Physics-Aware LLM-Based Probabilistic Wind Power Scenario Generation under Extreme Icing Conditions

Lei Wang,Ying Zhang
School of Electrical and Computer Engineering
Oklahoma State University,
Stillwater, OK, U.S.
{leiwang;y.zhang}@okstate.edu

Di Shi
Klipsch School of Electrical and Computer Engineering
New Mexico State University,
Las Cruces, NM, U.S.
dshi@nmsu.edu

Fei Ding
National Renewable Energy Laboratory
Golden, CO, U.S.
fei.ding@nrel.gov

***Abstract*—Accurately characterizing wind power uncertainty under icing and post-disaster conditions remains a critical challenge for resilient power system operation. To address this issue, this paper proposes a physics-aware large language model (LLM) framework for probabilistic wind power scenario generation under extreme icing conditions. The proposed framework integrates supervisory control and data acquisition (SCADA)-based physical modeling, multimodal tokenization, and a causal Transformer architecture trained in an autoregressive manner. A physics-aware decoding scheme effectively enforces rated power limits and ramping constraints on the generated trajectories while preserving stochastic diversity. Case studies using real wind turbine data show that the proposed method reproduces icing-induced power degradation and temporal variability observed during extreme weather. The resulting scenarios are physically consistent and high-fidelity, thereby significantly enhancing resilience assessment and recovery planning in renewable-integrated power systems.**

***Index Terms*—Physics-aware machine learning, icing conditions, large language model, probabilistic scenario generation, power system resilience, extreme weather, GenAI.**

I. Introduction

EXTREME weather events such as icing, blizzards, and storms increasingly threaten the reliable operation of renewable-integrated power systems. In cold-climate regions, wind farms are particularly vulnerable to blade icing, which causes aerodynamic degradation, mechanical stress, and severe power losses [1]. When such events occur on a large scale, the resulting generation uncertainty propagates across the grid's temporal and spatial scales, leading to voltage instability, load shedding, and delayed system recovery [2]. Following such disasters, resilient operations and restoration rely critically on accurate modeling of renewable generation uncertainty [3]. Hence, realistic post-disaster wind power scenarios provide a foundation for resilience assessment, restoration optimization, and secure operation planning in renewable-integrated power systems [4].

Existing scenario generation approaches can be broadly classified into explicit and implicit data-driven approaches. Explicit methods rely on predefined probabilistic models to characterize renewable generation uncertainty [5]–[7]. Although interpretable and computationally efficient, these methods usually assume linearity, stationarity, or fixed distributions, limiting their ability to capture nonlinear and non-stationary wind power characteristics under extreme conditions.

This research was supported by the National Science Foundation through award numbers 2418359 and 2346213.

In contrast, implicit methods directly learn the probability distribution of wind power from historical data without relying on explicit parametric assumptions [8]–[11]. Despite their strong expressive power, most existing data-driven methods are designed for normal operating conditions and neglect the physical impacts of disasters such as icing. As a result, generated scenarios may violate physical constraints or deviate from realistic post-disaster behaviors.

Despite these advances, most data-driven methods are designed for normal operating conditions and neglect the physical impacts of disasters such as icing or storms. These events alter turbine aerodynamics and mechanical behavior, introducing nonlinear and degradation-driven dynamics that purely statistical models fail to represent. As a result, the generated scenarios may violate physical constraints or deviate from realistic post-disaster behaviors, reducing their value for resilience-oriented applications.

