Computers, Materials & Continua

Tech Science Press

https://doi.org/10.32604/cmc.2026.087308

ARTICLE

Data-Driven Conditional Diffusion Generation Method for Anisotropic Mechanical Metamaterial Unit Cells

Hao Sun, Xiaohong Ding*, Min Xiong and Heng Zhang

School of Mechanical Engineering, University of Shanghai for Science and Technology, Shanghai, China

*Corresponding Author: Xiaohong Ding. Email: dingxhsh021@126.com

Received: 16 June 2026; Accepted: 03 August 2026; Published: 15 September 2026

ABSTRACT: Designing two-dimensional anisotropic mechanical metamaterial unit cells from prescribed effective properties remains a challenging inverse problem, particularly when directional stiffness and material usage need to be controlled simultaneously. In this work, a data-driven conditional diffusion framework is developed for generating unit-cell structures with target effective elastic moduli and volume fractions. A structure–property database containing 57,000 binary unit-cell images is first established through a random target-property-driven inverse homogenization method. The effective elastic moduli in the x and y directions, together with the volume fraction, are used as conditional labels, denoted as (E_x, E_y, V). A conditional denoising diffusion probabilistic model (DDPM) with a U-Net denoising backbone, referred to as DDPM-UNet, is then trained to generate unit-cell layouts under prescribed property conditions. To improve the consistency between generated geometries and target properties, a convolutional neural network (CNN) surrogate model is introduced. The CNN provides a property-consistency loss during diffusion model fine-tuning and is also used as a fast evaluator for candidate screening during inference. Compared with DDPM-UNet, the DDPM-UNet-CNN framework produces more valid candidates under 5%, 10%, and 20% relative error thresholds. FEA-based re-homogenization further shows that the selected candidates generally approach the prescribed effective properties. The developed framework provides a practical generative route for identifying candidate anisotropic mechanical metamaterial unit cells under prescribed property conditions.

KEYWORDS: Inverse generative design; conditional diffusion model; mechanical metamaterials; anisotropic unit cells

1 Introduction

Engineering applications increasingly require metamaterial structures with tailored combinations of mechanical properties. Mechanical metamaterials meet this demand by tuning the effective macroscopic response through the geometric design of microstructures at the unit-cell scale [1,2]. Unlike conventional homogeneous structures, their main mechanical parameters, such as elastic modulus [3], Poisson's ratio [4,5], and loss factor [6], are governed not only by the intrinsic properties of the base material, but also by structural features including unit-cell geometry, porosity, connectivity, and boundary conditions. Therefore, even when the same constituent material is used, changing the microstructural configuration can result in markedly different mechanical responses [7]. This geometry-dependent behavior makes it possible to tailor material properties in a controlled manner. With properly designed periodic or aperiodic microstructures, mechanical metamaterials can exhibit direction-dependent elastic modulus, negative Poisson's ratio, prescribed stress–strain responses, and high specific strength [8,9]. Consequently, an important task in metamaterial design is to efficiently identify unit-cell layouts that can satisfy prescribed mechanical-property requirements.

Because of their anisotropic responses, mechanical metamaterials have been used in a wide range of engineering applications. For anisotropic unit cells, describing the mechanical behavior with two directional effective moduli, E_x in the x direction and E_y in the y direction, provides a more direct representation than using a single homogenized parameter. Previous studies have shown that the mechanical performance of metamaterials can be improved by modifying unit-cell symmetry and topological connectivity [10]. Karathanasopoulos et al. [11] showed that both the macroscopic effective elastic modulus and the directional response of metamaterials can be regulated through the design of unit-cell architectures and their spatial arrangements. Kumar et al. [12] further reported that spinodoid metamaterials can balance directional stiffness and mass fraction by controlling micro- and nanoscale structural parameters, which is relevant to functionally graded materials and biomimetic bone tissue engineering. These studies suggest that, in unit-cell design, the combined consideration of E_x, E_y, and V is important for achieving directional mechanical performance while reducing material usage and structural weight, where V denotes the volume fraction, i.e., the ratio of the solid material area to the total unit-cell area for two-dimensional binary unit-cell images. Topology optimization and homogenization methods have long been used in metamaterial design [13]. Sigmund [14] proposed an inverse homogenization method for generating periodic microstructures with prescribed constitutive properties, and Andreassen and Andreasen [15] developed a numerical homogenization method for evaluating the effective elastic tensor of periodic composite materials. These methods provide a theoretical and numerical basis for linking microscopic configurations with macroscopic properties. However, when a specific target property is required, repeated optimization and parameter adjustment are usually needed, which increases the computational cost. The obtained structures may also depend strongly on the initial design, volume-fraction constraint, and filtering parameters. For high-resolution binary unit-cell images, additional issues such as structural connectivity, periodic-boundary consistency, and manufacturability must be checked after optimization. As a result, conventional topology-optimization-based design often requires a considerable amount of post-processing, screening, and correction work [16].

In recent years, with the rapid development of artificial intelligence, particularly deep learning, metamaterial design has gradually shifted from experience-based approaches to data-driven design paradigms [17]. In the field of metamaterials, deep learning has been widely applied to property prediction, geometric parameter optimization, and inverse design. Huang et al. [18] proposed an automated co-design method combining convolutional neural networks (CNNs) and generative adversarial networks (GANs). By randomly generating a large number of geometrically diverse samples and validating the results experimentally, they demonstrated that deep learning can effectively capture the features of complex and disordered geometries. Bastek et al. [19] developed an inverse design tool based on structure–property mapping to obtain metamaterial structures with prescribed anisotropic properties. Ha et al. [20] further demonstrated a machine-learning-based inverse design strategy for architected materials with prescribed stress–strain responses. Recent studies have further extended data-driven design to a wider range of mechanical metamaterial systems. For example, machine-learning- and generative-model-based inverse design methods have been used to generate mechanical metamaterials with prescribed nonlinear stress–strain responses and finite-strain mechanical behaviors [21,22]. Diffusion-model-based methods have also been applied to energy-absorbing metamaterials, where target mechanical responses, such as stress–strain curves or energy-absorption characteristics, are used to guide microstructure generation [23]. In addition, data-driven approaches have been used for composite lattice structures and programmable mechanical metamaterials with tunable effective properties or on-demand mechanical responses [24,25]. These studies show that data-driven structural design is becoming an effective tool for exploring large metamaterial design spaces.

Among current generative approaches, the denoising diffusion probabilistic model (DDPM) has become increasingly used in the inverse design of microstructures and mechanical metamaterials because it offers stable training behavior and can generate diverse samples [26,27]. Lyu and Ren [28] applied a diffusion-based deep generative model to reconstruct two-dimensional and three-dimensional random material microstructures, showing that DDPM can capture the distribution features of complex microstructural images. Wang et al. [29] further demonstrated that DDPM can generate energy-absorbing metamaterial microstructures from prescribed mechanical properties, which highlights its potential for inverse design. Du et al. [30] extended this idea to metamaterial unit-cell design and used DDPM to improve the diversity and functionality of the generated topologies. Taken together, these studies suggest that DDPM is well suited for probabilistic modeling and generation of complex microstructures.

Despite these advances, diffusion-based inverse design of two-dimensional anisotropic unit cells with continuous target properties remains limited. Most existing workflows use diffusion models as conditional generators, while surrogate models are mainly used after generation for fast property prediction or candidate screening. In this work, the CNN surrogate is instead introduced into the DDPM-UNet fine-tuning stage as a frozen differentiable property predictor. Its parameters are not updated during this stage, and it is used only to calculate a property-consistency loss for E_x, E_y and V. This allows the diffusion model to receive property-related feedback during fine-tuning, rather than relying only on post-generation filtering.

The proposed framework combines a coverage-oriented anisotropic unit-cell database, continuous property-conditioned DDPM generation, frozen CNN-guided fine-tuning, and FEA-based re-homogenization. The database is constructed using inverse-homogenization topology optimization, followed by re-homogenization, connectivity filtering, and interval-based sampling. The DDPM-UNet model is trained with E_x, E_y and V as continuous conditional variables, and the frozen CNN surrogate is used during fine-tuning to improve property consistency. The generated candidates are finally screened and re-evaluated by FEA-based homogenization. The aim of this framework is to increase the probability of obtaining valid candidates under prescribed error thresholds, rather than to guarantee exact inverse design for every generated structure.

2 Methods

2.1 Problem Description

For a given volume fraction and a set of target mechanical properties, the task is to generate a two-dimensional unit-cell topology whose effective response matches the prescribed requirements. To reduce the mismatch between the generated topology and the target properties, a property-consistency constraint is introduced into the generation process.

The original unit-cell structure is represented as a two-dimensional binary image composed of 0 and 1:

$$x_0 \in \{0,1\}^{H\times W} \tag{1}$$

where $x_0 = 1$ denotes the solid material phase, while $x_0 = 0$ denotes the void phase. H and W represent the height and width of the image, respectively. In this study, the unit-cell image resolution is set to 96×96 pixels.

For each unit-cell structure, the corresponding conditional label is defined as:

$$c = [E_x, E_y, V]^{\mathrm{T}} \in \mathbf{R}^3 \tag{2}$$

where E_x and E_y denote the effective elastic moduli of the unit cell along the x and y directions, respectively, and V represents the volume fraction of the unit cell.

