding 1.1 times the rated capacity P_{nom}, with $\lambda = 0.01$ as the weight for this part. When $\hat{y}_i > 1.1P_{\text{nom}}$, the loss is calculated as $\max(0, \hat{y}_i - 1.1P_{\text{nom}}) = \hat{y}_i - 1.1P_{\text{nom}}$; when $\hat{y}_i \leq 1.1P_{\text{nom}}$, the loss for this part is zero.

Both terms are normalized by $\frac{1}{N}$ to ensure dimensional consistency, making the physical constraint loss comparable in scale regardless of batch size.

2.4 Dynamic Weight Fusion Mechanism Based on Four-Dimensional Environmental Perception

2.4.1 Multi-dimensional Confidence Assessment System

Traditional fusion methods employ fixed weights. We propose a dynamic adjustment of model weights based on confidence assessment across four key dimensions. The threshold values in the following confidence functions are determined based on established photovoltaic operating characteristics and validated through sensitivity analysis on the validation set.

Irradiance confidence:

$$f_{\text{irrad}}(G) = \begin{cases} 1.0, & \text{if } G > 800\,\text{W/m}^2 \\ 0.3 + 0.7 \times \dfrac{G - 300}{500}, & \text{if } 300 < G \leq 800\,\text{W/m}^2 \\ 0.1 + 0.2 \times \dfrac{G - 50}{250}, & \text{if } 50 < G \leq 300\,\text{W/m}^2 \\ 0.05, & \text{if } G \leq 50\,\text{W/m}^2 \end{cases} \tag{47}$$

Time Confidence:

$$f_{\text{time}}(h) = \begin{cases} 0.8 + 0.2 \times \exp\left(-\dfrac{(h-12)^2}{32}\right), & \text{if } 8 \leq h \leq 16 \\ 0.5, & \text{if } 6 \leq h < 8 \text{ or } 16 < h \leq 19 \\ 0.1, & \text{otherwise} \end{cases} \tag{48}$$

Temperature Confidence:

$$f_{\text{temp}}(T) = \begin{cases} 1.0, & \text{if } 15 \leq T \leq 35 \\ 0.85, & \text{if } 10 \leq T < 15 \text{ or } 35 < T \leq 40 \\ 0.7, & \text{otherwise} \end{cases} \tag{49}$$

Weather Stability Confidence:

$$f_{\text{weather}}(v) = \begin{cases} 1.0, & \text{if } v < 3\,\text{m/s} \\ 0.8, & \text{if } 3 \leq v < 8\,\text{m/s} \\ 0.6, & \text{if } v \geq 8\,\text{m/s} \end{cases} \tag{50}$$

Composite Confidence (Weighted Geometric Mean):

The four factors influencing the data-driven model—irradiance, time, temperature, and weather stability—are assigned four-dimensional confidence weights (0.4, 0.3, 0.15, 0.15) based on feature importance analysis using SHAP (SHapley Additive exPlanations) values computed on the validation set. The SHAP analysis revealed that irradiance exhibits the highest sensitivity to prediction performance (42%), followed by time (29%). Temperature and weather stability contribute similarly (around 15% each). Therefore, weights are normalized according to relative importance to ensure the composite confidence reflects dominant factors in the real environment. The composite confidence C_{data} for the data-driven model is calculated using a weighted geometric mean.

$$C_{\text{data}} = \prod_{j=1}^{4} f_j^{w_j} = f_{\text{irrad}}^{0.4} \times f_{\text{time}}^{0.3} \times f_{\text{temp}}^{0.15} \times f_{\text{weather}}^{0.15} \tag{51}$$

The calculated C_{data} may yield overly extreme values (such as being too close to 0 or 1), so the clip function is used to constrain C_{data} between 0.05 and 0.95.

$$C_{\text{data}}^* = \text{clip}(C_{\text{data}}, 0.05, 0.95) \tag{52}$$

Confidence level of the physical model:

$$C_{\text{physics}} = 1 - C_{\text{data}}^* \tag{53}$$

Confidence Fusion Strategy:

$$\widehat{P}_{\text{confidence}} = C_{\text{data}}^* \cdot \widehat{P}_{\text{data}} + C_{\text{physics}} \cdot \widehat{P}_{\text{physics}} \tag{54}$$

2.4.2 Error-weighted fusion

Validation Set Error Calculation:

$$e_{\text{data}} = \frac{1}{N_{\text{val}}} \sum_{i=1}^{N_{\text{val}}} \left| y_i^{\text{val}} - \widehat{P}_{\text{data}}^{(i)} \right| \tag{55}$$

On the validation set, compute the average absolute error between the data-driven model's predicted value $p_{\text{di}}^{\text{val}}$ and the true value $v_{\text{di}}^{\text{val}}$. Here, N_{val} represents the number of samples in the validation set. By summing the absolute error for each sample and then taking the average, the overall error of the data-driven model on the validation set is obtained.

$$e_{\text{physics}} = \frac{1}{N_{\text{val}}} \sum_{i=1}^{N_{\text{val}}} \left| y_i^{\text{val}} - \widehat{P}_{\text{physics}}^{(i)} \right| \tag{56}$$

Calculate the average absolute error between the physical model prediction $\hat{P}_{\text{physics}}^{(i)}$ and the true value y_i^{val}. This follows the same principle as the data-driven model error calculation and measures the overall error of the physical model on the validation set.

