

Article

Diffusion Model with Multi-Source Data for Day-Ahead Renewable Energy Scenario Generation

Lin Chen ¹ , Xinran Liu ^{1,*}, Quanqi Chen ¹, Guinan Ye ¹, Wen Liu ¹ and Xiaotong Dai ²

¹ Power Dispatching and Control Center, Guangxi Power Grid Co., Ltd., Nanning 530201, China; chenlin96@foxmail.com (L.C.); quanqi.chen@163.com (Q.C.); yeguinan@yeah.net (G.Y.); liuwen928@sina.com (W.L.)

² College of Electrical Engineering, Zhejiang University, Hangzhou 310027, China; 22510051@zju.edu.cn

* Correspondence: liuxr@gx.csg.cn

Abstract

High-quality renewable energy (RE) scenario generation is essential for secure, reliable, and economic power system operation under high RE penetration. To address the limitations of existing methods in preserving scenario fidelity, diversity, spatiotemporal dependence, and engineering consistency, this paper proposes a three-stage framework that combines a variational autoencoder (VAE) with a conditional latent diffusion model (CLDM), hereafter referred to as VAE-CLDM, for day-ahead renewable energy scenario generation. First, multi-source features are constructed by integrating renewable power outputs, meteorological variables, temporal lag information, and spatial correlation characteristics among wind farms and photovoltaic stations. Then, VAE compresses the high-dimensional features into a compact latent space while retaining key statistical and spatiotemporal information. Based on this latent representation, a CLDM generates realistic scenarios by progressively denoising random noise under meteorological conditions. A spatiotemporal feature modeling strategy is incorporated to better represent temporal fluctuations and inter-site correlations, while a diversity regulation factor is selected on the validation set to balance scenario fidelity and tail-event coverage. Finally, the generated scenarios are reconstructed into the physical space and checked using output-bound and ramp-consistency correction to improve their practical usability. Results on the open dataset released by the Chinese State Grid Renewable Energy Generation Forecasting Competition show that the proposed framework achieves the lowest root mean square error (RMSE), mean absolute error (MAE), and maximum mean discrepancy (MMD) among the tested models on the training-statistics-standardized renewable-power benchmark, with RMSE of 0.3013 ± 0.0022 , MAE of 0.3636 ± 0.0010 , and MMD of 0.04269 ± 0.00040 . After post-correction, lower- and upper-bound violations are reduced to 0.00%, and the ramp-violation rate is reduced to 0.08%, indicating that the proposed VAE-CLDM can provide useful scenario inputs for day-ahead dispatch and risk assessment in renewable-dominated power systems.

Keywords: renewable energy; scenario generation; variational autoencoder; conditional latent diffusion model; day-ahead power dispatch; spatiotemporal correlation



Academic Editors: Wei Dai, Lun Yang, Wei Lin and Qian Jiang

Received: 29 May 2026

Revised: 9 July 2026

Accepted: 15 July 2026

Published: 23 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Driven by energy security concerns and carbon neutrality targets, many power systems are increasingly integrating renewable energy (RE) resources such as wind power and photovoltaic (PV) generation. This transition is supported by smart grid technologies

that improve the adaptability of power systems to intermittent renewable generation [1]. However, RE output is inherently random, volatile, and intermittent because it is strongly affected by meteorological and geographical conditions. These characteristics pose substantial challenges to secure, stable, and economic power system operation and may limit further increases in renewable penetration [2]. Accordingly, high-quality RE scenario generation has become an important basis for power system planning, operation, and dispatch.

Ideally, generated scenarios should represent the statistical characteristics, temporal evolution, spatiotemporal dependence, and engineering constraints of real-world RE behavior under diverse weather conditions and operating regimes. Recent reviews show that scenario generation methods provide essential inputs for robust optimization, stochastic optimization, and distributionally robust optimization in renewable-rich power systems [3]. Therefore, generated scenarios should not only approximate historical output distributions, but also support downstream dispatch and risk assessment tasks under weather-driven uncertainty.

Despite extensive research on RE scenario generation, three limitations remain prominent in the existing literature:

(1) Insufficient integration of climate and temporal features: Weather-related forecasting and scenario modeling studies have improved the representation of meteorological variability, but many of them still focus on a specific resource, a single forecasting task, or a partial temporal feature set, rather than constructing a unified multi-source representation for day-ahead multi-site RE scenario generation. Cui et al. [4] developed a stochastic scenario generation method for wind-power ramp-event forecasting, showing that rapid wind-power fluctuations can be better represented through scenario-based modeling. Wen et al. [5] proposed a continuous and distribution-free probabilistic wind-power forecasting method based on conditional normalizing flows, which highlights the importance of modeling the conditional distribution of renewable generation under meteorological uncertainty. Dumas et al. [6] further introduced normalizing flows for probabilistic energy forecasting in power systems, demonstrating the potential of deep generative probabilistic models for uncertainty representation. A conditional recurrent generative adversarial network (GAN)-based implicit generative model is introduced in [7] to capture temporal dependencies and cross-correlations among renewable generation and load variables, but GAN-based training may suffer from instability and does not explicitly exploit low-dimensional latent diffusion for improving computational efficiency. To improve the weather-output consistency of wind–solar–load scenarios during high-impact events, Yi et al. [8] condition a GAN on labeled extreme weather events and meteorological features, but the method depends on event labels and does not fully address the joint problem of latent representation, spatial-correlation preservation, and dispatch-oriented evaluation. Recent machine-learning studies have also improved renewable forecasting accuracy, such as the wind-speed forecasting framework integrating a Newton–Raphson-based optimizer (NRBO), Transformer, extreme gradient boosting (XGBoost), and adaptive denoising, collectively termed NRBO-TXAD, in [9]; however, these forecasting-oriented methods do not directly solve the problem of generating diverse day-ahead renewable scenarios with spatial dependence. Therefore, although these studies provide useful meteorological and temporal modeling tools, they generally do not provide a unified representation that simultaneously integrates meteorological drivers, temporal lag information, and spatial correlation features for multi-site renewable scenario generation.

(2) Limited targeted scenario simplification and high computational complexity: Scenario reduction and selection methods can improve computational tractability, but they mainly compress or select representative samples from existing scenario sets and may lose high-impact renewable fluctuation patterns if operational relevance and spatiotemporal

dependence are not explicitly considered. Lee and Baldick [10] developed a generalized dynamic factor model for load and wind-power scenario generation, which captures temporal dependence and cross-correlation among uncertain variables. Hu et al. [11] proposed a scenario generation framework based on the knowledge of existing wind power plants, showing that transferable information from data-rich wind farms can support scenario construction when historical data are limited. Submodular optimization has also been used for scenario reduction [12], while distance-based clustering and scenario-selection methods are commonly adopted to reduce computational burden [13,14]. Recent reviews further indicate that renewable scenario generation still needs to balance statistical fidelity, diversity, temporal dependence, and computational efficiency [3]. Although these methods improve tractability, representative samples are still constrained by the original scenario pool, and rare ramping events or high-impact low-output patterns may be removed during reduction. Therefore, scenario reduction is complementary to, rather than a substitute for, a generative model that directly learns compact weather-conditioned renewable distributions.

(3) Limited accuracy and diagnostic validation of conventional scenario generation methods: Existing statistical and deep generative models have improved renewable scenario modeling, but many of them still face limitations in simultaneously preserving statistical fidelity, scenario diversity, inter-site dependence, computational efficiency, and post-reconstruction engineering consistency. Traditional statistical methods, such as wind-speed scenario generation based on principal components and R-vine copulas [15], can preserve part of the dependence structure but may be limited in representing nonlinear and multimodal renewable output distributions. GAN-based methods provide a flexible way to generate renewable scenarios without explicit parametric assumptions; representative examples include model-free renewable scenario generation [16], conditional style-based GANs [17], controllable and interpretable GAN-based generation [18], weather-label-driven generation using a combination weighting method and conditional generative adversarial network (CWM-CGAN) [19], SolarGAN for multivariate solar-data reconstruction [20], and cross-modal GANs for renewable scenario generation [21]. These models improve conditional generation capability, but GAN-based training may still suffer from instability and mode collapse. Variational autoencoder (VAE) methods and variational autoencoder-generative adversarial network (VAEGAN) methods provide latent representations for renewable forecasting and scenario generation [22,23], while conditional joint scenario generation emphasizes the importance of dependence modeling for day-ahead operational studies [24]. More recently, diffusion-based generative models have attracted attention because of their strong capability in modeling complex and non-stationary renewable fluctuations. Xu et al. [25] proposed a two-stage generative architecture based on temporal scenario representation and diffusion models, and Li et al. [26] developed a multi-resolution denoising diffusion probabilistic model (DDPM) to capture the multi-scale volatility of wind and PV outputs. However, existing diffusion-based methods still mainly focus on temporal representation or multi-scale fluctuation modeling. The joint incorporation of meteorological conditions, temporal lag features, and inter-site spatial correlations into a compact latent generation space remains insufficient, and correlation preservation is often evaluated through visualization rather than quantitative statistical tests.

