

Article

Generation of Non-Gaussian Rough Surfaces Using a PSD-Amplitude-Constrained Phase C-VAE

Jinyuan Wang ¹, Weilin Zhu ¹, Xiaoli Zhao ¹, Xiansong He ¹, Meile Wang ², Bo Yu ², Taowen Xiao ² and Jianyong Yao ^{1,*}

¹ School of Mechanical Engineering, Nanjing University of Science and Technology, Nanjing 210094, China; wangjny@njust.edu.cn (J.W.); zwl_njust@163.com (W.Z.); xlzhao@njust.edu.cn (X.Z.); brownsong@njust.edu.cn (X.H.)

² Hengli Precision Industry Co., Ltd., Changzhou 213177, China; lisa@henglihydraulics.com (M.W.); dean@henglihydraulics.com (B.Y.); xiaotaowen@henglihydraulics.com (T.X.)

* Correspondence: yaojianyong@njust.edu.cn

Abstract

The non-Gaussian height distribution and power spectral density (PSD) characteristics of rough surfaces have significant effects on the real contact area, local pressure distribution, oil-film formation, and friction and wear behavior of lubricated contact interfaces in mechanical components. Conventional methods for generating non-Gaussian rough surfaces commonly rely on iterative correction under explicit statistical constraints, which limits their computational efficiency in large-scale sample generation. To address this issue, this study proposes a PSD-amplitude-constrained phase conditional variational autoencoder (phase C-VAE) for generating non-Gaussian rough surfaces. Unlike conventional constructive methods that repeatedly correct surface samples under explicit statistical constraints, the proposed method learns the conditional distribution of the Fourier phase, while the spectral amplitude used for reconstruction is directly determined from the prescribed PSD. By taking the target skewness, kurtosis, and PSD as conditional inputs, the proposed method achieves joint control of higher-order statistical characteristics and spectral characteristics within a unified generative framework. Under target conditions derived from measured surfaces, the generated non-Gaussian rough surface samples achieved mean absolute relative errors of 0.056% and 0.044% for skewness and kurtosis, respectively, with a generation time of 24.62 s. These results indicate that the proposed method can effectively match the target skewness and kurtosis while maintaining good consistency between the generated surfaces and the target PSD. The proposed method alleviates the efficiency limitation of conventional constructive methods in the large-scale generation of non-Gaussian rough surface samples and provides an effective machine-learning-based generative approach for rapid batch modeling of rough surfaces in lubrication, friction, and contact analyses.



Academic Editor: Kazumasa Kawasaki

Received: 24 June 2026

Revised: 24 July 2026

Accepted: 1 August 2026

Published: 3 August 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.

Keywords: non-Gaussian rough surfaces; power spectral density; PSD amplitude constraint; conditional variational autoencoder; skewness; kurtosis

1. Introduction

Rough surfaces are widely present in functional transmission components such as rolling linear guides, planetary roller screws, and ball screws. Compared with ideal smooth surfaces, real rough-surface contact is generally characterized by localized multi-asperity contact. The peak–valley structure not only affects interfacial load transfer and stress

concentration, but also changes lubricant storage, flow, and oil-film formation, thereby influencing contact, friction, and wear analyses [1–4]. In addition, machined surfaces exhibit pronounced randomness. Even under the same machining process and with similar statistical parameters, the local peak–valley morphology may differ among different locations or samples. Rough surfaces can generally be described by Gaussian or non-Gaussian random processes. Engineering surfaces of practical transmission components, especially those produced by grinding and turning, commonly exhibit non-Gaussian characteristics [5]. Their height distributions are governed not only by spectral features but also by higher-order statistical moments such as skewness and kurtosis. Compared with Gaussian rough surfaces, non-Gaussian rough surfaces can more accurately represent the peaks, valleys, and tail-distribution characteristics of real surfaces, and therefore have greater theoretical and engineering significance in tribology and lubrication analysis.

Existing methods for modeling non-Gaussian rough surfaces can be broadly divided into statistical-feature-based methods and fractal-feature-based methods. Fractal-feature-based methods are generally limited in their ability to directly control skewness and kurtosis, and are therefore not well suited for generating non-Gaussian rough surfaces with prescribed higher-order statistical characteristics [6]. Among statistical-feature-based approaches, autoregressive moving average (ARMA) time-series methods and spectral representation methods have been widely used for rough surface modeling. Davies et al. [7] established the relationship between the skewness and kurtosis of an output non-Gaussian random process and those of the corresponding non-Gaussian random noise using a standard ARMA model. Watson et al. [8,9] further derived the inverse relationship between the output sequence and the input noise. Based on this relationship, the noise distribution can be inferred from the target rough surface distribution, and the corresponding non-Gaussian noise can be obtained using the Johnson transformation [10], thereby completing the generation of non-Gaussian rough surfaces. However, because the autoregressive (AR) parameters are difficult to determine, the AR model is commonly converted into an equivalent moving average (MA) representation [11]. Patir et al. [12] proposed a method for directly generating rough surfaces using an MA model, in which an exponential autocorrelation function (ACF) was adopted and the model coefficients were determined using the Newton–Raphson method. Based on this work, Hu et al. [13] developed a two-dimensional digital filtering method, in which the PSD is calculated using the fast Fourier transform (FFT), and the model coefficients are subsequently obtained through the inverse Fourier transform. This method has been widely applied in rough surface modeling. To improve the accuracy and efficiency of coefficient determination, Bakolas et al. [14] used a nonlinear conjugate gradient method and showed that the ACF truncation length affects the generated rough surface. Liao et al. [15] adopted a nonlinear least-squares method to solve the model coefficients and indicated that the truncation length should range from 2 to 20 times the correlation length. Minet et al. [16] used the Newton–Raphson method to determine the model coefficients and investigated the influence of ACF truncation length on simulated surfaces. On this basis, Francisco et al. [17] further proposed a high-accuracy method for generating non-Gaussian rough surfaces over a broad range of skewness and kurtosis values.

The spectral representation method treats a rough surface as a sampled realization of a random process. Wu [18,19] proposed an FFT-based iterative method, in which an iterative procedure consisting of initial sequence generation, input distribution calculation based on target statistical characteristics, phase determination using the Johnson transformation, and new sequence construction was established. This method enables mapping between the spatial and frequency domains of rough surfaces. By fixing the PSD in the frequency domain and iteratively correcting the phase, it can simultaneously control the spectral

characteristics and non-Gaussian statistical features of rough surfaces. Wang et al. [20] introduced the translation random process theory into rough surface generation. In this method, a Gaussian random surface is first generated using FFT and then transformed into a target non-Gaussian random field through cumulative distribution function mapping. This approach converts the non-Gaussian surface generation problem into a nonlinear mapping problem based on a Gaussian random process, thereby avoiding the complex solution of frequency-domain phases. As a result, it improves computational efficiency while enhancing the stability and applicability of the method.

Although existing statistical-feature-based modeling methods for non-Gaussian rough surfaces can control the PSD, skewness, and kurtosis, they remain constructive methods that rely on explicit statistical constraints and sample-by-sample iterative correction. Therefore, when multiple random surface samples are required under the same target conditions, their computational efficiency and batch-generation capability remain limited. To address this issue, this study proposes a PSD-amplitude-constrained phase C-VAE. Unlike the conventional spectral representation method and the genetic-algorithm-optimized PSD-based generation (GA-PSD) method, which construct each surface sample by iteratively correcting its skewness and kurtosis, the proposed method learns the conditional distribution of the Fourier phase, while the spectral amplitude used for reconstruction is determined from the prescribed PSD. The target PSD, skewness, and kurtosis are jointly introduced as conditional inputs, and the higher-order statistical constraint loss and mode-seeking loss are incorporated to improve statistical matching accuracy and maintain sample diversity. Therefore, the main novelty of this study lies in integrating phase-based conditional generation with an explicit PSD amplitude constraint, thereby enabling the rapid generation of multiple non-Gaussian rough surface samples under prescribed statistical conditions.