Error-inverse weighting:

$$w_{\text{data}} = \frac{1/e_{\text{data}}}{1/e_{\text{data}} + 1/e_{\text{physics}}} \tag{57}$$

$$w_{\text{physics}} = \frac{1/e_{\text{physics}}}{1/e_{\text{data}} + 1/e_{\text{physics}}} \tag{58}$$

Error-weighted fusion prediction:

$$\widehat{P}_{\text{error}} = w_{\text{data}} \cdot \widehat{P}_{\text{data}} + w_{\text{physics}} \cdot \widehat{P}_{\text{physics}} \tag{59}$$

The purpose of this approach is to assign a higher weight to models with smaller errors in the fusion results, thereby improving the overall prediction accuracy.

2.4.3 Ultimate Integration

The final prediction is obtained by weighting the sum of confidence-based fusion and error-weighted fusion. The initial fusion weights ($\alpha = 0.7$, $\beta = 0.3$) were determined through grid search on the validation set, achieving minimal model error (RMSE $\approx$ 6.4 kW) within the $\alpha \in [0.6, 0.8]$ range. This indicates that the data-driven branch contributes more significantly under high-irradiance scenarios. To balance accuracy and robustness, the final initial fusion ratio was selected as $\alpha = 0.7$ and $\beta = 0.3$. This yields a more comprehensive and accurate final prediction result:

$$\widehat{P}_{\text{final}} = 0.7 \times \widehat{P}_{\text{confidence}} + 0.3 \times \widehat{P}_{\text{error}} \tag{60}$$

3 Experimental Setup

3.1 Dataset

This study utilizes empirical data collected from a photovoltaic power station in a prefecture-level city in China. The dataset spans an entire year across multiple seasons, with a 15-minute sampling interval, totaling 34,835 samples. It encompasses various typical weather conditions, including clear, cloudy, and overcast days, and consists of 23 variables covering meteorological parameters, electrical parameters, and temporal information.

3.2 Training Configuration

During the model training phase, the data is initially partitioned using a rigorous time-series segmentation approach to prevent future information leaking. The training set comprises 70%, while the validation and test sets each constitute 15%. The input sequence length is established at 96 time steps, encompassing a complete 24-hour cycle, with 37-dimensional input features, including meteorological, electrical, physical derivatives, and temporal loops. The Adam optimizer is employed, with an initial learning rate of 0.0005 and a learning rate scheduling strategy that reduces the rate to 0.7 times the original every 25 epochs to enhance convergence stability. Additionally, gradient clipping (threshold 1.0) is implemented during training to avert gradient explosion, and L2 regularization (coefficient 0.0005) is applied to the network weights to mitigate overfitting and ensure the model's generalization performance in photovoltaic power prediction tasks. The model training was finalized using an Intel i7-13700K CPU and an NVIDIA RTX 4060 GPU.

3.3 Evaluation Metrics

To comprehensively evaluate model performance, we employ the following four metrics.

Mean Absolute Error (MAE):

$$\mathrm{MAE} = \frac{1}{N}\sum_{i=1}^{N}|y_i - \widehat{y}_i| \tag{61}$$

MAE quantifies the average prediction error magnitude, expressed in kW for intuitive interpretation.

Root Mean Square Error (RMSE):

$$\mathrm{RMSE} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(y_i - \widehat{y}_i)^2} \tag{62}$$

RMSE is more sensitive to large errors and reflects prediction stability.

Mean Absolute Percentage Error (MAPE):

$$\mathrm{MAPE} = \frac{100}{N'}\sum_{i:y_i>\theta}\frac{|y_i - \widehat{y}_i|}{y_i} \tag{63}$$

where θ=10 kW is introduced to exclude low-power periods and prevent percentage distortion.

Coefficient of Determination (R^2):

$$R^2 = 1 - \frac{\sum_{i=1}^{N}(y_i - \widehat{y}_i)^2}{\sum_{i=1}^{N}(y_i - \bar{y})^2} \tag{64}$$

R^2 measures the model's ability to explain data variability, with values ranging from 0 to 1.

4 Results and Analysis

4.1 Overall Performance

The performance comparison between the proposed dual-driver prediction system and the single-model approach is presented in Table 5 and Fig. 7. The experimental results demonstrate that the proposed dual-driver prediction system significantly outperforms the single-model approach in photovoltaic power forecasting. To verify the statistical significance of the proposed method, a paired t-test was conducted to compare the prediction errors between the proposed method and the purely data-driven approach on the test set. The results indicate $p < 0.001$, confirming that the improvement is statistically significant.

The experimental results presented in Table 5 and Fig. 7 demonstrate that the proposed dual-driver prediction system significantly outperforms single-model methods in photovoltaic power forecasting.

The physical model, which relies exclusively on mechanism-based equations incorporating radiation and temperature parameters, fails to adequately account for environmental noise and equipment operational

Table 5: Prediction Results

Model	MAE (kW)	RMSE (kW)	MAPE (%)	R^2
Physical Model	15.585	32.95	66.77	0.4157
Neural Network	6.654	10.173	26.54	0.9443
Proposed	**3.571**	**6.384**	**22.62**	**0.9781**

variations. Consequently, it exhibits substantial prediction errors: the Mean Absolute Error (MAE) is 15.585 kW, the Root Mean Square Error (RMSE) is 32.95 kW, the Mean Absolute Percentage Error (MAPE) is 66.77%, and the coefficient of determination (R^2) is only 0.4157, indicating a limited capacity to accurately characterize actual power generation.