To address these limitations, there is a growing need for physics-aware, data-driven frameworks that integrate physical knowledge with learning-based generative models. Recent developments in large language models (LLMs) provide a promising foundation, offering strong capabilities in temporal reasoning and probabilistic sequence generation. Compared with variational autoencoder (VAE)-, conditional generative adversarial network (CGAN)-, or diffusion-based generators, an autoregressive LLM models wind power as a causal sequence, which better matches long-horizon, non-stationary icing dynamics. Its stepwise decoding also makes it straightforward to impose operational limits (e.g., rated-power and ramping constraints) during sampling without sacrificing diversity. By combining LLM-based generative modeling with physics-informed constraints, this study aims to produce realistic, interpretable, and physically consistent wind power scenarios under disaster conditions. The main contributions of this study are summarized as follows:

- A physics-aware LLM framework is proposed to generate probabilistic wind power scenarios under icing conditions. By integrating physical priors—such as rated power limits, ramping constraints, and icing-induced degrada-

979-8-3315-8138-1

tion behaviors—into the generative process, the framework ensures the physical consistency of the generated trajectories, enabling reliable use for resilience analysis and restoration planning.

- A causal Transformer architecture with multimodal tokenization is designed to capture nonlinear, stochastic, and context-dependent features in wind power data. This mechanism enhances the interpretability and reliability of scenario generation, outperforming conventional data-driven approaches such as VAE-, CGAN-, and diffusion-based models in modeling uncertainty under extreme weather conditions.

II. Problem Formulation

Let x_t denote the wind power output at time step t. The historical wind power sequence of length L is represented as $x_{t-L+1:t} = \{x_{t-L+1}, \ldots, x_t\}$. At each time step, we denote the exogenous feature vector as $\mathbf{f}_t$, which includes supervisory control and data acquisition (SCADA) measurements (e.g., pitch angle, rotor speed, nacelle internal temperature) and meteorological variables (e.g., wind speed, ambient temperature). Accordingly, $\mathbf{f}_{t-L+1:t}$ denotes the historical feature sequence aligned with $x_{t-L+1:t}$.

In this work, we focus exclusively on icing disaster scenarios, where the historical wind power and feature sequences are collected under icing conditions. This ensures that the model explicitly learns the temporal dynamics of wind power degradation caused by blade icing.

Given the historical power sequence $x_{t-L+1:t}$ and the associated historical features $\mathbf{f}_{t-L+1:t}$, the task of probabilistic scenario generation is to produce a set of M plausible future trajectories of wind power:

$$\left\{ \tilde{x}^{(m)}_{t+1:t+H} \right\}_{m=1}^{M} \sim P\left(x_{t+1:t+H} \mid x_{t-L+1:t},\, \mathbf{f}_{t-L+1:t}\right), \quad (1)$$

where $\tilde{x}^{(m)}_{t+1:t+H} = \left\{ \tilde{x}^{(m)}_{t+1}, \tilde{x}^{(m)}_{t+2}, \ldots, \tilde{x}^{(m)}_{t+H} \right\}$ denotes the m-th generated scenario. By concentrating on icing- affected data, the model captures the degradation patterns and enhanced variability induced by icing. The generated scenarios are post-processed with turbine physical limits (e.g., rated power cap and ramping constraints) to ensure physical plausibility.

III. Proposed Physics-Aware LLM Method

The proposed framework, illustrated in Fig. 1, generates probabilistic wind power scenarios under icing conditions by capturing temporal dependencies in power dynamics and enforcing physics-based constraints. It processes historical wind power, SCADA, and meteorological data, with the feature set optimized through correlation analysis with the target output. The workflow includes: i) Feature Selection: Retaining SCADA and meteorological variables strongly correlated with wind power to reduce noise and improve efficiency; ii) Physical Model Identification: Fitting a wind speed–power curve to non-icing data via least squares to establish a normal-operation baseline that constrains subsequent generation [12]; iii) Tokenization: Discretizing the selected features into tokens

