

Article

Not peer-reviewed version

HyperDiff: An Inverse Design Framework for Hyperelastic Microstructures Based on a Conditional Diffusion Model

[Yu Zhang](#)[†], [Lijie Liu](#)[†], Shukai Li, [K.I. Elkhodary](#), [Zongliang Du](#)^{*}, [Shan Tang](#)^{*}

Posted Date: 24 November 2025

doi: 10.20944/preprints202511.0619.v2

Keywords: hyperelastic microstructures; inverse design; conditional diffusion models; large deformation; generative modeling



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

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.

Article

HyperDiff: An Inverse Design Framework for Hyperelastic Microstructures Based on a Conditional Diffusion Model

Yu Zhang ^{1,2,†}, Lijie Liu ^{1,2,†}, Shukai Li ^{1,2}, K.I. Elkhodary ³, Zongliang Du ^{1,2,*}, Shan Tang ^{1,2,*}

¹ State Key Laboratory of Structural Analysis, Optimization and CAE Software for Industrial Equipment, Department of Engineering Mechanics, Dalian University of Technology, Dalian 116024, China

² International Research Center for Computational Mechanics, Dalian University of Technology, Dalian 116024, China

³ Department of Mechanical Engineering, The American University in Cairo, 11835 New Cairo, Egypt

* Correspondence: zldu@dlut.edu.cn (Z.D.); shantang@dlut.edu.cn (S.T.)

† These authors contributed equally to this work.

Abstract

Designing hyperelastic porous microstructures under finite strain is challenging because bending, buckling, contact, and densification interact to produce nonconvex and one-to-many relations between topology and response. We present *HyperDiff*, a conditional diffusion framework that reformulates inverse design as probabilistic sampling rather than deterministic regression. A compact B-spline encoding of the target force–displacement curve captures the system's energy–evolution trend, providing temporal and mechanical context that guides the denoising process toward physically consistent configurations with the desired multi-stage deformation behavior. The workflow integrates Gaussian random field (GRF)-based topology generation, constitutive calibration, large-deformation finite-element simulations, and quasi-static compression experiments. Across held-out and interpolated targets, the generated microstructures accurately reproduce sequential deformation stages (bending–buckling–densification) and global responses, with deviations typically below 10%, while preserving manufacturability and one-to-many design diversity. The current implementation focuses on two-dimensional unit cells under quasi-static compression, yet the framework is extensible to 3D, multi-resolution, and multi-physics systems. By combining physics-aware conditioning with generative sampling, *HyperDiff* establishes a practical front end for mechanics-based design workflows, applicable to programmable soft actuators, impact-energy absorbers with tunable plateaus, and rapid exploration of nonlinear architected materials for soft and deformable systems.

Keywords: hyperelastic microstructures; inverse design; conditional diffusion models; large deformation; generative modeling

1. Introduction

Mechanical metamaterials are architected solids whose effective properties are governed primarily by geometry rather than base material composition. These engineered structures enable functionalities such as lightweight energy absorption, programmable deformation, multistability, and shape morphing, with applications in aerospace structures, biomedical devices, soft robotics, and protective systems [1–7]. Recent work has shown that controlled buckling, snap-through, and phase-transition-like transformations can be used to realize tunable force–displacement or stress–strain responses on demand, rather than relying on the intrinsic response of a homogeneous material [8,9]. As a result, the design of metamaterials is increasingly viewed not only as structural layout design, but as direct control over nonlinear mechanical behavior for specific functions [10–13].

Although many classes of high-performance metamaterials have been demonstrated, their design remains challenging. Conventional “forward” workflows typically begin with a candidate topology

inspired by natural templates such as cellular, porous, or origami-like patterns, and then iterate through simulation and manual refinement [8,14,15]. This approach has produced notable successes, including ultralight lattices and multistable snapping systems, but it is time-consuming, problem-specific, and scales poorly when the goal is to match a prescribed mechanical response under large deformation [5,7]. In modern applications, designers increasingly require materials whose global response is staged or switchable — for example, an initially compliant regime for impact mitigation, followed by a stiffening or locking regime for load-bearing [4,6,9]. Achieving such behavior by manual tuning is difficult because performance, manufacturability, and stability must be balanced simultaneously [10,12,16].

Inverse design aims to address this challenge by reversing the traditional workflow: instead of starting from geometry and predicting behavior, it starts from the desired mechanical response and searches for a compatible microstructure [17–19]. This paradigm is particularly attractive for metamaterials, where the mapping from geometry to response is often nonlinear, history-dependent, and governed by local instabilities such as buckling and self-contact [20]. With recent advances in machine learning, inverse design is increasingly seen as a path toward performance-by-design structures in both structural and functional materials [10–13,21].

Current machine-learning-based inverse design strategies can be broadly grouped into two categories. The first category is *parametric* or low-dimensional design. In these approaches, the metamaterial is represented using a finite set of geometric parameters, and surrogate models or neural networks are trained to map those parameters to mechanical performance [22–25]. These models can then be coupled with heuristic or evolutionary optimizers to efficiently explore the parameter space and produce customized load–displacement curves for specific structure families [22,26]. Although computationally efficient and interpretable, such methods are constrained by the assumed parameterization. They typically operate within a predefined family of unit cells (e.g., shells, graded lattices, or spinoid structures), and may struggle to capture drastic topological changes or complex post-buckling behavior [24–26].

The second category is *non-parametric* generative design, in which the microstructure is treated directly as an image-, voxel-, or implicit-surface field rather than a hand-crafted geometric template [27–30]. This class of methods leverages deep generative models to learn high-dimensional structure–property relationships from data. Convolutional neural networks and generative adversarial networks (GANs) have been used to synthesize topologies with target stiffness or target bulk modulus, and to encode the geometry of cellular materials in an implicit representation [27,28,31,32]. Variational autoencoders (VAEs) take a related approach by embedding complex microstructures into a smooth latent space, which can then be searched or optimized for target properties [26,33,34].

More recently, diffusion models have emerged as a powerful alternative for generative inverse design. These models iteratively transform random noise into structured outputs by learning a reverse denoising process [35–38]. Unlike GANs, diffusion models are known for their sampling stability and controllability. They have been applied to reconstruct complex microstructures, generate energy-absorbing metamaterials tailored to specified loading curves, and model time-dependent deformation fields in nonlinear metamaterials under load [39–44]. These advances suggest that generative modeling can move beyond “shape suggestion” and toward direct synthesis of structures with prescribed mechanical behavior.

However, important challenges remain for the inverse design of *hyperelastic* metamaterials that undergo large strain, geometric nonlinearity, and path-dependent softening. Under compression, such microstructures often exhibit multiple sequential regimes: an initially compliant phase dominated by bending, followed by buckling and stress softening, and eventually densification accompanied by contact and rapid stiffening [9,20,45,46]. Capturing these regimes is essential for applications such as impact mitigation and reusable energy absorption, yet the mapping from topology to force–displacement behavior is strongly nonconvex and not easily parameterized [22,24,26]. While recent diffusion-based studies have demonstrated controllable generation for elastoplastic or energy-absorbing systems [39,40,44], they differ substantially in how the mechanical response is represented and conditioned.

Vlassis and Sun [39] use a denoising diffusion probabilistic model (DDPM) guided by embedded feature vectors extracted from target hyperelastic energy functionals in the Mechanical Modified National Institute of Standards and Technology (MNIST) dataset to generate microstructures whose effective responses match prescribed nonlinear energy–stretch curves. Bastek and Kochmann [40] cast the inverse design of nonlinear mechanical metamaterials as a video denoising diffusion problem, training on full-field stress and displacement histories of periodic cellular structures made of an elastoplastic thermoplastic resin under large-strain compression so that a target stress–strain curve yields both a candidate design and its predicted deformation path. Wang et al. [44] propose DiffMat, a conditional diffusion framework trained on homogenization-based topology-optimization data to map desired energy-absorbing stress–strain curves to diverse pixelated unit-cell designs for energy-absorbing metamaterials. Nevertheless, there is still limited work on learning a direct mapping from a target nonlinear force–displacement curve to a manufacturable hyperelastic porous topology that remains physically realistic under large deformation, including buckling and post-buckling contact.

A key difficulty in this field lies in the intrinsic ill-posedness of the inverse mapping between geometry and mechanical response. For hyperelastic porous microstructures, the mapping from a target force–displacement curve to geometry is inherently ill-posed for four coupled reasons: (i) non-uniqueness, where distinct topologies can realize nearly identical curves while similar patterns may diverge; (ii) nonconvexity and mode switching introduced by buckling, contact, and densification; (iii) high sensitivity of the response to small geometric perturbations (e.g., manufacturing tolerances); and (iv) the limited descriptive power of low-dimensional features such as volume fraction. These challenges motivate the adoption of a probabilistic, condition-guided generative formulation that can naturally represent one-to-many mappings instead of enforcing a single deterministic inverse. Representative examples illustrating these phenomena are provided in Section 3.

To address this gap, we propose *HyperDiff*, a conditional diffusion-based inverse design framework for hyperelastic porous microstructures with nonlinear mechanical responses, as outlined in Figure 1. The core idea is to treat the desired global response, here expressed as a force–displacement curve, as a conditioning signal for a denoising diffusion probabilistic model (DDPM) [35,36]. The diffusion model is implemented using a U-Net backbone with residual connections and attention-based conditioning [28,37,39,47,48], which enables it to generate candidate microstructure topologies directly in image form. To supply meaningful training data, we construct an integrated pipeline that (i) procedurally generates porous microstructures, (ii) calibrates a hyperelastic constitutive model for the base material using tensile experiments, (iii) simulates large-deformation compression via finite element analysis to obtain full force–displacement curves, and (iv) encodes those curves in a compact form suitable for conditional generation [10,16,45,46,49]. The resulting framework links data generation, inverse design, forward simulation, and experimental validation.

In summary, this work demonstrates that conditional diffusion modeling can be used to generate diverse and physically consistent microstructures that reproduce prescribed nonlinear force–displacement behavior, including regimes dominated by bending, buckling, softening, and densification. By focusing on hyperelastic, large-strain responses (a setting where topology, instability, and contact play central roles), this study aims to contribute to the practical integration of generative AI tools into the design workflow of nonlinear architected materials [9,12,13,21,44].

The remainder of this article is organized as follows. Section 2 presents the theoretical basis of denoising diffusion probabilistic models and details the proposed *HyperDiff* framework. Section 3 describes the dataset generation pipeline, including GRF-based topology synthesis, material characterization, and finite element (FE) simulation of force–displacement responses. Section 4 explains the model training strategy, including conditioning, data normalization, and key hyperparameters. Section 5 reports numerical results on both test and interpolated targets to assess performance and generalization. Section 6 presents experimental verification via quasi-static compression tests and compares measurements with simulations. Finally, Section 8 summarizes the main findings, discusses limitations, and outlines future directions.