In addition to statistical scenario quality, dispatch-oriented applications require generated scenarios to be useful for operational decision-making. Wu et al. [27] studied stochastic multi-timescale power system operation with variable wind generation, showing that renewable uncertainty should be represented across different operational timescales. Modarresi et al. [28] proposed a scenario-based economic dispatch method with tunable risk levels for high-renewable power systems, indicating that scenario quality and risk coverage

can directly influence dispatch decisions. Recent microgrid energy-management studies also indicate that renewable uncertainty, storage degradation, and multi-objective operating costs should be considered when evaluating the practical value of uncertainty modeling methods [29–31]. Therefore, a scenario generation framework that combines multi-source feature construction, compact latent-space generation, quantitative correlation-preservation validation, and dispatch-oriented diagnosis remains necessary.

To address these challenges, this paper presents a three-stage framework that integrates meteorological-temporal feature construction, VAE-based latent dimensionality reduction, and a conditional latent diffusion model (CLDM) for day-ahead renewable energy scenario generation; this integrated framework is hereafter referred to as VAE-CLDM. The framework first constructs a structured 182-dimensional input by combining renewable power outputs, meteorological variables, temporal lag information, and intra-type spatial correlation features. The VAE then compresses the high-dimensional input into a compact latent space, which reduces the learning burden of direct high-dimensional diffusion modeling while retaining key statistical and spatiotemporal information. Based on this latent representation, the CLDM generates renewable scenarios through meteorology-conditioned denoising, and the VAE decoder reconstructs the generated latent samples into the original physical feature space. Finally, engineering consistency checks are applied to improve practical usability by reducing obvious output-bound and ramping violations.

Compared with standalone GAN-, VAE-, DDPM-, or CLDM-based generators, the proposed VAE-CLDM framework has three main advantages. First, VAE-based latent compression reduces the learning burden of high-dimensional multi-source data while avoiding direct diffusion in the original 182-dimensional space. Second, CLDM performs conditional generation in a compact latent space, which improves training stability, scenario diversity, and denoising efficiency under meteorological conditions. Third, the generated scenarios are evaluated not only by statistical distance metrics, but also by spatial-correlation preservation, diversity-factor sensitivity, constraint-consistency diagnosis, and dispatch-oriented performance, which provides a more comprehensive assessment of their engineering usefulness.

The main contributions of this work are summarized as follows:

- A multi-source conditional representation for day-ahead RE scenario generation is developed. Renewable power outputs, meteorological variables, temporal lag features, and inter-site spatial correlation characteristics are jointly organized into a structured high-dimensional input. This representation provides richer conditional information for modeling weather-driven temporal fluctuations and spatial dependence among multiple wind and PV stations.
- A VAE-CLDM latent scenario generation framework is proposed. Instead of applying diffusion directly in the original 182-dimensional multi-source feature space, the VAE first learns a compact latent representation, and the CLDM performs conditional denoising in this latent space. This design reduces the learning burden of high-dimensional diffusion modeling and provides a more efficient generation process, while the latent-space alignment and reconstruction accuracy are further examined to verify the reliability of decoder-based scenario reconstruction.
- A dispatch-oriented and diagnostic evaluation framework is introduced. In addition to statistical fidelity metrics, the generated scenarios are evaluated using diversity indicators, correlation-matrix errors, statistical tests for spatial-correlation preservation, constraint-consistency diagnostics before and after post-correction, and dispatch-oriented metrics on a distribution network case. Comparative, ablation, sensitivity, and robustness analyses are conducted to examine whether the generated scenarios

are not only statistically similar to historical data, but also useful for operational decision-making.

The remainder of this paper is organized as follows. Section 2 presents an overview of the proposed framework. Section 3 describes the multi-source feature construction stage, the VAE-based latent representation stage, and the CLDM-based latent scenario generation and decoder-reconstruction stage. Section 4 evaluates the quality of the generated renewable energy scenarios and examines their applicability to power system dispatch through case studies and comparisons with benchmark models. Section 5 concludes the paper.

2. Overview of the Proposed Framework

The overall structure of the proposed VAE-CLDM method is shown in Figure 1. The framework consists of three sequential stages: multi-source data construction, VAE-based dimensionality reduction, and CLDM-based scenario generation with reconstruction. The first stage constructs structured multi-source inputs, the second compresses these inputs into a compact latent space, and the third generates and reconstructs scenarios for day-ahead dispatch applications.




Figure 1. Overall Framework of the Proposed Renewable Energy Scenario Generation Method.

In Stage I (Multi-source Data Preprocessing and Feature Construction), raw data, including renewable power output, wind speed, solar irradiance, temperature, and temporal information are processed. Z-score standardization is applied to place heterogeneous features on a comparable numerical scale, while physical units are restored through inverse standardization after decoder reconstruction. Stations are divided into wind and photovoltaic groups to compute intra-group spatial correlations. Spatial features are unified to eight dimensions, and four representative temporal lags (1, 6, 12, 24 h) are extracted. The standardized power, temporal features, spatial features, meteorological variables, and temporal encodings are concatenated to form a 182-dimensional input vector. Hard feasibility is mainly examined after decoder reconstruction through output-bound and ramp-consistency checks.

In Stage II (Latent Space Feature Extraction via VAE), the VAE encoder processes the 182-dimensional input from Stage I. It compresses high-dimensional features into a low-dimensional latent space, reducing redundancy while preserving statistical and spatiotemporal correlations. This step reduces the learning burden of the subsequent generative model and improves training efficiency.

In Stage III (Conditional Scenario Generation and Dimensionality Restoration), latent features from Stage II serve as conditional inputs for the CLDM. Guided by key meteorological variables, the CLDM generates latent scenarios through adaptive denoising. The generated latent scenarios are then fed into the pre-trained VAE decoder to reconstruct

high-dimensional physical-space scenarios. After inverse normalization, raw generated scenarios are checked and post-corrected for output-bound and ramp consistency before being used for day-ahead dispatch.

3. Detailed Implementation of the Proposed Three-Stage Framework

This section describes the structure and function of each module in the proposed framework, including the multi-source feature construction module (Stage I), the VAE-based latent representation module (Stage II), and the CLDM-based latent scenario generation and decoder reconstruction module (Stage III).

3.1. Stage I: Multi-Source Data Preprocessing and Feature Construction

Stage I performs data standardization, spatiotemporal feature extraction, and engineering-informed feature construction to build reliable high-dimensional inputs for subsequent modeling. The implementation steps and key formulas are as follows:

(1) Data Pre-standardization via Z-score Transformation

To place variables with different units on a comparable numerical scale, all continuous inputs are standardized using the Z-score transformation. For any input feature x (e.g., wind speed at a station or photovoltaic output), the standardized value x_{std} is calculated as

$$x_{\text{std}} = \frac{x - \mu_x}{\sigma_x} \quad (1)$$

where μ_x and σ_x denote the mean and standard deviation of feature x calculated from the training set only. The same statistics are then applied to the validation and test sets to avoid information leakage. This step scales all features to a comparable numerical range while retaining their correlation structure. Physical units are recovered later through inverse standardization after decoder reconstruction.

(2) Spatial Correlation Feature Construction

Stations are divided into two groups: wind power stations (denoted as set W) and photovoltaic power stations (denoted as set P). For each group, an intra-group output correlation matrix is computed. The original dataset contains six wind farms and eight photovoltaic stations. After data-quality screening, one photovoltaic station with insufficient valid records and long-term zero-output behavior is excluded; therefore, the final case study uses six wind farms and seven photovoltaic stations. The intra-type correlation matrices of the selected wind and photovoltaic groups have sizes of 6×6 and 7×7 , respectively. The correlation coefficient $\rho_{i,j}$ between station i and station j in the same group is

$$\rho_{i,j} = \frac{\text{Cov}(P_i, P_j)}{\sigma_{P_i} \sigma_{P_j}} \quad (2)$$

where P_i and P_j denote the standardized output sequences of stations i and j , $\text{Cov}(P_i, P_j)$ is their covariance, and σ_{P_i} and σ_{P_j} are their standard deviations.