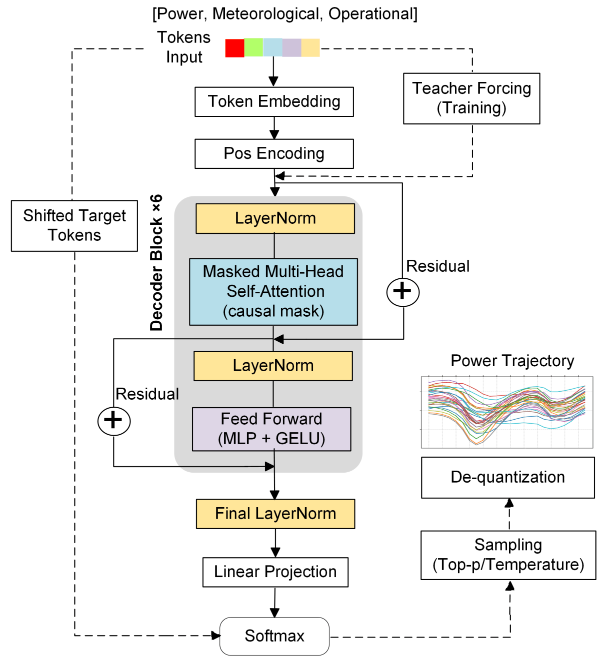

Fig. 1: LLM-based Causal Transformer architecture for wind power scenario generation under icing conditions.

according to their physical ranges and concatenating them into multimodal input sequences; iv) LLM-based Modeling: Feeding tokenized sequences into a decoder-only causal Transformer trained autoregressively with teacher forcing to predict future power tokens; v) Physics-aware Training and Decoding: Incorporating physics-based penalties—such as rated power caps, baseline bounds, ramp-rate limits, and smoothness—while applying constrained sampling to enforce physical feasibility; and vi) Scenario Generation: De-quantizing sampled tokens into continuous trajectories and producing multiple scenarios via top-p and temperature sampling that reflect icing-induced degradation with a lower mean and higher variability.

A. Tokenization and LLM Architecture

The goal of tokenization is to transform continuous turbine measurements into discrete tokens suitable for autoregressive modeling with LLMs [13]. In this study, we specifically focus on icing scenarios, and therefore, only on turbine operation under iced conditions is considered. Accordingly, SCADA and meteorological features are normalized and discretized into fixed bins, as summarized in Table I.

Each feature is first scaled to the interval [0, 1] according to its physical or operational range. The wind power output, normalized by rated capacity, is quantized into 256 μ-law bins to emphasize detail in the low-production regime while compressing the high-power region. Meteorological variables are discretized at different resolutions: wind speed into 64 bins, and temperature into 16 bins. Operational variables, such as pitch angle, yaw position, and generator speed, are quantized into 16 bins to encode control dynamics under icing conditions.

979-8-3315-8138-1

TABLE I: Tokenization Settings for Features

Feature type	Variables	Discretization (bins)
Power	Active power	μ-law 256 (μ = 120)
Meteorological	Wind speed, temp.	Quantile 64 (wind), 16 (temp.)
Operational	Pitch angle, yaw, generator speed	Quantile 16

At every time step, the tokenized features are concatenated into a multimodal tuple [Power, Meteorological, Operational]. The temporal sequence of such tuples then forms the model input, enabling the Transformer to autoregressively generate wind power trajectories that reflect both environmental and operational contexts.

At time step t, the tokenized input sequence is denoted as

$$\mathbf{x}_t = \{x_t^{(p)},\, x_t^{(m)},\, x_t^{(o)}\}, \quad t = 1, \ldots, L, \tag{2}$$

where $x_t^{(p)}$, $x_t^{(m)}$, and $x_t^{(o)}$ represent power, meteorological, and operational tokens, respectively. Each discrete token index is mapped into a dense representation by a learnable embedding matrix $\mathbf{E} \in \mathbb{R}^{V\times d}$, with V the vocabulary size and d the embedding dimension. Positional encodings $\mathbf{p}_{t,k} \in \mathbb{R}^d$ are added to preserve temporal order, yielding

$$\mathbf{z}_{t,k} = \mathbf{E}[x_t^{(k)}] + \mathbf{p}_{t,k}, \tag{3}$$

where $k \in \{p, m, o\}$ indexes the modalities corresponding to power, meteorological, and operational tokens, respectively.