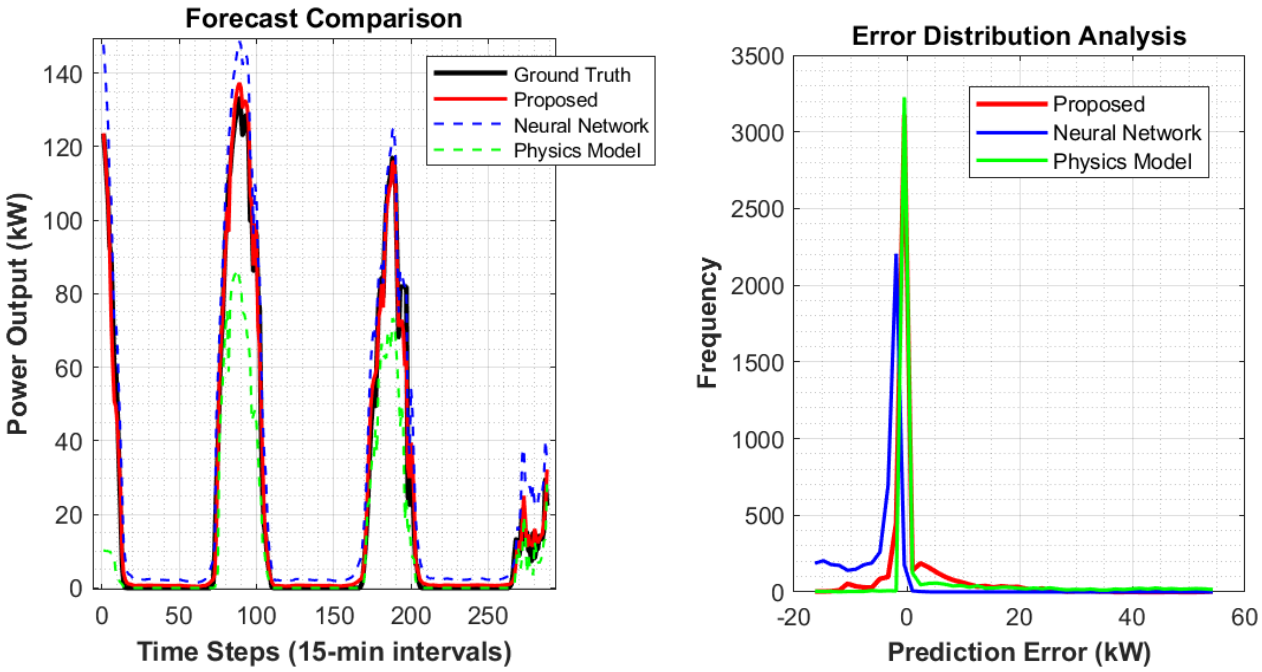


Figure 7: Prediction Results: (a) Prediction Comparison; (b) Error Distribution Analysis

The data-driven model markedly improves predictive accuracy by capturing complex nonlinear relationships through deep neural networks. The MAE decreases to 6.654 kW, the RMSE to 10.173 kW, the MAPE to 26.54%, and the R^2 increases to 0.9443, demonstrating a strong capability to identify temporal patterns and correlations across multidimensional inputs. However, this approach still exhibits insufficient adaptability to abrupt changes under adverse weather conditions.

The proposed ensemble fusion model integrates the mechanism-based constraints of physical models with the feature extraction capabilities of data-driven models, yielding substantial performance improvements. The final MAE is reduced to 3.571 kW, RMSE to 6.384 kW, MAPE to 22.62%, and R^2 reaches 0.9781, approaching optimal fitting performance. These results indicate that the dual-driver methodology not only enhances predictive accuracy but also strengthens the model's generalization capability under complex and variable meteorological conditions, highlighting the synergistic benefits of integrating data-driven and physics-based approaches.

4.2 Ablation Experiment Analysis

We performed systematic ablation experiments to validate the efficacy of each module. Figure 8 illustrates the outcomes of the ablation experiment, with further data provided in Table 6.

Table 6: Ablation Experiment Results

Configuration	MAE	RMSE	MAPE	R^2	ΔMAE
Full Model (Proposed)	**3.571**	**6.384**	**22.62**	**0.9781**	—
w/o Physics Loss	3.892	6.930	24.78	0.9724	+9.0%
w/o Multi-Scale CNN	4.156	7.450	26.91	0.9683	+16.4%
w/o BiLSTM	4.523	8.090	28.84	0.9632	+26.7%
w/o Fusion	6.654	10.21	40.10	0.9443	+86.4%
w/o Physical Features	5.238	9.310	32.42	0.9538	+46.7%
w/o Time Features	4.687	8.12	30.54	0.9601	+31.2%
Base Features Only	7.823	13.02	48.71	0.9215	+119.1%

Ablation tests were conducted to evaluate the necessity and contribution of each component by progressively removing core modules, including temporal features, physical features, the fusion module, the BiLSTM layer, the Multi-scale CNN module, and the physics-informed loss function. The full model served as the baseline for performance comparison, with evaluation metrics including MAE, RMSE, MAPE, and R^2. The experimental results are presented in Table 6 and Fig. 8.