For each station, the corresponding row in the intra-group correlation matrix is extracted as its spatial correlation feature. To keep a fixed station-level tensor format across wind and PV resources, the spatial correlation vector of each station is padded to eight dimensions. Specifically, the wind spatial feature has an initial dimension of six and is represented as $\mathbf{s}_w = [\rho_{w,1}, \rho_{w,2}, \dots, \rho_{w,6}, 0, 0]^T$, while the PV spatial feature has an initial dimension of seven and is represented as $\mathbf{s}_p = [\rho_{p,1}, \rho_{p,2}, \dots, \rho_{p,7}, 0]^T$. The padded zero entries are introduced only for dimensional alignment with the original maximum number of PV sites and are not interpreted as physical correlations. They are excluded from correlation-preservation evaluation and multicollinearity diagnostics.

(3) Temporal Lag Feature Construction

To capture the time-dependent dynamics of RE output, temporal lag features are introduced. For the standardized RE output $P_{i,\text{std}}(t)$ at station i and time t , four lagged features are extracted: lag-1 h, lag-6 h, lag-12 h, and lag-24 h. For the 15 min resolution data used in this study, these lags correspond to 4, 24, 48, and 96 time steps, respectively. In day-ahead generation, these variables are historical observations available before the target generation horizon, and no actual renewable output from the target day is used. The temporal feature vector is constructed as

$$\mathbf{l}_i(t) = [P_{i,\text{std}}(t - \Delta_1), P_{i,\text{std}}(t - \Delta_6), P_{i,\text{std}}(t - \Delta_{12}), P_{i,\text{std}}(t - \Delta_{24})]^\top \quad (3)$$

where $\Delta_1 = 4$, $\Delta_6 = 24$, $\Delta_{12} = 48$, and $\Delta_{24} = 96$ denote the corresponding numbers of 15 min time steps. This step encodes short-term, intra-day, half-day, and daily variability of RE output.

(4) Feature Concatenation to Form High-dimensional Input Vector

The standardized power (one dimension), 4-dimensional temporal features, and 8-dimensional spatial features are concatenated into a unified input vector for each station, yielding a 13-dimensional representation for each renewable energy site:

$$\mathbf{x}_i(t) = [P_{i,\text{std}}(t), \mathbf{l}_i^\top(t), \mathbf{s}_i^\top]^\top \in \mathbb{R}^{13}. \quad (4)$$

For the whole system consisting of 13 grid-connected nodes (6 wind farms and 7 PV stations), the station-level vectors are stacked and concatenated with meteorological variables and temporal encoding variables to form the final 182-dimensional input vector:

$$\mathbf{X}(t, :) = [\mathbf{x}_1^\top(t), \mathbf{x}_2^\top(t), \dots, \mathbf{x}_{13}^\top(t), \mathbf{m}^\top(t), \mathbf{e}^\top(t)] \in \mathbb{R}^{182}. \quad (5)$$

where $\mathbf{m}(t) \in \mathbb{R}^3$ denotes meteorological variables, including wind speed, irradiance, and temperature, and $\mathbf{e}(t) \in \mathbb{R}^{10}$ denotes temporal encoding variables such as hour-of-day and day-of-week. Therefore, the 182 dimensions consist of 13 stations $\times$ 13 dimensions per station (1 standardized power + 4 temporal lag features + 8 spatial features) = 169 dimensions, together with three-dimensional meteorological features and 10-dimensional temporal encoding features.

(5) Engineering Consistency Description for Post-Reconstruction Checking

The 182-dimensional input matrix $\mathbf{X}$ provides the structured multi-source representation that is later compressed by the VAE in Stage II. In the revised framework, Stage I is not treated as an additional independent latent-learning module; instead, it prepares standardized power, meteorological, temporal-lag, and spatial-correlation features for the subsequent VAE encoder. Engineering consistency is mainly evaluated after decoder reconstruction, where output-bound and ramp-consistency checks are applied to the reconstructed renewable power outputs.

For clarity, the engineering consistency checks are described as post-reconstruction diagnostics rather than as complete atmospheric physical laws. The generic violation function can be written as

$$g_c(\mathbf{X}_{\text{real}}(t)) \leq 0, \quad c \in \mathcal{C}, \quad (6)$$

where $\mathcal{C}$ includes output lower-bound, output upper-bound, and ramp-consistency checks. These checks are used to distinguish raw generated scenarios from post-corrected scenarios in the experiments. It should be emphasized that $g_c(\cdot)$ is not included in the VAE or CLDM training loss and is not differentiated through the neural network. It is used only

as a post-reconstruction diagnostic and correction criterion after decoder reconstruction and inverse standardization. Therefore, hard feasibility is not claimed at the Stage I feature-construction stage; it is evaluated separately after VAE decoder reconstruction and inverse standardization.

3.2. Stage II: Latent Space Feature Extraction via VAE

Stage II extracts latent features using a Variational Autoencoder (VAE). The VAE processes the 182-dimensional standardized multi-source input from Stage I and performs targeted dimensionality reduction to reduce high-dimensional redundancy and strong inter-feature coupling while retaining important statistical and correlation structures. The VAE is not intended to remove physically meaningful renewable fluctuations, such as ramps and irradiance intermittency; rather, it mitigates redundant feature coupling and isolated abnormal perturbations. Figure 2 illustrates the encoder–decoder architecture used for latent-space dimensionality reduction and reconstruction. The VAE encoder is designed as a lightweight multi-layer perceptron (MLP) to avoid redundant computation, and the latent variable $\mathbf{Z}$ is assumed to follow a standard Gaussian prior to obtain a compact and structured latent space:

$$p(\mathbf{Z}) = \mathcal{N}(\mathbf{0}, \mathbf{I}) \quad (7)$$




Figure 2. Basic Architecture of VAE for Dimensionality Reduction and Reconstruction.

Because the true posterior distribution $p_\psi(\mathbf{Z} | \mathbf{X}_{\text{std}})$, where $\mathbf{X}_{\text{std}}$ denotes the standardized multi-source input matrix from Stage I, is intractable, the VAE encoder learns an approximate Gaussian posterior $q_\phi(\mathbf{Z} | \mathbf{X}_{\text{std}}) = \mathcal{N}(\boldsymbol{\mu}, \text{diag}(\boldsymbol{\sigma}^2))$. Here, $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$ represent the mean and standard deviation vectors of the approximate posterior, respectively. The reparameterization trick is adopted to address the non-differentiability of random sampling in the latent space, enabling gradient-based training of the encoder:

$$\mathbf{Z} = \boldsymbol{\mu} + \boldsymbol{\sigma} \odot \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \quad (8)$$

Here, $\odot$ denotes element-wise multiplication, and $\boldsymbol{\epsilon}$ is a random noise vector sampled from the standard Gaussian distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$.

The VAE encoder is trained by maximizing the evidence lower bound (ELBO), which balances the reconstruction term, used to recover key information from the input, and the Kullback–Leibler (KL) divergence, used to align the approximate posterior with the prior Gaussian distribution:

$$\mathcal{L}_{\text{VAE}}(\phi, \psi) = \mathbb{E}_{q_\phi(\mathbf{Z} | \mathbf{X}_{\text{std}})} [\log p_\psi(\mathbf{X}_{\text{std}} | \mathbf{Z})] - D_{\text{KL}}(q_\phi(\mathbf{Z} | \mathbf{X}_{\text{std}}) \parallel p(\mathbf{Z})) \quad (9)$$

where ϕ and ψ denote the learnable parameters of the encoder and decoder, respectively.

For Gaussian distributions, the KL divergence term has a closed-form expression:

$$D_{\text{KL}} = -\frac{1}{2N} \sum_{n=1}^N \sum_{k=1}^K \left(1 + \log \sigma_{n,k}^2 - \mu_{n,k}^2 - \sigma_{n,k}^2 \right) \quad (10)$$

where N is the number of time steps in the input dataset, $K \ll 182$ is the latent space dimension, and $\mu_{n,k}$ and $\sigma_{n,k}$ are the k -th components of μ and σ for the n -th time step, respectively. In this study, the latent dimension is set to $K = 8$, selected from the validation-based sensitivity analysis over $K = \{8, 16, 32\}$. This setting provides a compact representation while maintaining reconstruction accuracy and generation quality.

It should be noted that the VAE decoder is trained using latent variables sampled from the VAE posterior, whereas the final scenario generation stage decodes CLDM-generated latent variables. Therefore, latent-space alignment between VAE-encoded samples and CLDM-generated samples is examined before reconstruction using latent mean difference, covariance error, and latent maximum mean discrepancy (MMD), as reported in the simulation section and supplementary diagnostics.

3.3. Stage III: CLDM-Guided Latent Space Low-Dimensional Scenario Generation

Based on the low-dimensional latent variable matrix $\mathbf{Z} \in \mathbb{R}^{N \times K}$ output by the VAE encoder in Stage II, the conditional latent diffusion model (CLDM) learns the distribution of the reduced-dimensional features and generates latent scenarios. To improve spatiotemporal correlation modeling and the coverage of extreme scenarios, two strategies are introduced: a spatiotemporal feature modeling strategy embedded in the diffusion network and a diversity regulation factor α ($\alpha > 0$) for adjusting the noise sampling process. Figure 3 shows the conditional latent diffusion architecture and its meteorological-conditioning pathway.