Therefore, the training dataset can be expressed as:

$$\mathcal{D} = \{x_0^n, c^n\}_{n=1}^{N} \quad (3)$$

where x_0^n is the n-th binary unit-cell image, and c^n is its corresponding conditional label.

In this study, two models are constructed for comparison:

DDPM-UNet: E_x, E_y, and V are used as conditional inputs to enable target-property-driven unit-cell generation.

DDPM-UNet-CNN: Based on conditional diffusion generation, a CNN-based property predictor is introduced. A property-consistency constraint is further incorporated to improve the agreement between the generated structures and the target mechanical properties.

2.2 Training Data

The unit-cell database was built by a random target-property-driven inverse-homogenization topology optimization method. Each unit cell was discretized into a 96 × 96 regular finite element grid, with each pixel corresponding to one finite element. The solid phase was Ti-6Al-4V, with an elastic modulus of E_0 = 110,000 MPa and a Poisson's ratio of ν = 0.34.

A large candidate pool was generated by continuous random target sampling and inverse-homogenization optimization. Sequential interval-based sampling was then applied to construct the final database. The candidates were grouped by the recalculated E_x values using 5 GPa intervals, with at most 5000 samples randomly retained in each interval. The remaining candidates were used to supplement underrepresented E_y and volume-fraction intervals, yielding 57,000 binary unit-cell images with improved coverage over the E_x, E_y, and V spaces.

For each optimization run, a target property set was sampled first. The main targets were the directional effective moduli E_x^{tar} and E_y^{tar}, and the target volume fraction V^{tar}. The target directional moduli were sampled within the stiffness range of 5–80 GPa, and the target volume fraction was uniformly sampled from 0.3 to 0.8. Four target-sampling modes were used to cover different regions of the E_x-E_y space: E_x dominant, E_y dominant, balanced, and mild-anisotropic modes. These modes were only used for target sampling during database generation. They were not used as discrete class labels in the CNN surrogate model or in the DDPM. The learning models were conditioned only on the continuous property vector (E_x, E_y, V). Transposition-based data augmentation was not applied in this work. If the x and y directions of a unit cell are exchanged, the corresponding labels should also be transformed by exchanging E_x and E_y, while V remains unchanged. Since the database already contains both E_x dominant and E_y dominant samples, the current study did not use this augmentation strategy, but it will be considered in future work to further improve data symmetry and coverage.

The representative unit-cell structures generated under different target-property categories are shown in Fig. 1. A main modulus *Emain* was sampled from the prescribed stiffness interval. In the E_x dominant mode, E_x^{tar} = *Emain*, and E_y^{tar} was assigned as a lower fraction of *Emain*. In the E_y dominant mode, the same rule was used after exchanging the two directions. In the balanced mode, E_x^{tar} and E_y^{tar} were both sampled near *Emain*, so the two directional moduli remained close. In the mild-anisotropic mode, a moderate difference between E_x^{tar} and E_y^{tar} was introduced without producing strongly directional targets. This sampling setting provided candidate structures with different degrees of anisotropy.

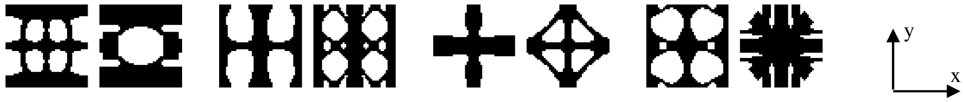

Figure 1: Representative unit-cell structures generated under four target-property categories.

Since E_x^{tar} and E_y^{tar} alone cannot determine the complete in-plane stiffness tensor of an anisotropic unit cell, two auxiliary targets were introduced during inverse homogenization. The auxiliary shear modulus target was defined as

$$G_{xy}^{tar} = \eta_g \sqrt{E_x^{tar} E_y^{tar}} \tag{4}$$

where η_g is the shear-stiffness adjustment coefficient. In this study, η_g was uniformly sampled from 0.20 to 0.42.

The auxiliary Poisson's ratio target ν_{xy}^{tar} was uniformly sampled from 0.05 to 0.30. This range was used to keep the auxiliary Poisson's-ratio target in a stable range and to prevent the Poisson's-ratio term from dominating the main objectives. The quantities G_{xy}^{tar} and ν_{xy}^{tar} were only used as weak auxiliary constraints in the topology optimization step. They were not used as conditional labels in the generative model.

The unit-cell topology was represented by a continuous density field. The initial density field was generated according to V^{tar}, and random density perturbations were added to produce different initial layouts. Several low-density circular regions were inserted as initial void seeds. A high-density cross-shaped skeleton was placed near the center of the unit cell to promote connected load-bearing paths in the early optimization stage. This skeleton was only used in the initial density field. It was not imposed as a required feature of the final binary topology. During optimization, it could disappear, deform, or merge into other solid paths.

The initial density field (ρ_0) was processed by biaxial mirroring and periodic-boundary correction. Specifically, the density field was averaged with its left–right and up–down mirrored copies, and opposite edges were then averaged to enforce boundary consistency. The mirroring operation improved the symmetry of the initial layout, while the periodic-boundary correction made the opposite edges compatible for repeated tiling.

The element stiffness was interpolated by the SIMP model:

$$E_e = E_{\min} + \rho_0^p (E_0 - E_{\min}) \tag{5}$$

where $E_{\min} = 10^{-6} E_0$, and the SIMP penalization factor was set to $p = 3.0$.

Finite element analysis was then performed to calculate the homogenized stiffness tensor $\boldsymbol{C}^H$ under periodic boundary conditions. The effective engineering constants were obtained from the compliance matrix $\boldsymbol{S} = \left(\boldsymbol{C}^{\mathrm{H}}\right)^{-1}$ as follows:

$$E_x = \frac{1}{S_{11}}, E_y = \frac{1}{S_{22}}, G_{xy} = \frac{1}{S_{33}}, \nu_{xy} = -\frac{S_{12}}{S_{11}} \tag{6}$$

The density field was updated by minimizing the following property-error objective:

$$J = w_x \left(\frac{E_x - E_x^{tar}}{E_x^{tar}}\right)^2 + w_y \left(\frac{E_y - E_y^{tar}}{E_y^{tar}}\right)^2 + w_g \left(\frac{G_{xy} - G_{xy}^{tar}}{G_{xy}^{tar}}\right)^2 + w_\nu \left(\frac{\nu_{xy} - \nu_{xy}^{tar}}{\nu_{xy}^{tar}}\right)^2 + \lambda_V (V - V^{tar})^2 \tag{7}$$

where w_x, w_y, w_g, and w_v are the weighting coefficients, and λ_V is the penalty coefficient for the volume fraction.

In this work, $w_g = 0.45$, $w_v = 0.03$, and $\lambda_V = 6.0$. The weights of E_x and E_y were adjusted according to the target-sampling mode. For the E_x dominant mode, $w_x = 1.35$ and $w_y = 0.80$; for the E_y dominant mode, $w_y = 1.35$ and $w_x = 0.80$; for the balanced mode, $w_x = w_y = 1.10$; otherwise, $w_x = w_y = 1.00$. The E_x, E_y, and V terms were the main objectives. The G_{xy} and ν_{xy} terms served as weak auxiliary constraints.

The density variables were updated using an Adam-based projected gradient descent method [31,32]. The learning rate was set to 0.035, and the maximum move limit of each design variable was 0.08 in one update step. The Adam parameters were $\beta 1 = 0.9$, $\beta 2 = 0.999$ and $\varepsilon = 10^{-8}$. During each iteration, a periodic density filter with a radius of 3.5 elements was applied to reduce checkerboard patterns and improve spatial smoothness. A smooth Heaviside projection was then used to promote a clearer solid–void distribution [33]:

$$\bar{\rho} = \frac{\tanh(\beta\eta) + \tanh(\beta(\rho - \eta))}{\tanh(\beta\eta) + \tanh(\beta(1 - \eta))} \tag{8}$$

where $\beta = 3.0$ and $\eta = 0.5$.

The optimization was terminated when the maximum number of iterations was reached. In this work, the maximum iteration number was set to 100. An early stopping criterion was also used after 30 iterations: the optimization was stopped when the relative errors of E_x, E_y, and G_{xy} were lower than 8%, 8%, and 15%, respectively.

After optimization, the best continuous density field was filtered and projected again, and then converted into a binary unit-cell image by adaptive thresholding. The threshold was determined from the density distribution so that the binary solid fraction was close to the desired volume fraction. A minimum binary solid fraction of 0.30 was imposed to avoid overly sparse and weakly connected structures. After binarization, opposite boundaries were corrected again to ensure periodic-boundary consistency. Disconnected solid components were repaired by adding thin solid bridges with a width of 2 pixels. Connected-component analysis was then used to check the final binary structure. Samples containing isolated solid islands or multiple disconnected main solid components were excluded during final database selection.

Finally, each retained binary unit cell was recalculated by periodic-boundary homogenization. The recalculated properties E_x, E_y, V were stored together with the corresponding binary image. The database selection was also based on these recalculated properties rather than the initially sampled target values. The overall database-generation procedure is illustrated in Fig. 2. In the following learning stage, only the recalculated E_x, E_y, and V were used as conditional labels for the CNN surrogate model and the DDPM.