Concatenating all modalities across the temporal dimension gives the model input sequence

$$\mathbf{Z}^{(0)} = \{z_{1,p},\, z_{1,m},\, z_{1,o},\, \ldots,\, z_{L,p},\, z_{L,m},\, z_{L,o}\}. \tag{4}$$

The embedded sequence is passed through a stack of N decoder blocks, each following a design where layer normalization is applied before every sub-layer and residual connections are preserved. Given the input representation $\mathbf{H}^{(l-1)} \in \mathbb{R}^{L\times d}$ to the l-th block, masked multi-head self-attention is first applied. For each attention head $h = 1, \ldots, H$, the query, key, and value matrices are computed as

$$\mathbf{Q}_h = \mathbf{X}\mathbf{W}_h^Q, \quad \mathbf{K}_h = \mathbf{X}\mathbf{W}_h^K, \quad \mathbf{V}_h = \mathbf{X}\mathbf{W}_h^V, \tag{5}$$

where $\mathbf{X} = \mathrm{LN}(\mathbf{H}^{(l-1)})$, $\mathbf{W}_h^Q, \mathbf{W}_h^K, \mathbf{W}_h^V \in \mathbb{R}^{d\times d_k}$ are learnable projection matrices, d is the embedding dimension, and $d_k = d/H$ is the dimension per head. A causal mask $\mathbf{M}_{\text{causal}}$ is applied to prevent attending to future tokens. The attention output for head h is

$$\mathrm{Attn}_h = \mathrm{Softmax}\left(\frac{\mathbf{Q}_h\mathbf{K}_h}{\sqrt{d_k}} + \mathbf{M}_{\text{causal}}\right)\mathbf{V}_h. \tag{6}$$

Outputs from all heads are concatenated and linearly projected, followed by a residual connection:

$$\mathbf{U}^{(l)} = \mathbf{H}^{(l-1)} + \mathrm{Concat}(\mathrm{Attn}_1, \ldots, \mathrm{Attn}_H)\mathbf{W}^O, \tag{7}$$

where $\mathbf{W}^O \in \mathbb{R}^{Hd_k\times d}$ is the output projection.

The second sub-layer is a position-wise feed-forward network with GELU activation,

$$\mathrm{FFN}(\mathbf{x}) = \mathbf{W}_2\,\sigma(\mathbf{W}_1\mathbf{x} + \mathbf{b}_1) + \mathbf{b}_2, \tag{8}$$

where $\mathbf{W}_1 \in \mathbb{R}^{d\times d_{\text{ff}}}$, $\mathbf{W}_2 \in \mathbb{R}^{d_{\text{ff}}\times d}$ are learnable weights, $\mathbf{b}_1, \mathbf{b}_2$ are biases, and d_{ff} is the hidden dimension of the feed-forward layer. After applying layer normalization, the output is added through another residual connection, yielding the block output

$$\mathbf{H}^{(l)} = \mathbf{U}^{(l)} + \mathrm{FFN}\Big(\mathrm{LN}(\mathbf{U}^{(l)})\Big). \tag{9}$$

After N stacked decoder blocks, the final hidden states are normalized and projected into the vocabulary space:

$$\mathbf{o}_t = \mathrm{Softmax}\Big(\mathbf{W}_o\,\mathrm{LN}(\mathbf{h}_t^{(N)})\Big), \quad \mathbf{o}_t \in \mathbb{R}^V \tag{10}$$

where $\mathbf{h}_t^{(N)} \in \mathbb{R}^d$ is the hidden representation of the t-th token after the last block, $\mathbf{W}_o \in \mathbb{R}^{V\times d}$ is the output projection matrix, and V is the vocabulary size determined by the discretization bins across all features. The output vector $\mathbf{o}_t$ is a probability distribution over the next possible tokens.