The ablation study systematically removes individual components to quantify their respective contributions to overall performance. Fig. 8 presents the MAE comparison across different configurations in a bar chart format. The baseline configuration "Base Features Only" exhibits the highest MAE of 7.823 kW, serving as a reference point for assessing component contributions.

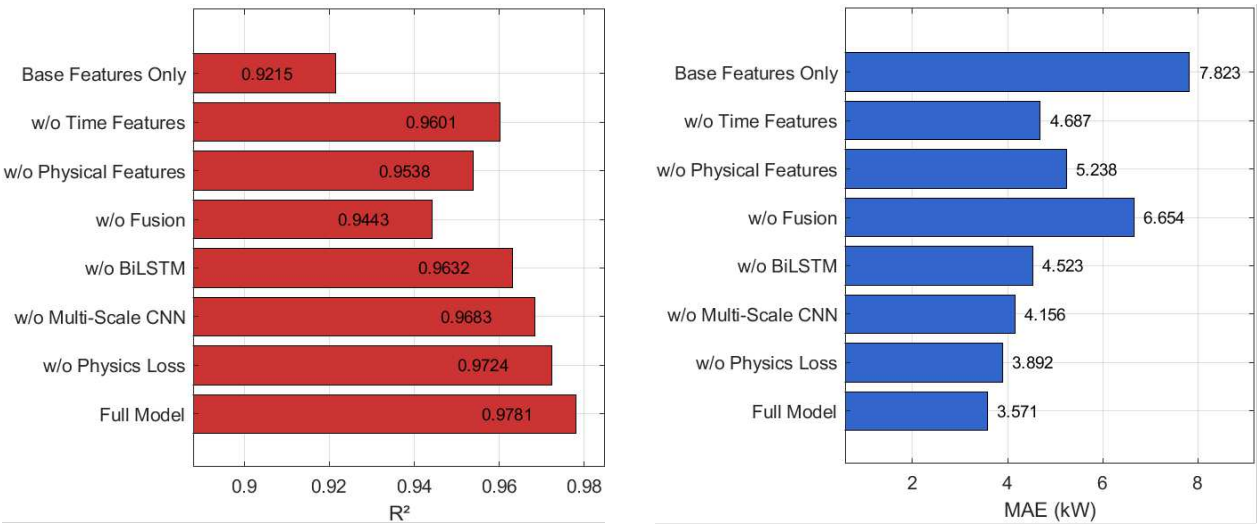


Figure 8: Ablation Experiment Results: (a) MAE Comparison; (b) R^2 Comparison

Removing temporal features (w/o Time Features) results in an MAE of 4.687 kW, demonstrating that temporal features are critical for capturing temporal dynamics and reducing prediction errors. Excluding physical features (w/o Physical Features) yields an MAE of 5.238 kW, indicating that physical features provide domain-specific prior constraints that substantially reduce prediction bias. The removal of the fusion module (w/o Fusion) causes the MAE to increase significantly to 6.654 kW, highlighting the essential role of the fusion mechanism in integrating multi-source information from both data-driven and physics-based components.

Eliminating the BiLSTM layer (w/o BiLSTM) results in an MAE of 4.523 kW, while removing the multi-scale CNN module (w/o Multi-Scale CNN) yields an MAE of 4.156 kW. These results confirm that both the sequential modeling capability of BiLSTM and the multi-scale feature extraction capability of CNN are integral to improving prediction accuracy. Excluding the physics-informed loss term (w/o Physics Loss) produces an MAE of 3.892 kW, whereas the full model achieves the lowest MAE of 3.571 kW.

These findings demonstrate that the progressive incorporation of each component consistently reduces

the MAE, confirming the synergistic contributions of temporal features, physical features, the fusion module, BiLSTM, Multi-scale CNN, and physics-informed loss to prediction accuracy. In summary, the ablation experiments comprehensively validate the effectiveness of the proposed integrated fusion framework. The synergistic interaction among functional modules ensures both high prediction accuracy and model robustness.

4.3 Comparative Experimental Analysis

To comprehensively evaluate the performance of the proposed method, we compare it with several mainstream forecasting approaches, including traditional statistical methods (moving average, SARIMA), machine learning methods (SVR, random forest), deep learning methods (LSTM, GRU, CNN-LSTM), and the latest Transformer architectures (Informer, Autoformer, PINN-Transformer). All experiments were conducted on the same dataset and hardware environment, employing 15-minute intervals for forecasting with a 3-day prediction window (288 time steps). Evaluation metrics encompassed Mean Absolute Error (MAE), Root Mean Square Error (RMSE), Mean Absolute Percentage Error (MAPE), Coefficient of Determination (R^2), and model training time to comprehensively assess each method's forecasting accuracy and computational efficiency.

The comparison results with 10 benchmark methods are shown in Table 7. The proposed method significantly outperforms the comparison methods across all metrics.