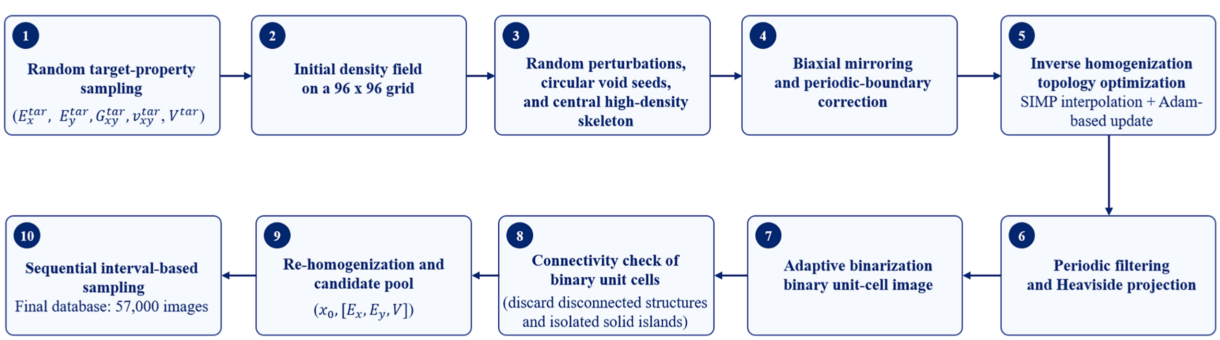

Figure 2: Workflow for constructing the anisotropic mechanical metamaterial unit cell database.

Fig. 3 summarizes the statistical characteristics of the constructed unit-cell database. Fig. 3a shows the sample-count distribution in the E_x-E_y property space. The horizontal and vertical axes denote the recalculated effective moduli E_x and E_y, respectively, and the color represents the number of samples in each modulus interval. The database covers a wide range of directional stiffness combinations, including E_x dominant, E_y dominant, and nearly balanced stiffness regions. The nonuniform sample density reflects both the feasibility of different anisotropic property combinations and the subsequent interval-based selection strategy.

Fig. 3b gives the sample-count distribution over different volume-fraction intervals. The samples are distributed over the prescribed volume-fraction range, providing unit cells with different solid material contents. Since the volume fraction directly affects both stiffness and material usage, this distribution ensures that the database contains structures with different lightweight design levels.

The anisotropy degree of the unit cells in the two principal directions is quantified by

$$\mathrm{R} = \max\left(\frac{E_x}{E_y}, \frac{E_y}{E_x}\right) \tag{9}$$

Fig. 3c shows the distribution of R in the database. When R approaches 1, the unit cell exhibits nearly balanced directional mechanical characteristics. As R increases, the anisotropy degree of the unit cell gradually becomes more pronounced. The database contains both nearly balanced structures with R close to 1 and clearly anisotropic structures with R ≥ 1.5. As shown in Fig. 3d, the horizontal and vertical axes represent E_x and E_y, respectively, and the value in each grid denotes the average V of the samples within the corresponding effective-modulus interval. Overall, V increases with increasing E_x and E_y. For two-dimensional binary unit cells, achieving higher bidirectional effective stiffness generally requires a larger fraction of solid material.

2.3 CNN Surrogate Model

To enhance the consistency between the generated structures and the target mechanical properties during diffusion model training, a convolutional neural network (CNN) is introduced as a surrogate property-prediction model [34,35]. The CNN is trained independently before the diffusion model training. After training, its network parameters are frozen, and the model is used only as a mechanical-property supervision module to participate in the forward propagation and loss calculation of the diffusion model.

Before model training, the original binary unit-cell image x_0 was preprocessed by a pixel-wise linear mapping. Since the binary image uses 0 and 1 to represent void and solid phases, respectively, each pixel value was mapped to [−1, 1] as

$$\tilde{x}_0 = 2x_0 - 1 \tag{10}$$

Thus, void pixels were mapped from 0 to −1, while solid pixels were mapped from 1 to 1. This operation was applied to each pixel of the image tensor, rather than using an image-wise normalization.

The property label $c = [E_x, E_y, V]$ was standardized dimension by dimension using the mean and standard deviation calculated from the training set:

$$\tilde{c}_j = \frac{c_j - \mu_j}{\sigma_j}, j \in \{E_x, E_y, V\} \tag{11}$$

where μ_j and σ_j denote the mean and standard deviation of the j-th property in the training set. The same statistics were used in training, validation, and generation. During evaluation, the predicted normalized properties were transformed back to the physical scale by inverse standardization.

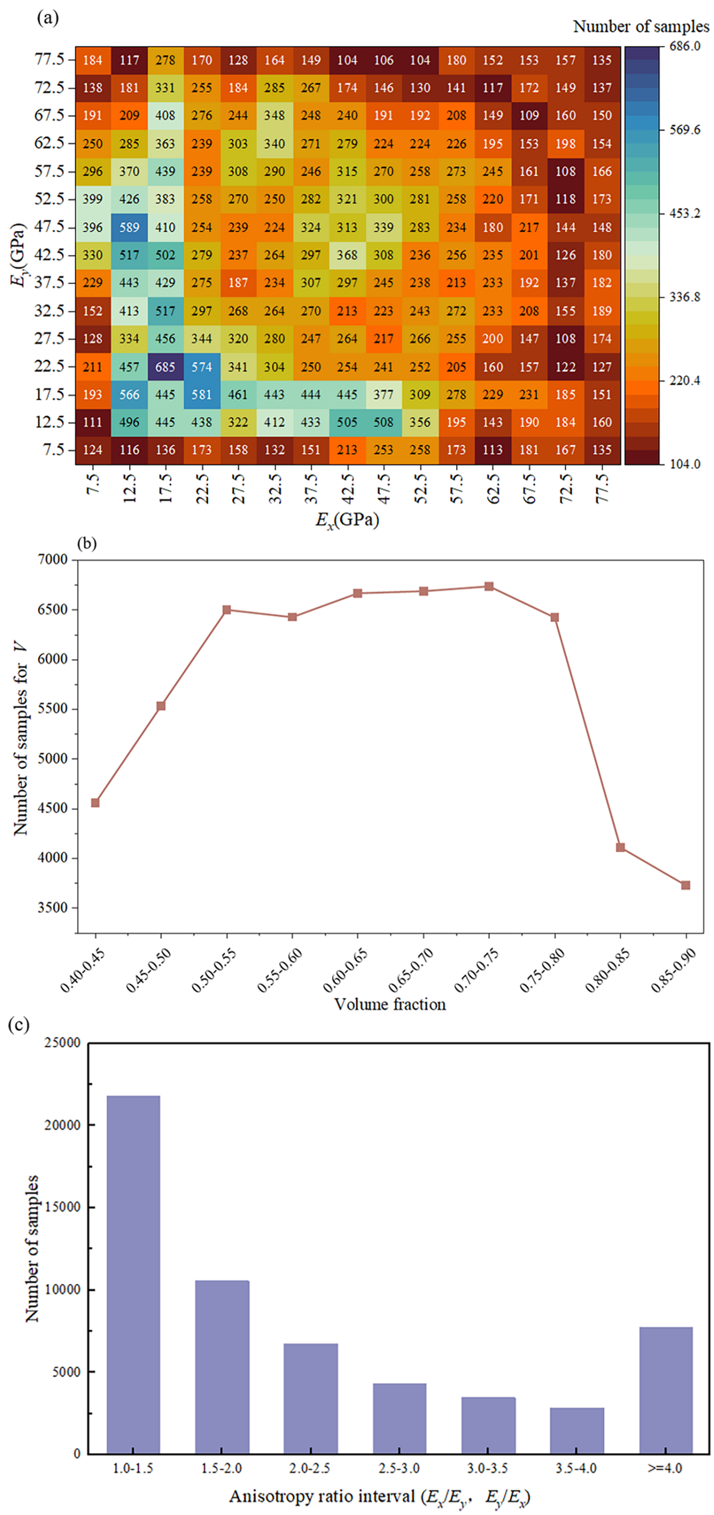

Figure 3: (Continued)

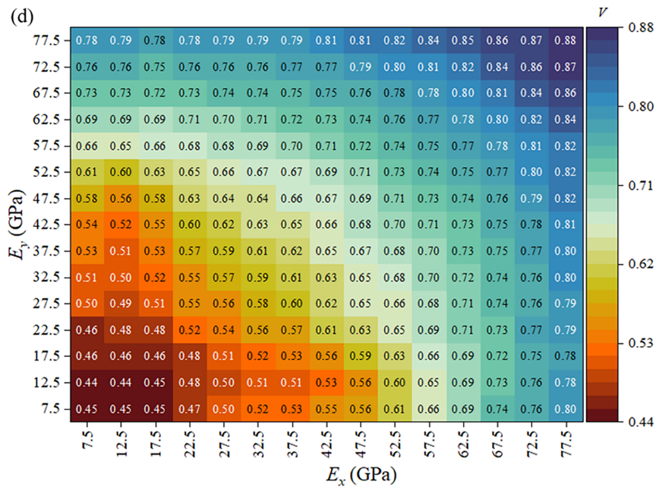

Figure 3: Statistical characteristics of the constructed unit-cell database: (**a**) sample-count heat map in the E_x-E_y property space; (**b**) sample-count distribution over volume-fraction intervals; (**c**) anisotropy-ratio distribution; and (**d**) average volume fraction in the E_x-E_y property space.