During training, the model is optimized with teacher forcing, where the shifted ground-truth sequence $\{x_2, \ldots, x_{L+1}\}$ serves as the prediction target. The training objective is the cross-entropy loss

$$\mathcal{L}_{\text{CE}} = -\sum_{t=1}^{L} \log \mathbf{o}_t[x_{t+1}], \tag{11}$$

where $\mathbf{o}_t[x_{t+1}]$ denotes the predicted probability assigned to the true next token x_{t+1}.

During inference, tokens are generated autoregressively as

$$x_{t+1} \sim \mathbf{o}_t, \tag{12}$$

using nucleus sampling with top-p filtering and temperature scaling. The resulting discrete sequence is finally de-quantized back into continuous wind power trajectories.

B. Physics-aware Scenario Generation

After autoregressive modeling in Section III-A, the model outputs a probability distribution $\mathbf{o}_t \in \mathbb{R}^V$ over possible next power tokens at each time step t. Let $q(\cdot)$ and $q^{-1}(\cdot)$ denote the quantization and de-quantization mappings, respectively, with $q^{-1} : \{1, \ldots, V_P\} \to [0, 1]$ mapping discrete tokens back to normalized power values. The generated token $\hat{x}_{t+1}$ is decoded into power as

$$\hat{p}_{t+1} = q^{-1}(\hat{x}_{t+1}) \cdot P_{\text{rated}}, \tag{13}$$

where P_{rated} is the rated capacity of the turbine. Meteorological and operational features are not predicted, but instead serve as conditioning variables to guide the generation process, ensuring that power trajectories reflect the actual operating environment under icing conditions.

To enforce physical plausibility, three classes of constraints are imposed. First, a power cap loss prevents predicted outputs from exceeding either the rated power or a scaled physical baseline $P_{\text{norm}}(v_t)$, obtained from our previous work [12]:

$$\begin{aligned}\mathcal{L}_{\text{cap}} &= \frac{1}{L}\sum_{t=1}^{L}\Big[\hat{p}_t - \min\big(P_{\text{rated}}, \alpha P_{\text{norm}}(v_t)\big)\Big]_+, \\ [z]_+ &= \max(z, 0),\end{aligned} \tag{14}$$

979-8-3315-8138-1

where $0 < \alpha \leq 1$ is a safety margin factor.

Second, a ramping loss penalizes sharp variations between consecutive steps:

$$\mathcal{L}_{\text{ramp}} = \frac{1}{L}\sum_{t=2}^{L}\Big[\,(\hat{p}_{t+1} - \hat{p}_t - R^{\uparrow})_+^2 + (\hat{p}_t - \hat{p}_{t+1} - R^{\downarrow})_+^2\,\Big] \tag{15}$$

where $R^{\uparrow}$ and $R^{\downarrow}$ are upward and downward ramp-rate limits derived from the historical data.

Third, a temporal smoothness penalty is included via a Huber-style total variation term:

$$\mathcal{L}_{\text{tv}} = \frac{1}{L}\sum_{t=2}^{L}\rho_\delta(\hat{p}_t - \hat{p}_{t-1}), \tag{16}$$

where ρ_δ denotes the Huber loss.

The overall objective combines data likelihood with physics-based regularization:

$$\mathcal{L} = \mathcal{L}_{\text{CE}} + \lambda_{\text{cap}}\mathcal{L}_{\text{cap}} + \lambda_{\text{ramp}}\mathcal{L}_{\text{ramp}} + \lambda_{\text{tv}}\mathcal{L}_{\text{tv}} \tag{17}$$

where λ_{cap}, λ_{ramp}, and λ_{tv} are balancing hyperparameters.