The present study focuses on the generation of one-dimensional non-Gaussian rough surface profiles with prescribed PSD, skewness, and kurtosis. The conventional spectral representation method and the GA-PSD method are adopted as comparison methods to evaluate the proposed method in terms of statistical feature matching accuracy, computational efficiency, and sample diversity. The remainder of this paper is organized as follows: Section 2 introduces the conventional spectral representation method and its genetic-algorithm-based improvement of the error-iteration process. Section 3 presents the proposed machine-learning-based rough surface generation method. Section 4 presents the experimental setup; compares the three methods in terms of statistical accuracy, height-distribution consistency, and computational efficiency; and further evaluates the repeatability, adaptability to different target statistical conditions, and sample diversity of the proposed method. Hyperparameter sensitivity and ablation analyses are also reported.

2. Conventional Non-Gaussian Rough Surface Generation Methods

2.1. Spectral Representation Method

The spatial profile of a one-dimensional rough surface can be represented as a stationary random process with statistical characteristics [21]. The spectral representation method is a random-process construction method based on Fourier spectral decomposition. In this method, the statistical characteristics of a rough surface are constrained by introducing the target PSD and ACF in the frequency domain, and the frequency-domain information is then transformed back into the spatial domain using the inverse Fourier transform to generate a rough surface with prescribed spectral characteristics. For non-Gaussian rough surfaces, higher-order statistical characteristics such as skewness and kurtosis need to be controlled. Therefore, the Johnson transformation system is introduced into the spectral representation method to jointly control the probability distribution and higher-order statistical moments of rough surfaces. Assume that a zero-mean discrete height sequence

of a one-dimensional non-Gaussian rough surface is given as $\{z_r \mid r = 0, 1, 2, \dots, N-1\}$. The discrete PSD of the one-dimensional non-Gaussian rough surface can be obtained by applying the FFT to the discretized ACF. The target discrete PSD is expressed as:

$$R_r = \frac{1}{N} \sum_{s=0}^{N-1} z_s z_{s+r}, \quad r = 0, 1, 2, \dots, N-1 \quad (1)$$

$$P_k = \frac{1}{N} \sum_{r=0}^{N-1} R_r \exp\left(-\frac{i2\pi kr}{N}\right), \quad k = 0, 1, 2, \dots, N-1 \quad (2)$$

where R_r and P_k denote the discrete autocorrelation function and the discrete PSD of the one-dimensional non-Gaussian rough surface, respectively. N is the number of sampling points of the non-Gaussian rough surface z_r . The autocorrelation function is used to describe the similarity between the rough surface and itself at different sampling points. In this paper, circular autocorrelation is adopted; that is, when $s+r > N$, $z_{s+r} = z_{s+r-N}$. In addition, to characterize the higher-order statistical features of the non-Gaussian rough surface, the target skewness Sk_t and kurtosis K_t are defined as follows:

$$Sk_t = \frac{\frac{1}{N} \sum_{r=0}^{N-1} z_r^3}{\sigma^3}, \quad K_t = \frac{\frac{1}{N} \sum_{r=0}^{N-1} z_r^4}{\sigma^4}. \quad (3)$$

where σ denotes the standard deviation of the rough surface height sequence. Because the height sequence is assumed to have zero mean, i.e., $\mu = 0$, the centered height $z_r - \mu$ reduces to z_r .

During spectral shaping, the higher-order moments of the Johnson-transformed non-Gaussian sequence may deviate from the prescribed skewness and kurtosis. Therefore, correction parameters are introduced for the target skewness and kurtosis before applying the Johnson transformation. Let $Sk_1 = Sk_t$ and $K_1 = K_t$, and map Sk_1 and K_1 to the skewness Sk_n and kurtosis K_n of a non-Gaussian random sequence:

$$Sk_n = \frac{\left(\sum_{r=0}^{m-1} h_r^2\right)^{3/2}}{\sum_{r=0}^{m-1} h_r^3} Sk_1, \quad K_n = \frac{K_1 \left[\left(\sum_{r=0}^{m-1} h_r^2\right)^2 - 6 \sum_{r=0}^{m-2} \sum_{p=r+1}^{m-1} h_r^2 h_p^2 \right]}{\sum_{r=0}^{m-1} h_r^4} \quad (4)$$

where h_r and h_p are the weighting coefficients for linear spectral shaping.

The Johnson system is a distribution construction method based on monotonic and invertible variable transformations. Its core idea is to establish a one-to-one relationship between the target random variable and a standard normal variable through a parameterized transformation, thereby generating a non-Gaussian random sequence with prescribed skewness and kurtosis. By taking Sk_n and K_n as the target skewness and kurtosis inputs of the Johnson transformation, respectively, the Johnson transformation system is used to obtain a one-dimensional non-Gaussian rough surface z_r that satisfies the target kurtosis and skewness.

$$\phi'_k = \tan^{-1} \left(\frac{\sum_{r=0}^{N-1} \eta'_r \sin(2\pi kr/N)}{\sum_{r=0}^{N-1} \eta'_r \cos(2\pi kr/N)} \right) \quad (5)$$

$$z_t = \sum_{k=0}^{N-1} \sqrt{P_k} \exp \left\{ i2\pi \left[\phi'_k + \frac{kr}{N} \right] \right\}, \quad r = 0, 1, 2, \dots, (N-1) \quad (6)$$

According to Equations (4)–(6), an updated rough surface z_t can be generated, and its skewness Sk_2 and kurtosis K_2 can be calculated. The errors between (Sk_2, K_2) and the initial target values (Sk_t, K_t) are then used to iteratively correct Sk_1 and K_1 as follows:

$$\begin{aligned} Sk_1 &= Sk_1 + \alpha_1(Sk_t - Sk_2), \\ K_1 &= K_1 + \alpha_2(K_t - K_2). \end{aligned} \quad (7)$$

The iteration is repeated until the prescribed convergence tolerance is satisfied. Once this criterion is met, z_t is regarded as the required one-dimensional random rough surface.

2.2. Spectral Representation Method with Genetic-Algorithm-Optimized Iteration

In the iterative stage of the spectral representation method, only a simple linear iteration is adopted, which leads to limitations in both efficiency and accuracy. To address this issue, a spectral representation method with genetic-algorithm-optimized iteration is introduced for non-Gaussian rough surface generation. Based on the original power-spectrum-based generation framework, the spectral representation method and the Johnson transformation are retained, while the simple linear iteration of skewness and kurtosis is replaced by a genetic algorithm search.

Taking Sk_1 and K_1 as the optimization variables, the deviations between the actual skewness and kurtosis of the generated sample and their target values are used as the optimization objective. The objective function is expressed as:

$$J(Sk_1, K_1) = w_1 \left(\frac{Sk_2 - Sk_t}{Sk_t + \varepsilon} \right)^2 + w_2 \left(\frac{K_2 - K_t}{K_t + \varepsilon} \right)^2 \quad (8)$$

where w_1 and w_2 are weighting coefficients, and ε is a small constant introduced to prevent division by zero. The optimization aims to minimize the objective function. A smaller objective function value indicates that the higher-order statistical characteristics of the generated sample corresponding to the current parameters are closer to the target values.

Assume that the i -th individual in the population of the g -th generation is expressed as:

$$(Sk_{1,i}^{(g)}, K_{1,i}^{(g)}), \quad i = 1, 2, \dots, N_p \quad (9)$$

Let the parent individuals in the genetic optimization process be defined as $(Sk_{1,a}, K_{1,a})$ and $(Sk_{1,b}, K_{1,b})$. After crossover, the offspring individual can be expressed as:

$$\begin{cases} Sk_{1,c} = \lambda Sk_{1,a} + (1 - \lambda) Sk_{1,b}, \\ K_{1,c} = \lambda K_{1,a} + (1 - \lambda) K_{1,b}. \end{cases} \quad (10)$$

Once the optimization objective is reached, the genetic search is terminated. The correction parameters Sk_1 and K_1 corresponding to the current best individual, together with the generated rough surface sample, are recorded, and this sample is taken as the final generated non-Gaussian rough surface.