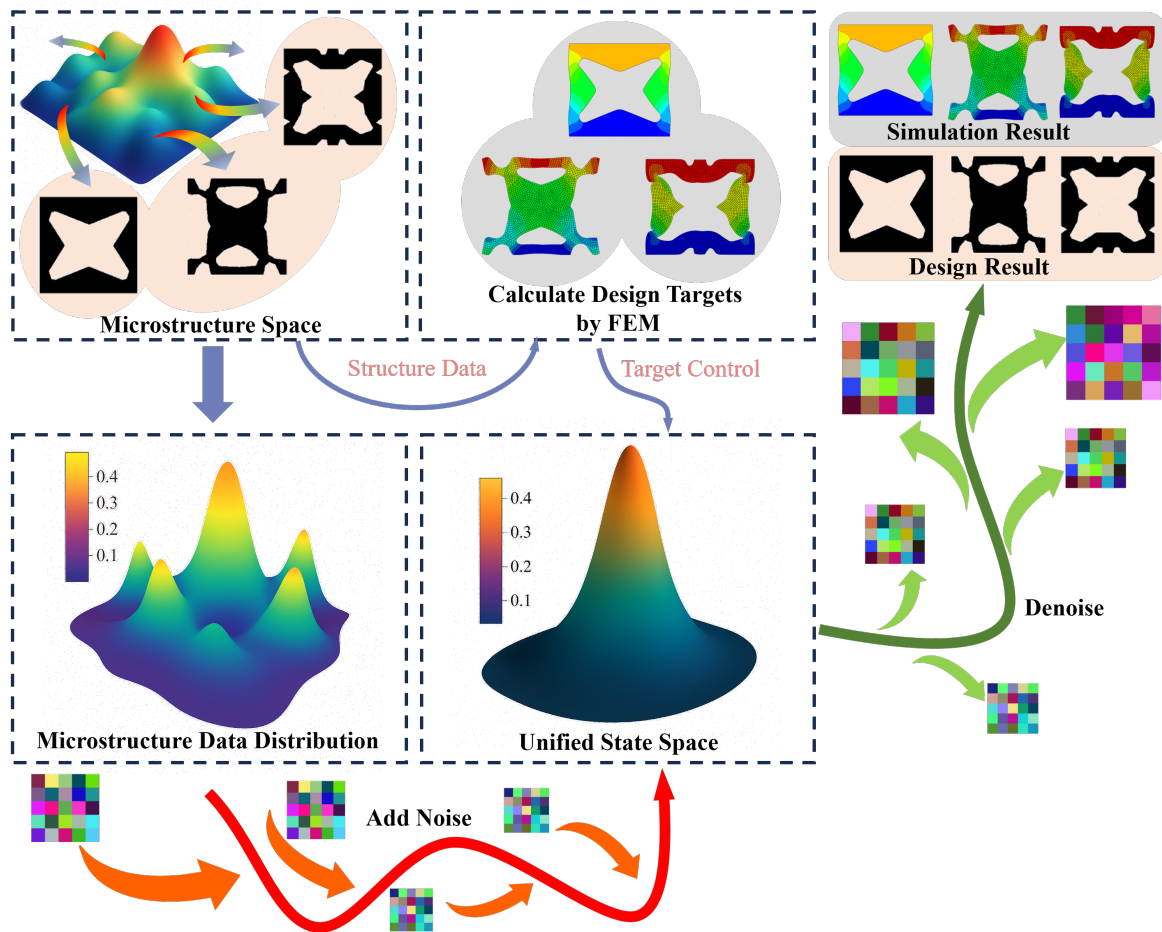


Figure 1. Overall technical roadmap of the proposed *HyperDiff* framework for the inverse design of hyperelastic porous microstructures. The workflow integrates data generation, conditional diffusion modeling, and multi-physics validation. (a) Diverse microstructure topologies are generated using Gaussian random fields (GRF) and analyzed via finite element simulations under quasi-static compression to obtain nonlinear force–displacement curves, forming a large structure–response dataset. (b) The desired target mechanical curve is encoded as a compact conditional vector that guides the diffusion model during denoising. (c) Starting from random noise, the model progressively reconstructs candidate microstructures consistent with the specified mechanical behavior. (d) The framework yields multiple physically admissible and topologically diverse solutions that reproduce the target nonlinear response, enabling one-to-many, data-driven inverse design of hyperelastic metamaterials.

2. Methodology

This section details the theoretical framework of *HyperDiff*. Section 2.1 outlines the principles of the denoising diffusion probabilistic model (DDPM), while Section 2.2 describes the custom neural network architecture developed for this research.

2.1. Generative Diffusion Model Theory

DDPMs comprise two processes: a forward process that incrementally adds noise to the data, and a reverse process that learns to denoise. The reverse process starts with pure Gaussian noise and iteratively refines it into a coherent data sample over a sequence of steps. At each step, a neural network acts as a denoiser, predicting and removing a small amount of noise to gradually recover the underlying data structure. Therefore, the reverse process can be understood as a sequence of denoising operations that collectively reconstruct data from pure noise.

2.1.1. Forward Process

The purpose of the forward process is to systematically and controllably corrupt an initial data sample \mathbf{x}_0 (e.g., a clear microstructure image) into pure, unstructured noise. This is achieved through a fixed Markov chain, where each state \mathbf{x}_t depends only on the preceding state \mathbf{x}_{t-1} , analogous to a time-stepping scheme in a physical simulation. This gradual corruption process can be conceptually related to stochastic relaxation in physical systems, where disorder increases step by step in the forward direction and order is progressively restored during the reverse diffusion. This process unfolds over T discrete time steps, generating a sequence of increasingly noisy latent variables $\mathbf{x}_1, \dots, \mathbf{x}_T$.

At each time index t , the preceding sample is perturbed by a small Gaussian increment. The core of this operation is described by the conditional probability:

$$q(\mathbf{x}_t | \mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t} \mathbf{x}_{t-1}, \beta_t \mathbf{I}), \quad (1)$$

where β_t is a small, positive constant defined by a variance schedule, with $0 < \beta_1 < \beta_2 < \dots < \beta_T < 1$. This schedule controls the magnitude of noise added at each step. The term $\sqrt{1 - \beta_t}$ acts as a scaling factor that slightly shrinks the previous state \mathbf{x}_{t-1} , while the variance $\beta_t \mathbf{I}$ defines the newly added noise.

A key property of this formulation, which makes it computationally tractable, is that we can sample \mathbf{x}_t at any arbitrary time step t directly from the initial data \mathbf{x}_0 , without having to iterate through all intermediate steps. This closed-form solution is obtained by defining $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$, which yields:

$$\mathbf{x}_t = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (2)$$

This equation can be understood as a weighted sum. The terms $\sqrt{\bar{\alpha}_t}$ and $\sqrt{1 - \bar{\alpha}_t}$ represent the scaling factors for the original signal (\mathbf{x}_0) and the noise ($\boldsymbol{\varepsilon}_t$), respectively. As t increases, $\bar{\alpha}_t$ decreases towards zero, causing the contribution from the original data to diminish while the contribution from the noise grows.

This process ensures that as t approaches the final time step T , the cumulative product $\bar{\alpha}_T$ becomes nearly zero. Consequently, the state \mathbf{x}_T loses all information about the original sample \mathbf{x}_0 and converges to an isotropic Gaussian distribution, $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. This provides a well-defined starting point of pure noise from which the reverse process can begin reconstruction.

2.1.2. Reverse Process

The generative mechanism of the model is driven by the reverse process. Its goal is to reverse the noising procedure by starting with pure Gaussian noise $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and iteratively denoising it, step by step, to recover a sample from the original data distribution $q(\mathbf{x}_0)$. This is modeled as a learned Markov chain parameterized by a neural network θ :

$$p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t) = \mathcal{N}(\mathbf{x}_{t-1}; \boldsymbol{\mu}_\theta(\mathbf{x}_t, t), \sigma_\theta^2(\mathbf{x}_t, t) \mathbf{I}), \quad (3)$$

where the neural network learns both the mean $\boldsymbol{\mu}_\theta$ and the scalar variance σ_θ^2 . However, Ho et al. [35] found that fixing the variance to a set of predefined constants and only learning the mean resulted in more stable training and higher-quality samples. Following this standard practice, we set $\sigma_\theta^2(\mathbf{x}_t, t)$ to a non-learned hyperparameter σ_t^2 (typically $\sigma_t^2 = \beta_t$ or $\sigma_t^2 = \tilde{\beta}_t$).¹ The network's task is therefore reduced to learning only the mean $\boldsymbol{\mu}_\theta$.

The primary challenge is that the true reverse conditional probability $q(\mathbf{x}_{t-1} | \mathbf{x}_t)$ is *intractable*. Calculating it would require marginalizing over the entire (unknown) data distribution, which is computationally infeasible. However, a crucial insight is that if we also condition this probability on

¹ The two standard choices for this fixed variance are β_t , which is the variance from the corresponding forward noising step, and $\tilde{\beta}_t$, which is the variance of the ideal posterior $q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0)$ shown in Equation (4). The model is robust to either choice.

the *initial data sample* \mathbf{x}_0 , the posterior becomes *tractable* and follows a Gaussian (normal) probability distribution:

$$q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_{t-1}; \tilde{\boldsymbol{\mu}}(\mathbf{x}_t, \mathbf{x}_0), \tilde{\boldsymbol{\beta}}_t \mathbf{I}), \quad (4)$$

where the mean $\tilde{\boldsymbol{\mu}}$ and variance $\tilde{\boldsymbol{\beta}}_t$ are known functions of the forward process variables $(\alpha_t, \bar{\alpha}_t)$. While \mathbf{x}_0 is unknown during inference, this tractable posterior is invaluable during training, as it provides a target for the neural network.

To make this insight practical for training, we use a *reparameterization trick*. Recall from the forward process (Equation (2)) that any noisy sample \mathbf{x}_t can be expressed in terms of the initial sample \mathbf{x}_0 and the noise $\boldsymbol{\varepsilon}_t$ that was added. By rearranging that equation, we can express \mathbf{x}_0 as:

$$\mathbf{x}_0 = \frac{1}{\sqrt{\bar{\alpha}_t}} \left(\mathbf{x}_t - \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\varepsilon}_t \right). \quad (5)$$

Substituting this expression for \mathbf{x}_0 into the equation for the tractable posterior mean $\tilde{\boldsymbol{\mu}}$ (Equation (4)) allows us to express that mean as a function of the observable state \mathbf{x}_t and the noise $\boldsymbol{\varepsilon}_t$:

$$\tilde{\boldsymbol{\mu}}(\mathbf{x}_t, t) = \frac{1}{\sqrt{\bar{\alpha}_t}} \left(\mathbf{x}_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \boldsymbol{\varepsilon}_t \right). \quad (6)$$

This reformulation reveals a critical simplification: learning the complex posterior mean $\boldsymbol{\mu}_\theta$ is equivalent to learning the noise vector $\boldsymbol{\varepsilon}_t$ that was added at step t . Predicting this noise is a more direct and stable objective for a neural network.

Therefore, instead of training the network to predict $\boldsymbol{\mu}_\theta$ directly, we train a network $\boldsymbol{\varepsilon}_\theta(\mathbf{x}_t, t)$ to predict the noise component $\boldsymbol{\varepsilon}_t$ from the noisy sample \mathbf{x}_t at time step t . The learning objective reduces to minimizing the mean squared error between the ground-truth noise and the predicted noise, a widely adopted and effective loss:

$$\mathcal{L}_{\text{simple}}(\theta) = \mathbb{E}_{t, \mathbf{x}_0, \boldsymbol{\varepsilon}_t} \left[\|\boldsymbol{\varepsilon}_t - \boldsymbol{\varepsilon}_\theta(\mathbf{x}_t, t)\|^2 \right], \quad (7)$$

where $\mathbb{E}_{t, \mathbf{x}_0, \boldsymbol{\varepsilon}_t}$ denotes the expectation over the distribution of data, time steps, and noise, which in practice is approximated by randomly sampling these components at each step of training. Accordingly, all experiments reported below optimize the L2 (mean-squared) loss in Equation (7), and the hyperparameter summary reflects the same objective.