During inference, physics-aware sampling is applied. At each step t, candidate power tokens are first drawn using nucleus sampling with temperature scaling. The candidate set $\mathcal{C}_t$ is then pruned by discarding those that violate cap or ramping limits:

$$\begin{aligned}\mathcal{C}_t \leftarrow \Big\{\, i \in \mathcal{C}_t \;\big|\; p(i) \leq \min\big(P_{\text{rated}},\, \alpha P_{\text{norm}}(v_t)\big),\\ \big|p(i) - \hat{p}_t\big| \leq \tilde{R}\Big\}\end{aligned} \tag{18}$$

where $\tilde{R}$ is a relaxed ramping tolerance. If no feasible candidate remains, the nearest valid token is projected back. Repeating this process over H steps produces a complete power trajectory.

Finally, the discrete sequence $\{\hat{x}_{L+1}, \dots, \hat{x}_{L+H}\}$ is de-quantized into continuous values $\{\hat{p}_{L+1}, \dots, \hat{p}_{L+H}\}$. Light smoothing (e.g., moving average within physical limits) is optionally applied to suppress spurious spikes. Sampling multiple trajectories yields a diverse yet physically consistent scenario set $\{\hat{p}^{(m)}_{L+1:L+H}\}_{m=1}^{M}$, aligning probabilistic generation with both data-driven patterns and turbine physics.

IV. Case Study

The proposed LLM-based stochastic scenario generation framework is implemented in Python 3.8.2 using the PyTorch 1.8.0 GPU backend. All experiments are conducted on a workstation equipped with an Intel Xeon E5-2680 v4 CPU (2.4 GHz) and an NVIDIA GeForce RTX 3080 Ti GPU.

The SCADA dataset from two wind turbines in North China [12] is adopted. Following standard preprocessing in [12], the data is cleaned and resampled to 1-min resolution, and 26 SCADA variables are retained and normalized to $[0, 1]$. Non-icing periods are used to fit a four-parameter physical power curve and estimate rated power and ramping limits, while icing periods are used for training and evaluation.

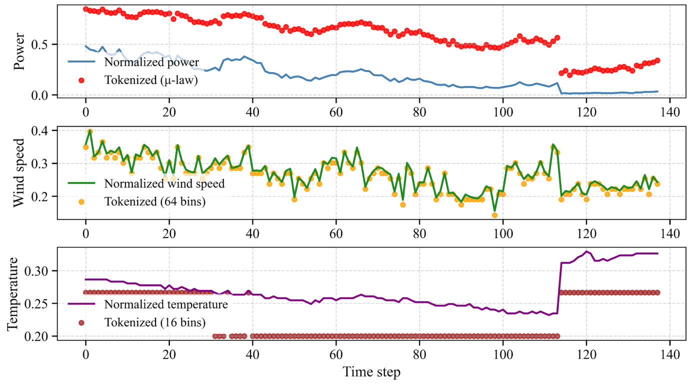

Fig. 2: Tokenization results for power, wind speed, and temperature.

The data are split chronologically into 70%/15%/15% for training/validation/testing.

The model performance is evaluated using three criteria: (i) the continuous ranked probability score (CRPS) [14] for probabilistic accuracy, (ii) the Kullback–Leibler divergence (KLD) [11] for statistical similarity, and (iii) the violation rate against physical constraints (rated power and ramping limits) for physical consistency.

A. Tokenization & Scenario Generation

To enable the LLM-based scenario generation, continuous SCADA features are transformed into discrete tokens. Each feature is normalized using the statistics of non-icing data to ensure consistent scaling between normal and icing conditions. Then, the normalized signals are discretized according to their physical or operational ranges, as summarized in Table I. The active power is quantized into 256 μ-law bins to emphasize variations in low-output regions, while wind speed and temperature are uniformly discretized into 64 and 16 bins, respectively. This feature-specific quantization preserves the physical dynamics of different variables while reducing dimensional complexity. The tokenized sequences of power, wind speed, and temperature (Fig. 2) clearly reflect the underlying temporal patterns and serve as structured symbolic inputs for the subsequent LLM-based scenario generation.