Table 7: Comparative Experimental Results

Methods	**MAE**	**RMSE**	**MAPE**	**R^2**	**Time (min)**
Proposed	**3.571**	**6.384**	**22.62**	**0.9781**	21.52
CNN-LSTM	5.54	7.64	27.60	0.9684	35.7
GRU	5.69	7.82	28.35	0.9670	28.4
LSTM	7.24	9.97	36.05	0.9463	31.2
Random Forest	10.07	18.12	43.28	0.8225	8.5
SVR	16.15	25.36	66.42	0.6522	15.3
SARIMA	17.94	29.48	72.15	0.5303	2.1
MA-24	18.60	30.25	74.82	0.5056	<0.1
Informer	6.728	10.100	32.95	0.9448	27.30
PINN-Transformer	7.321	8.526	25.93	0.9606	30.50
Autoformer	10.787	12.332	33.55	0.9176	33.60

As shown in Fig. 9 and Table 7, the proposed method achieves the best performance across all evaluation metrics. The proposed method has a mean absolute error (MAE) of 3.571 kW, which is 35.5% lower than the second-best method (CNN-LSTM) and 50.7% lower than the standard LSTM method. The RMSE is 6.384 kW, which is lower than the RMSE of other deep learning methods. The MAPE is 22.62%, indicating a clear advantage in relative error control. The R^2 value reaches 0.9781, demonstrating the model's strong ability to capture load variations.

From a methodological perspective, deep learning approaches generally outperform traditional statistical and machine learning methods. SARIMA and moving-average methods yield MAEs of 17.94 kW and 18.60 kW, respectively, indicating substantial errors and failing to capture nonlinear load characteristics. Support vector regression (SVR) and random forests perform marginally better than statistical methods but still exhibit considerable errors, with random forests achieving an MAE of 10.07 kW and SVR achieving 16.15 kW.

Among deep learning approaches, recurrent neural network-based methods demonstrate moderate performance, with GRU achieving an MAE of 5.69 kW and LSTM achieving 7.24 kW. CNN-LSTM, which combines convolutional and recurrent structures, achieves improved performance with an MAE of 5.54 kW. Among Transformer-based architectures, Informer achieves an MAE of 6.728 kW, PINN-Transformer achieves 7.321 kW, and Autoformer achieves 10.787 kW; none of these methods surpasses the proposed approach.

Furthermore, the proposed method demonstrates notable computational efficiency with a training time of 21.52 minutes, which is the shortest among all deep learning approaches. In comparison, Informer requires 27.30 minutes, GRU requires 28.4 minutes, and CNN-LSTM requires 35.7 minutes. This indicates that the proposed method achieves superior prediction accuracy without incurring additional computational overhead.

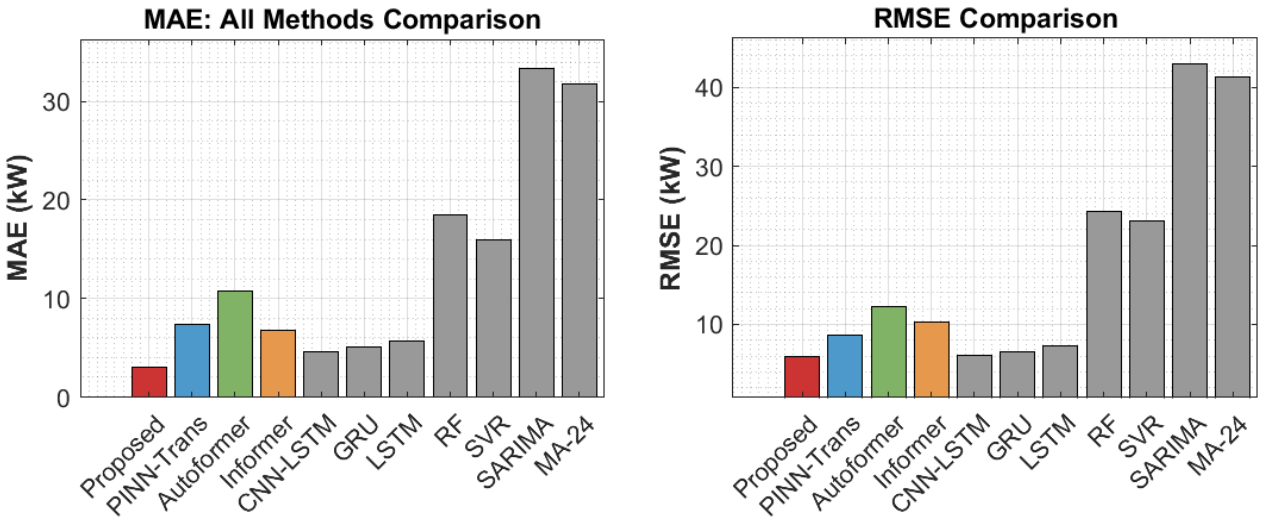


Figure 9: Comparison of MAE and RMSE across methods

Fig. 10 compares the prediction curves of the proposed method with three representative Transformer-based approaches over a 3-day prediction horizon. To clearly illustrate prediction details, four top-performing deep learning methods are visualized: the proposed method, PINN-Transformer, Autoformer, and Informer. Traditional and machine learning methods, which exhibit substantial prediction errors (MAE $>$ 10 kW), show pronounced deviations from actual values, rendering detailed analysis impractical; therefore, they are excluded from the figure. Other deep learning methods (LSTM, GRU, and CNN-LSTM) exhibit prediction curve patterns similar to those of the selected methods and are omitted to avoid visual clutter.