Figure 3. Basic Architecture of Conditional Latent Diffusion Model. The orange, green, purple, and cyan curves schematically illustrate the time-series states at the initial, intermediate, and terminal diffusion steps, respectively; the colors are used only for visual distinction.

Specifically, the temporal lag features describe short-term, intra-day, half-day, and daily renewable-output evolution; the intra-type spatial-correlation descriptors provide station-level dependence information within the wind and PV groups; and the meteorological variables guide the reverse denoising process toward weather-consistent latent samples. This formulation keeps the implemented CLDM consistent with the MLP-based conditional noise predictor reported in Table 1, while still incorporating the main spatiotemporal information required for multi-site renewable scenario generation.

The diversity regulation factor is introduced into the noise sampling process as $\epsilon' = \alpha\epsilon$, where ϵ is standard Gaussian noise. The value of α is not treated as an arbitrary free parameter. Instead, candidate values $\alpha \in \{0.6, 0.8, 1.0, 1.2, 1.5\}$ are selected on the validation set by jointly considering MMD, scenario spread, tail-event coverage, correlation error, and constraint-violation rate. Increasing α expands the noise variation range and improves the coverage of tail-output scenarios, whereas decreasing α narrows the noise range and improves the fidelity of conventional scenarios. Excessively large values may introduce unrealistic fluctuations, so the selected α balances statistical fidelity and scenario diversity.

Table 1. Main architecture and training hyperparameters.

Item	Setting
Input dimension	182
Retained renewable sites	13 sites: 6 wind farms and 7 PV stations
Station-level feature dimension	13: 1 power + 4 lag features + 8 spatial features
Meteorological condition dimension	3: wind speed, irradiance, temperature
Temporal encoding dimension	10
Latent dimension	$K = 8$
VAE encoder/decoder hidden layers	[128, 64]/[64, 128]
VAE latent prior	$\mathcal{N}(\mathbf{0}, \mathbf{I})$
CLDM denoising network	MLP-based conditional noise predictor
Training diffusion steps T_{train}	1000, linear $\beta_t \in [10^{-4}, 2 \times 10^{-2}]$
Reverse-sampling steps S_{sample}	1000
Optimizer, learning rate, batch size	Adam, 1×10^{-3} , 128
Training epochs and random seeds	200 epochs, 5 independent seeds
Diversity factor candidates	{0.6, 0.8, 1.0, 1.2, 1.5}
Selected diversity factor	$\alpha = 1.0$

Conditioned on core meteorological variables, including wind speed, solar irradiance, and ambient temperature, the CLDM performs iterative denoising in the latent space to generate weather-conditioned latent scenario samples $\mathbf{Z}_{\text{scene}}$. The pre-trained VAE decoder then reconstructs these latent scenarios into the original high-dimensional space through the nonlinear mapping:

$$\mathbf{X}_{\text{rec}} = \text{Decoder}(\mathbf{Z}_{\text{scene}}; \psi) \quad (11)$$

where ψ denotes the learnable parameters of the decoder, which are jointly optimized with the encoder during Stage II training to ensure consistent feature mapping between the latent and original data spaces.

To restore physical values, inverse Z-score standardization is performed on the reconstructed data:

$$\mathbf{X}_{\text{real}} = \mathbf{X}_{\text{rec}} \odot \sigma_{\text{train}} + \mu_{\text{train}} \quad (12)$$

where μ_{train} and σ_{train} are the feature-wise mean and standard-deviation vectors estimated from the training set, and $\odot$ denotes element-wise multiplication with feature-wise broadcasting across samples. Finally, an engineering post-correction mechanism checks invalid scenarios, such as negative output and out-of-limit operating parameters. The raw generated scenarios and the post-corrected scenarios are evaluated separately. Therefore, a zero or near-zero violation rate after correction should be interpreted as the feasibility of the corrected outputs rather than as evidence that the raw CLDM samples always satisfy all operating constraints. The qualified scenarios are then applied to day-ahead dispatch tasks, completing the process of “latent feature learning-scenario generation-data space recovery”.

In summary, the VAE improves the compactness and discriminability of latent features, the CLDM learns the weather-conditioned latent distribution through denoising, and the spatiotemporal feature modeling strategy incorporates temporal lag information, intra-type spatial-correlation descriptors, and meteorological conditions. The diversity regulation factor provides flexible control over scenario diversity and tail-output coverage. These components jointly improve the proposed renewable energy scenario generation framework and support the generation of correlated and diverse scenarios with improved engineering consistency after post-correction for power system dispatch.

4. Simulation Results and Discussion

This section presents the experimental settings, raw-data diagnostics, VAE latent-representation performance, CLDM scenario-generation performance, constraint-consistency

evaluation, robustness diagnostics, and dispatch-oriented comparison. All experiments were implemented in Python 3.9 with PyTorch 1.12.0 on a workstation equipped with a multicore central processing unit (CPU) with 24 logical threads, an artificial intelligence (AI) processing unit delivering approximately 280 tera floating-point operations per second (TFLOPS) of theoretical peak performance at 16-bit floating-point precision (FP16), and 64 gigabytes (GB) of system memory.

Table 1 summarizes the main architecture and training hyperparameters used in this study. The settings include the input dimension, latent dimension, denoising-network structure, diffusion schedule, optimizer, training epochs, diversity-factor candidates, and random-seed configuration.

The training diffusion steps $T_{\text{train}} = 1000$ are used as a fixed noise schedule for reproducibility rather than as a claimed theoretically necessary or globally optimal setting. The final reported experiments use the corresponding $S_{\text{sample}} = 1000$ reverse-sampling configuration to keep the benchmark evaluation consistent with the retained training noise schedule. Therefore, this setting is treated as an implementation choice rather than as a theoretical optimum.

4.1. Raw Data Description

The experimental data used in this study are derived from the open dataset released by the Chinese State Grid Renewable Energy Generation Forecasting Competition. The dataset was published by Chen and Xu [32] and is publicly available at <https://doi.org/10.1038/s41597-022-01696-6> (accessed on 28 January 2026). The full dataset contains field measurements from six wind farms and eight photovoltaic (PV) power stations in China from 2019 to 2020, with a temporal resolution of 15 min. The dataset includes historical renewable power records and meteorological variables, and therefore provides a realistic basis for training and validating day-ahead renewable energy scenario generation models.

After data screening and quality control, the case study uses measurements from 13 grid-connected renewable sites, including six wind farms and seven PV stations. One PV station is excluded because it contains long-term zero-output behavior and insufficient valid records, which would otherwise distort the PV output distribution and spatial-correlation structure. Detailed station-screening statistics are reported in Supplementary Table S1. The core variables include wind speed (0–26.9 m/s), solar irradiance (0–1467 W/m²), ambient temperature (−24.5–37.6 °C), and renewable power output.

Missing values lasting fewer than 10 consecutive time steps are filled by linear interpolation, while samples with a missing rate exceeding 20% are removed to ensure data integrity.

The effect of short-gap interpolation is quantitatively checked by comparing temporal statistics before and after interpolation. The lag-1 autocorrelation, lag-4 autocorrelation, and ramp standard deviation change by only 0.003, 0.003, and 0.002, respectively, indicating that the interpolation procedure for short missing gaps does not materially distort the temporal dependence structure. The detailed interpolation diagnostic is provided in Supplementary Table S2.

To avoid information leakage, all normalization statistics are calculated using the training set only. Specifically, the mean and standard deviation used in Z-score normalization are estimated from the training set and then applied unchanged to the validation and test sets. The intra-type spatial correlation matrices used for feature construction are also computed only from training-set renewable output sequences. The dataset is split in chronological order rather than randomly, so that the validation and test periods are strictly later than the training period.

Figure 4 presents raw data visualizations from a representative wind farm and PV station. The temporal curves show the fluctuation characteristics of wind and PV generation, while the marginal distributions illustrate the non-Gaussian and weather-dependent nature of the renewable variables. These characteristics motivate the use of multi-source feature construction and generative scenario modeling.

A multicollinearity diagnostic is further conducted on the training-set features after excluding constant padding columns. The temporal lag features have a maximum variance inflation factor (VIF) of 8.53 and a condition number of 14.06, while meteorological and static spatial descriptors show stronger collinearity due to shared regional weather drivers and repeated station-level descriptors. Therefore, the VAE is not introduced under an independence assumption for the original 182-dimensional feature space. Instead, it is used to compress correlated multi-source inputs into a compact latent representation before CLDM-based generation. The full multicollinearity diagnostic and the temporal/seasonal split results are reported in Supplementary Tables S3 and S4, respectively.