Based on the spectral representation method, a genetic algorithm is introduced to iteratively optimize the correction parameters for skewness and kurtosis, which helps improve the stability and adaptability of the skewness and kurtosis matching process. However, this method is still essentially a constructive rough surface generation method based on explicit statistical constraints. Its core procedure is to construct or iteratively solve surface samples that satisfy the prescribed target statistical conditions. Therefore, when

a large number of rough surface samples need to be rapidly generated under the same statistical conditions, this method still has certain limitations in terms of computational efficiency, conditional mapping flexibility, and batch generation capability.

3. Improved Phase C-VAE with PSD Amplitude Constraint for Non-Gaussian Rough Surface Generation

To address the limited computational efficiency of explicit constructive methods for non-Gaussian rough surface generation in large-scale sample generation scenarios, this study proposes a phase C-VAE method with a PSD amplitude constraint for generating non-Gaussian rough surfaces. The proposed method combines an explicit target PSD amplitude constraint with phase-based conditional generation. On the one hand, the frequency-domain amplitude is determined from the target PSD, thereby explicitly constraining the spectral structure of the generated surface. On the other hand, the target skewness, kurtosis, and normalized PSD are used as conditional inputs, and a C-VAE with a conditional prior is employed to learn the frequency-domain phase generation pattern under prescribed statistical conditions. In the generation stage, the model first samples the latent variable according to the target statistical conditions and generates the frequency-domain phase. The generated phase is then combined with the target PSD amplitude to construct the complex spectrum, which is finally transformed using the inverse fast Fourier transform (IFFT) to obtain the spatial-domain non-Gaussian rough surface. The generation workflows of the GA-PSD method and the proposed phase C-VAE method are shown in Figure 1.




Figure 1. Flowcharts of the GA-PSD method and the phase C-VAE method for generating non-Gaussian rough surfaces.

3.1. Network Structure of the Phase C-VAE with a Conditional Prior

Let x denote the standardized rough surface sequence sample; c denote the condition vector composed of the target skewness, kurtosis, and PSD features; and z denote the latent variable. The condition vector is constructed as follows:

$$c = [Sk_t, K_t, P_t] \quad (11)$$

where P_t denotes the standardized target PSD feature vector obtained through FFT-based discretization.

Considering the significant differences in dimensionality and physical meaning between skewness, kurtosis, and PSD features, a branched conditional encoding structure is adopted in this study to represent the condition vector. Specifically, skewness and kurtosis are encoded as low-dimensional statistical features, while the PSD features are encoded through a spectral feature branch. The skewness and kurtosis are fed into the low-dimensional statistical feature encoding branch, whereas the PSD features are fed into the spectral feature encoding branch. The encoded features from the two branches are then fused to obtain a unified conditional embedding:

$$e_c = E_c(c) \quad (12)$$

where $E_c(\cdot)$ denotes the condition encoder, and e_c is the conditional embedding vector. This design helps enhance the model's representation capability for heterogeneous statistical conditions and provides a unified input for subsequent posterior encoding and conditional prior modeling.

In the training stage, the sample x and the conditional embedding e_c are jointly fed into the posterior encoder to estimate the approximate posterior distribution of the latent variable:

$$q_\psi(z | x, e_c) \quad (13)$$

where ψ denotes the network parameters of the posterior encoder. The approximate posterior distribution is used to extract latent representations related to the random local morphological variations from real rough surface samples. It should be noted that the latent variable in this study is not directly used for spatial-domain height sequence regression. Instead, it is used to control the random variations in the frequency-domain phase distribution under the same target skewness, kurtosis, and PSD conditions, thereby further affecting the spatial-domain rough surface morphology obtained after IFFT.

To enable the model to generate samples without relying on real sample inputs during the generation stage, a conditional prior network is further introduced in this study to learn the conditional prior distribution of the latent variable from the conditional embedding:

$$p_\eta(z | e_c) \quad (14)$$

where η denotes the parameters of the conditional prior network. The conditional prior network enables the latent-variable distribution to adaptively vary with the target skewness, kurtosis, and PSD conditions. During training, the consistency between the approximate posterior distribution and the conditional prior distribution is constrained, allowing the trained model to directly sample latent variables from the target statistical conditions and thereby achieve conditional generation.

To make the random sampling process differentiable for gradient backpropagation, the reparameterization trick is adopted to sample the latent variable during training:

$$z_q = \mu_q(x, e_c) + \sigma_q(x, e_c) \odot \zeta, \quad \zeta \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \quad (15)$$

where ζ is a random noise vector following the standard normal distribution, and $\odot$ denotes element-wise multiplication. Correspondingly, in the generation stage, the latent variable is sampled from the conditional prior distribution:

$$z_p = \mu_p(e_c) + \sigma_p(e_c) \odot \zeta, \quad \zeta \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \quad (16)$$

During training, the model learns the local random morphological information of real samples through the posterior path, whereas during generation, latent-variable sampling is

directly performed according to the target statistical conditions through the conditional prior path.

The input to the phase decoder consists of the latent variable and the encoded conditional information. In the detailed implementation, the encoded conditional information includes the fused conditional embedding and the 32-dimensional skewness–kurtosis feature described in Table 1. For concise notation, these condition-related inputs are collectively denoted by e_c in the decoder mapping. In this study, the C-VAE decoder is defined as a frequency-domain phase generator. Its direct output is not the spatial-domain rough surface height sequence, but the frequency-domain phase under the target PSD amplitude constraint. To avoid confusion between the roles of the conditional prior network and the decoder, the conditional prior network is denoted as $p_\eta(z | e_c)$, which describes the distribution of the latent variable under given conditions, whereas the decoder is represented as a deterministic phase mapping:

$$\hat{\phi} = D_\theta(z, e_c) \quad (17)$$

During training, the decoder can receive the posterior-sampled latent variable z_q to form the posterior reconstruction path. During generation, the decoder receives the conditional-prior-sampled latent variable z_p to form the conditional generation path.

Based on the generated phase and the frequency-domain amplitude determined by the target PSD, the complex spectrum is constructed as follows:

$$\hat{X}_k = A_k \exp(i\hat{\phi}_k) \quad (18)$$

where $A_k = \sqrt{P_k}$ is the amplitude of the k -th frequency component determined by the target PSD, and $\hat{\phi}_k$ is the phase of the k -th frequency component generated by the decoder. To ensure that the rough surface obtained after IFFT is real-valued, the complex spectrum should satisfy the conjugate symmetry condition:

$$\hat{X}_{N-k} = \hat{X}_k^* \quad (19)$$

where $(\cdot)^*$ denotes the complex conjugate. Accordingly, the complete complex spectrum $\hat{X}$ is constructed and then transformed by IFFT to obtain the spatial-domain rough surface:

$$\hat{x} = \text{Re}\{\text{IFFT}(\hat{X})\} \quad (20)$$

To address the inconsistency between the representation forms of the generated samples and the training samples, as well as the interference caused by differences in absolute amplitude scale during model learning, the spatial-domain rough surface obtained after IFFT is normalized to have zero mean and unit standard deviation:

$$\tilde{x} = \frac{\hat{x} - \mu_{\hat{x}}}{\sigma_{\hat{x}} + \varepsilon} \quad (21)$$

Zero-mean and unit-standard-deviation normalization ensures consistency in representation between the generated samples and the training samples. It weakens the influence of the absolute amplitude scale of rough surfaces on network training, enabling the model to focus more on the PSD structure and higher-order statistical features themselves, thereby improving its ability to learn and represent the target statistical constraints.

In this study, the decoder directly outputs the frequency-domain phase, whereas the spatial-domain rough surface is the final output obtained by combining the generated phase with the target PSD amplitude in the frequency domain followed by IFFT. The subsequent reconstruction loss, higher-order statistical constraint loss, and mode-seeking loss do not act directly on the frequency-domain phase, but instead act on the normalized spatial-domain

rough surface. To address the problem that the decoder may rely too strongly on the conditional input, leading to an insufficient role of the latent variable in representing local morphological randomness, a dropout mechanism is introduced into the conditional embedding. By randomly deactivating part of the conditional embedding, the model is forced to use the latent variable for phase generation when part of the conditional information is suppressed. This helps alleviate the problem of latent-variable collapse and improves the ability of the model to express sample diversity under the same statistical conditions.