2.1.3. Conditional Generation with Classifier-Free Guidance

To guide the generation process with auxiliary information \mathbf{y} , such as force–displacement data, we employ the *classifier-free guidance* (CFG) strategy [36]. This technique allows a single model to learn both conditional and unconditional generation pathways, enabling robust control over the output at inference time.

The core idea is to modify the training process. With a fixed probability (e.g., 10–20% of the time), the conditioning vector \mathbf{y} is replaced with a learned null token, \emptyset . This forces the network to learn to predict the noise both with and without conditioning information. The training objective is therefore to predict the noise $\boldsymbol{\varepsilon}_t$ from the noisy state \mathbf{x}_t , conditioned on either the true label \mathbf{y} or the null token \emptyset :

$$\mathcal{L}_{\text{CFG}}(\theta) = \mathbb{E}_{t, \mathbf{x}_0, \boldsymbol{\varepsilon}_t, \mathbf{y}} \left[\|\boldsymbol{\varepsilon}_t - \boldsymbol{\varepsilon}_\theta(\mathbf{x}_t, \mathbf{y}, t)\|^2 \right], \quad (8)$$

where \mathbf{y} is randomly replaced by \emptyset for a subset of the training batch.

The true power of CFG is realized during inference. At each denoising step t , the model makes two predictions for the noise: one conditional, $\boldsymbol{\varepsilon}_\theta(\mathbf{x}_t, \mathbf{y}, t)$, and one unconditional, $\boldsymbol{\varepsilon}_\theta(\mathbf{x}_t, \emptyset, t)$. These are combined to form a guided prediction $\hat{\boldsymbol{\varepsilon}}_t$:

$$\hat{\boldsymbol{\varepsilon}}_t(\mathbf{x}_t, \mathbf{y}, t) = \boldsymbol{\varepsilon}_\theta(\mathbf{x}_t, \emptyset, t) + w(\boldsymbol{\varepsilon}_\theta(\mathbf{x}_t, \mathbf{y}, t) - \boldsymbol{\varepsilon}_\theta(\mathbf{x}_t, \emptyset, t)), \quad (9)$$

where w is the *guidance scale*, a hyperparameter that controls the strength of the conditioning. When $w = 0$, generation is purely unconditional. When $w > 1$, the model is pushed to generate samples that more strongly adhere to the condition \mathbf{y} . This guided noise estimate $\hat{\epsilon}_t$ is then used in the reverse process equations (e.g., Equation (6)) to compute the mean of the denoised sample \mathbf{x}_{t-1} .

2.2. HyperDiff Architecture

To achieve the inverse design objectives of the *HyperDiff* framework, the neural network must predict the noise term ϵ_t at each reverse diffusion step, conditioned on the target force–displacement curve \mathbf{y} . This requires the model to extract features across multiple spatial scales while coupling local geometric patterns with global mechanical constraints. A U-Net backbone, with its symmetric encoder–decoder structure and skip connections, provides a suitable foundation for this purpose.

The proposed architecture builds upon the conventional diffusion-model U-Net [35,50] and introduces two key modules that enhance its representational and physical consistency: residual blocks and spatial transformers. As illustrated in Figure 2, panel (a) depicts the overall denoising process with a multi-scale U-Net backbone, which progressively refines the noisy microstructure representation $x_T \rightarrow x_0$. The encoder extracts hierarchical features from fine to coarse scales, while the decoder reconstructs the topology by combining deep semantic and shallow geometric information through skip connections.

Panel (b) shows the residual block structure. Each residual block acts as a local adaptive filter that stabilizes gradient flow and captures deformation-sensitive features. The diffusion time-step embedding is injected into the block using FiLM-style modulation, allowing the network to dynamically adjust its denoising behavior according to the current diffusion state.

Panel (c) presents the spatial transformer module, which captures long-range correlations and integrates the conditioning mechanical information. Through self-attention, the model learns global spatial dependencies within the feature map, while cross-attention propagates the target force–displacement condition \mathbf{y} across the entire field, ensuring that every spatial region adheres to the prescribed macroscopic mechanical behavior.

By jointly employing residual blocks for localized learning and spatial transformers for global constraint propagation, *HyperDiff* achieves a balanced local–global coupling mechanism. This design ensures that the generated microstructures are both topologically coherent and mechanically consistent with the desired target responses.

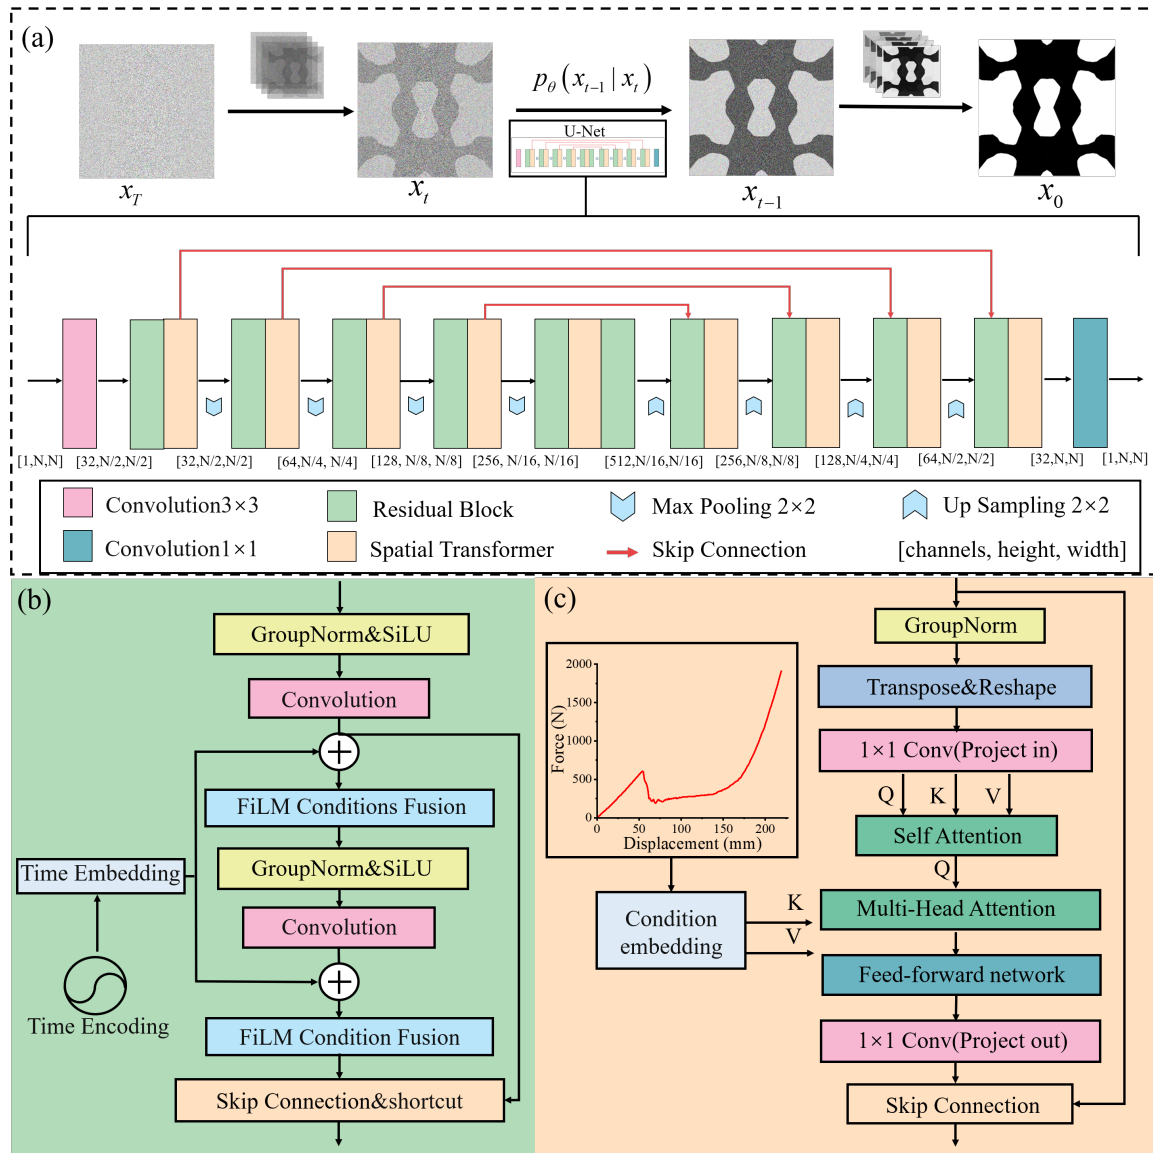


Figure 2. Network overview and core modules of *HyperDiff*. (a) *Denoising diffusion with a U-Net backbone.* Starting from pure Gaussian noise x_T , the model iteratively predicts $p_\theta(x_{t-1} | x_t)$ and reconstructs a clear microstructure x_0 . The lower row shows the multi-scale U-Net with encoder-decoder symmetry and skip connections (red). Channel sizes and spatial resolutions are indicated as [channels, height, width]. Green blocks denote residual blocks; orange blocks denote spatial transformers; blue markers indicate 2×2 max pooling / up sampling; 1×1 convolutions perform channel projection. (b) *Residual block (green region).* Each block uses GroupNorm and SiLU activations with two convolutions, and injects the diffusion time-step via FiLM-style feature-wise modulation (time embedding \rightarrow scale/shift). A residual shortcut preserves stable gradients and enables dynamic, time-aware local denoising. (c) *Spatial transformer (orange region).* Features are normalized, reshaped, and linearly projected, followed by self-attention and multi-head attention. The target force-displacement curve is encoded as a condition vector and fused as keys/values, enabling cross-attention guidance that propagates global mechanical constraints over the spatial field. A feed-forward network and output projection (1×1 conv) complete the block, with residual connections for stability.

3. Dataset Generation

A comprehensive dataset mapping microstructural topologies to their nonlinear mechanical responses is essential for training the *HyperDiff* model. This section details the three-stage pipeline developed to generate this dataset: (1) procedural generation of diverse microstructure topologies using a Gaussian random field (GRF) method, (2) experimental characterization of the hyperelastic

material properties, and (3) large-scale finite element simulation to compute the force–displacement response for each topology.

3.1. Microstructure Topology Generation

To build a physically grounded dataset of porous microstructures with diverse geometric and mechanical characteristics, a two-dimensional Gaussian random field (GRF) sampling approach is employed, as illustrated in Figure 3(a). A continuous GRF is first generated over a square domain and then thresholded to obtain a binary solid–void distribution. To ensure structural integrity and manufacturability, connectivity and coverage thresholds are applied, and a four-side symmetric topology is reconstructed by mirroring the binary pattern. Contour detection, sampling, and smoothing operations are subsequently performed to remove irregularities and obtain well-defined smooth boundaries suitable for finite element analysis. This process produces geometrically diverse yet physically consistent unit cells that serve as the input domain of the *HyperDiff* framework. The corresponding material calibration and finite element simulation procedures, shown in Figure 3(b,c), provide the ground-truth mechanical responses used to supervise the conditional diffusion model.