Figure 4. Raw data visualization of a typical wind farm and PV station. (a) Temporal curves of wind speed and wind power; (b) temporal curves of solar irradiance and PV power; (c) statistical distribution of wind speed; (d) statistical distribution of solar irradiance.

4.2. VAE Latent Representation and Correlation-Preservation Diagnostics

The VAE encoder compresses the standardized 182-dimensional multi-source input matrix into a compact latent representation. The latent dimension is set to $K = 8$ according to validation-based sensitivity analysis over $K = \{8, 16, 32\}$. Among the tested settings, $K = 8$ provides the best overall trade-off among reconstruction accuracy, distributional similarity, latent-space compactness, and runtime. This setting is therefore adopted in the subsequent CLDM generation stage. The complete latent-dimension sensitivity table is provided in Supplementary Table S5.

The adopted temporal lag configuration $\{1, 6, 12, 24\}$ h is also checked against alternative lag configurations and a learned temporal embedding. The diagnostic results show that the adopted lag set performs close to the best explicit lag settings and clearly outperforms

the no-lag configuration. Although adding a learned temporal embedding slightly reduces root mean square error (RMSE) and mean absolute error (MAE) in the diagnostic proxy, the improvement is small, whereas the explicit $\{1, 6, 12, 24\}$ h lag set remains simpler, more interpretable, and does not require an additional trainable temporal-embedding module. The full lag-set diagnostic is reported in Supplementary Table S8.

To examine whether the generated latent samples can be reliably decoded by the pretrained VAE decoder, latent-space alignment is evaluated between VAE-encoded samples and CLDM-generated samples. The CLDM-generated latent samples remain close to the VAE-encoded latent distribution, with a latent mean difference of 0.021, standard deviation difference of 0.034, latent MMD of 0.1434, and covariance error of 0.058. These results reduce the risk that the decoder is applied to out-of-distribution latent samples. The detailed latent-alignment metrics are listed in Supplementary Table S6.

Decoder reconstruction accuracy is evaluated in both normalized and original physical scales. The normalized RMSE values of wind speed, solar irradiance, ambient temperature, and renewable power are 0.091, 0.088, 0.076, and 0.104, respectively. These values correspond approximately to 0.36 m/s, 21.56 W/m², 0.65 °C, and 5.72 MW in the original physical scale. Therefore, the decoder can reconstruct the main physical variables with acceptable accuracy before the generated scenarios are used in subsequent constraint-consistency and dispatch-oriented evaluations. The complete normalized and original-scale reconstruction errors are reported in Supplementary Table S7.

Spatial dependence is first examined through real and generated correlation matrices. Let $\mathbf{R}^{\text{real}}$ and $\mathbf{R}^{\text{gen}}$ denote the correlation matrices calculated from the real and generated renewable output sequences, respectively. The correlation error matrix is defined as:

$$\Delta \mathbf{R} = \mathbf{R}^{\text{gen}} - \mathbf{R}^{\text{real}}. \quad (13)$$

The padded zero entries introduced only for dimensional alignment are excluded from all correlation-preservation calculations.

Figure 5 compares the real correlation matrix, generated correlation matrix, and absolute correlation error matrix. The generated scenarios preserve the dominant intra-type correlation patterns, especially among PV stations, while wind–PV cross-type dependence remains more difficult to reproduce.

As shown in Table 2, the overall spatial-correlation MAE is 0.049, and the Wilcoxon test does not show a statistically significant systematic shift at the 0.05 level. The wind–PV group has a larger error than the intra-type groups, indicating that cross-type dependence remains more difficult to preserve when intra-type correlation features are used as the main spatial descriptors.




Figure 5. Quantitative spatial-correlation preservation analysis. (a) Real correlation matrix; (b) generated correlation matrix; (c) absolute correlation error matrix.

Table 2. Statistical evaluation of spatial-correlation preservation.

Group	Pairs	Corr. MAE	Corr. RMSE	Pres. Corr.	<i>p</i> -Value
Wind–Wind	15	0.041	0.055	0.912	0.184
PV–PV	21	0.028	0.036	0.887	0.238
Wind–PV	42	0.063	0.079	0.742	0.058
Overall	78	0.049	0.064	0.836	0.102

4.3. Performance of CLDM Scenario Generation and Constraint Consistency

The CLDM is trained in the VAE latent space rather than directly in the original 182-dimensional feature space. This design reduces the learning burden of the diffusion model while preserving the main statistical and spatiotemporal characteristics encoded by the VAE. Meteorological variables, including wind speed, solar irradiance, and ambient temperature, are used as explicit conditioning information during reverse denoising.

The diversity regulation factor α is introduced by scaling the reverse-sampling noise term. Its value is selected on the validation set rather than manually fixed. Candidate values are evaluated using statistical fidelity, scenario spread, tail-event coverage, correlation-preservation error, and constraint-violation rate.

Table 3 shows that increasing α improves tail-event coverage but may also increase statistical errors and constraint violations. Based on the validation results, $\alpha = 1.0$ is adopted as the default value because it achieves the lowest MMD and provides a balanced trade-off among statistical fidelity, tail-event coverage, and constraint consistency.

Table 3. Sensitivity analysis of the diversity regulation factor α on the original renewable-power scale.

α	RMSE (MW)	MAE (MW)	MMD	Tail-Event Coverage	Violation (%)
0.6	37.90	27.86	0.168	58.3	0.12
0.8	36.18	26.52	0.151	71.6	0.18
1.0	35.27	25.95	0.143	82.4	0.26
1.2	35.84	26.21	0.149	91.7	0.43
1.5	38.42	28.11	0.176	96.8	0.91

To clarify the source of constraint consistency, raw generated scenarios and post-corrected scenarios are evaluated separately. The raw scenarios are checked for negative output, output exceeding installed capacity, and excessive ramping. The corrected scenarios are evaluated using the same criteria after output-bound and ramp-consistency correction.

Table 4 shows that the raw generated scenarios still contain a small number of lower-bound, upper-bound, and ramping violations. After post-correction, output-bound violations are removed and the ramp-violation rate is reduced to 0.08%. Therefore, the near-zero violation rate should be interpreted as the feasibility of the post-corrected scenarios rather than as evidence that the raw CLDM outputs always satisfy all operational constraints.

Table 4. Constraint-violation diagnosis before and after post-correction.

Scenario	Lower (%)	Upper (%)	Ramp (%)	Corrected (%)	Avg. Corr. (MW)
Raw generated	0.84	0.31	2.17	–	–
Post-corrected	0.00	0.00	0.08	2.73	0.86

The sensitivity to imperfect meteorological conditioning is also examined by perturbing wind speed, solar irradiance, and ambient temperature inputs by 5–20%. As the

perturbation level increases from 0% to 20%, RMSE increases from 6.7539 to 7.2915, MAE increases from 3.2995 to 3.6802, and the constraint-violation rate increases from 0.233419 to 0.244073. These results indicate gradual degradation rather than abrupt failure under controlled meteorological-condition errors. The MMD and correlation-error indicators vary only slightly in this diagnostic, so the result is not overstated as strong forecast-error robustness. Since the open dataset contains historical measurements rather than operational numerical weather prediction products, this test is interpreted as a controlled forecast-error proxy rather than real numerical weather prediction (NWP) validation. The detailed perturbation results are provided in Supplementary Table S9.

To examine whether scenario realism is retained when 15 min scenarios are converted to hourly intervals, four consecutive 15 min outputs are averaged into hourly scenarios. The MMD changes from 0.143 to 0.137, and the correlation error changes from 0.128 to 0.132. These results indicate that the generated scenarios retain comparable distributional similarity and correlation-preservation performance under hourly downsampling. The complete hourly diagnostic is provided in Supplementary Table S10.

4.4. Comprehensive Performance Comparison

To evaluate the proposed VAE-CLDM framework, two aspects are considered: scenario generation quality and dispatch-oriented diagnostic applicability. Scenario generation quality is assessed using distance-based metrics, including RMSE, MAE, and MMD, while the dispatch-oriented diagnostic is evaluated using network-loss, renewable-absorption, dispatch-cost, cost-volatility, and voltage-overlimit indicators.

4.4.1. Scenario Generation Performance

The benchmark comparison evaluates whether the proposed VAE-CLDM improves the statistical quality of generated renewable-power scenarios relative to representative generative baselines. The compared configurations are defined as follows:

- Standalone VAE: a reconstruction-based baseline that performs dimensionality reduction and basic scenario reconstruction without the diffusion generation module.
- Wasserstein generative adversarial network (WGAN) and DDPM: representative deep generative baselines for renewable scenario generation.
- Standalone CLDM: a direct-space conditional diffusion baseline operating in the original 182-dimensional feature space without VAE-based latent compression or decoder reconstruction.
- VAE-CLDM: the proposed framework integrating multi-source feature construction, VAE-based latent dimensionality reduction, CLDM-based conditional scenario generation, decoder reconstruction, and diversity regulation.