All neural-network modules in the proposed phase C-VAE were implemented using fully connected layers. For a standardized surface profile of length $L = 1025$, the real-valued fast Fourier transform contains $M = L/2 + 1 = 513$ frequency components. Therefore, the condition vector consisted of the target skewness, target kurtosis, and a 513-dimensional standardized PSD feature vector, resulting in a total dimension of 515. Because the higher-order statistical parameters and PSD features have different dimensions and numerical scales, they were first encoded separately and then fused into a 256-dimensional conditional embedding. Both the posterior and conditional prior distributions were parameterized using a 16-dimensional latent variable, and the decoder produced a 513-dimensional phase vector. During training, dropout with a rate of 0.20 was applied to the conditional embedding before it was provided to the phase decoder. The numbers of layers and neurons and the activation functions of the main network components are summarized in Table 1.

Table 1. Detailed network architecture of the proposed phase C-VAE.

Main Component	Layers and Numbers of Neurons	Activation Function
Condition encoder	The two-dimensional skewness–kurtosis input is processed by two fully connected hidden layers with 64 neurons each; the PSD input is processed by two fully connected hidden layers with 256 neurons each; the extracted features are then fused by one fully connected layer with 256 neurons.	ReLU for all hidden layers
Posterior encoder	The surface profile is first processed by one fully connected hidden layer with 128 neurons; after fusion with the conditional embedding, it is processed by one fully connected hidden layer with 256 neurons; two independent linear output layers with 16 neurons each subsequently produce the latent mean and log-variance.	ReLU for the hidden layers and linear activation for the output layers
Conditional prior network	The conditional embedding is processed by one fully connected hidden layer with 256 neurons, followed by two independent linear output layers with 16 neurons each that produce the mean and log-variance of the conditional prior.	ReLU for the hidden layer and linear activation for the output layers
Phase decoder	The target skewness and kurtosis are first projected into a 32-dimensional feature. After concatenation with the latent variable and conditional embedding, the resulting feature is processed by two fully connected hidden layers with 512 neurons each, followed by a 513-neuron output layer.	ReLU for the hidden layers and $\pi \tanh(\cdot)$ for the output layer

3.2. Multi-Objective Loss Function and Training–Generation Strategy

To meet the training requirements of the phase C-VAE with a PSD amplitude constraint for statistically constrained generation of non-Gaussian rough surfaces, a multi-objective loss function is constructed in this study. The loss function consists of a reconstruction loss, a conditional prior alignment loss, a higher-order statistical constraint loss, and a mode-seeking loss. A progressive training strategy is also adopted to improve training stability. Since the frequency-domain structure of the generated surface is mainly achieved

through the explicit constraint imposed by the target PSD amplitude, PSD is not included as an independent optimization loss term during training. This study focuses on evaluating the matching accuracy of higher-order statistical features, generation efficiency, and sample diversity under the same statistical conditions, while the PSD constraint is ensured at the methodological level through the frequency-domain amplitude construction process.

Let $\tilde{x}_{\text{post}}$ denote the reconstructed surface obtained by applying IFFT to the complex spectrum. The reconstruction loss does not directly constrain the frequency-domain phase output by the decoder, but instead constrains the consistency between the spatial-domain reconstructed surface generated from the phase and the input standardized surface. It is defined as follows:

$$L_{\text{rec}} = \sum_{i=1}^N (x_i - \tilde{x}_{\text{post},i})^2. \quad (22)$$

where N is the sequence length of the rough surface sample, and x_i and $\tilde{x}_{\text{post},i}$ denote the values of the input sample and the reconstructed sample at the i -th point, respectively. This loss is used to constrain the model to recover the basic morphological features of the input rough surface in the spatial domain.

To maintain consistency between the approximate posterior distribution learned during training and the conditional prior distribution used during generation, a conditional prior alignment loss is introduced in this study. By constraining the discrepancy between the approximate posterior distribution and the conditional prior distribution, the model can directly sample latent variables from the conditional prior after training, without relying on real sample inputs. This loss is expressed using the Kullback–Leibler (KL) divergence:

$$L_{\text{KL}} = D_{\text{KL}}[q_{\psi}(z | \mathbf{x}, \mathbf{e}_c) \| p_{\eta}(z | \mathbf{e}_c)] \quad (23)$$

The reconstruction loss and conditional prior alignment loss mainly constrain the sample reconstruction capability and the consistency of the latent-space distribution, but they cannot directly ensure that the higher-order statistical features of the generated surfaces satisfy the target skewness and kurtosis constraints. Therefore, a higher-order statistical constraint loss is further introduced for the posterior path to explicitly constrain the consistency between the higher-order statistical features of the posterior reconstructed samples and the target statistical conditions. Let w_{sk} and w_k denote the weighting coefficients of the skewness and kurtosis terms, respectively, and let s_k denote the normalization scale for the kurtosis error. The higher-order statistical constraint loss for the posterior path can be expressed as:

$$L_{\text{mom}}^{\text{post}} = w_{sk} \left(\frac{\text{Sk}(\tilde{x}_{\text{post}}) - Sk_t}{|Sk_t| + \varepsilon} \right)^2 + w_k \left(\frac{\text{K}(\tilde{x}_{\text{post}}) - K_t}{s_k} \right)^2 \quad (24)$$

To improve the ability of the model to directly generate samples satisfying the target statistical constraints by sampling from the conditional prior during the generation stage, a higher-order statistical constraint loss is further introduced into the conditional prior path. Under the same conditional embedding, two sets of latent variables, $\mathbf{z}_{p,1}$ and $\mathbf{z}_{p,2}$, are sampled from the conditional prior distribution. After decoder-based phase generation, PSD amplitude synthesis, IFFT, and normalization, two generated surfaces from the prior path, $\tilde{x}_{\text{prior},1}$ and $\tilde{x}_{\text{prior},2}$, are obtained. The higher-order statistical constraint loss for the prior path is defined as:

$$L_{\text{mom}}^{\text{prior}} = \frac{1}{2} \sum_{m=1}^2 \left[w_{sk} \left(\frac{\text{Sk}(\tilde{x}_{\text{prior},m}) - Sk_t}{|Sk_t| + \varepsilon} \right)^2 + w_k \left(\frac{\text{K}(\tilde{x}_{\text{prior},m}) - K_t}{s_k} \right)^2 \right] \quad (25)$$

By combining the constraint losses from the posterior path and the prior path, the overall higher-order statistical constraint loss is constructed as follows:

$$L_{\text{mom}} = \frac{1}{2} (L_{\text{mom}}^{\text{post}} + L_{\text{mom}}^{\text{prior}}) \quad (26)$$

This loss simultaneously constrains the higher-order statistical features of the posterior reconstructed samples and the conditional-prior-generated samples, which helps improve the responsiveness of the model to the target skewness and kurtosis.

To prevent the decoder from ignoring random latent perturbations and consequently generating overly similar surfaces under the same statistical conditions, a mode-seeking loss is introduced. This loss follows the basic idea of mode-seeking regularization [22], which increases the difference between generated outputs induced by a unit random latent perturbation. Under the same conditional embedding e_c , two independent standard Gaussian noise vectors, ζ_1 and ζ_2 , are sampled, and

$$\begin{aligned} \zeta_k &= T_p \zeta_k, \\ z_{p,k} &= \mu_p(e_c) + \sigma_p(e_c) \odot \zeta_k, \quad k = 1, 2, \end{aligned} \quad (27)$$

where $T_p = 0.8$ is the prior-sampling temperature used during training. The two conditional-prior latent variables are processed through phase decoding, target-PSD amplitude synthesis, IFFT, and normalization to obtain two generated surfaces, $\tilde{x}_{\text{prior},1}$ and $\tilde{x}_{\text{prior},2}$. The random-perturbation distance and generated-surface distance are respectively defined as

$$\begin{aligned} D_{\zeta} &= \frac{1}{d} \sum_{j=1}^d |\zeta_{1,j} - \zeta_{2,j}|, \\ D_x &= \frac{1}{N} \sum_{i=1}^N |\tilde{x}_{\text{prior},1,i} - \tilde{x}_{\text{prior},2,i}|. \end{aligned} \quad (28)$$

Because the desired objective is to increase D_x/D_{ζ} , whereas model training is performed through loss minimization, its reciprocal is used to construct the mode-seeking loss:

$$L_{\text{ms}} = \mathbb{E} \left[\frac{D_{\zeta}}{\max(D_x, \epsilon_{\text{ms}})} \right], \quad \epsilon_{\text{ms}} = 10^{-6}. \quad (29)$$

where d and N denote the latent dimension and the number of surface sampling points, respectively. The mean L_1 distances eliminate the influence of the different dimensions on the distance magnitudes. For a given pair of random perturbations, minimizing L_{ms} encourages an increase in D_x , thereby preventing different random perturbations from producing nearly identical surfaces. Because the two surfaces share the same statistical conditions and target PSD amplitudes, their differences mainly correspond to changes in local peak–valley positions and morphological fluctuations induced by phase variations. Therefore, this loss improves local morphological diversity under the same statistical conditions, while the higher-order statistical constraints and PSD amplitude constraint ensure that the generated surfaces continue to satisfy the prescribed statistical features.