The CNN surrogate takes the normalized unit-cell image $\tilde{x}_0$ as input and outputs the normalized property vector $\tilde{c}_{pred}$:

$$\tilde{c}_{pred} = f_{cnn}(\tilde{x}_0) \tag{12}$$

where $f_{cnn}(\cdot)$ denotes the surrogate property-prediction network. The feature extractor consists of four convolutional blocks with 32, 64, 128, and 256 channels, respectively. Each convolutional layer uses a 3 × 3 kernel and is followed by batch normalization and ReLU activation. Max pooling is applied after the first three convolutional blocks. The extracted feature map is then processed by adaptive average pooling and passed to a fully connected regression head with dimensions 256→128→32→3. A dropout layer with a rate of 0.2 is used in the regression head.

During the training stage of the surrogate network, the loss function is defined as:

$$L_{cnn} = \mathcal{L}(\tilde{c}_{pred}, \tilde{c}) \tag{13}$$

where $\mathcal{L}(\cdot)$ denotes the regression loss used for CNN training.

The dataset was split into training, validation, and test sets with a ratio of 8:1:1 using a fixed random seed. The CNN surrogate was trained for 200 epochs with a batch size of 32. AdamW was used as the optimizer, with a learning rate of 1×10^{-3} and a weight decay of 1×10^{-4}. The model with the lowest validation loss was saved and used in the subsequent DDPM fine-tuning stage. No random geometric data augmentation, such as rotation, flipping, or cropping, was used during CNN training.

After training, the CNN parameters are frozen. During the subsequent training of the diffusion model, the pretrained CNN serves as a differentiable surrogate model for property evaluation. Specifically, it predicts the mechanical properties of the denoised structure estimated by the DDPM-UNet, while allowing the property-consistency loss to be back-propagated to the generative model. In this way, target-property constraints are imposed on the generation process.

2.4 Conditional Diffusion Generative Model

In this study, a conditional DDPM-UNet is adopted as the structural generative model [26,36]. The noisy structural image x_t, the time step t, and the target property condition $\tilde{c} = [E_x, E_y, V]^{\mathrm{T}}$ are used as the model inputs. In the decoding stage of the U-Net, the conditional information and time-step information are mapped into feature vectors and fused with the image features [37,38], as shown in Fig. 4.

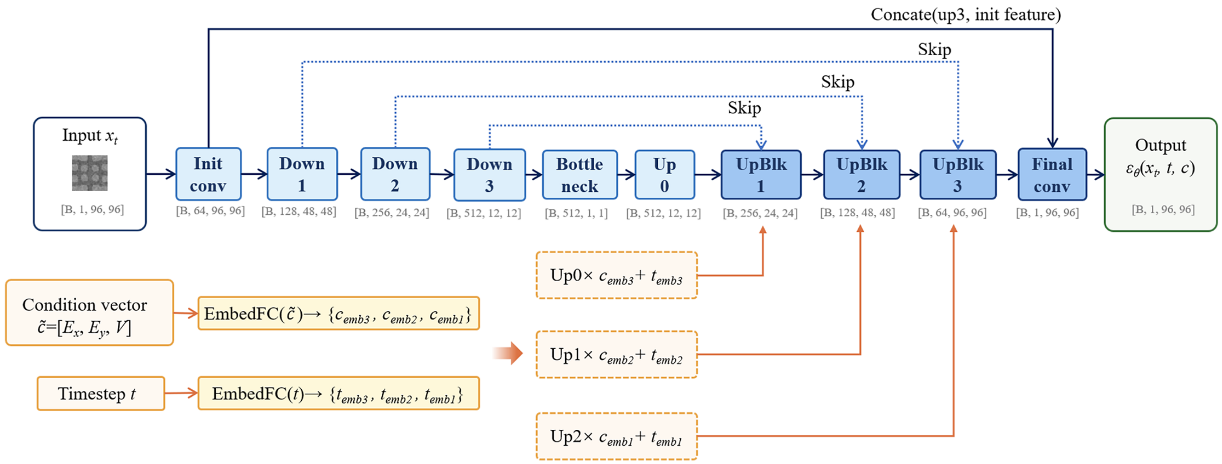

Figure 4: Architecture of the conditional DDPM-UNet for unit-cell generation.

During training, the time step t is randomly sampled, and the noisy sample x_t is constructed according to the forward diffusion process:

$$x_t = \sqrt{\overline{\alpha}_t}x_0 + \sqrt{1-\overline{\alpha}_t}\varepsilon \quad \varepsilon \sim \mathcal{N}(0, \boldsymbol{I}) \tag{14}$$

where $\overline{\alpha}_t$ is the cumulative noise-scheduling coefficient in the diffusion process, and $\boldsymbol{I}$ denotes the identity matrix. The conditional U-Net is used to predict the noise at the current time step:

$$\varepsilon_\theta = \varepsilon_\theta(x_t, t, \tilde{c}) \tag{15}$$

The basic diffusion loss is defined as the mean squared error between the predicted noise and the true noise:

$$\mathrm{L}_{\mathrm{ddpm}} = \left\|\varepsilon - \varepsilon_\theta(x_t, t, \tilde{c})\right\|_2^2 \tag{16}$$

When only the diffusion loss is used, the model mainly learns the morphological features of structures from the data distribution, and the consistency between the generated results and the target mechanical properties may be insufficient. Therefore, after the DDPM-UNet training is completed, a frozen CNN-based surrogate property-prediction network is introduced to impose a property-consistency constraint on the generated structures, and the DDPM-UNet is further fine-tuned.

Based on the predicted noise, the denoised structural estimate corresponding to the current time step can be inferred as:

$$\hat{x}_0 = \frac{x_t - \sqrt{1-\overline{\alpha}_t}\varepsilon_\theta\left(x_t, t, \tilde{c}\right)}{\sqrt{\overline{\alpha}_t}} \tag{17}$$

Subsequently, $\hat{x}_0$ is fed into the frozen CNN-based surrogate property-prediction network to obtain the predicted properties of the generated structure:

$$\hat{c}_{pred} = f_{cnn}(\hat{x}_0) \tag{18}$$

To enforce consistency between the properties of the inferred structure and the target condition, the property-consistency loss is constructed as:

$$\mathrm{L_{phys}} = \mathcal{L}(\hat{c}_{pred}, \tilde{c}) \tag{19}$$

Finally, the total loss function of the diffusion model is defined as:

$$\mathrm{L} = \mathrm{L_{ddpm}} + w_p \mathrm{L_{phys}} \tag{20}$$

where w_p is the weighting coefficient of the property-consistency loss.

In the structure generation stage, given a target property condition $c*$, multiple candidate unit-cell structures are generated through the reverse diffusion process. The frozen CNN is then used for rapid property prediction and error-based screening of the candidate structures, so that the final selected candidates are closer to the prescribed target properties.

Based on the above procedures, including database construction, CNN surrogate pretraining, and conditional diffusion model training, the proposed DDPM-UNet-CNN unit-cell generation framework is illustrated in Fig. 5. The framework consists of three stages. First, a unit-cell database with true effective property labels is generated using the inverse homogenization method, and the CNN-based surrogate property-prediction model is pretrained. Second, during the training of the conditional DDPM-UNet, the frozen CNN is introduced to construct the property-consistency loss. Finally, in the inference stage, the trained CNN is used to rapidly evaluate the mechanical properties of the generated candidate structures and perform error-based screening. It should be noted that the inverse design problem is generally nonunique, since different unit-cell topologies may have similar effective properties. Therefore, the proposed model does not learn a deterministic mapping from c to a single topology. Instead, it learns the conditional distribution $p(x_0|c)$. For the same target property vector, multiple candidate structures can be generated from different initial Gaussian noises and then screened by the CNN surrogate model and FEA verification.

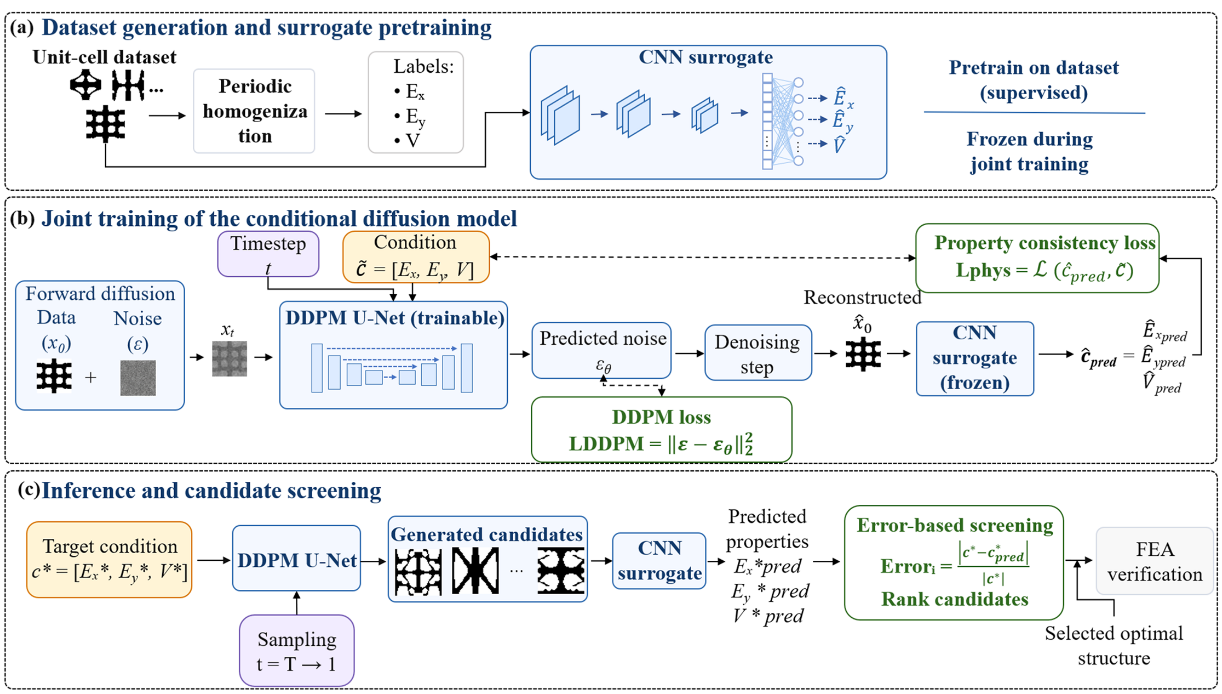

Figure 5: Flowchart of the DDPM-UNet-CNN training and inference process.