All models are trained under the same chronological dataset split, batch size, number of training epochs, and hardware environment. To account for initialization and sampling variability, each stochastic configuration is repeated over five independent random seeds, and Table 5 reports the mean values and standard deviations on the normalized benchmark scale.

As shown in Table 5, VAE-CLDM achieves the lowest mean RMSE (0.3013 ± 0.0022), MAE (0.3636 ± 0.0010), and MMD (0.04269 ± 0.00040) on the training-statistics-standardized renewable-power benchmark. These normalized-scale generation-quality results should be compared within Table 5 and should not be mixed directly with physical-scale sensitivity results, dispatch diagnostics, or the separate CPU timing-scope diagnostic reported in Supplementary Table S12.

Table 5. Multi-seed performance comparison of scenario generation models on the training-statistics-standardized renewable-power benchmark.

Model	RMSE	MAE	MMD	Sampling Time (s/1000)
VAE	0.7942 ± 0.0054	0.6179 ± 0.0070	0.29773 ± 0.00489	0.073 ± 0.0007
CLDM	0.5101 ± 0.0005	0.4795 ± 0.0007	0.04873 ± 0.00042	7.2740 ± 0.1471
DDPM	0.5176 ± 0.0006	0.3897 ± 0.0005	0.08124 ± 0.00043	2.8092 ± 0.1878
WGAN	0.5814 ± 0.0023	0.5366 ± 0.0010	0.09211 ± 0.00216	0.062 ± 0.0007
VAE-CLDM	0.3013 ± 0.0022	0.3636 ± 0.0010	0.04269 ± 0.00040	9.8892 ± 0.6483

For the normalized-scale MMD comparison between VAE-CLDM and standalone CLDM, a Wilcoxon signed-rank diagnostic is additionally performed using the same five random seeds. The one-sided test corresponding to the directional hypothesis that VAE-CLDM reduces MMD gives $p = 0.03125$, while the two-sided exact test gives $p = 0.0625$. Because only five seeds are available, this result is interpreted as diagnostic evidence rather than a definitive statistical conclusion. The corresponding normalized-scale MMD summary and Wilcoxon diagnostic are provided in Supplementary Table S13.

The relative improvements are calculated directly from Table 5 and can be summarized as follows:

- Compared with the standalone VAE model, VAE-CLDM reduces RMSE by 62.1%, MAE by 41.2%, and MMD by 85.7% on the training-statistics-standardized renewable-power benchmark.
- Compared with standalone CLDM, VAE-CLDM reduces RMSE by 40.9%, MAE by 24.2%, and MMD by 12.4%. This supports the normalized-scale generation-quality advantage of VAE-CLDM in the implemented benchmark. However, the sampling time of VAE-CLDM in this benchmark is larger than that of standalone CLDM, so quality improvement and sampling cost are discussed as separate aspects rather than combined into a single superiority claim.
- Compared with WGAN and DDPM, VAE-CLDM improves the normalized-scale generation-quality metrics, with MMD reduced by 53.6% relative to WGAN and by 47.5% relative to DDPM.

To directly address the contribution of the multi-source feature-construction stage, a feature-level ablation diagnostic is further conducted using the available feature-ablation evaluation scale. This diagnostic does not treat Stage I as an independently trainable physical-constraint-learning module; instead, it isolates the effects of the constructed temporal-lag features, spatial-correlation descriptors, and meteorological conditioning pathways.

Table 6 shows that the full VAE-CLDM configuration clearly outperforms the power-only input baseline and the settings without temporal-lag features or meteorological conditioning. In particular, the MMD increases from 0.0203 ± 0.0025 to 0.0514 ± 0.0092 when temporal-lag features are removed, and to 0.0276 ± 0.0021 when meteorological conditioning is removed. These results indicate that the multi-source feature-construction stage is useful, especially the temporal-lag features and meteorological conditioning. Removing the spatial-correlation descriptors also leads to a modest degradation: RMSE increases from 10.452 ± 0.184 to 10.574 ± 0.162 , MAE increases from 6.028 ± 0.067 to 6.124 ± 0.127 , MMD increases from 0.0203 ± 0.0025 to 0.0236 ± 0.0046 , and Corr. MAE increases from 0.0617 ± 0.0051 to 0.0623 ± 0.0018 . Therefore, the spatial-correlation descriptors are retained as station-level dependence descriptors with a modest but consistent contribution, while the dominant gains in this ablation come from temporal-lag features and meteorological conditioning. Overall, this ablation supports the implemented feature-construction de-

sign while avoiding the unsupported claim that Stage I itself is an independently learned physical-constraint module.

Table 6. Feature-level ablation diagnostic for the multi-source feature-construction stage.

Configuration	RMSE	MAE	MMD	Corr. MAE
Full VAE-CLDM	10.452 ± 0.184	6.028 ± 0.067	0.0203 ± 0.0025	0.0617 ± 0.0051
w/o temporal-lag features	13.250 ± 0.313	7.654 ± 0.219	0.0514 ± 0.0092	0.0786 ± 0.0070
w/o spatial-correlation descriptors	10.574 ± 0.162	6.124 ± 0.127	0.0236 ± 0.0046	0.0623 ± 0.0018
w/o meteorological conditioning	11.031 ± 0.192	6.572 ± 0.075	0.0276 ± 0.0021	0.0700 ± 0.0040
Power-only input	18.763 ± 0.503	12.697 ± 0.310	0.1976 ± 0.0151	0.3154 ± 0.0216

4.4.2. Grid Dispatch Performance Evaluation

To further evaluate the practical value of the scenarios generated by the VAE-CLDM model, grid dispatch simulations are conducted on the Institute of Electrical and Electronics Engineers (IEEE) 33-node distribution network. The dispatch study is used as a diagnostic evaluation rather than a full market-clearing model. The objective minimizes the sum of grid-purchase cost, renewable-curtailment cost, network-loss cost, and reactive-compensation cost over 96 15 min intervals. The constraints include nodal power balance, voltage limits, line-flow limits, renewable-output bounds, and power-flow feasibility checks.

The core dispatch parameters are as follows:

- Cost parameters: renewable-energy curtailment cost \$0.8/kWh, grid-purchase cost \$1.2/kWh, network-loss cost \$1.0/kWh, and reactive-power compensation cost \$0.05/kVArh.
- Time resolution: 15 min intervals, corresponding to 96 time slots for 24 h dispatch.
- Renewable placement: buses 4, 6, 8, 10, 12, 14, 17, 20, 22, 24, 26, 29, and 31 in the IEEE 33-node system.
- Evaluation metrics: average network loss (MW), renewable energy absorption rate (%), average 24 h dispatch cost (\$), cost standard deviation (\$), and voltage over-limit rate (%).

Additional dispatch-model settings, including the power-flow approximation, voltage limits, renewable placement, load profile, scenario treatment, and cost terms, are summarized in Supplementary Table S14. The absolute dispatch reference cases, including the perfect-forecast and deterministic/no-scenario baselines, are reported in Supplementary Table S15.

Figure 6 compares five dispatch metrics across all models. The values are normalized to the interval 0–1, and lower values indicate better performance. The VAE-CLDM model achieves normalized values of 0.07–0.13 across all indicators. Specifically, for average network loss (0.12), average dispatch cost (0.07), cost standard deviation (0.13), and voltage overlimit rate (0.13), VAE-CLDM performs better than VAE (0.90), WGAN (0.68–0.79), DDPM (0.41–0.59), and CLDM (0.25–0.40). For renewable energy absorption rate, which is reciprocally normalized for consistent interpretation, VAE-CLDM obtains the lowest normalized value (0.09), corresponding to the highest actual absorption rate of 95.4%. For reference, Supplementary Table S15 reports that the perfect-forecast case reaches 96.4% absorption, the deterministic/no-scenario case reaches 92.8%, and the VAE-CLDM case reaches 95.4% under the same diagnostic dispatch setting.




Figure 6. Comparison of dispatch metrics across benchmark models.

Figure 7 shows the 24 h network-loss curves of different models on a typical day. VAE-CLDM exhibits the smoothest loss profile, with a fluctuation range close to 0.03–0.031 MW, whereas the peak network loss of the VAE model approaches 0.039 MW.

Figure 8 compares the 24 h dispatch-cost trajectories of different models on the same typical day. VAE-CLDM maintains the most stable dispatch-cost curve at approximately \$248 per 15 min interval, with a lower fluctuation amplitude than VAE, WGAN, DDPM, and CLDM.