Based on the above loss terms, the overall loss function is constructed as follows:

$$L = \lambda_{\text{rec}} L_{\text{rec}} + \beta L_{\text{KL}} + \lambda_{\text{mom}} L_{\text{mom}} + \lambda_{\text{ms}} L_{\text{ms}} \quad (30)$$

where λ_{rec} , β , λ_{mom} , and λ_{ms} denote the coefficients of the corresponding loss terms, respectively.

According to the roles of the individual loss terms during training, their corresponding weights were scheduled over different training stages. Both the conditional prior alignment

weight β and the higher-order statistical constraint weight λ_{mom} were linearly increased to 1.0 over the first 30 epochs, thereby progressively strengthening the latent-distribution regularization and statistical-moment constraints. The reconstruction-loss weight λ_{rec} was maintained at 1.0 during the first 20 epochs and then gradually reduced to 0.20 using cosine annealing between epochs 20 and 80, after which it remained unchanged. A fixed weight of $\lambda_{\text{ms}} = 0.10$ was assigned to the mode-seeking loss to improve the diversity of the generated samples without allowing this term to dominate the overall optimization. These weight settings enabled the model to balance surface reconstruction, latent-space regularization, higher-order statistical feature preservation, and generated-sample diversity.

4. Experimental Validation and Results Analysis

4.1. Experimental Setup and Evaluation Metrics

The complete dataset consisted of 10,000 numerically generated rough surface samples and 100 measured rough surface profile samples. The numerical samples were obtained using the non-Gaussian rough surface generation method proposed by Chen et al. [23], whereas the measured samples were extracted from ground-surface profiles acquired from rolling linear guide pairs. The 100 measured samples were fixed-length one-dimensional profile samples obtained by segmenting the original measured profiles using a unified sampling length. The numerical samples were mainly used to increase the training-data volume and extend the range of statistical conditions, whereas the measured samples introduced the morphological characteristics of actual ground surfaces. For each sample, the skewness, kurtosis, and PSD were calculated, and the height sequence was standardized to have zero mean and unit standard deviation. The standardized PSD was used as the spectral condition input, whereas the corresponding unstandardized PSD was used to determine the target frequency-domain amplitude during generation. After preprocessing, the numerical and measured samples were separately divided into training, validation, and test sets at a ratio of 8:1:1, with segmented samples originating from the same measured profile assigned to the same data subset. The training set was used for model-parameter optimization, the validation set was used for model selection, and the test set was used only for final performance evaluation.

Based on the constructed dataset, the proposed method was trained. The model inputs consisted of normalized rough surface samples and their corresponding statistical condition vectors. The condition vector was first transformed into a conditional embedding through the branched condition encoder. The posterior encoder then sampled the latent variable based on x and e_c , and the decoder output the frequency-domain phase. Subsequently, the generated phase was combined with the frequency-domain amplitude determined by the target PSD to construct the complex spectrum, which was further processed by IFFT and normalization to obtain the reconstructed surface. The model was trained for 100 epochs using the Adam optimizer with a batch size of 256. The complete training hyperparameters and loss-weight settings are summarized in Table 2. Model training and inference were performed on a computing platform equipped with an Intel Core i7-13700KF central processing unit (CPU) and an NVIDIA GeForce RTX 4060 Ti graphics processing unit (GPU).

To quantitatively evaluate the performance of the proposed method in generating non-Gaussian rough surfaces, an evaluation framework is established from four aspects: higher-order statistical feature matching accuracy, height-distribution similarity, computational efficiency, and sample diversity. The relative errors of skewness and kurtosis are used to characterize the differences in higher-order statistical features between the generated samples and the target values. To further evaluate the consistency of the overall height probability distributions, the first-order Wasserstein distance W_1 is introduced. Let $\tilde{h}_{(i)}^t$ and

$\tilde{h}_{(i)}^g$ denote the i -th order statistics of the standardized height samples of the target and generated surfaces, respectively. For the target and generated surfaces, each containing N samples, the first-order Wasserstein distance is defined as:

$$W_1 = \frac{1}{N} \sum_{i=1}^N \left| \tilde{h}_{(i)}^g - \tilde{h}_{(i)}^t \right|. \quad (31)$$

Table 2. Values and Settings of Network Training Hyperparameters.

Parameter Name	Symbol	Value or Setting
Optimizer	—	Adam, with $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1.0 \times 10^{-8}$, and a weight decay of 0
Learning rate	l_r	Initial value of 1.0×10^{-3} ; multiplied by 0.3 after epochs 70 and 90
Training epochs	epochs	100
Batch size	batch_size	256
Gradient clipping	$g_{\max}$	Maximum L_2 norm of 5.0
Reconstruction-loss weight	λ_{rec}	1.0 for the first 20 epochs; cosine-annealed to 0.20 between epochs 20 and 80; maintained at 0.20 thereafter
Conditional prior alignment weight	β	Linearly increased to 1.0 over the first 30 epochs
Skewness and kurtosis error weights	w_{sk}/w_k	8.0/4.0
Higher-order statistical constraint weight	λ_{mom}	Linearly increased to 1.0 over the first 30 epochs
Mode-seeking loss weight	λ_{ms}	Fixed at 0.10

A smaller W_1 indicates that the height probability distribution of the generated surface is closer to that of the target surface. The sample generation time is recorded to evaluate the computational efficiency of different methods, while sample diversity is evaluated using the pointwise standard deviation and pairwise RMSE in Section 4.3. Under the same evaluation framework, the conventional spectral representation method and the GA-PSD method are adopted as comparison methods. In the GA-PSD method, the error-iteration procedure is improved using a genetic algorithm.

4.2. Comparison and Evaluation of Generation Performance

To verify the effectiveness and advantages of the proposed method in generating non-Gaussian rough surfaces, a representative rough surface profile was selected from the measured test set that was not involved in either model training or model selection, and its statistical features were used as the target conditions. Under the same target skewness, kurtosis, and PSD conditions, the conventional spectral representation method, the GA-PSD method, and the phase C-VAE method with a PSD amplitude constraint were each used to generate one rough surface sample. The results are shown in Figure 2.




Figure 2. Comparison of the measured target surface and the rough surface samples generated by the three methods.

To further evaluate the height-distribution similarity of the generated surfaces, the empirical cumulative distribution functions of the standardized heights of the measured target surface and the three generated surfaces were compared, as shown in Figure 3. Under the same target conditions, the W_1 values of the conventional spectral representation method, GA-PSD method, and phase C-VAE method with a PSD amplitude constraint were 0.0627, 0.0640, and 0.0660, respectively. The results show that the empirical CDF curves of the three generated surfaces are close to that of the target surface and that the corresponding W_1 values are similar, indicating that all three methods can effectively preserve the target height probability distribution.




Figure 3. Empirical cumulative distribution functions of the standardized heights of the measured target surface and the rough surfaces generated by the three methods.