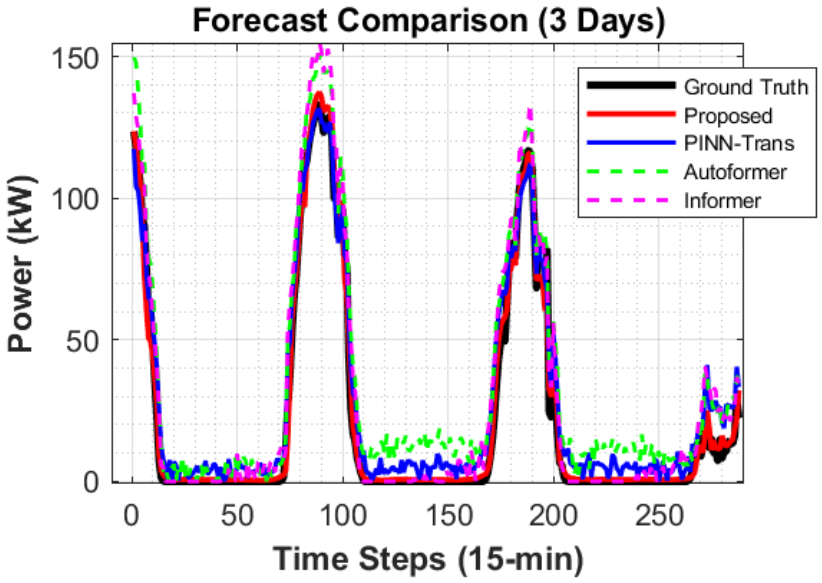


Figure 10: Prediction Comparison Curve

As shown in Fig. 10, the proposed method (red solid line) closely matches the ground truth (black solid line), accurately capturing both peak-to-trough variations and daily cycle characteristics. During peak load periods (approximately 150 kW), the prediction curve of the proposed method nearly overlaps with the ground truth. In contrast, PINN-Transformer (blue solid line) and Autoformer (green dashed line) exhibit notable underestimation at peak values, with prediction errors reaching 15–25 kW. During low-load periods (approximately 10–30 kW), the proposed method also demonstrates excellent performance,

with predictions remaining stable near the ground truth. In contrast, Informer (pink dashed line) exhibits noticeable fluctuations during low-load periods, with deviations exceeding 10 kW at certain time steps, indicating poor stability. Notably, during the load transition phase from time steps 250 to 288, the proposed method demonstrates strong tracking capability and rapid response to both rising and falling load trends. Other methods exhibit noticeable lag, with a delay of approximately 2–3 time steps (30–45 minutes), which could result in delayed decision-making in practical power dispatch scenarios.

4.4 Performance Under Different Weather Conditions

Table 8: Performance Under Different Weather Conditions

Weather	Samples	MAE (kW)	RMSE (kW)	R^2
Sunny	2,190 (41.8%)	2.34	4.12	0.9856
Cloudy	1,682 (32.2%)	3.78	6.45	0.9742
Overcast	1,029 (19.7%)	5.12	8.67	0.9623
Rainy	324 (6.2%)	6.89	11.23	0.9478
Overall	**5,225**	**3.57**	**6.38**	**0.9781**

Table 8 indicates that the proposed model exhibits significant variations in predictive performance across different weather scenarios. Under clear-sky conditions, the model demonstrates the best performance, achieving a coefficient of determination (R^2) of 0.9856, with MAE and RMSE as low as 2.34 kW and 4.12 kW, respectively. This reflects strong adaptability to stable illumination conditions. As weather complexity increases (cloudy → overcast → rainy), the model's prediction errors progressively rise. Under rainy conditions, MAE and RMSE increase to 6.89 kW and 11.23 kW, respectively, while R^2 decreases to 0.9478. Although prediction accuracy exhibits some degradation under complex weather scenarios with significant illumination fluctuations, the overall R^2 remains at a high level of 0.9781. This validates the robustness and practical applicability of the proposed method under diverse weather conditions, demonstrating its capability to meet the requirements of real-world PV power forecasting.

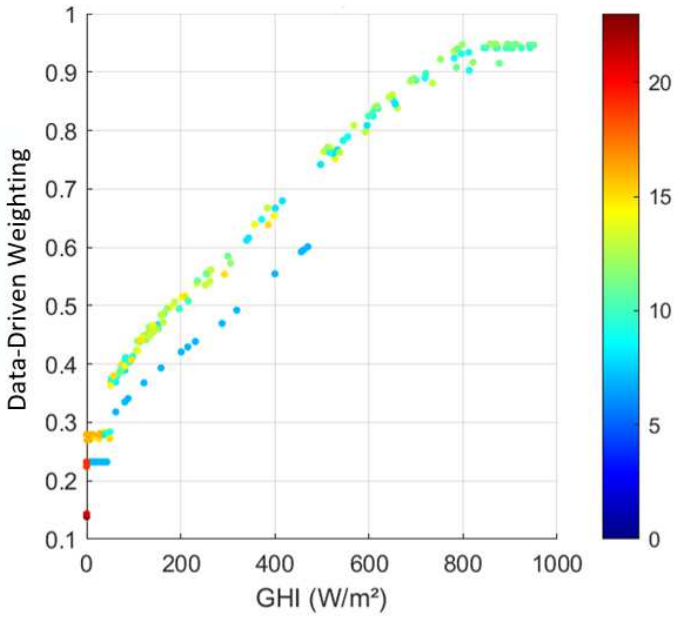


Figure 11: Dynamic Weight Allocation

4.5 Dynamic Weight Fusion Mechanism

The dynamic weight allocation mechanism is illustrated in Fig. 11. Experimental results indicate that the weights assigned to the data-driven model are positively correlated with GHI. Under high irradiance conditions (GHI $>$ 600 W/m^2), the data-driven weights reach 0.75–0.98; under moderate irradiance conditions (200 $<$ GHI $<$ 600 W/m^2), weights range from 0.35 to 0.75; under low irradiance conditions (GHI $<$ 200 W/m^2), weights decrease to 0.15–0.45.