To evaluate the stochastic scenario generation capability and the impact of physics-aware constraints, 50 wind power trajectories are generated conditioned on the predicted normalized power. Fig.3 (a) shows the results without physical constraints, where the black solid line denotes the true power, the light blue curves represent unconstrained LLM-generated scenarios, and the blue dashed line shows their mean. The unconstrained results accurately capture the overall temporal evolution and nonlinear variations of the true curve, including the abrupt drop around step 40, but exhibit occasional unrealistic spikes and excessive fluctuations due to the absence of physical regularization. In contrast, Fig.3(b) presents the results with physics-aware constraints including the power-cap loss, ramp-rate penalty, and temporal smoothness regularization. These constraints effectively suppress non-physical excursions while preserving the stochastic diversity of the generated samples. The trajectories exhibit improved temporal consistency and

979-8-3315-8138-1

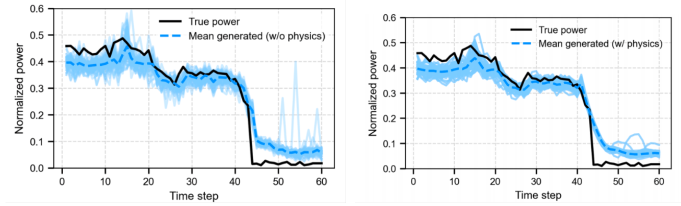

(a) (b)

Fig. 3: Comparison of LLM-generated wind power scenarios without and with physics-aware constraints.

TABLE II: Performance Comparison of Different Scenario Generation Models

Model	CRPS	KLD	VR (%)
CGAN [9]	0.162	0.061	7.05
VAE [10]	0.189	0.076	7.72
Diffusion [11]	0.145	0.052	6.58
Proposed LLM	**0.121**	**0.038**	**2.85**

physical plausibility, particularly in the post-event region, confirming that the proposed physics-aware sampling enforces realistic operational dynamics without degrading probabilistic variability. Overall, these results demonstrate that the proposed LLM-based framework can serve as a powerful probabilistic yet physically consistent generator for uncertainty quantification and resilient operation analysis in power systems.

B. Performance Benchmarking and Baseline Comparison

To evaluate the effectiveness of the proposed LLM-based scenario generator, we compare it with three representative deep generative baselines, including the CGAN [9]-, VAE [10]-, and diffusion [11]-based methods. All models are assessed by the CRPS, KLD, and violation rate of the rated-power and ramping constraints. As summarized in Table II, the proposed LLM-based method achieves the lowest CRPS and KLD among other methods, indicating improved probabilistic fidelity and better alignment with the statistical characteristics of ground-truth wind power trajectories. Moreover, its violation rate is significantly lower (2.85%), demonstrating enhanced compliance with physical limits. These results collectively show that the LLM captures nonlinear temporal patterns more effectively and produces diverse yet physically credible scenarios, outperforming other generative approaches.

To assess if the projection-back fallback in (18) reduces stochastic diversity during volatile icing events, a decoding sensitivity study without retraining is performed. We select high-volatility icing windows from the test set and generate 50 scenarios per window under three settings: Unconstrained, Default (paper), and Stricter (tighter ramp/cap). The overall violation rate (same as Table II), a compact diversity score (average cross-scenario variability over the horizon), and the projection rate (fraction of steps triggering projection-back) are reported in Table III. It can be shown that Default achieves a low violation rate with non-trivial diversity and a low projection rate, while stricter constraints further reduce the violation rate at the cost of reduced diversity and more frequent fallback.