Based on the rough surface samples generated by the three methods, the generation time was recorded, and the skewness and kurtosis of the generated samples were calculated. The three rough surface generation methods were compared in terms of higher-order

statistical feature matching and generation efficiency. Under the same target conditions, the relative skewness errors of the conventional spectral representation method, GA-PSD method, and phase C-VAE method with a PSD amplitude constraint were 0.11%, 0.012%, and 0.0014%, respectively. The corresponding relative kurtosis errors were 0.026%, 0.019%, and 0.020%, respectively, and the generation times were 98.26 s, 797.8 s, and 24.62 s, respectively. The corresponding W_1 results are also listed in Table 3, while the relative skewness and kurtosis errors are compared in Figure 4.

The results show that the skewness and kurtosis errors of the proposed phase C-VAE method remain at low levels. Compared with the conventional spectral representation method, the skewness and kurtosis errors are reduced by approximately 98.73% and 23.08%, respectively. Compared with the GA-PSD method, the skewness error is reduced by 88.33%, whereas the kurtosis error increases slightly by 5.26%. These results indicate that the proposed phase C-VAE method can effectively preserve the target higher-order statistical features. Although the W_1 value of the phase C-VAE method is slightly higher than those of the two comparison methods, the differences among the three methods are small. Combined with the high consistency among the empirical CDF curves, this result indicates that the proposed method achieves height-distribution matching performance comparable to that of the two comparison methods. Meanwhile, its generation time is only 24.62 s, which is shorter than those of the conventional spectral representation method and the GA-PSD method, demonstrating a clear advantage in computational efficiency. Overall, the proposed method achieves high statistical feature matching accuracy while enabling the rapid generation of non-Gaussian rough surfaces, making it suitable for friction, lubrication, and contact analyses that require the batch generation of rough surface samples.

Table 3. Higher-order statistical errors, first-order Wasserstein distances, and generation times of rough surface samples generated by the three methods.

Metric	Measured	Power Spectrum Method	GA-PSD Method	Phase C-VAE Method
Skewness	−0.38185	−0.38144	−0.3818	−0.3818
Relative error of skewness	0	0.11%	0.012%	0.0014%
Kurtosis	3.7092	3.7082	3.7085	3.7084
Relative error of kurtosis	0	0.026%	0.019%	0.020%
First-order Wasserstein distance W_1	0	0.0627	0.0640	0.0660
Computation time (s)	/	98.26	797.8	24.62

To further evaluate the repeatability of the proposed stochastic generative model under different random seeds, five independent random seeds were used to generate five non-Gaussian rough surface samples while keeping the target PSD, skewness, kurtosis, and trained model parameters unchanged. For each generated sample, the relative skewness error, relative kurtosis error, and first-order Wasserstein distance were calculated, and the results were summarized in the form of mean $\pm$ standard deviation. As shown in Figure 5, the absolute relative skewness error over the five independent generations was $(0.00257 \pm 0.00093)\%$, the absolute relative kurtosis error was $(0.02109 \pm 0.00286)\%$, and the first-order Wasserstein distance was 0.07368 ± 0.00833 . These results indicate that the proposed phase C-VAE method can consistently maintain low higher-order statistical errors and stable height-distribution matching performance under different random seeds, demonstrating good stochastic generation repeatability.




Figure 4. Comparison of higher-order statistical errors for rough surfaces generated by the three methods.




Figure 5. Repeatability evaluation of the phase C-VAE method over five independent random seeds.

To further evaluate the adaptability of the proposed model to different higher-order statistical conditions, Pawlus et al. reported that, for the ground surface topographies examined in their study, the skewness S_{sk} ranged from -0.85 to -0.10 , while the kurtosis S_{ku} ranged from 3.2 to 6.0 , with a strong inverse correlation between these two parameters ($R = -0.82$) [24]. In the present study, these reported areal-parameter ranges were used as reference bounds for defining the statistical conditions of one-dimensional profiles. Four target conditions were selected along the reported joint variation trend of skewness and kurtosis: $(-0.85, 6.0)$, $(-0.60, 5.0)$, $(-0.35, 4.0)$, and $(-0.10, 3.2)$. In all four tests, the model architecture and parameters were kept unchanged, without retraining or parameter adjustment. Only the target skewness and kurtosis were varied in each test, and one rough surface profile was generated for each condition. The generated results and their relative errors are summarized in Table 4.

Table 4. Skewness, kurtosis, and corresponding relative errors of rough surface profiles generated by the phase C-VAE under different target conditions.

Target Skewness	Target Kurtosis	Generated Skewness	Relative Error of Skewness	Generated Kurtosis	Relative Error of Kurtosis
-0.85	6.0	-0.85029	-0.034%	6.0017	0.028%
-0.60	5.0	-0.59977	$+0.038\%$	4.9978	-0.045%
-0.35	4.0	-0.34991	$+0.025\%$	3.9979	-0.052%
-0.10	3.2	-0.10003	-0.030%	3.2090	0.281%

As shown in Table 4, when the target skewness varies from -0.85 to -0.10 and the target kurtosis varies from 6.0 to 3.2 , the proposed phase C-VAE generates rough surface

profiles whose higher-order statistical features remain close to their target values under all four substantially different statistical conditions. The tested conditions cover the variation from strong negative skewness and high kurtosis to weak negative skewness and low kurtosis. The absolute relative skewness errors range from 0.025% to 0.038%, while the absolute relative kurtosis errors range from 0.028% to 0.281%. Although the kurtosis error is relatively larger for the target condition $(-0.10, 3.2)$, it remains below 0.3%. These results demonstrate that, within the tested literature-based reference range, the proposed model maintains high higher-order statistical feature-matching accuracy despite substantial variations in the target skewness and kurtosis, indicating good adaptability to different statistical conditions.

4.3. Generation Capability and Diversity Analysis Under the Same Statistical Conditions

To meet the demand for multiple rough surface samples in friction and contact analyses of non-Gaussian rough surfaces, a single measured surface was used as the target surface, with its PSD, skewness, and kurtosis specified as the target statistical conditions. Using the trained phase C-VAE model described above, eight rough surface samples were generated. The skewness and kurtosis of each generated sample were calculated to evaluate the capability of the proposed method to generate diverse samples under the same statistical conditions. The generation time was also recorded to evaluate statistical-feature preservation, sample diversity, and computational efficiency. The statistical results and generated profiles are presented in Table 5 and Figure 6, respectively.

Table 5. Statistical results of eight rough surface samples generated by the phase C-VAE method under the same target conditions.

Sample	Skewness	Relative Error of Skewness	Kurtosis	Relative Error of Kurtosis
Measured	−0.38185	/	3.7092	/
Sample 1	−0.38184	+0.0014%	3.7084	−0.020%
Sample 2	−0.38146	+0.1018%	3.7099	+0.020%
Sample 3	−0.38207	−0.057%	3.7076	−0.043%
Sample 4	−0.38179	+0.016%	3.7061	−0.081%
Sample 5	−0.38147	+0.099%	3.7119	+0.073%
Sample 6	−0.38156	+0.074%	3.7121	+0.079%
Sample 7	−0.38174	+0.029%	3.7103	+0.030%
Sample 8	−0.38210	−0.066%	3.7089	−0.008%
Mean absolute relative error	/	0.056%	/	0.044%
Time (s)	24.62			

As shown in Figure 6, the eight generated profiles are consistent with the measured target surface in terms of the overall amplitude range and fluctuation characteristics. However, differences can be observed among the profiles in the local peak–valley positions, peak amplitudes, and detailed fluctuations, indicating that the proposed method has a strong capability to generate morphologically diverse samples under the same statistical constraints. As shown in Table 5, the skewness and kurtosis of the measured surface are −0.38185 and 3.7092, respectively. The skewness values of the eight generated samples are all close to the target value, and their kurtosis values remain within a small fluctuation range. The maximum absolute relative errors of skewness and kurtosis among the generated samples are 0.1018% and 0.081%, respectively. The mean absolute relative errors of the eight samples are 0.056% and 0.044%, respectively. Meanwhile, the time required to complete the generation and statistical analysis of the eight samples is 24.62 s. Because the phase C-VAE generates multiple surface profiles through batch inference, its multi-sample generation

time does not increase linearly with the number of generated samples. By comparison, the generation times of the conventional spectral representation method and the GA-PSD method are approximately 800 s and 6400 s, respectively. These results demonstrate that the proposed method can generate multiple rough surface samples while preserving the target higher-order statistical features. It exhibits good statistical consistency and greatly improves the generation efficiency in multi-sample generation scenarios.