The GRF is generated in the frequency domain to control its spatial characteristics via a prescribed power spectrum. First, a complex-valued field $z(\mathbf{k})$ is constructed from independent standard normal variables z_1 and z_2 for each wave vector \mathbf{k} :

$$z(\mathbf{k}) = z_1(\mathbf{k}) + iz_2(\mathbf{k}). \quad (10)$$

A power spectrum filter is then applied to shape the field's correlations:

$$F(\mathbf{k}) = z(\mathbf{k}) (|\mathbf{k}|^2 + \epsilon)^{-\alpha/2}, \quad (11)$$

where the parameter α controls the spatial correlation of the field, and ϵ is a small constant to prevent singularities. The real-space representation $f(\mathbf{x})$ is obtained via the inverse Fourier transform:

$$f(\mathbf{x}) = \mathcal{F}^{-1}[F(\mathbf{k})]. \quad (12)$$

To ensure $f(\mathbf{x})$ remains strictly real before thresholding, we enforce Hermitian symmetry by constructing $F(-\mathbf{k}) = \overline{F(\mathbf{k})}$ (with consistent random phases) prior to the inverse FFT, so the recovered field is physically meaningful for subsequent binarization.

After normalization and Gaussian smoothing, the field is binarized using a threshold t to define the material phase $\chi(\mathbf{x})$:

$$\chi(\mathbf{x}) = \begin{cases} 1, & \text{if } f(\mathbf{x}) > t, \\ 0, & \text{otherwise.} \end{cases} \quad (13)$$

Finally, to form a fully symmetric structure, the base microstructure quadrant $\chi(x, y)$ is mirrored across both axes:

$$\chi_{\text{full}}(x, y) = \begin{bmatrix} \chi(x, y) & \chi(x, N - y) \\ \chi(N - x, y) & \chi(N - x, N - y) \end{bmatrix}, \quad (14)$$

where (x, y) denotes the spatial coordinates and N is the grid size. This GRF-based method efficiently generates a diverse set of porous microstructures with rich, multiscale features suitable for training a robust generative model.

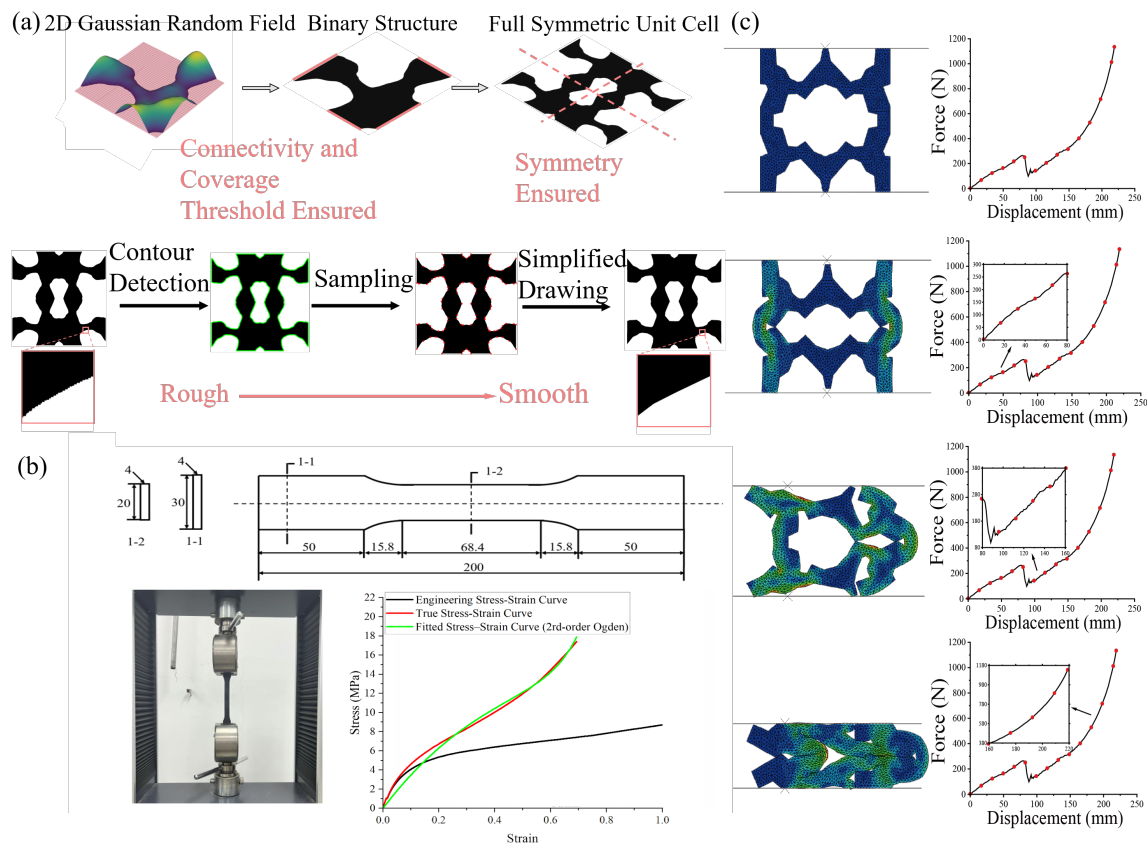


Figure 3. Dataset construction and physical simulation framework for *HyperDiff*. (a) *Microstructure generation based on a 2D Gaussian random field (GRF)*. A continuous GRF is thresholded to produce binary topologies, followed by connectivity filtering, symmetry reconstruction, and contour smoothing from rough to smooth edges. This process yields physically consistent and geometrically diverse unit cells. (b) *Material calibration using uniaxial tensile tests*. The TPU specimen geometry and testing setup are shown, together with the engineering, true, and fitted stress–strain curves. A second-order Ogden model accurately captures the nonlinear hyperelastic behavior. (c) *Finite element simulations and force–displacement responses*. Different generated microstructures are compressed under displacement control. The von Mises stress contours and corresponding force–displacement curves reveal characteristic deformation stages of elastic response, buckling, and densification. These results form a physically grounded dataset for training the conditional diffusion model.

3.2. Material Characterization and Hyperelastic Model Calibration

To characterize the nonlinear mechanical behavior of the base material used in the generated microstructures, uniaxial tensile tests were conducted on thermoplastic polyurethane (TPU) specimens in accordance with the GB/T 228.1–2021 standard. The experimental setup, specimen geometry, and resulting stress–strain curves are shown in Figure 3(b). The tests were performed using a TFW-100S universal testing machine at a constant displacement rate of 50 mm/min, corresponding to a nominal strain rate of $4.17 \times 10^{-3} \text{ s}^{-1}$, which satisfies the quasi-static loading condition. Both engineering and true stress–strain curves were recorded, and the data were fitted using a second-order Ogden hyperelastic model to capture the material’s large-deformation behavior. The fitted curve shows excellent agreement with the experimental data, confirming that the Ogden model can accurately describe the nonlinear, strain-hardening response of TPU. The calibrated parameters were subsequently applied to all finite element simulations in Figure 3(c), ensuring consistency between the experimental characterization and numerical modeling stages of the dataset.

Assuming material incompressibility, the engineering data were converted to true stress (σ_t) and logarithmic strain (ε_t):

$$\sigma_t = \frac{F}{A_0} (\varepsilon_{\text{eng}} + 1), \quad (15)$$

$$\varepsilon_t = \ln(\varepsilon_{\text{eng}} + 1), \quad (16)$$

where F is the applied load, A_0 is the initial cross-sectional area, and ε_{eng} is the engineering strain. To enable finite element simulations, the true stress–strain data were fitted using the Ogden hyperelastic model [49], which provided the best fit. Its strain energy density function is expressed as:

$$U = \sum_{i=1}^N \frac{2\mu_i}{\alpha_i^2} (\bar{\lambda}_1^{\alpha_i} + \bar{\lambda}_2^{\alpha_i} + \bar{\lambda}_3^{\alpha_i} - 3) + \sum_{i=1}^N \frac{1}{D_i} (J_{\text{el}} - 1)^{2i}, \quad (17)$$

where $\bar{\lambda}_i = J^{-1/3} \lambda_i$ are the deviatoric principal stretches, and μ_i , α_i , and D_i are material parameters. The Ogden model is particularly suitable for capturing the nonlinear stiffening behavior of hyperelastic materials. The fitted parameters obtained from Abaqus are listed in Table 1. When a reported parameter satisfies $\alpha_i \approx 0$ (as for the first term in Table 1), the implementation evaluates the $\alpha_i \rightarrow 0$ limit, which is equivalent to inserting a neo-Hookean term with modulus $2\mu_i$ and avoids the apparent singularity in the $2\mu_i/\alpha_i^2$ prefactor.

Table 1. Fitted Ogden model parameters for the TPU material.

Parameter	Ogden model
μ_1	12.7464
α_1	0.0000
D_1	0.0000
μ_2	0.0108
α_2	12.3711
D_2	0.0000

3.3. Finite Element Simulation and Force–Displacement Responses

To evaluate the mechanical responses of the generated microstructures and establish the ground-truth data for model training, finite element simulations were performed on each unit cell using the calibrated hyperelastic parameters obtained from Figure 3(b). In metamaterial design, full assemblies typically require 5–10 repeating unit cells to achieve homogenization, which considerably increases computational cost. To balance efficiency and fidelity, the simulations in this study focused on individual unit cells, which sufficiently capture the intrinsic nonlinear mechanisms relevant to the inverse-design task.

As illustrated in Figure 3(c), each microstructure was placed between two rigid plates and subjected to uniaxial compression under quasi-static loading conditions, and the corresponding von Mises stress contours were used to visualize the evolution of local stress concentrations during deformation. The simulations were conducted in a two-dimensional plane-stress setting with appropriate symmetry and boundary constraints to represent periodic behavior. The computed force–displacement curves exhibit the characteristic three-stage compressive response of porous hyperelastic materials: (i) an initial elastic regime dominated by frame bending and stretching, (ii) a nonlinear instability regime associated with progressive buckling and softening, and (iii) a final densification regime where pore collapse leads to a rapid increase in stiffness. These stages correspond closely to the deformation and stress-localization mechanisms observed in quasi-static compression experiments, ensuring that the dataset remains both physically interpretable and experimentally consistent.

All simulations were performed in ABAQUS/Explicit using the Ogden hyperelastic model with parameters listed in Table 1. The upper rigid plate was coupled to a reference point for displacement

control, while the lower plate was fully constrained. Contact definitions included hard contact with a tangential friction coefficient of 0.18 between the plates and the microstructure, as well as general contact for internal self-interaction to prevent element interpenetration during large deformation. The microstructure was meshed with CPS3 (plane-stress) triangular elements (seed size 2.5 mm), and the rigid plates were meshed with R2D2 elements (seed size 5.0 mm). A smooth-step amplitude curve was used to impose 60% compressive strain quasi-statically. The explicit central-difference integration scheme was adopted for its robustness in handling large deformations and complex contact evolution without convergence issues.

A fully automated workflow was implemented through a Python script based on the ABAQUS API. The script automatically performs geometry construction, Boolean operations, contact definition, material assignment, mesh generation, and batch job submission, enabling efficient multi-core parallel simulations of approximately 20,000 unique microstructures. This large-scale dataset captures not only the canonical three-stage responses but also non-classical two- or four-stage behaviors caused by local instabilities and complex topological patterns. Notably, microstructures with similar volume fractions can exhibit markedly different force–displacement responses, while distinct topologies may produce nearly identical curves. This non-unique structure–property mapping is visualized in Figures 4 and 5, which highlight the diversity, nonconvexity, and weak identifiability of the dataset—factors that make the inverse problem intrinsically challenging and motivate the adoption of data-driven generative frameworks such as *HyperDiff*.