Figure 7. Twenty-four-hour network loss comparison on a typical day.




Figure 8. Twenty-four-hour dispatch cost comparison on a typical day.

The dispatch-oriented diagnostic results indicate that the proposed VAE-CLDM method can provide useful scenario data for renewable energy grid integration and dispatch analysis. These results should be interpreted as application-oriented diagnostic evidence under the adopted IEEE 33-node setting, rather than as a universal guarantee of operational superiority.

4.5. Performance Summary

The comprehensive evaluation results indicate that the proposed VAE-CLDM framework improves normalized-scale scenario generation quality, supports engineering consistency after post-correction, and provides useful diagnostic scenario inputs for dispatch-oriented analysis. The main observations are summarized as follows:

1. **Scenario generation quality:** The VAE-CLDM model achieves the lowest RMSE, MAE, and MMD among the tested benchmark models in the training-statistics-standardized renewable-power benchmark. In particular, the MMD of VAE-CLDM is 0.04269 ± 0.00040 , lower than standalone CLDM (0.04873 ± 0.00042), DDPM (0.08124 ± 0.00043), WGAN (0.09211 ± 0.00216), and standalone VAE (0.29773 ± 0.00489). This indicates that the generated scenarios better match the statistical distribution of real renewable power output in the implemented benchmark.
2. **Spatiotemporal representation:** The explicit $\{1, 6, 12, 24\}$ h lag set provides a simple and interpretable multi-scale temporal representation. The feature-level ablation further shows that temporal-lag features and meteorological conditioning are the most important parts of the constructed input. Intra-type spatial-correlation descriptors are retained for interpretable station-level dependence representation, and the updated feature-level ablation shows a modest but consistent contribution of these descriptors. The overall spatial-correlation MAE is 0.049, and the Wilcoxon test does not indicate a systematic shift at the 0.05 level in Table 2.
3. **Constraint consistency:** Raw generated scenarios and post-corrected scenarios are evaluated separately. The output-bound violations are removed after post-correction, and the ramp-violation rate is reduced to 0.08%. Therefore, feasibility is attributed to the corrected scenario outputs rather than overstated as a hard guarantee of the raw generative model.
4. **Robustness and applicability:** Under controlled meteorological perturbations, RMSE and MAE increase gradually, indicating degradation rather than abrupt failure. This diagnostic is a controlled forecast-error proxy rather than real NWP validation, and it is not used to claim strong forecast-error robustness. Hourly downsampling does not materially degrade distributional similarity or correlation preservation. The dispatch-oriented diagnostic further shows that VAE-CLDM provides useful scenario inputs for renewable-rich distribution-network operation studies, but it should be interpreted as application-oriented diagnostic evidence rather than a universal guarantee of operational superiority.

Overall, the key methodological advantage of VAE-CLDM lies in combining VAE latent compression with CLDM conditional denoising. This design improves normalized-scale generation quality and provides a practical diagnostic workflow for evaluating scenario fidelity, feature contribution, post-correction consistency, and dispatch applicability under clearly stated evaluation scopes.

5. Conclusions

This paper proposes a three-stage VAE-CLDM framework for day-ahead renewable energy scenario generation. The framework integrates multi-source feature construction, VAE-based latent representation, conditional latent diffusion, decoder-based reconstruction,

and engineering post-correction. By performing diffusion-based generation in a compact latent space rather than directly in the original 182-dimensional feature space, the proposed method reduces the learning burden of high-dimensional generation while retaining key meteorological, temporal, and spatial-correlation information.

The simulation results show that the proposed framework improves statistical fidelity and engineering consistency after post-correction. In the training-statistics-standardized renewable-power benchmark, VAE-CLDM achieves the lowest RMSE, MAE, and MMD among the tested models, with RMSE of 0.3013 ± 0.0022 , MAE of 0.3636 ± 0.0010 , and MMD of 0.04269 ± 0.00040 . The feature-level ablation confirms that the multi-source feature-construction stage is useful, especially the temporal-lag features and meteorological conditioning; it also shows a modest but consistent contribution of spatial-correlation descriptors to station-level dependence representation. The overall spatial-correlation MAE is 0.049, and the Wilcoxon signed-rank test does not indicate a statistically significant systematic shift at the 0.05 level in the reported spatial-correlation diagnostic. The raw generated scenarios still contain a small number of output-bound and ramping violations, but after post-correction, lower- and upper-bound violations are reduced to 0.00%, and the ramp-violation rate is reduced to 0.08%.

The dispatch-oriented diagnostic study on the IEEE 33-node distribution network indicates that the generated scenarios can provide useful inputs for day-ahead dispatch analysis under renewable uncertainty, while the dispatch results should be interpreted as diagnostic evidence under the adopted setting rather than as a universal operational-performance guarantee. The present study is based on a single public regional dataset and uses controlled perturbation rather than real operational NWP forecasts. Therefore, cross-climate-zone generalization and real forecast-driven validation are not claimed. Future work will further incorporate key grid operation constraints into the generation process, introduce cross-type graph conditioning for wind–PV dependence, extend decoder reconstruction to probabilistic or ensemble forms to propagate latent and decoder uncertainty, and validate the framework on additional regions and operational weather forecasts.

Supplementary Materials: The following supporting information can be downloaded at: <https://www.mdpi.com/article/doi/s1>, Table S1: Data-quality screening statistics of renewable stations; Table S2: Interpolation-effect diagnostic for short missing gaps; Table S3: Multicollinearity diagnostic on training-set features after excluding constant padding columns; Table S4: Temporal and seasonal split diagnostics within the same regional dataset; Table S5: Sensitivity analysis of latent dimension K ; Table S6: Latent-space alignment between CLDM-generated and VAE-encoded latent samples; Table S7: Decoder reconstruction errors in normalized and original physical scales; Table S8: Lag-set sensitivity diagnostic; Table S9: Controlled meteorological perturbation diagnostic; Table S10: Scenario realism under different temporal resolutions; Table S11: Multi-seed performance comparison on the training-statistics-standardized renewable-power benchmark; Table S12: Full-pipeline computational accounting and timing-scope MMD diagnostic; Table S13: Seed-level normalized-scale MMD diagnostic for VAE-CLDM versus standalone CLDM; Table S14: Dispatch-model settings for the IEEE 33-node diagnostic case; Table S15: Absolute dispatch reference cases; Table S16: Station-level diagnostic for W1.

Author Contributions: Conceptualization, L.C., X.L. and G.Y.; Methodology, L.C., X.L., Q.C. and X.D.; Software, L.C., X.L. and W.L.; Validation, L.C. and X.L.; Formal analysis, L.C., X.L., Q.C. and X.D.; Investigation, L.C., X.L., W.L. and X.D.; Resources, L.C., X.L., Q.C., W.L. and X.D.; Data curation, L.C., X.L., Q.C., G.Y. and X.D.; Writing—original draft, L.C., X.L. and X.D.; Writing—review & editing, L.C., X.L., G.Y., W.L. and X.D.; Visualization, L.C., X.L., G.Y., W.L. and X.D. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the Science and Technology Project of Guangxi Power Grid Co., Ltd. (Project Name: Research on Key Technologies of Intelligent Agents for Primary-

Secondary Integration Analysis and Decision-Making in Power Dispatching and Operation, Project No.: GXKJXM20240240).

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The publicly available dataset analyzed in this study can be accessed at <https://doi.org/10.1038/s41597-022-01696-6> (accessed on 28 January 2026). The processed results supporting the conclusions of this article are included within the article and its Supplementary Materials.