3 Results

3.1 CNN Surrogate Validation

To evaluate the predictive capability of the CNN surrogate model, the predicted and true values of E_x, E_y, and the volume fraction V were compared on the test set. As shown in Fig. 6, the scatter points of the three parameters are distributed close to the ideal prediction line, indicating that the CNN surrogate model can effectively learn the mapping relationship between unit-cell geometric features and the corresponding property parameters.

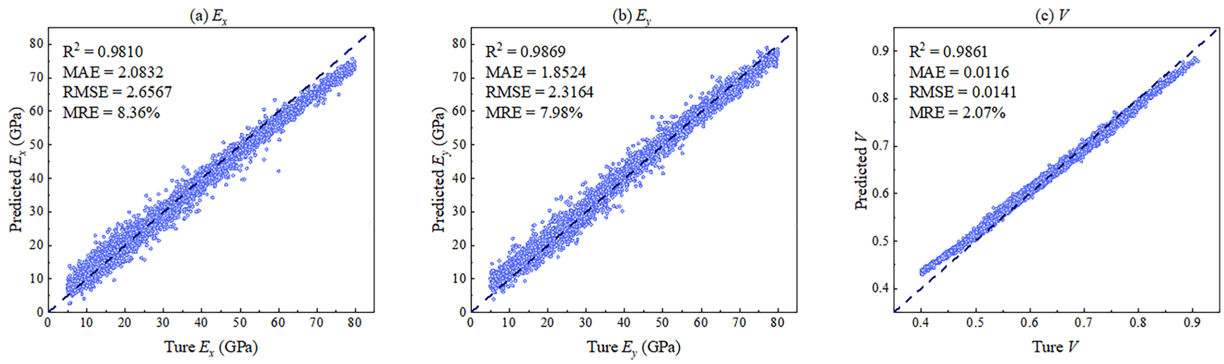

Figure 6: Predictive performance of the CNN surrogate model on the test set: (**a**) E_x; (**b**) E_y; (**c**) V.

In addition, the R_2 values for E_x and E_y are 0.9810 and 0.9869, respectively, with mean relative errors of 8.36% and 7.98%. These results indicate that the model achieves high accuracy in predicting the effective mechanical properties. For V, the R_2 value reaches 0.9861, and the mean relative error is only 2.07%. This is because V is directly related to the proportion of the solid phase in the unit cell and can be determined more explicitly. In contrast, E_x and E_y are strongly affected by the topological configuration, structural connectivity, and stress–strain distribution, leading to a more complex structure–property relationship and a higher prediction difficulty. Overall, the CNN surrogate model exhibits good accuracy in predicting E_x, E_y, and V of the unit cells. Therefore, it can be used as a property-supervision module in the subsequent fine-tuning stage of the DDPM-UNet and for rapid evaluation and screening of generated candidate structures.

3.2 Comparative Analysis of Generation Results

Fig. 7 records the loss evolution during the two-stage training process. The left dashed region corresponds to the initial DDPM-UNet training stage, while the right dashed region represents the CNN-assisted fine-tuning stage with the property-consistency constraint. In the first stage, the diffusion loss drops rapidly and then approaches a stable level, showing that the DDPM-UNet has learned the main geometric patterns of the unit cells in the training set. After the CNN surrogate model is introduced, the total loss, CNN-based property loss, and diffusion loss decrease gradually during fine-tuning. This trend shows that the property-consistency term can be coupled with the diffusion training process without causing unstable optimization.

During this fine-tuning stage, the CNN surrogate model was kept frozen, and its parameters were not updated. It was only used as a fixed property predictor to calculate the property-consistency loss between the generated structures and the target property labels. At this stage, the diffusion loss remains relatively smooth, which means that the CNN constraint does not weaken the original structure-generation capability of the diffusion model. Instead, it helps guide the generated structures toward the prescribed property conditions. After approximately 120–140 epochs, both the property loss and the total loss decrease more slowly and begin to level off. At this stage, the model has nearly reached convergence, and further training brings only limited improvement. The final model was selected based on the target-property errors of the generated structures estimated by the frozen CNN surrogate model. FEA-based re-homogenization was only performed after generation to verify the selected candidate unit cells and was not involved in model training or model selection.

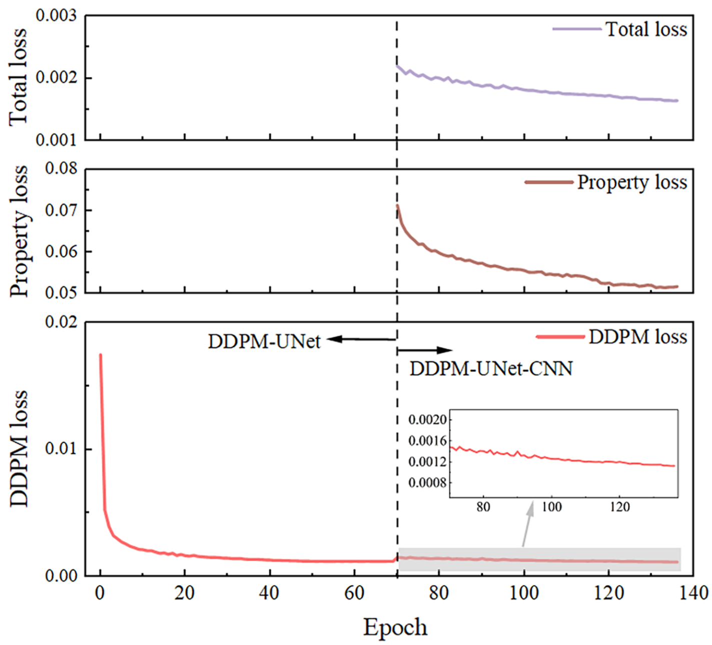

Figure 7: Training loss curves of the DDPM-UNet and DDPM-UNet-CNN stages.

The target-property-driven generation performance of DDPM-UNet-CNN was evaluated using four representative target categories: E_x dominant, E_y dominant, balanced, and mild anisotropic cases. The corresponding target conditions were set to [45, 20, 0.6], [20, 45, 0.6], [30, 30, 0.6], and [40, 30, 0.6], respectively.

Table 1 summarizes the number of valid generated structures under 5%, 10%, and 20% relative error thresholds. For each target type, each model generated 1000 candidate structures, and the generated samples were evaluated using the same frozen CNN-based property predictor. A generated structure was regarded as valid when the relative errors between the CNN-predicted properties and the target properties were lower than the prescribed threshold.

Table 1: Comparison of valid generated samples under 5%, 10%, and 20% relative error thresholds.

Threshold	5%		10%		20%	
Target Types	DDPM-UNet-CNN	DDPM-UNet	DDPM-UNet-CNN	DDPM-UNet	DDPM-UNet-CNN	DDPM-UNet
E_x dominant	185	156	210	198	537	425
E_y dominant	156	133	186	153	513	399
Balanced	167	138	215	182	465	326
Mild anisotropic	254	201	354	296	793	693

As expected, the number of valid samples decreases as the error threshold becomes stricter. Compared with the DDPM-UNet model, the DDPM-UNet-CNN model produces more valid candidates under all three thresholds and across all four target types. This indicates that the CNN-assisted property consistency constraint improves the correspondence between generated structures and target properties. At the 5% threshold, the valid sample ratio increases by 2.3%–5.3%. At the 10% threshold, the increase is 1.2%–5.8%. When the threshold is relaxed to 20%, the improvement becomes more evident, with an increase of 10.0%–13.9%. These results indicate that CNN-assisted property-consistency fine-tuning improves the probability of obtaining property-consistent candidates, although the improvement should be interpreted as an increase in valid-candidate yield rather than a guarantee of exact property matching for every generated structure.

Since DDPM-UNet-CNN achieved higher valid ratios, the subsequent FEA-based verification was performed on this final model. For each target type, candidates satisfying the 5% CNN-predicted error criterion were first selected, and the top 100 candidates with the smallest CNN-predicted errors were saved for FEA-based re-homogenization.

Fig. 8 shows the consistency between the target elastic moduli and the FEA-recalculated elastic moduli of the selected DDPM-UNet-CNN candidates. The blue dashed line denotes ideal agreement ($y = x$), and the two black solid lines represent the ±5% relative error bounds. Most E_x and E_y values are located close to the ideal agreement line, showing good agreement between the generated structures and the target elastic moduli. The 5% threshold was used for CNN-based candidate screening, not as a strict FEA-based error bound. Due to the prediction discrepancy between the CNN surrogate model and FEA-based homogenization, some selected candidates may exhibit FEA-recalculated errors larger than 5%.