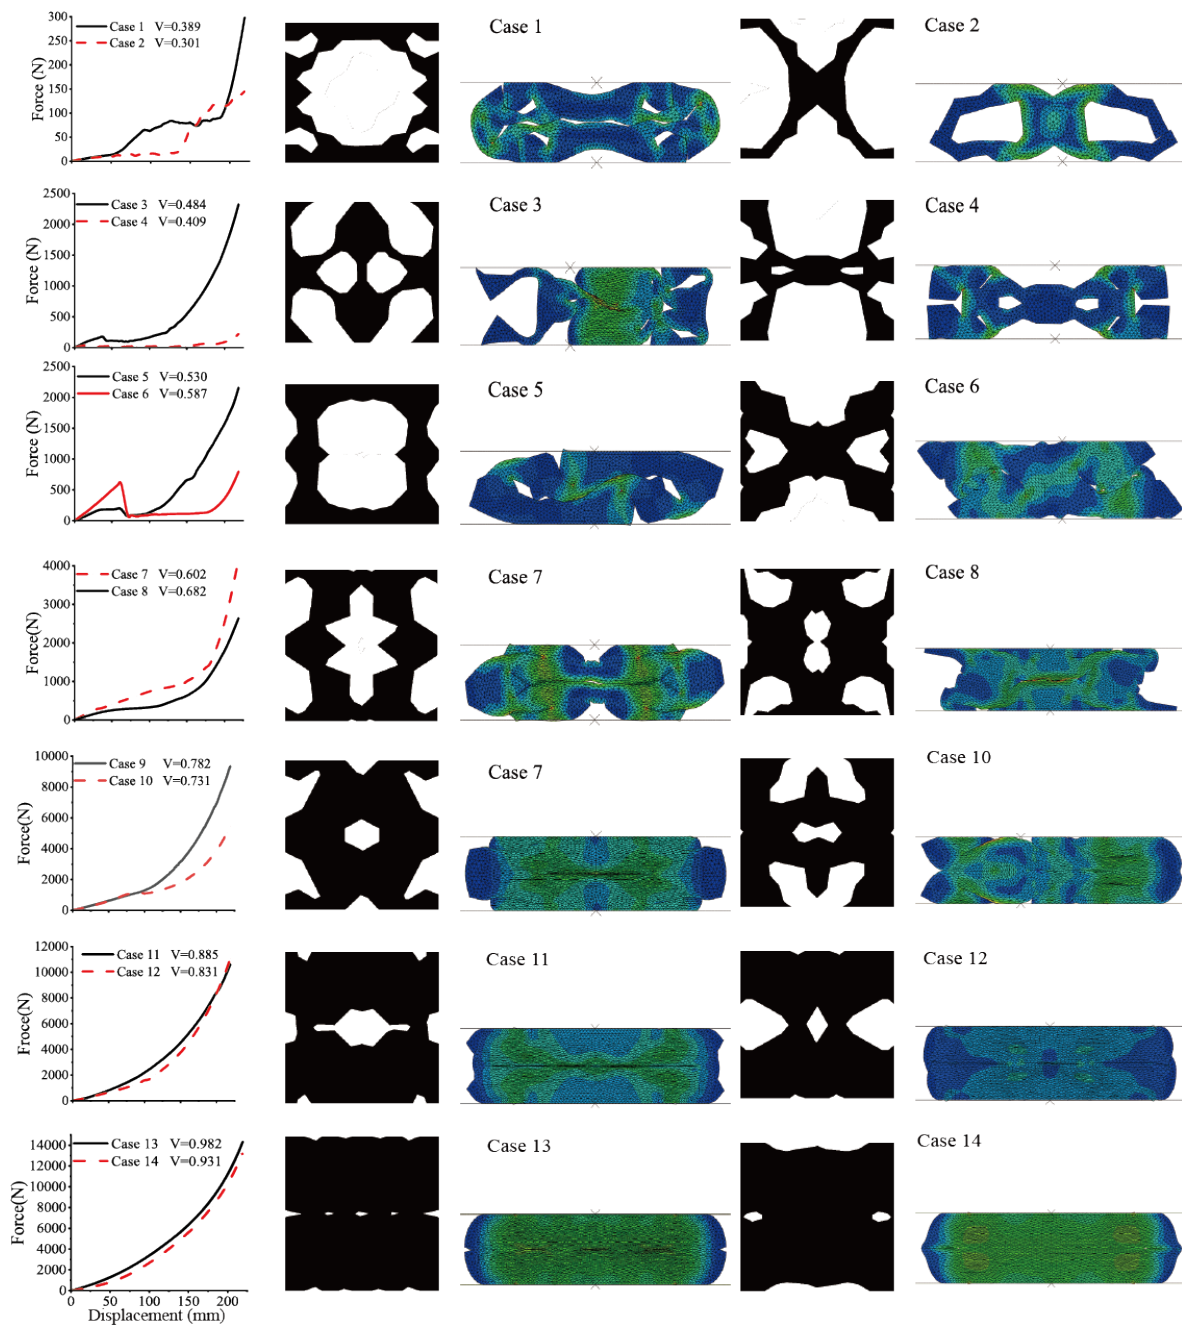


Figure 4. Representative FE samples illustrating multi-stage deformation and response diversity (Cases 1–14). Each row shows a unique topology, its von Mises stress field under large deformation, and the corresponding force–displacement curve. The dataset spans bending-dominated compliance, buckling-induced softening (plateaus), and contact-driven densification, revealing a strongly nonconvex structure–response landscape.




Figure 5. Non-unique and weakly identifiable structure–property mappings (Cases 15–18). *Top:* distinct topologies with the same volume fraction yield markedly different force–displacement curves (mode switching under similar global descriptors). *Bottom:* topologies with different porosities produce nearly identical responses (one-to-many mapping). These comparisons highlight non-uniqueness, nonconvexity, and the limits of low-dimensional parametrizations, motivating a probabilistic inverse formulation.

4. Model Training Strategy

The primary objective of our training strategy is to establish a reliable mapping from a target force–displacement curve to a corresponding microstructural design. A common approach in related literature, such as that by [39], involves conditioning a model on simplified mechanical responses like stress–strain curves. However, such a setup can be ineffective for designing microstructures that exhibit complex behaviors such as structural instability and buckling. The limitation lies in the indirect nature of this mapping, where the model learns from static geometry without sufficient information about the dynamic mechanical response.

To provide a compact and robust conditioning signal for the diffusion model, the force–displacement curves are (i) smoothed and de-noised, then (ii) parameterized by B-spline.

B-spline encoding and feature vector

To obtain a compact conditioning signal, each target force–displacement curve is first approximated by a B-spline,

$$P(u) = \sum_{i=0}^n P_i B_{i,k}(u), \quad u \in [u_{k-1}, u_{n+1}], \quad (18)$$

where P_i are control-point coefficients and $B_{i,k}(u)$ are k -th order basis functions defined by a nondecreasing knot vector $\{u_j\}_{j=0}^m$. To make the coefficients numerically stable and scale-invariant, we apply min–max normalization

$$P'_i = \frac{P_i - P_{\min}}{P_{\max} - P_{\min}}, \quad (19)$$

and retain (P_{\min}, P_{\max}) so that the original curve can be reconstructed at inference. Finally, the one-dimensional conditioning vector is formed by concatenating the normalized coefficients, the stored extrema, and the knot locations (each group in ascending order):

$$V = [P'_0, P'_1, \dots, P'_n, P_{\min}, P_{\max}, u_0, u_1, \dots, u_m]. \quad (20)$$

This compact representation preserves the essential curve shape while improving training efficiency, numerical stability, and generalization.

Hyperparameters

The key hyperparameter settings employed in our training process are detailed in Table 2. This configuration was optimized to balance model capacity, training stability, and computational efficiency.

This training strategy provides a stable conditioning signal for the diffusion model, supporting accurate and interpretable inverse design predictions under complex deformation mechanisms such as buckling and densification.

Table 2. Key hyperparameters used for model training.

Hyperparameter	Value
Batch size	64
Learning rate	1×10^{-4}
Decay steps	25
Decay factor	0.99
Input channels	1
Output channels	1
Base channels	32
Residual blocks	1
Attention layers enabled	[False, True, False, True, False]
Channel multipliers	[1, 1, 2, 2, 2]
Attention heads	1
Transformer layers	1
Condition embedding dimension	[48]
Noise scale range (β_t)	($1e-4$, 0.02)
DDPM timesteps	1000
Loss function	L2 loss (MSE)
Image resolution	64×64 px
Training iterations	100,000

5. Results and Validation

We evaluated *HyperDiff*'s performance through systematic numerical experiments. We selected four reference force–displacement curves from the test set as targets. For each target, we generated 50 microstructures by sampling different random noise vectors, thereby assessing the model's ability to produce diverse yet mechanically consistent designs.

All generated microstructures were analyzed using the same ABAQUS finite element setup described in Section 3.3. To quantify performance, we used the normalized root-mean-square error (NRMSE), ε , between the generated force response F and the target response F_{target} :

$$\varepsilon(F, F_{\text{target}}) = \frac{1}{F_{\text{ref}}} \sqrt{\frac{1}{n} \sum_{i=1}^n (F_i - F_{\text{target},i})^2}, \quad (21)$$

where $F_{\text{ref}} = \max_i |F_{\text{target},i}| + 10^{-6}$ provides a single, curve-wide normalization factor that is strictly positive even when the target force crosses or equals zero. This definition keeps the metric scale-invariant without incurring divisions by $F_{\text{target},i}$ at the unloaded state. So the entire curve, including the initial flat segment, contributes smoothly to the reported error. Consequently, a 10% deviation on a low-force (soft) curve is weighted just as heavily as a 10% deviation on a high-force (stiff) curve, ensuring the model is evaluated fairly across a diverse range of mechanical responses.

5.1. Performance on Test Curves

To assess *HyperDiff*'s generalization to unseen data, we selected four representative force–displacement curves from the test set, covering both weakly and strongly nonlinear behaviors. Figure 6 (left block) summarizes the results: the top row reports per-case profiles of the normalized point-wise relative error along the displacement axis, the middle row shows NRMSE distributions over 50 stochastic generations, and the bottom row presents a representative force–displacement overlay for Case 1. For three of the four cases, the median error was below 7%, and Case 4 achieved a median error of approximately 0.4%. Additional per-target overlays and the corresponding topologies for the remaining test cases are provided in Figure 7, which further illustrate consistent shape matching and one-to-many topology diversity.

As shown in Figure 6, the generated microstructures and their corresponding force–displacement curves for the minimum- and median-error samples closely match their targets. Notably, the generated

microstructures exhibit significant topological diversity while achieving nearly identical mechanical performance. For example, the best- and median-performing structures for Case 1 share a similar overall layout but differ in their internal connectivity. This one-to-many mapping capability is a key advantage of our generative approach, providing designers with a range of valid solutions for a single performance target.