Figure 6. Eight rough surface samples generated by the phase C-VAE method under the same target conditions.

To further characterize sample diversity, Figure 7 shows the mean profile and pointwise standard-deviation band of the generated samples, whereas Figure 8 presents the corresponding pairwise RMSE matrix.




Figure 7. Mean profile and standard deviation band of the generated samples.




Figure 8. RMSE heatmap among the eight rough surface samples.

4.4. Ablation and Sensitivity Analysis

To investigate the influence of important hyperparameters on the performance of the proposed phase C-VAE, a one-factor-at-a-time sensitivity analysis was first conducted. The final model configuration was used as the baseline, and only one hyperparameter was varied in each experiment while the training dataset, number of training epochs, optimizer settings, and generation procedure were kept unchanged. The investigated hyperparameters included the latent dimension d_z , the higher-order statistical constraint weight λ_{mom} , the mode-seeking loss weight λ_{ms} , and the dropout probability of the conditional embedding p_{drop} . For each trained model, eight surfaces were generated under the same target conditions. The mean absolute relative errors of skewness and kurtosis, the mean first-order Wasserstein distance W_1 , and the mean pairwise RMSE of the eight samples were used as evaluation metrics. The results are summarized in Table 6.

As shown in Table 6, the baseline configuration provides a favorable balance between higher-order statistical accuracy and sample diversity. When the latent dimension is reduced to $d_z = 8$ or increased to $d_z = 32$, the skewness and kurtosis errors increase compared with those of the baseline model, indicating that an excessively small latent space limits phase-distribution representation, whereas an excessively large latent space does not further improve generation performance under the present training conditions. For the higher-order statistical constraint weight, $\lambda_{\text{mom}} = 0.5$ leads to larger skewness and kurtosis errors, suggesting that the statistical constraint is insufficient. Increasing this weight to $\lambda_{\text{mom}} = 2.0$ maintains similar skewness accuracy but reduces the mean pairwise RMSE, indicating a decrease in generated-sample diversity.

Table 6. Hyperparameter sensitivity analysis of the proposed phase C-VAE method.

Hyperparameter Setting	Mean Absolute Relative Skewness Error	Mean Absolute Relative Kurtosis Error	W_1	Mean Pairwise RMSE
Baseline	0.0557%	0.0442%	0.0698	0.7343
$d_z = 8$	0.0906%	0.0668%	0.0820	0.7231
$d_z = 32$	0.0985%	0.1367%	0.0834	0.6391
$\lambda_{\text{mom}} = 0.5$	0.1104%	0.1171%	0.0711	0.8115
$\lambda_{\text{mom}} = 2.0$	0.0557%	0.0515%	0.0705	0.6053
$\lambda_{\text{ms}} = 0$	0.0157%	0.0163%	0.0698	0.0313
$\lambda_{\text{ms}} = 0.05$	0.0504%	0.0422%	0.0744	0.6379
$\lambda_{\text{ms}} = 0.20$	0.0795%	0.0753%	0.0852	0.8072
$p_{\text{drop}} = 0$	0.0772%	0.0732%	0.0815	0.7153
$p_{\text{drop}} = 0.30$	0.0979%	0.0819%	0.0739	0.6521

The mode-seeking loss weight mainly affects the diversity of generated samples. When $\lambda_{\text{ms}} = 0$, the skewness and kurtosis errors remain small, but the mean pairwise RMSE decreases sharply to 0.0313, indicating that the generated samples tend to become highly similar. Increasing λ_{ms} improves morphological diversity, but an excessively large value, such as $\lambda_{\text{ms}} = 0.20$, increases both the statistical errors and W_1 . For the conditional dropout probability, both $p_{\text{drop}} = 0$ and $p_{\text{drop}} = 0.30$ result in larger statistical errors than the baseline setting, suggesting that a moderate dropout probability helps stabilize conditional-feature learning. Overall, the baseline setting with $d_z = 16$, $\lambda_{\text{mom}} = 1.0$, $\lambda_{\text{ms}} = 0.10$, and $p_{\text{drop}} = 0.20$ achieves a balanced performance in terms of statistical-feature matching, height-distribution consistency, and generated-sample diversity.

To evaluate the effects of individual loss terms on the generation performance of the proposed phase C-VAE and analyze the balance between statistical feature matching and sample diversity, an ablation study was conducted. With the network architecture, training parameters, and generation parameters kept unchanged, the mode-seeking loss, higher-order statistical constraint loss, conditional prior alignment loss, and reconstruction loss were removed separately. The resulting changes in skewness error, kurtosis error, and generated-sample diversity were then compared among the models. This analysis clarifies the contribution of each loss term and verifies the necessity and rationality of the final loss-function design.

The quantitative ablation results are summarized in Table 7. The higher-order statistical errors, diversity metrics, and trade-off between statistical accuracy and diversity are presented in Figures 9, 10, and 11, respectively.

Table 7. Mean absolute relative higher-order statistical errors and mean pairwise RMSE of generated rough surface samples in each ablation group.

Group	Mean Absolute Relative Skewness Error	Mean Absolute Relative Kurtosis Error	Mean Pairwise RMSE
Full model	0.056%	0.044%	0.73
Without mode-seeking loss	0.016%	0.016%	0.031
Without higher-order statistical constraint loss	2.79%	3.50%	1.32
Without conditional prior alignment loss	0.092%	0.076%	0.78
Without reconstruction loss	0.087%	0.073%	0.70




Figure 9. Mean absolute relative errors of higher-order statistics of generated rough surface samples.




Figure 10. Diversity of generated rough surface samples.

As shown in Table 7, the full model achieves a good balance between higher-order statistical matching accuracy and sample diversity. Its mean absolute relative errors of skewness and kurtosis are 0.056% and 0.044%, respectively, and the mean pairwise RMSE is 0.73. After removing the mode-seeking loss, the skewness and kurtosis errors remain at low levels, but the mean pairwise RMSE decreases to 0.031, indicating that the generated samples are highly similar to each other and that the model exhibits obvious diversity degradation. After removing the higher-order statistical constraint loss, the mean pairwise RMSE increases to 1.32, whereas the mean absolute relative errors of skewness and kurtosis increase to 2.79% and 3.50%, respectively. This indicates that although the generated samples show greater morphological differences, their higher-order statistical features fail to satisfy the target requirements. After removing the conditional prior alignment loss or the reconstruction loss, the model can still generate rough surfaces satisfying the basic statistical constraints, but both the skewness and kurtosis errors increase compared with those of the full model. This suggests that these two loss terms play auxiliary roles in maintaining latent-variable distribution consistency and improving phase-structure representation capability, respectively.




Figure 11. Trade-off between higher-order statistical accuracy and diversity of generated rough surface samples. The mean statistical error is defined as the average of the mean absolute relative skewness and kurtosis errors.

In summary, the higher-order statistical constraint loss is the key term for ensuring skewness and kurtosis matching accuracy, whereas the mode-seeking loss is the key term for maintaining sample diversity under the same conditions. The conditional prior alignment loss and reconstruction loss further improve model stability and generation accuracy. The ablation study verifies the rationality and necessity of the constructed multi-objective loss function.