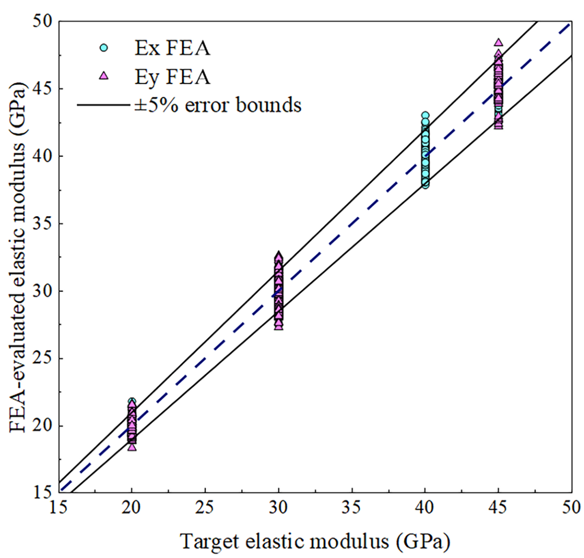

Figure 8: FEA-based consistency scatter plot of selected DDPM-UNet-CNN candidates.

The FEA-based relative errors of E_x, E_y, and V were further calculated for the 100 selected candidates in each target type. Table 2 reports the mean error, maximum error, and standard deviation. The mean errors of E_x and E_y are generally within 2%–4%, while the mean error of V remains close to or below 1% for most target types. These results confirm that the selected candidates preserve good property consistency after FEA-based verification.

Table 2: FEA-based error statistics of selected DDPM-UNet-CNN candidates.

Target Type	E_x Dominant	E_y Dominant	Balanced	Mild Anisotropic
N	100	100	100	100
E_x mean (%)	2.375	2.370	3.591	2.636
E_x max (%)	5.866	9.038	8.378	7.617
E_x std (%)	1.345	1.665	1.996	1.555
E_y mean (%)	2.655	2.083	3.253	3.264
E_y max (%)	8.294	7.553	8.785	8.957
E_y std (%)	1.842	1.578	1.937	1.996
V mean (%)	0.451	0.836	0.883	1.155
V max (%)	1.042	3.067	3.299	5.035
V std (%)	0.314	0.834	0.849	1.272

It should be noted that the CNN surrogate was used both in the property-consistency loss during DDPM fine-tuning and in candidate screening. This may introduce surrogate-model bias, because the CNN is only an approximation of FEA-based homogenization. Therefore, a low CNN-predicted error does not necessarily guarantee a low FEA-recalculated error. For this reason, the CNN-based results were used only for preliminary screening, while the final evaluation was based on FEA-based re-homogenization. The mean error, maximum error, and standard deviation reported in Table 2 were calculated from FEA-recalculated properties rather than CNN predictions.

Fig. 9 shows representative unit-cell structures generated by DDPM-UNet-CNN under four target categories, together with their FEA-recalculated effective properties. Clear differences in topology can be seen under different target conditions. For the E_x dominant case, the generated structures develop continuous load-bearing paths mainly along the horizontal direction, which leads to a relatively high E_x. For the E_y dominant case, better connectivity appears in the vertical direction, resulting in a relatively high E_y. Under the balanced target condition, the structures exhibit comparable connectivity in both the x and y directions, so that E_x and E_y remain close to each other. For the mild anisotropic case, the generated structures retain bidirectional load-bearing capability while still showing a directional preference with E_x slightly higher than E_y. For DDPM-UNet-CNN, the representative FEA-recalculated results under the E_x dominant, E_y dominant, balanced, and mild anisotropic targets are [43.6, 20.3, 0.6], [20.5, 45.9, 0.6], [33.9, 32.3, 0.6], and [39.3, 32.8, 0.6], respectively, showing good agreement with the prescribed target properties.

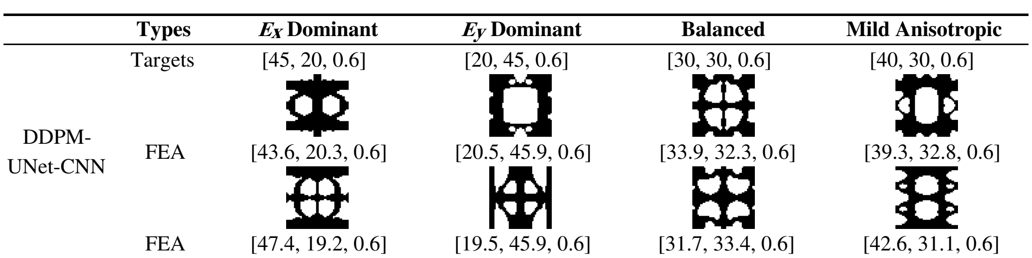

Figure 9: Representative unit-cell structures generated by DDPM-UNet-CNN under four target conditions.

3.3 Controllable Structure Generation under Multiple Target Conditions

To test whether DDPM-UNet-CNN can respond to a wider range of property inputs, 18 target conditions were further designed. These targets were formed by combining six groups of effective elastic moduli with three volume-fraction levels. The six modulus groups represent low-modulus, E_x dominant, E_y dominant, balanced-modulus, mild anisotropic, and high-modulus cases. For each modulus group, low, medium, and high volume fractions were assigned to examine whether the model could adjust the generated topology according to both stiffness and material-usage requirements. The complete target settings are given in Table 3.

Table 3: Target property settings for DDPM-UNet-CNN generation evaluation.

Target Types	E_x (GPa)	E_y (GPa)	$V1$	$V2$	$V3$
Low modulus	20	20	0.45	0.50	0.55
E_x dominant	45	20	0.50	0.55	0.60
E_y dominant	20	45	0.50	0.55	0.60
Balanced	40	40	0.55	0.60	0.65
Mild anisotropic	50	40	0.65	0.70	0.75
High modulus	65	65	0.75	0.80	0.85

Fig. 10 shows several representative unit-cell structures generated by DDPM-UNet-CNN under different target conditions, together with their FEA-recalculated effective properties. The generated topologies vary clearly with the prescribed property combinations. For low-modulus targets, the unit cells contain larger void regions and only a limited number of continuous load-bearing paths. As the target volume fraction and effective moduli increase, the generated structures become more compact, with more connected solid regions. When the target stiffness is dominated by a single principal direction, either E_x or E_y, the material distribution tends to form stronger load-bearing paths along the corresponding direction. Under balanced-modulus targets, the connectivity in the x and y directions becomes more comparable. For mild anisotropic targets, the generated structures still maintain bidirectional connectivity, but a certain directional preference can be observed. These FEA-recalculated results show that DDPM-UNet-CNN can generate unit-cell structures that respond reasonably to different $[E_x, E_y, V]$ conditions.

A nearest-neighbor morphological comparison was further conducted to examine whether the generated structures were directly copied from the training database. In Fig. 10, "Nearest DB" denotes the database sample with the closest morphology to the generated structure, and "Difference Map" gives the pixel-wise difference between them. Black regions in the difference map correspond to different pixels, whereas white regions indicate identical pixels. The generated structures share similar global shapes and main skeletal features with their nearest database samples, which suggests that they remain within the structural distribution learned from the database. However, the local differences shown in the difference maps also indicate that the generated structures are not direct reproductions of existing samples. Instead, DDPM-UNet-CNN produces new local topological variations while preserving the main structural characteristics learned from the training data.

Although connected-component analysis and periodic-boundary correction were used in this work, practical manufacturability was not fully enforced in the current generation process. Minimum feature size, stress concentration, and additive-manufacturing constraints were not included as explicit hard constraints. Therefore, the generated structures should be regarded as candidate designs that still require further manufacturability checking and FEA-based mechanical verification before practical application.