TABLE III: Decoding Sensitivity on Volatile Icing Windows

Setting	$\mathrm{VR}_{\mathrm{total}}$ (%) ↓	Diversity ↑	Projection (%) ↓
Unconstrained	7.0	0.092	0.0
Default	3.3	0.079	1.2
Stricter	2.4	0.066	3.8

V. Conclusions

This paper proposes a physics-aware LLM framework for probabilistic wind power scenario generation under icing conditions. The proposed method captures nonlinear temporal dependencies and icing-induced degradation while maintaining physical consistency through physics-aware decoding. Case studies demonstrate superior probabilistic accuracy, statistical consistency, and physical realism compared with the VAE, CGAN, and diffusion-based baselines. Beyond producing high-fidelity probabilistic scenarios, the framework provides a generative foundation for post-disaster resilience analysis, supporting data-driven evaluation of recovery strategies, uncertainty-aware restoration optimization, and resilient operation planning under extreme weather.

References

[1] L. Wang, Y. He, Y. Zhou *et al.*, “A novel approach to wind turbine blade icing detection with limited sensor data via spatiotemporal attention siamese network,” *IEEE Trans. Ind. Informat.*, vol. 20, no. 6, pp. 8993–9005, 2024.

[2] M. H. Athari and Z. Wang, “Impacts of wind power uncertainty on grid vulnerability to cascading overload failures,” *IEEE Trans. Sustain. Energy*, vol. 9, no. 1, pp. 128–137, 2017.

[3] L. Xu, K. Feng, N. Lin *et al.*, “Resilience of renewable power systems under climate risks,” *Nat. Rev. Electr. Eng.*, vol. 1, no. 1, pp. 53–66, 2024.

[4] J. Li, J. Zhou, and B. Chen, “Review of wind power scenario generation methods for optimal operation of renewable energy systems,” *Appl. Energy*, vol. 280, p. 115992, 2020.

[5] M. Rayati, M. Bozorg, M. Carpita *et al.*, “Stochastic optimization and markov chain-based scenario generation for exploiting the underlying flexibilities of an active distribution network,” *Sustain. Energy Grids Netw.*, vol. 34, p. 100999, 2023.

[6] A. B. Krishna and A. R. Abhyankar, “Time-coupled day-ahead wind power scenario generation: A combined regular vine copula and variance reduction method,” *Energy*, vol. 265, p. 126173, 2023.

[7] Y. Chen, J. Wen, and S. Cheng, “Probabilistic load flow method based on nataf transformation and latin hypercube sampling,” *IEEE Trans. Sustain. Energy*, vol. 4, no. 2, pp. 294–301, 2012.

[8] Y. Chen, Y. Wang, D. Kirschen *et al.*, “Model-free renewable scenario generation using generative adversarial networks,” *IEEE Trans. Power Syst.*, vol. 33, no. 3, pp. 3265–3275, 2018.

[9] R. Yuan, B. Wang, Y. Sun *et al.*, “Conditional style-based generative adversarial networks for renewable scenario generation,” *IEEE Trans. Power Syst.*, vol. 38, no. 2, pp. 1281–1296, 2022.

[10] Z. Li, X. Peng, W. Cui *et al.*, “A novel scenario generation method of renewable energy using improved vaegan with controllable interpretable features,” *Appl. Energy*, vol. 363, p. 122905, 2024.

[11] X. Dong, Y. Sun, Y. Yang *et al.*, “Controllable renewable energy scenario generation based on pattern-guided diffusion models,” *Appl. Energy*, vol. 398, p. 126446, 2025.

[12] L. Wang, Y. He, Y. He *et al.*, “Wind turbine blade icing risk assessment considering power output predictions based on scso-ifcm clustering algorithm,” *Renew. Energy*, vol. 223, p. 119969, 2024.

[13] X. Zhang, R. R. Chowdhury, R. K. Gupta *et al.*, “Large language models for time series: A survey,” *arXiv preprint arXiv:2402.01801*, 2024.

[14] C. Gilbert, J. Browell, and D. McMillan, “Leveraging turbine-level data for improved probabilistic wind power forecasting,” *IEEE Trans. Sustain. Energy*, vol. 11, no. 3, pp. 1152–1160, 2019.

979-8-3315-8138-1