5. Conclusions

To meet the demand for rapid multi-sample modeling of non-Gaussian rough surfaces in friction, lubrication, and contact analyses, this study proposed a phase C-VAE method with a PSD amplitude constraint and validated it against the conventional spectral representation method and the GA-PSD method. The experimental results show that the proposed method can rapidly generate diverse non-Gaussian rough surfaces while preserving the target PSD, skewness, and kurtosis constraints. Under the statistical conditions of the same measured surface, the relative skewness error, relative kurtosis error, and generation time of the sample generated by the phase C-VAE method were 0.0014%, 0.020%, and 24.62 s, respectively, demonstrating higher generation efficiency than the conventional spectral representation method and the GA-PSD method. For the eight rough surface samples generated under the same target conditions, the mean absolute relative errors of skewness and kurtosis were 0.056% and 0.044%, respectively, and the mean pairwise RMSE was 0.73, indicating that the proposed method can preserve the consistency of higher-order statistical features while maintaining local morphological differences among samples. The ablation study further shows that the higher-order statistical constraint loss plays a key role in skewness and kurtosis matching accuracy, whereas the mode-seeking loss is important for maintaining sample diversity under the same conditions. Overall, the proposed method alleviates the reliance of conventional constructive methods on sample-by-sample iterative solution and the associated efficiency limitations in batch generation, providing an effective machine-learning-based generation method for non-Gaussian rough surface sample augmentation and batch modeling in lubrication, friction, and contact analyses.

The present study focuses on the conditional generation of one-dimensional non-Gaussian rough surface profiles with prescribed PSD, skewness, and kurtosis. Extension of the proposed method to two-dimensional rough surface generation will be investigated in future work.

Author Contributions: Conceptualization, J.W., X.Z. and J.Y.; funding acquisition, J.Y., W.Z., M.W., B.Y. and T.X.; investigation, M.W., B.Y. and T.X.; methodology, J.W., W.Z., X.Z. and X.H.; software, J.W., X.H. and T.X.; supervision, X.Z. and J.Y.; writing—original draft, J.W., W.Z. and X.H.; writing—review & editing, J.W., X.Z. and J.Y. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported in part by the National Science and Technology Major Project of the Ministry of Science and Technology of China (No. G202401).

Data Availability Statement: The datasets presented in this article are not readily available because the full suite of data and software is not yet ready for release. Insufficient safeguards and documentation currently make it unsuitable for public use. Requests to access the datasets should be directed to the author Jinyuan Wang.

Conflicts of Interest: Authors Meile Wang, Bo Yu, and Taowen Xiao were employed by the company Hengli Precision Industry 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.

Abbreviations

The following abbreviations are used in this manuscript:

ACF	Autocorrelation Function
AR	Autoregressive
ARMA	Autoregressive Moving Average
C-VAE	Conditional Variational Autoencoder
CPU	Central Processing Unit
FFT	Fast Fourier Transform
GA	Genetic Algorithm
GA-PSD	Genetic-Algorithm-Optimized PSD-Based Generation Method
GPU	Graphics Processing Unit
IFFT	Inverse Fast Fourier Transform
KL	Kullback–Leibler
MA	Moving Average
PSD	Power Spectral Density
CDF	Cumulative Distribution Function
ReLU	Rectified Linear Unit
RMSE	Root Mean Square Error

References

1. Šperka, P.; Knotek, J.; Omasta, M.; Křupka, I.; Polach, P.; Hartl, M. Film Thickness and Friction of Textured Surfaces in Hydrodynamic Inclined and Parallel Gaps—An Experimental Study. *Lubricants* **2026**, *14*, 26. [\[CrossRef\]](#)
2. Cordier, M.; Diab, Y.; Cavolet, J.; Majdoub, F.; Changenet, C.; Ville, F. An experimental study on the impact of roughness orientation on the friction coefficient in EHL contact. *Lubricants* **2025**, *13*, 340. [\[CrossRef\]](#)
3. He, T.; Chen, W.; Liu, Z.; Gong, Z.; Du, S.; Zhang, Y. The impact of surface roughness on the friction and wear performance of GCr15 bearing steel. *Lubricants* **2025**, *13*, 187. [\[CrossRef\]](#)
4. Tang, Y.; Huang, L.; Wu, L.; Meng, X. Simulation of lubrication on rough surfaces with multiscale lubrication neural networks. *Sci. China Technol. Sci.* **2025**, *68*, 1320303. [\[CrossRef\]](#)
5. Ghorbani, H.; Shi, B.; Attia, H. Experimental Investigation and Modelling of High-Speed Turn-Milling of H13 Tool Steel: Surface Roughness and Tool Wear. *Lubricants* **2025**, *13*, 444. [\[CrossRef\]](#)

6. Wang, Y.; Liu, Y. Preparation of measured engineering surfaces for modeling tribological systems, part I: Characterization and reconstruction. *Tribol. Int.* **2024**, *199*, 110026. [[CrossRef](#)]
7. Davies, N.; Spedding, T.; Watson, W. Autoregressive moving average processes with non-normal residuals. *J. Time Ser. Anal.* **1980**, *1*, 103–109. [[CrossRef](#)]
8. Watson, W.; King, T.; Spedding, T.A.; Stout, K. The machined surface—time series modelling. *Wear* **1979**, *57*, 195–205. [[CrossRef](#)]
9. Watson, W.; Spedding, T.A. The time series modelling of non-Gaussian engineering processes. *Wear* **1982**, *83*, 215–231. [[CrossRef](#)]
10. Hill, I.; Hill, R.; Holder, R. Algorithm AS 99: Fitting Johnson curves by moments. *J. R. Stat. Society. Ser. C Appl. Stat.* **1976**, *25*, 180–189. [[CrossRef](#)]
11. Uchidate, M.; Shimizu, T.; Iwabuchi, A.; Yanagi, K. Generation of reference data of 3D surface texture using the non-causal 2D AR model. *Wear* **2004**, *257*, 1288–1295. [[CrossRef](#)]
12. Patir, N. A numerical procedure for random generation of rough surfaces. *Wear* **1978**, *47*, 263–277. [[CrossRef](#)]
13. Hu, Y.; Tonder, K. Simulation of 3-D random rough surface by 2-D digital filter and Fourier analysis. *Int. J. Mach. Tools Manuf.* **1992**, *32*, 83–90. [[CrossRef](#)]
14. Bakolas, V. Numerical generation of arbitrarily oriented non-Gaussian three-dimensional rough surfaces. *Wear* **2003**, *254*, 546–554. [[CrossRef](#)]
15. Liao, D.; Shao, W.; Tang, J.; Li, J. An improved rough surface modeling method based on linear transformation technique. *Tribol. Int.* **2018**, *119*, 786–794. [[CrossRef](#)]
16. Minet, C.; Brunetiere, N.; Tournier, B.; Fribourg, D. Analysis and modeling of the topography of mechanical seal faces. *Tribol. Trans.* **2010**, *53*, 799–815. [[CrossRef](#)]
17. Francisco, A.; Brunetière, N. A hybrid method for fast and efficient rough surface generation. *Proc. Inst. Mech. Eng. Part J J. Eng. Tribol.* **2016**, *230*, 747–768. [[CrossRef](#)]
18. Wu, J.J. Simulation of rough surfaces with FFT. *Tribol. Int.* **2000**, *33*, 47–58. [[CrossRef](#)]
19. Wu, J.J. Simulation of non-Gaussian surfaces with FFT. *Tribol. Int.* **2004**, *37*, 339–346. [[CrossRef](#)]
20. Wang, Y.; Liu, Y.; Zhang, G.; Wang, Y. A simulation method for non-Gaussian rough surfaces using fast Fourier transform and translation process theory. *J. Tribol.* **2018**, *140*, 021403. [[CrossRef](#)]
21. Wang, Y.; Azam, A.; Wilson, M.C.; Neville, A.; Morina, A. A comparative study for selecting and using simulation methods of Gaussian random surfaces. *Tribol. Int.* **2022**, *166*, 107347. [[CrossRef](#)]
22. Mao, Q.; Lee, H.Y.; Tseng, H.Y.; Ma, S.; Yang, M.H. Mode seeking generative adversarial networks for diverse image synthesis. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 15–20 June 2019; pp. 1429–1437.
23. Chen, J.; Zang, F.; Zhao, X.; Li, H.; Tong, Z.; Yuan, K.; Zhu, L. Generating non-Gaussian rough surfaces using analytical functions and spectral representation method with an iterative algorithm. *Appl. Math. Model.* **2025**, *137*, 115665. [[CrossRef](#)]
24. Pawlus, P.; Reizer, R.; Wiczorowski, M.; Królczyk, G. Parametric description of one-process surface texture. *Measurement* **2022**, *204*, 112066. [[CrossRef](#)]

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.