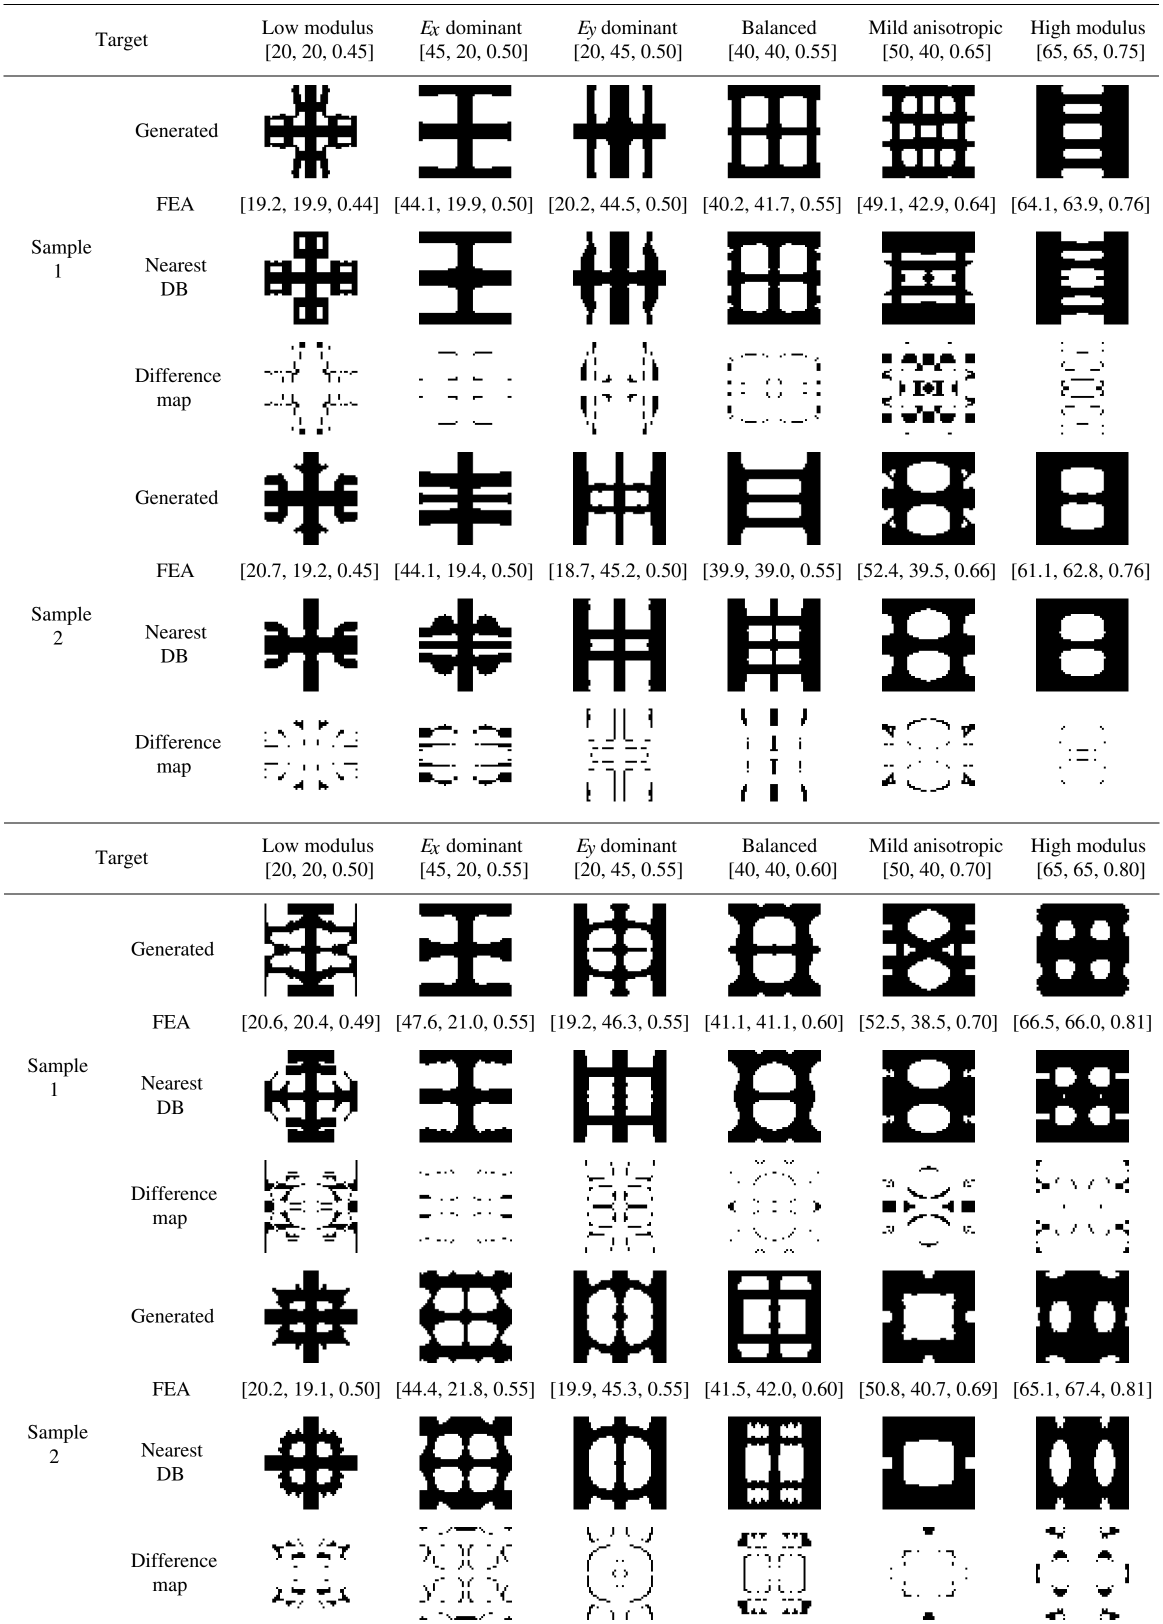

Figure 10: (Continued)

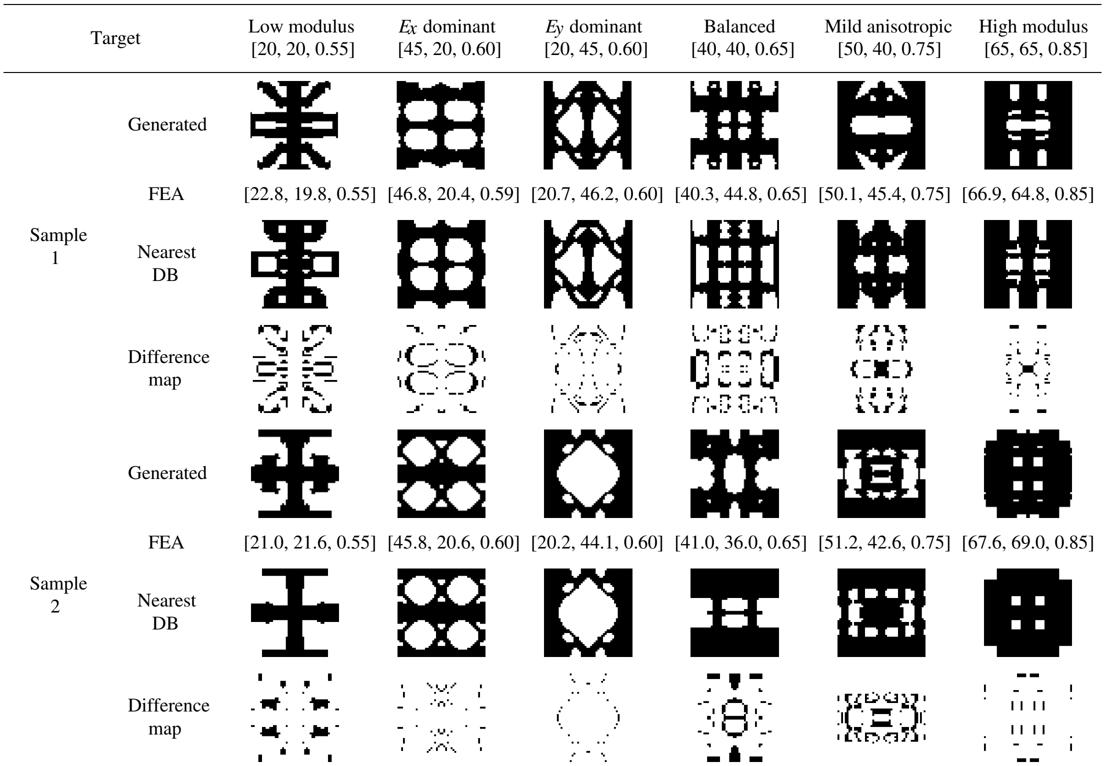

Figure 10: Generated structures, FEA-recalculated effective properties, and nearest-neighbor database comparisons under different target conditions using DDPM-UNet-CNN.

4 Conclusions

In this work, a conditional diffusion-based inverse design framework was developed for two-dimensional anisotropic mechanical metamaterial unit cells. The framework combines DDPM-UNet with a CNN surrogate model, so that the generation process can be guided by prescribed effective elastic moduli and volume fraction. A topology-diverse structure–property database was first constructed, and the CNN surrogate was then used to introduce property supervision during model fine-tuning and to support candidate evaluation during inference. The main conclusions are as follows.

(1) A database containing 57,000 two-dimensional binary unit cells was generated using a random target-property-driven inverse homogenization topology optimization method. The samples cover different effective moduli, volume fractions, and anisotropy levels, providing a sufficiently diverse dataset for conditional diffusion model training.

(2) By introducing the frozen CNN surrogate model into DDPM-UNet, a differentiable property-consistency term was added during fine-tuning. Compared with the original DDPM-UNet, the fine-tuned DDPM-UNet-CNN model produced more valid candidates under all tested relative error thresholds. The valid sample ratio increased by 2.3%–5.3% at the 5% threshold, 1.2%–5.8% at the 10% threshold, and 10.0%–13.9% at the 20% threshold. This result suggests that CNN-assisted fine-tuning improves the probability of obtaining property-consistent candidates from the generated sample pool.

(3) The results under multiple target conditions show that DDPM-UNet-CNN can adjust the generated topology according to different E_x, E_y, and V inputs. With increasing target volume fraction and effective moduli, the generated unit cells tend to become denser. When the target condition is dominated by E_x or E_y, clearer load-bearing paths are formed along the corresponding direction.

These results indicate that DDPM-UNet-CNN can generate two-dimensional anisotropic mechanical metamaterial unit cells under prescribed property conditions and improve the efficiency of target-property matching in microstructural inverse design.

Acknowledgement: The authors acknowledge the support provided by the School of Mechanical Engineering, University of Shanghai for Science and Technology, Shanghai, China.