This mechanism effectively leverages the complementary strengths of both models: the data-driven model excels under stable high-irradiance conditions due to abundant training samples and clear patterns, while the physical model provides reliable constraints under complex low-irradiance conditions where data-driven predictions may be less stable. The smooth transition of weights across different irradiance levels demonstrates the adaptive capability of the proposed fusion strategy.

Furthermore, the physical constraints embedded in the model ensure that all predictions remain within physically reasonable bounds, preventing non-physical results such as negative power values or outputs exceeding theoretical limits. These results validate the effectiveness of the physics-informed neural network approach in photovoltaic power forecasting.

5 Discussion

The experimental results demonstrate that the proposed dual-driven framework achieves improvements in prediction accuracy, stability, and interpretability. The physical model ensures consistency with fundamental principles, while the dynamic weighting mechanism enhances adaptability by leveraging data-driven precision under clear-sky conditions and physical stability during weather fluctuations.

Several limitations warrant acknowledgment. The model depends on high-quality historical data and may deteriorate with data gaps or noise. Hyperparameter selection requires adaptive optimization for different configurations, and the deep network component may constrain real-time efficiency. Additionally, the current implementation focuses on single-site modeling without capturing spatial correlations.

Future research may explore graph neural networks for multi-site prediction, uncertainty quantification for probabilistic forecasts, online learning for improved adaptability, and lightweight architectures for reduced deployment costs.

6 Conclusion

This paper proposes a dual-driven photovoltaic power prediction framework that integrates physical prior knowledge with deep learning methodologies. The physical model ensures prediction consistency with fundamental principles, while the data-driven component captures complex nonlinear patterns. A dynamic weighting mechanism enables adaptive fusion based on environmental conditions.

Experimental validation on real-world datasets confirms that the proposed framework outperforms conventional single-paradigm approaches in both accuracy and reliability under variable meteorological conditions. The framework achieves a balance between interpretability and predictive performance, providing a promising direction for photovoltaic power forecasting research.

Acknowledgments

This research was supported by the National Natural Science Foundation of China (Project No. U24B6004).

References

[1] Kalair A, Abas N, Saleem MS, Kalair AR, Khan N (2020) Role of energy storage systems in energy transition from fossil fuels to renewables. Energy Storage 3(1):e135. https://doi.org/10.1002/est2.135

[2] Pan C, Liu Y, Oh Y et al (2024) Short-term photovoltaic power forecasting using PV data and sky images in an auto cross modal correlation attention multimodal framework. Energies 17(24):6378. https://doi.org/10.3390/en17246378

[3] International Energy Agency (2025) Global Energy Review 2025. IEA Publications, Paris. https://www.iea.org/reports/global-energy-review-2025

[4] Du XD, Zhao JL, Liu KY et al (2022) Digital twin early warning method study for overload risk of distribution network with a high proportion of photovoltaic access. Power Syst Prot Control 50(9):136–144

[5] Cao YS, Liu G, Luo DH, Bavirisetti DP, Xiao G (2023) Multi-timescale photovoltaic power forecasting using an improved Stacking ensemble algorithm based LSTM-Informer model. Energy 283:128669. https://doi.org/10.1016/j.energy.2023.128669

[6] Ahmed R, Sreeram V, Mishra Y, Arif M (2020) A review and evaluation of the state-of-the-art in PV solar power forecasting: techniques and optimization. Renew Sustain Energy Rev 124:109792. https://doi.org/10.1016/j.rser.2020.109792

[7] Corti F, Laudani A, Lozito GM, Palermo M, Quercio M, Pattini F, Rampino S (2023) Dynamic analysis of a supercapacitor DC-link in photovoltaic conversion applications. Energies 16(16):5864. https://doi.org/10.3390/en16165864

[8] Jia S, Gao ZW (2025) Extended modeling for wind turbines with application to hybrid renewable energy systems. Eng Sci Technol Int J 70:102168. https://doi.org/10.1016/j.jestch.2025.102168

[9] Asghar R, Fulginei FR, Quercio M, Mahrouch A (2024) Artificial neural networks for photovoltaic power forecasting: a review of five promising models. IEEE Access 12:90461–90485. https://doi.org/10.1109/ACCESS.2024.3420693

[10] Meenal R, Binu D, Ramya KC, Michael PA, Kumar KV, Rajasekaran E, Sangeetha B (2022) Weather forecasting for renewable energy system: a review. Arch Comput Methods Eng 29(5):2875–2891. https://doi.org/10.1007/s11831-021-09695-3

[11] Wang F, Lu X, Mei S, Su Y, Zhen Z, Zou Z et al (2022) A satellite image data based ultra-short-term solar PV power forecasting method considering cloud information from neighboring plant. Energy 238:121946. https://doi.org/10.1016/j.energy.2021.121946