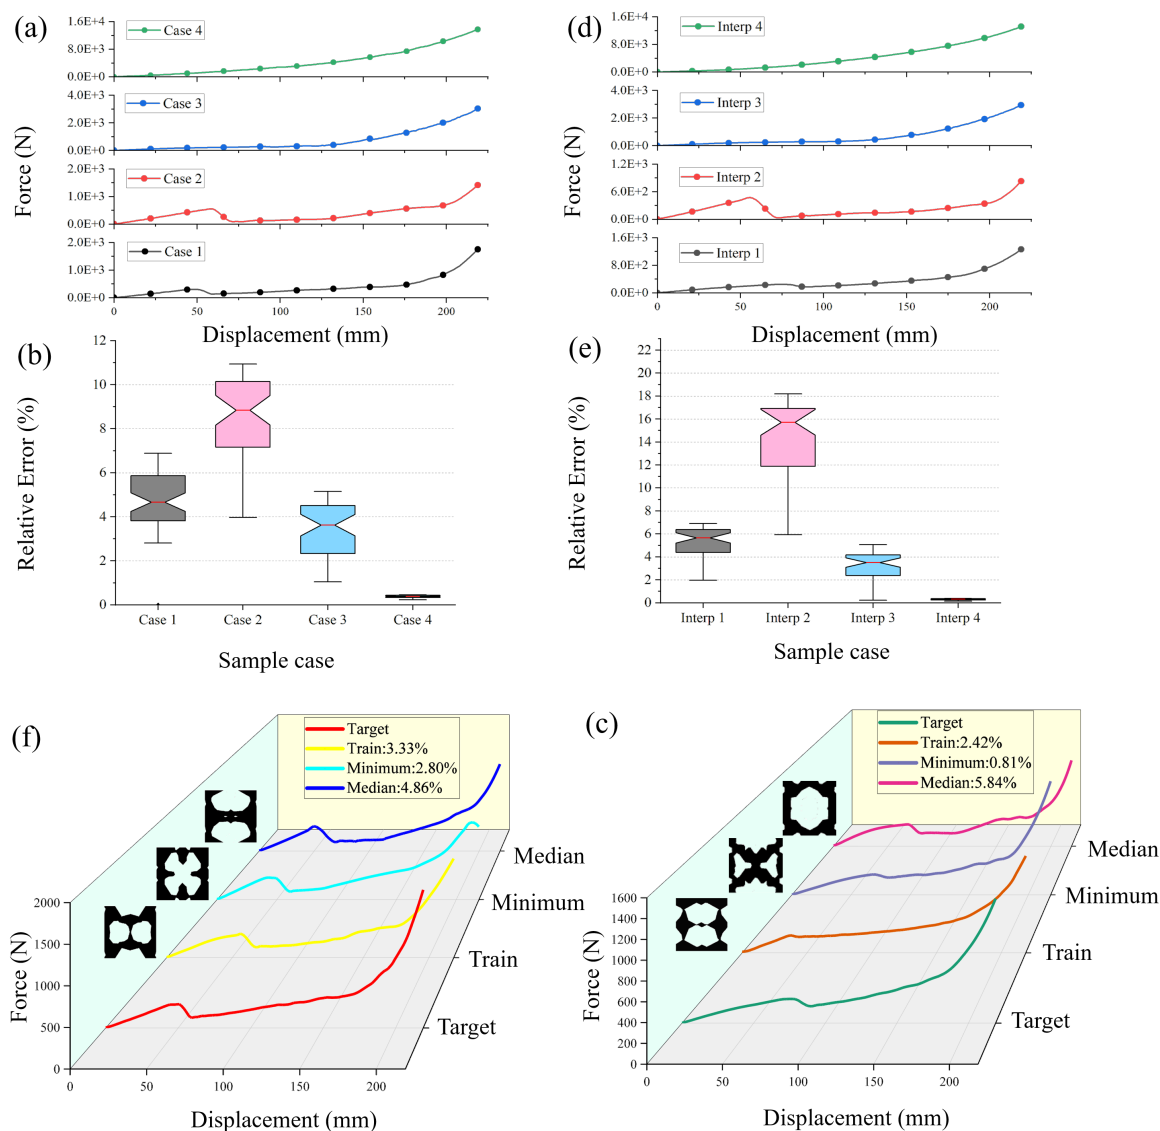


Figure 6. Comprehensive evaluation of *HyperDiff* on test targets (left) and interpolated targets (right). *Top row:* per-case profiles of the normalized pointwise relative error along the displacement axis for four test cases (left) and four interpolated cases (right). *Middle row:* per-case NRMSE% distributions over 50 stochastic generations (center line = median; box = interquartile range; whiskers = 10–90th percentile). *Bottom row:* representative force–displacement overlays for Case 1 (left) and Interp. 1 (right); the target curve is compared to its nearest neighbor in the training set (NN-train) and the generated minimum/median samples. Legends report the corresponding percentage errors. Small topology thumbnails highlight one-to-many design diversity. Additional cases are detailed in Figure 7.

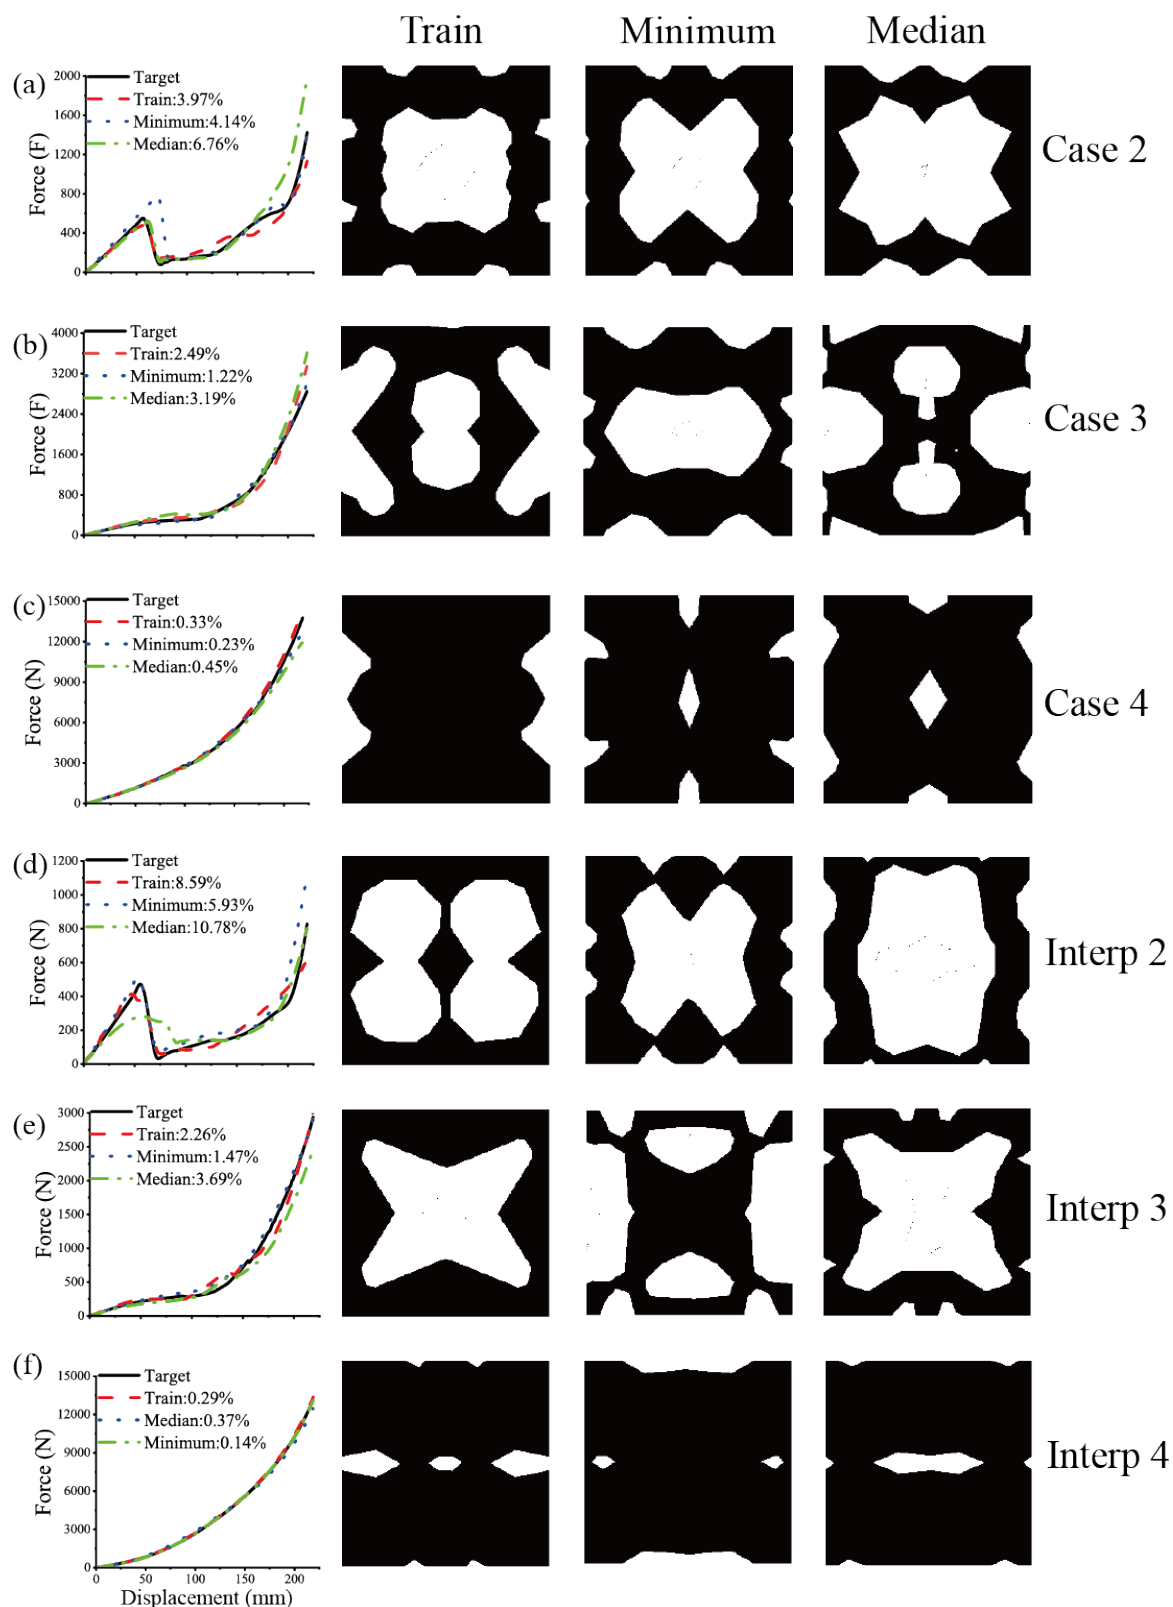


Figure 7. Detailed overlays and corresponding topologies for the remaining test and interpolated targets. For each subfigure, the left panel shows the force–displacement overlay (Target, NN-train, Minimum, Median) with overall NRMSE%, and the three right panels show the corresponding topologies (Train, Minimum, Median).

5.2. Performance on Interpolated Curves

To further test *HyperDiff*'s generalization, we evaluated its ability to generate microstructures for interpolated force–displacement curves—targets that lie between samples in the training data

distribution. We created four sets of interpolated targets by linearly combining pairs of curves from the training set (see Figure 6, right block).