Conflicts of Interest: Authors Lin Chen, Xinran Liu, Quanqi Chen, Guinan Ye, Wen Liu were employed by the company Guangxi Power Grid Co., Ltd. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest. The funders (Guangxi Power Grid Co., Ltd.) had no role in the design of the study; in the collection, analysis, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

References

1. Maihemuti, S.; Wang, W.; Wu, J.; Wang, H.; Muhedaner, M. New Energy Power System Static Security and Stability Region Calculation Research Based on IPSO-RLS Hybrid Algorithm. *Energies* **2022**, *15*, 9655. [\[CrossRef\]](#)
2. Liu, Z.; Guo, H.; Zhang, Y.; Zuo, Z. A Comprehensive Review of Wind Power Prediction Based on Machine Learning: Models, Applications, and Challenges. *Energies* **2025**, *18*, 350. [\[CrossRef\]](#)
3. Zheng, K.; Sun, Z.; Song, Y.; Zhang, C.; Zhang, C.; Chang, F.; Yang, D.; Fu, X. Stochastic Scenario Generation Methods for Uncertainty in Wind and Photovoltaic Power Outputs: A Comprehensive Review. *Energies* **2025**, *18*, 503. [\[CrossRef\]](#)
4. Cui, M.; Ke, D.; Sun, Y.; Gan, D.; Zhang, J.; Hodge, B.M. Wind Power Ramp Event Forecasting Using a Stochastic Scenario Generation Method. *IEEE Trans. Sustain. Energy* **2015**, *6*, 422–433. [\[CrossRef\]](#)
5. Wen, H.; Pinson, P.; Ma, J.; Gu, J.; Jin, Z. Continuous and Distribution-Free Probabilistic Wind Power Forecasting: A Conditional Normalizing Flow Approach. *IEEE Trans. Sustain. Energy* **2022**, *13*, 2250–2263. [\[CrossRef\]](#)
6. Dumas, J.; Wehenkel, A.; Lanaspeze, D.; Cornélusse, B.; Sutura, A. A Deep Generative Model for Probabilistic Energy Forecasting in Power Systems: Normalizing Flows. *Appl. Energy* **2022**, *305*, 117871. [\[CrossRef\]](#)
7. Dalal, D.; Bilal, M.; Shah, H.; Sifat, A.I.; Pal, A.; Augustin, P. Cross-Correlated Scenario Generation for Renewable-Rich Power Systems Using Implicit Generative Models. *Energies* **2023**, *16*, 1636. [\[CrossRef\]](#)
8. Yi, D.; Yu, M.; Wang, Q.; Tian, H.; Wang, L.; Yan, Y.; Wu, C.; Hu, B.; Li, C. Method for Wind–Solar–Load Extreme Scenario Generation Based on an Improved InfoGAN. *Appl. Sci.* **2024**, *14*, 9163. [\[CrossRef\]](#)
9. Hou, Z.; Liu, J.; Shao, Z.; Ma, Q.; Liu, W. Machine Learning Innovations in Renewable Energy Systems with Integrated NRBO-TXAD for Enhanced Wind Speed Forecasting Accuracy. *Electronics* **2025**, *14*, 2329. [\[CrossRef\]](#)
10. Lee, D.; Baldick, R. Load and Wind Power Scenario Generation Through the Generalized Dynamic Factor Model. *IEEE Trans. Power Syst.* **2017**, *32*, 400–410. [\[CrossRef\]](#)
11. Hu, J.; Li, H.; Liu, Z. A Novel Scenario Generation Framework Based on the Knowledge of Existing Wind Power Plants. *IEEE Trans. Sustain. Energy* **2021**, *12*, 1229–1241. [\[CrossRef\]](#)
12. Wang, Y.; Liu, Y.; Kirschen, D.S. Scenario Reduction with Submodular Optimization. *IEEE Trans. Power Syst.* **2017**, *32*, 2479–2480. [\[CrossRef\]](#)
13. Kang, B.; Kim, S.; Jung, H.; Choe, J.; Lee, K. Efficient Assessment of Reservoir Uncertainty Using Distance-Based Clustering: A Review. *Energies* **2019**, *12*, 1859. [\[CrossRef\]](#)
14. Akhavadegan, F.; Wang, L.; McCalley, J. Scenario Selection for Iterative Stochastic Transmission Expansion Planning. *Energies* **2020**, *13*, 1203. [\[CrossRef\]](#)
15. Goh, H.H.; Peng, G.; Zhang, D.; Dai, W.; Kurniawan, T.A.; Goh, K.C.; Cham, C.L. A New Wind Speed Scenario Generation Method Based on Principal Component and R-Vine Copula Theories. *Energies* **2022**, *15*, 2698. [\[CrossRef\]](#)
16. Chen, Y.; Wang, Y.; Kirschen, D.S.; Zhang, B. Model-Free Renewable Scenario Generation Using Generative Adversarial Networks. *IEEE Trans. Power Syst.* **2018**, *33*, 3265–3275. [\[CrossRef\]](#)
17. Yuan, R.; Wang, B.; Sun, Y.; Song, X.; Watada, J. Conditional Style-Based Generative Adversarial Networks for Renewable Scenario Generation. *IEEE Trans. Power Syst.* **2023**, *38*, 1281–1296. [\[CrossRef\]](#)
18. Dong, W.; Chen, X.; Yang, Q. Data-Driven Scenario Generation of Renewable Energy Production Based on Controllable Generative Adversarial Networks with Interpretability. *Appl. Energy* **2022**, *308*, 118387. [\[CrossRef\]](#)

19. He, G.; Liu, K.; Wang, S.; Lei, Y.; Li, J. CWM-CGAN Method for Renewable Energy Scenario Generation Based on Weather Label Multi-Factor Definition. *Processes* **2022**, *10*, 470. [[CrossRef](#)]
20. Zhang, W.; Luo, Y.; Zhang, Y.; Srinivasan, D. SolarGAN: Multivariate Solar Data Imputation Using Generative Adversarial Network. *IEEE Trans. Sustain. Energy* **2021**, *12*, 743–746. [[CrossRef](#)]
21. Kang, M.; Zhu, R.; Chen, D.; Li, C.; Gu, W.; Qian, X.; Yu, W. A Cross-Modal Generative Adversarial Network for Scenarios Generation of Renewable Energy. *IEEE Trans. Power Syst.* **2024**, *39*, 2630–2640. [[CrossRef](#)]
22. Khan, M.; Naeem, M.R.; Al-Ammar, E.A.; Ko, W.; Vettikalladi, H.; Ahmad, I. Power Forecasting of Regional Wind Farms via Variational Auto-Encoder and Deep Hybrid Transfer Learning. *Electronics* **2022**, *11*, 206. [[CrossRef](#)]
23. Li, Z.; Peng, X.; Cui, W.; Xu, Y.; Liu, J.; Yuan, H.; Lai, C.S.; Lai, L.L. A Novel Scenario Generation Method of Renewable Energy Using Improved VAEGAN with Controllable Interpretable Features. *Appl. Energy* **2024**, *363*, 122905. [[CrossRef](#)]
24. Krishna, A.B.; Abhyankar, A.R. An Efficient Data-Driven Conditional Joint Wind Power Scenario Generation for Day-Ahead Power System Operations Planning. *IEEE Trans. Power Syst.* **2024**, *39*, 3105–3117. [[CrossRef](#)]
25. Xu, C.; Xu, P.; Dai, Y.; Su, S.; Zhang, L.; Zhang, J.; Bai, Y.; Gao, T.; Xie, Q.; Shang, L.; et al. A Two-Stage Generative Architecture for Renewable Scenario Generation Based on Temporal Scenario Representation and Diffusion Models. *Energies* **2025**, *18*, 1275. [[CrossRef](#)]
26. Li, D.; Zhao, X.; Xu, W.; Ge, C.; Li, C. A Novel Renewable Energy Scenario Generation Method Based on Multi-Resolution Denoising Diffusion Probabilistic Models. *Energies* **2025**, *18*, 3781. [[CrossRef](#)]
27. Wu, H.; Krad, I.; Florita, A.; Hodge, B.M.; Ela, E.; Ibanez, E.; Zhang, J. Stochastic Multi-Timescale Power System Operations with Variable Wind Generation. *IEEE Trans. Power Syst.* **2017**, *32*, 3325–3337. [[CrossRef](#)]
28. Modarresi, M.S.; Xie, L.; Campi, M.C.; Garatti, S.; Carè, A.; Thatte, A.A.; Kumar, P.R. Scenario-Based Economic Dispatch with Tunable Risk Levels in High-Renewable Power Systems. *IEEE Trans. Power Syst.* **2019**, *34*, 5103–5114. [[CrossRef](#)]
29. Alanazi, M.; Alanazi, A.; Memon, Z.A.; Awan, A.B.; Deriche, M. Multi-Objective Energy Management in Microgrids: Improved Honey Badger Algorithm with Fuzzy Decision-Making and Battery Aging Considerations. *Energies* **2024**, *17*, 4373. [[CrossRef](#)]
30. Zhao, Y.; Li, H.; Wan, C.; Du, D.; Chen, B. Microgrid Energy Management Considering Energy Storage Degradation Cost. *Batteries* **2025**, *11*, 169. [[CrossRef](#)]
31. Arévalo, P.; Benavides, D.; Ochoa-Correa, D.; Ríos, A.; Torres, D.; Villanueva-Machado, C.W. Smart Microgrid Management and Optimization: A Systematic Review Towards the Proposal of Smart Management Models. *Algorithms* **2025**, *18*, 429. [[CrossRef](#)]
32. Chen, Y.; Xu, J. Solar and Wind Power Data from the Chinese State Grid Renewable Energy Generation Forecasting Competition. *Sci. Data* **2022**, *9*, 577. [[CrossRef](#)] [[PubMed](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.