Funding Statement: This study was supported by the National Natural Science Foundation of China (Grant No. 52375257) and the Shanghai Pujuang Program (24PJD073).

Author Contributions: Hao Sun: Software, data curation, and writing—original draft. Xiaohong Ding: Conceptualization, methodology, funding acquisition, and project administration. Min Xiong: Resources and investigation. Heng Zhang: Software guidance and algorithm design supervision. All authors reviewed and approved the final version of the manuscript.

Availability of Data and Materials: The datasets generated during the current study are available from the corresponding author upon reasonable request. The codes used for database generation, model training, and analysis are also available from the corresponding author upon reasonable request.

Ethics Approval: Not applicable. This study did not involve human participants, animal experiments, clinical data, or personal information.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Surjadi JU, Gao L, Du H, Li X, Xiong X, Fang NX, et al. Mechanical metamaterials and their engineering applications. Adv Eng Mater. 2019;21(3):1800864. doi:10.1002/adem.201800864.
2. Bertoldi K, Vitelli V, Christensen J, van Hecke M. Flexible mechanical metamaterials. Nat Rev Mater. 2017;2(11):17066. doi:10.1038/natrevmats.2017.66.
3. Zheng X, Smith W, Jackson J, Moran B, Cui H, Chen D, et al. Multiscale metallic metamaterials. Nat Mater. 2016;15(10):1100–6. doi:10.1038/nmat4694.
4. Clausen A, Wang F, Jensen JS, Sigmund O, Lewis JA. Topology optimized architectures with programmable Poisson's ratio over large deformations. Adv Mater. 2015;27(37):5523–7. doi:10.1002/adma.201502485.
5. Zeng Z, Han C, Sun CN, Kandukuri S, Zhu Y, Huang J, et al. Tunable isotropy on the mechanical properties of wavy hexachiral metamaterials: numerical simulation and additive manufacturing. Smart Manuf. 2022;1:2250001. doi:10.1142/s2737549822500016.
6. Shan S, Kang SH, Zhao Z, Fang L, Bertoldi K. Design of planar isotropic negative Poisson's ratio structures. Extrem Mech Lett. 2015;4:96–102. doi:10.1016/j.eml.2015.05.002.
7. Lu C, Hsieh M, Huang Z, Zhang C, Lin Y, Shen Q, et al. Architectural design and additive manufacturing of mechanical metamaterials: a review. Engineering. 2022;17:44–63. doi:10.1016/j.eng.2021.12.023.
8. Lakes R. Foam structures with a negative Poisson's ratio. Science. 1987;235(4792):1038–40. doi:10.1126/science.235.4792.1038.
9. Zheng X, Lee H, Weisgraber TH, Shusteff M, DeOtte J, Duoss EB, et al. Ultralight, ultrastiff mechanical metamaterials. Science. 2014;344(6190):1373–7. doi:10.1126/science.1252291.
10. Liu Y, Wang Y, Ren H, Meng Z, Chen X, Li Z, et al. Ultrastiff metamaterials generated through a multilayer strategy and topology optimization. Nat Commun. 2024;15(1):2984. doi:10.1038/s41467-024-47089-8.
11. Karathanasopoulos N, Reda H, Ganghoffer JF. Designing two-dimensional metamaterials of controlled static and dynamic properties. Comput Mater Sci. 2017;138:323–32. doi:10.1016/j.commatsci.2017.06.035.
12. Kumar S, Tan S, Zheng L, Kochmann DM. Inverse-designed spinodoid metamaterials. npj Comput Mater. 2020;6:73. doi:10.1038/s41524-020-0341-6.

13. Bensoussan A, Lions JL, Papanicolaou G. Asymptotic analysis for periodic structures. Providence, RI, USA: American Mathematical Society; 2011. doi:10.1090/chel/374.
14. Sigmund O. Materials with prescribed constitutive parameters: an inverse homogenization problem. Int J Solids Struct. 1994;31(17):2313–29. doi:10.1016/0020-7683(94)90154-6.
15. Andreassen E, Andreasen CS. How to determine composite material properties using numerical homogenization. Comput Mater Sci. 2014;83:488–95. doi:10.1016/j.commatsci.2013.09.006.
16. Lee D, Chen WW, Wang L, Chan YC, Chen W. Data-driven design for metamaterials and multiscale systems: a review. Adv Mater. 2024;36(8):e2305254. doi:10.1002/adma.202305254.
17. Zheng X, Zhang X, Chen TT, Watanabe I. Deep learning in mechanical metamaterials: from prediction and generation to inverse design. Adv Mater. 2023;35(45):2302530. doi:10.1002/adma.202302530.
18. Huang J, Mai H, Wan H, Chen J, Chen R, He T. Deep learning-based intelligent inverse design of cylindrical randomized plate acoustic metamaterials. Appl Acoust. 2026;243:111136. doi:10.1016/j.apacoust.2025.111136.
19. Bastek JH, Kumar S, Telgen B, Glaesener RN, Kochmann DM. Inverting the structure-property map of truss metamaterials by deep learning. Proc Natl Acad Sci U S A. 2022;119(1):e2111505119. doi:10.1073/pnas.2111505119.
20. Ha CS, Yao D, Xu Z, Liu C, Liu H, Elkins D, et al. Rapid inverse design of metamaterials based on prescribed mechanical behavior through machine learning. Nat Commun. 2023;14:5765. doi:10.1038/s41467-023-40854-1.
21. Bastek JH, Kochmann DM. Inverse design of nonlinear mechanical metamaterials via video denoising diffusion models. Nat Mach Intell. 2023;5(12):1466–75. doi:10.1038/s42256-023-00762-x.
22. Thakolkaran P, Espinal M, Dhulipala S, Kumar S, Portela CM. Experiment-informed finite-strain inverse design of spinodal metamaterials. Extrem Mech Lett. 2025;74:102274. doi:10.1016/j.eml.2024.102274.
23. Du Z, Bian T, Ren X, Jia Y, Tang S, Cui T, et al. Inverse design of mechanical metamaterial achieving a prescribed constitutive curve. Theor Appl Mech Lett. 2024;14(1):100486. doi:10.1016/j.taml.2023.100486.
24. Peng XL, Xu BX. Data-driven inverse design of composite triangular lattice structures. Int J Mech Sci. 2024;265:108900. doi:10.1016/j.ijmecsci.2023.108900.
25. Chai Z, Zong Z, Yong H, Ke X, Zhu J, Ding H, et al. Tailoring stress-strain curves of flexible snapping mechanical metamaterial for on-demand mechanical responses via data-driven inverse design. Adv Mater. 2024;36(33):2404369. doi:10.1002/adma.202404369.
26. Ho J, Jain A, Abbeel P. Denoising diffusion probabilistic models. Adv Neural Inf Process Syst. 2020;33:6840–51.
27. Nichol AQ, Dhariwal P. Improved denoising diffusion probabilistic models. In: Proceedings of the 38th International Conference on Machine Learning; 2021 Jul 18–24; Online. p. 8162–71.
28. Lyu X, Ren X. Microstructure reconstruction of 2D/3D random materials via diffusion-based deep generative models. Sci Rep. 2024;14(1):5041. doi:10.1038/s41598-024-54861-9.
29. Wang H, Du Z, Feng F, Kang Z, Tang S, Guo X. DiffMat: data-driven inverse design of energy-absorbing metamaterials using diffusion model. Comput Meth Appl Mech Eng. 2024;432:117440. doi:10.1016/j.cma.2024.117440.
30. Du L, Hu J, Wang S, Jiang Y, Lei N, He Y, et al. Topo-GenMeta: generative design of metamaterials based on diffusion model with attention to topology. Comput Aided Des. 2026;190:103977. doi:10.1016/j.cad.2025.103977.
31. Kingma DP, Ba J. Adam: a method for stochastic optimization. In: Proceedings of the International Conference on Learning Representations; 2015 May 7–9; San Diego, CA, USA. p. 1–13.
32. Tavakoli R, Zhang H. A nonmonotone spectral projected gradient method for large-scale topology optimization problems. Numer Algebra Control Optim. 2012;2(2):395–412. doi:10.3934/naco.2012.2.395.
33. Guest JK, Prévost JH, Belytschko T. Achieving minimum length scale in topology optimization using nodal design variables and projection functions. Numer Meth Eng. 2004;61(2):238–54. doi:10.1002/nme.1064.
34. LeCun Y, Bottou L, Bengio Y, Haffner P. Gradient-based learning applied to document recognition. Proc IEEE. 1998;86(11):2278–324. doi:10.1109/5.726791.
35. Krizhevsky A, Sutskever I, Hinton GE. ImageNet classification with deep convolutional neural networks. Commun ACM. 2017;60(6):84–90. doi:10.1145/3065386.

36. Sohl-Dickstein J, Weiss EA, Maheswaranathan N, Ganguli S. Deep unsupervised learning using nonequilibrium thermodynamics. In: Proceedings of the 32nd International Conference on Machine Learning; 2015 Jul 7–9; Lille, France. p. 2256–65.
37. Li G, Liu D. Inverse design of composite materials with desired mechanical behaviors based on conditional diffusion model. Eng Struct. 2025;345:121417. doi:10.1016/j.engstruct.2025.121417.
38. Ronneberger O, Fischer P, Brox T. U-Net: convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015. Cham, Switzerland: Springer; 2015. p. 234–41. doi:10.1007/978-3-319-24574-4_28.