The results indicate that *HyperDiff* exhibits strong interpolation capability. As shown in Figure 6, the median error was below 8% in three of the four interpolation cases, while Interp. 4 achieved a notably low median error of approximately 0.4%. In contrast, Interp. 2 showed a higher median error (around 16%), which can be attributed to its pronounced nonlinear characteristics, making it a more challenging target for accurate design.

The generated microstructures in Figure 6 successfully match the interpolated mechanical targets, with minimum errors below 6% in all cases. Again, the model produces topologically diverse solutions for each target. For instance, the minimal-error and median-error structures for Interp. 3 both feature a primary “X”-shaped load path but differ in their central region. This confirms *HyperDiff*'s ability to generalize beyond the training set and explore the design space to find novel and effective microstructures. Complementary per-target overlays and the corresponding topologies for the remaining interpolated cases are presented in Figure 7.

6. Experimental Verification

To validate the predictive accuracy and physical consistency of the proposed *HyperDiff* framework, a series of quasi-static compression experiments were performed on representative energy-absorbing metamaterial specimens. Each specimen was a prismatic cubic block with overall dimensions of 50 mm × 50 mm × 50 mm, obtained by extruding the 2D unit-cell designs along the out-of-plane direction. The thickness is therefore on the same order as the in-plane dimensions, effectively suppressing global out-of-plane buckling and preserving the deformation mechanisms (pore closure, local buckling, and contact-driven densification) captured by the plane-strain FE simulations.

The samples were fabricated via **3D printing (additive manufacturing)** and tested using a universal testing machine under a constant compression rate of 2 mm/min, ensuring quasi-static loading conditions. During the tests, force and displacement data were synchronously recorded to accurately capture the mechanical responses. The experimental setup enabled continuous observation of deformation patterns and ensured repeatability of the results. A comprehensive comparison between experimental and simulation results is presented in Figure 8.

As shown in Figure 8, the left column presents the force–displacement curves obtained from the target, numerical simulation, and experimental measurements (black dashed, red dotted, and blue dash–dot lines, respectively). The normalized root-mean-square error (NRMSE) was used to evaluate prediction accuracy. For the representative case shown, the NRMSE between the experimental and target curves was 9.41%, between the simulated and target curves 6.10%, and between the experimental and simulated curves 4.42%. These results confirm that the *HyperDiff* model can accurately reproduce the nonlinear mechanical behavior of the designed metamaterials.

The right-hand side of Figure 8 illustrates the deformation sequences (stages I–VI) under progressive compression. The top row shows the simulated von Mises stress contours, while the bottom row presents the corresponding experimental deformation images. As strain increased, small internal pores closed first (stages II–III). When the displacement reached approximately 15 mm (stage IV), the central pore was nearly closed. Further loading caused contact between the upper and lower beams and the compression plate, leading to a rapid increase in tangent stiffness. The regions of high stress in the simulations correspond closely to the zones of visible large-strain localization and contact-induced densification in the experiments, indicating that the model captures the essential physical mechanisms rather than merely fitting data correlations. This agreement reinforces the interpretability and physical reliability of the proposed framework.

A detailed analysis of the first case in Figure 8 is presented as a representative example. The other tested specimens exhibited consistent deformation modes and response trends, demonstrating the robustness of both the model and experimental validation process.

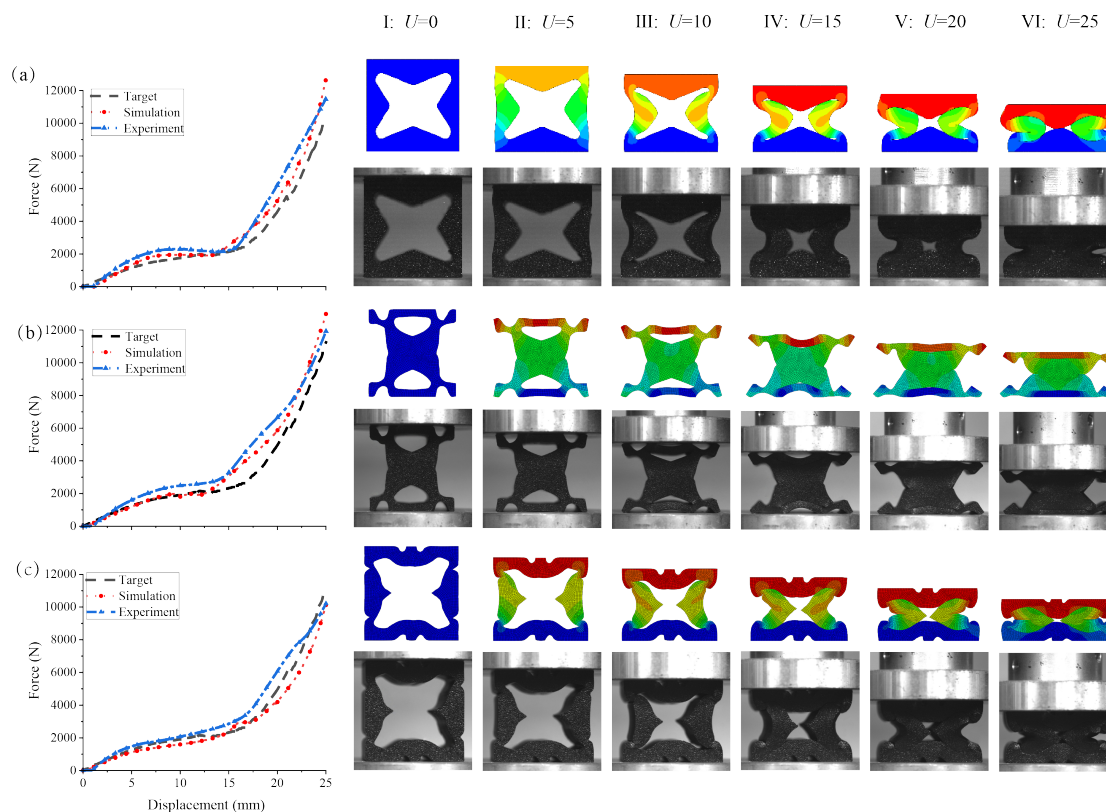


Figure 8. Experimental verification of energy-absorbing metamaterials generated by *HyperDiff*. Comparison between simulation and experiment under quasi-static compression. The *left panels* show force–displacement curves (black dashed = target, red dotted = simulation, blue dash–dot = experiment). Vertical ticks on each curve mark six representative loading stages (I–VI) corresponding to global compression displacements $U = 0, 5, 10, 15, 20,$ and 25 mm. The *right panels* display the corresponding deformation sequences, where the *upper rows* are simulated von Mises stress contours and the *lower rows* are experimental images. Regions of high simulated stress spatially coincide with the zones of visible large-strain localization and contact-induced densification observed in the experiments, indicating that the proposed model captures the essential deformation mechanisms and nonlinear responses of hyperelastic porous metamaterials. Here, U denotes the imposed displacement of the upper platen in millimeters (mm).

7. Discussion

The results in Sections 5 and 6 show that *HyperDiff* can recover the main features of nonlinear force–displacement responses for hyperelastic porous microstructures across both numerical simulations and experiments. This section provides a mechanics-aware interpretation of these results, explains how the conditional diffusion process organizes structure–response relations, and outlines the scope and limitations of the present framework.

7.1. From a Difficult Inverse Problem to a Tractable Generative View

Inverse design in the large-deformation hyperelastic regime is challenging because the macroscopic response arises from multiple sequential mechanisms—early-stage bending compliance, instability-induced softening, contact-mediated load transfer, and late-stage densification. These mechanisms lead to a highly nonconvex, inherently one-to-many mapping between topology and response. Small geometric perturbations can trigger a completely different instability path, while distinct microstructures may yield similar global responses. Under such conditions, parametric regression or gradient-based optimization often becomes trapped in local minima and lacks generality across design families.

The present framework addresses this difficulty by shifting from “solve-for-the-one” to “sample-from-the-set.” Instead of pursuing a deterministic inverse map, *HyperDiff* learns the conditional *distribution* of admissible topologies given a target response. This generative perspective is consistent

with the physical nature of the problem: for a prescribed force–displacement curve, multiple feasible microstructures usually exist. Sampling diverse candidates and retaining those verified by forward analysis thus provides a pragmatic and robust solution.

7.2. A Mechanics-Aware Reading of the Diffusion Process

Although diffusion models are typically introduced as statistical denoisers, the present results suggest a mechanics-aware interpretation in this context. The U-Net backbone with residual connections captures local void morphologies, while spatial transformers ensure global coherence consistent with compression load paths. During denoising, structural features emerge in a physically ordered sequence: early iterations determine initial compliance (void slenderness and ligament curvature), mid-stage updates control buckling onset and amplitude, and later iterations form load-bearing bridges that sustain densification.

Hence, the iterative sampling process does not merely “reconstruct an image,” but instead relaxes toward configurations that are both statistically plausible and mechanically admissible under the given conditioning signal. This explains why the generated microstructures reproduce not only the overall force–displacement trends but also the observed deformation stages in both simulation and experiment.

7.3. Conditioning via Force–Displacement Curves as Energy-Trend Guidance

A distinctive element of *HyperDiff* lies in its conditioning strategy based on a compact B-spline representation of the target force–displacement curve. Although no explicit energy integration is enforced, the curve inherently reflects the cumulative exchange of mechanical work during loading; the corresponding B-spline coefficients thus serve as a concise descriptor of the system’s *energy-evolution trend*. Conditioning on this vector provides temporal (stage-wise) and mechanical (stiffness/softening/densification) context to the generator. Empirically, this guidance enables the model to converge toward structures that exhibit the correct *sequence* of deformation mechanisms, not merely the right scalar magnitudes.

This interpretation also clarifies the difference from encodings that treat the curve as a generic signal: here, the conditioning is designed to carry physically meaningful information. Extending the conditioning vector with explicit cumulative-energy or stage-indicator terms is straightforward if finer phase control is desired in future studies.

7.4. Applicability, Robustness, and Boundaries

The current implementation focuses on two-dimensional unit cells at a fixed image resolution and on quasi-static compression of a calibrated hyperelastic material. Within this scope, the method is robust: it achieves low deviations on held-out targets, maintains accuracy on interpolated curves, and yields printed samples whose deformation patterns match simulations. These results support that the learned topology–response relationships are physically meaningful rather than numerical coincidences.

At the same time, clear boundaries exist. Out-of-plane effects are excluded by design; rate dependence, viscoelasticity, and thermal coupling are not considered; and the minimum printable feature size is constrained by the image resolution. Diffusion-based sampling is computationally more demanding than direct regression, but the gain in solution diversity is valuable in nonconvex design spaces. These factors should be kept in mind when generalizing beyond the current dataset.

7.5. Relation to Prior Inverse Design Studies

Compared with inverse design frameworks developed for elastoplastic or energy-absorbing systems, the present study targets the finite-strain hyperelastic regime, where geometric and constitutive nonlinearities interact. Methodologically, the main differences are: (i) a generative, distributional formulation of the inverse problem; (ii) a physically meaningful conditioning vector encoding the energy-evolution trend; and (iii) validation of both global responses and stage-wise deformation

modes through numerical and experimental means. These aspects explain the consistent performance observed for highly nonlinear targets that typically challenge deterministic regressors.