[12] Korkmaz D (2021) SolarNet: a hybrid reliable model based on convolutional neural network and variational mode decomposition for hourly photovoltaic power forecasting. Appl Energy 300:117410. https://doi.org/10.1016/j.apenergy.2021.117410

[13] Liu L, Zhao Y, Chang D, Xie J, Ma Z, Sun Q, Yin H, Wennersten R (2018) Prediction of short-term PV power output and uncertainty analysis. Appl Energy 228:700–711. https://doi.org/10.1016/j.apenergy.2018.06.112

[14] Gaboitaolelwe J, Zungeru AM, Yahya A, Lebekwe CK, Vinod DN, Salau AO (2023) Machine learning based solar photovoltaic power forecasting: a review and comparison. IEEE Access 11:40820–40845. https://doi.org/10.1109/ACCESS.2023.3270041

[15] Markovics D, Mayer MJ (2022) Comparison of machine learning methods for photovoltaic power forecasting based on numerical weather prediction. Renew Sustain Energy Rev 161:112364. https://doi.org/10.1016/j.rser.2022.112364

[16] Tang YG, Yang K, Zhang SJ, Zhang Z (2024) Photovoltaic power forecasting: a dual-attention gated recurrent unit framework incorporating weather clustering and transfer learning strategy. Eng Appl Artif Intell 130:107691. https://doi.org/10.1016/j.engappai.2023.107691

[17] Ghimire S, Nguyen-Huy T, Prasad R et al (2023) Hybrid convolutional neural network-multilayer perceptron model for solar radiation prediction. Cogn Comput 15:645–671. https://doi.org/10.1007/s12559-022-10070-y

[18] Chen DY, Shi XD, Jiang MK et al (2025) Selecting effective NWP integration approaches for PV power forecasting with deep learning. Sol Energy 301:113939. https://doi.org/10.1016/j.solener.2025.113939

[19] Liu WC, Mao ZZ (2024) Short-term photovoltaic power forecasting with feature extraction and attention mechanisms. Renew Energy 226:120437. https://doi.org/10.1016/j.renene.2024.120437

[20] Palermo M, Forconi F, Belloni E, Quercio M, Lozito GM, Fulginei FR (2023) Optimization of a feedforward neural network's architecture for an HVAC system problem. In: Proc 3rd Int Conf Electr Comput Commun Mechatronics Eng (ICECCME), Tenerife, Spain, pp 1–6. https://doi.org/10.1109/ICECCME57830.2023.10252568

[21] Belloni E, Grasso F, Lozito GM, Poli D, Fulginei FR, Talluri G (2023) Neural-assisted HVACs optimal scheduling for renewable energy communities. Energy Build 301:113658. https://doi.org/10.1016/j.enbuild.2023.113658

[22] Tao KJ, Zhao JH, Tao Y et al (2024) Operational day-ahead photovoltaic power forecasting based on transformer variant. Appl Energy 373:123825. https://doi.org/10.1016/j.apenergy.2024.123825

[23] Wang L, Mao MX, Xie J, Liao Z et al (2023) Accurate solar PV power prediction interval method based on frequency-domain decomposition and LSTM model. Energy 262:125592. https://doi.org/10.1016/j.energy.2022.125592

[24] Piantadosi G, Dutto S, Galli A et al (2024) Photovoltaic power forecasting: a Transformer based framework. Energy AI 18:100444. https://doi.org/10.1016/j.egyai.2024.100444

[25] Yuan L, Wang XT, Sun Y et al (2025) Multistep photovoltaic power forecasting based on multi-timescale fluctuation aggregation attention mechanism and contrastive learning. Int J Electr Power Energy Syst 164:110389. https://doi.org/10.1016/j.ijepes.2024.110389

[26] Xiang X, Li X, Zhang Y et al (2024) A short-term forecasting method for photovoltaic power generation based on the TCN-ECANet-GRU hybrid model. Sci Rep 14:6744. https://doi.org/10.1038/s41598-024-56751-6

[27] Zhou H, Wang J, Ouyang F, Cui C, Li X (2023) A two-stage method for ultra-short-term PV power forecasting based on data-driven. IEEE Access 11:41175–41189. https://doi.org/10.1109/ACCESS.2023.3267515

[28] Ren X, Zhang F, Yan J, Liu Y (2024) A novel convolutional neural net architecture based on incorporating meteorological variable inputs into ultra-short-term photovoltaic power forecasting. Sustainability 16:2786. https://doi.org/10.3390/su16072786

[29] Thaker J, Höller R (2024) Hybrid model for intra-day probabilistic PV power forecast. Renew Energy 232:121057. https://doi.org/10.1016/j.renene.2024.121057

[30] Chen R, Gao S, Zhao Y, Li D, Lin S (2024) A hybrid model based on the photovoltaic conversion model and artificial neural network model for short-term photovoltaic power forecasting. Front Energy Res 12:1446422. https://doi.org/10.3389/fenrg.2024.1446422

[31] Liao R, Liu Y, Xu X, Li Z, Chen Y, Shen X, Liu J (2024) Enhanced photovoltaic power generation forecasting for newly-built plants via physics-infused transfer learning with domain adversarial neural networks. Energy Convers Manag 322:119114. https://doi.org/10.1016/j.enconman.2024.119114