7.6. Practical Implications

In practice, *HyperDiff* can be used in two complementary modes. First, as a *design exploration tool*, it rapidly proposes multiple distinct microstructures for a single target, allowing post-selection under additional criteria such as manufacturability or stress localization. Second, as a *front-end generator*, it provides physically consistent initial designs for subsequent physics-based optimization. Both usages exploit the model's diversity while maintaining mechanical consistency.

Summary. When conditioned on an energy-trend encoding of the desired response, the diffusion model acts as a probabilistic sampler over physically feasible topologies in a nonconvex design space. This perspective reconciles data-driven generation with mechanics-based understanding and explains the cross-consistency observed between simulations and experiments.

8. Conclusions and Future Work

This work introduced *HyperDiff*, a conditional diffusion framework for the inverse design of hyperelastic porous microstructures undergoing large deformation, buckling, softening, and densification. By modeling the distribution of admissible topologies conditioned on a physically meaningful encoding of the target force–displacement path, the framework bridges response specification, generative synthesis, and physical validation.

Main contributions.

- **Generative formulation of a nonconvex inverse problem.** The inverse design problem is reformulated as conditional sampling rather than deterministic inversion, enabling one-to-many solutions consistent with the inherent non-uniqueness of the finite-strain hyperelastic regime. This improves robustness for targets where local minima or mode switching hinder parametric regressors.
- **Physics-aware conditioning via B-spline encoding.** The conditioning vector compactly represents the energy–evolution trend along the loading path, providing temporal and mechanical context that guides the denoising process toward configurations that reproduce the correct sequence of deformation stages. This enhances interpretability and facilitates faithful reproduction of bending, buckling, and densification behavior.
- **Integrated numerical and experimental validation.** Large-scale finite element simulations and quasi-static compression experiments confirm that generated designs match both global responses and local deformation stages, typically within 10% deviation. This supports that the learned topology–response relationships are physically realizable rather than purely statistical correlations.

Limitations.

The present framework considers two-dimensional unit cells at a fixed resolution and a single constitutive calibration under quasi-static compression. Out-of-plane mechanisms, rate effects, viscoelasticity, and thermo-mechanical coupling are beyond the current scope. Diffusion-based sampling incurs higher computational cost than direct regression, trading runtime for solution diversity.

Future work.

Future extensions include: (i) **Volumetric generalization:** extending to three-dimensional architected materials using volumetric diffusion and 3D U-Nets to preserve geometric detail; (ii) **Multi-scale and multi-resolution coupling:** linking unit-cell generation with homogenization-informed global constraints to improve stress localization while maintaining load paths; (iii) **Physics- and process-aware guidance:** embedding explicit physical priors and manufacturability constraints during sampling

to ensure feasible, print-ready designs; (iv) **Broader materials and loadings:** expanding datasets to rate-dependent, viscoelastic, or thermo-mechanical systems and testing robustness under varied conditions.

Outlook.

Beyond this specific application, the results highlight a generative, physics-aware route for inverse design in nonlinear mechanics. Treating the target response as an energy-trend guidance and employing conditional diffusion as a probabilistic sampler provides a principled way to unify data-driven diversity with mechanical admissibility. This paradigm can extend to other classes of architected materials and integrate naturally with downstream optimization and verification workflows.

Data Availability Statement: The data that support the findings of this study are available from the corresponding author upon reasonable request.

Acknowledgments: This work was supported by the Liaoning Provincial Science and Technology Major Special Project (Grant No. 2024JH1/11700046) and the National Natural Science Foundation of China (Grant No. 12372194).

References

1. M. E. Davis, Ordered porous materials for emerging applications, *Nature* 417 (6891) (2002) 813–821.
2. Y. Gogotsi, A. Nikitin, H. Ye, W. Zhou, J. E. Fischer, B. Yi, H. C. Foley, M. W. Barsoum, Nanoporous carbide-derived carbon with tunable pore size, *Nature materials* 2 (9) (2003) 591–594.
3. M. L. K. Hoa, M. Lu, Y. Zhang, Preparation of porous materials with ordered hole structure, *Advances in colloid and interface science* 121 (1-3) (2006) 9–23.
4. K. Bertoldi, V. Vitelli, J. Christensen, M. Van Hecke, Flexible mechanical metamaterials, *Nature Reviews Materials* 2 (11) (2017) 1–11.
5. S. Shan, S. H. Kang, J. R. Raney, P. Wang, L. Fang, F. Candido, J. A. Lewis, K. Bertoldi, Multistable architected materials for trapping elastic strain energy, *Adv. Mater* 27 (29) (2015) 4296–4301.
6. A. Rafsanjani, K. Bertoldi, A. R. Studart, Programming soft robots with flexible mechanical metamaterials, *Science Robotics* 4 (29) (2019) eaav7874.
7. Y. Zhong, W. Tang, H. Xu, K. Qin, D. Yan, X. Fan, Y. Qu, Z. Li, Z. Jiao, H. Yang, et al., Phase-transforming mechanical metamaterials with dynamically controllable shape-locking performance, *National Science Review* 10 (9) (2023) nwad192.
8. J. T. Overvelde, T. A. De Jong, Y. Shevchenko, S. A. Becerra, G. M. Whitesides, J. C. Weaver, C. Hoberman, K. Bertoldi, A three-dimensional actuated origami-inspired transformable metamaterial with multiple degrees of freedom, *Nature communications* 7 (1) (2016) 10929.
9. Z. Chai, Z. Zong, H. Yong, X. Ke, J. Zhu, H. Ding, C. F. Guo, Z. Wu, Tailoring stress–strain curves of flexible snapping mechanical metamaterial for on-demand mechanical responses via data-driven inverse design, *Advanced Materials* 36 (33) (2024) 2404369.
10. D. Lee, W. Chen, L. Wang, Y.-C. Chan, W. Chen, Data-driven design for metamaterials and multiscale systems: a review, *Advanced Materials* 36 (8) (2024) 2305254.
11. K. Guo, Z. Yang, C.-H. Yu, M. J. Buehler, Artificial intelligence and machine learning in design of mechanical materials, *Materials Horizons* 8 (4) (2021) 1153–1172.
12. X.-Q. Han, X.-D. Wang, M.-Y. Xu, Z. Feng, B.-W. Yao, P.-J. Guo, Z.-F. Gao, Z.-Y. Lu, Ai-driven inverse design of materials: Past, present, and future, *Chinese Physics Letters* 42 (2) (2025) 027403.
13. M. Shafiq, K. Thakre, R. Pandurangan, R. V. S. Lalitha, Generative ai designs the next generation of smart materials from pixels to products, *The International Journal of Advanced Manufacturing Technology* (2025) 1–12.
14. J. Mei, T. Liao, H. Peng, Z. Sun, Bioinspired materials for energy storage, *Small Methods* 6 (2) (2022) 2101076.
15. M. Liu, S. Wang, L. Jiang, Nature-inspired superwettability systems, *Nature Reviews Materials* 2 (7) (2017) 1–17.
16. S. Bonfanti, S. Hiemer, R. Zulkarnain, R. Guerra, M. Zaiser, S. Zapperi, Computational design of mechanical metamaterials, *Nature Computational Science* 4 (8) (2024) 574–583.
17. O. Sigmund, Materials with prescribed constitutive parameters: an inverse homogenization problem, *International Journal of Solids and Structures* 31 (17) (1994) 2313–2329.

18. M. P. Bendsoe, O. Sigmund, *Topology optimization: theory, methods, and applications*, Springer Science & Business Media, 2013.
19. K. V. Wong, A. Hernandez, A review of additive manufacturing, *International scholarly research notices* 2012 (1) (2012) 208760.
20. P. Yang, Z. Guo, N. Hu, W. Sun, Y. Chen, Hyperelastic behaviors of closed-cell porous materials at a wide porosity range, *Composite Structures* 294 (2022) 115792.
21. H. Park, Z. Li, A. Walsh, Has generative artificial intelligence solved inverse materials design?, *Matter* 7 (7) (2024) 2355–2367.
22. Y. Wang, Q. Zeng, J. Wang, Y. Li, D. Fang, Inverse design of shell-based mechanical metamaterial with customized loading curves based on machine learning and genetic algorithm, *Computer Methods in Applied Mechanics and Engineering* 401 (2022) 115571.
23. Q. Zhou, A. Zhao, H. Wang, C. Liu, Machine learning guided design of mechanically efficient metamaterials with auxeticity, *Materials Today Communications* 39 (2024) 108944.
24. S. Kumar, S. Tan, L. Zheng, D. M. Kochmann, Inverse-designed spinodoid metamaterials, *npj Computational Materials* 6 (1) (2020) 73.
25. Q. Zeng, Z. Zhao, H. Lei, P. Wang, A deep learning approach for inverse design of gradient mechanical metamaterials, *International Journal of Mechanical Sciences* 240 (2023) 107920.
26. C. S. Ha, D. Yao, Z. Xu, C. Liu, H. Liu, D. Elkins, M. Kile, V. Deshpande, Z. Kong, M. Bauchy, et al., Rapid inverse design of metamaterials based on prescribed mechanical behavior through machine learning, *Nature Communications* 14 (1) (2023) 5765.
27. H. T. Kollmann, D. W. Abueidda, S. Koric, E. Guleryuz, N. A. Sobh, Deep learning for topology optimization of 2d metamaterials, *Materials & Design* 196 (2020) 109098.
28. J. Wang, W. W. Chen, D. Da, M. Fuge, R. Rai, Ih-gan: A conditional generative model for implicit surface-based inverse design of cellular structures, *Computer Methods in Applied Mechanics and Engineering* 396 (2022) 115060.
29. I. Sosnovik, I. Oseledets, Neural networks for topology optimization, *Russian Journal of Numerical Analysis and Mathematical Modelling* 34 (4) (2019) 215–223.
30. X. Zheng, X. Zhang, T.-T. Chen, I. Watanabe, Deep learning in mechanical metamaterials: from prediction and generation to inverse design, *Advanced Materials* 35 (45) (2023) 2302530.
31. I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, Generative adversarial nets, *Advances in neural information processing systems* 27 (2014).
32. S. Arora, R. Ge, Y. Liang, T. Ma, Y. Zhang, Generalization and equilibrium in generative adversarial nets (gans), in: *International conference on machine learning*, PMLR, 2017, pp. 224–232.
33. D. P. Kingma, M. Welling, Auto-encoding variational bayes, *arXiv preprint arXiv:1312.6114* (2013).
34. X. Y. Lee, J. R. Waite, C.-H. Yang, B. S. S. Pokuri, A. Joshi, A. Balu, C. Hegde, B. Ganapathysubramanian, S. Sarkar, Fast inverse design of microstructures via generative invariance networks, *Nature Computational Science* 1 (3) (2021) 229–238.
35. J. Ho, A. Jain, P. Abbeel, Denoising diffusion probabilistic models, *Advances in neural information processing systems* 33 (2020) 6840–6851.
36. J. Ho, T. Salimans, Classifier-free diffusion guidance, *arXiv preprint arXiv:2207.12598* (2022).
37. R. Rombach, A. Blattmann, D. Lorenz, P. Esser, B. Ommer, High-resolution image synthesis with latent diffusion models, in: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 10684–10695.
38. G. Daras, H. Chung, C.-H. Lai, Y. Mitsufuji, J. C. Ye, P. Milanfar, A. G. Dimakis, M. Delbracio, A survey on diffusion models for inverse problems, *arXiv preprint arXiv:2410.00083* (2024).
39. N. N. Vlassis, W. Sun, Denoising diffusion algorithm for inverse design of microstructures with fine-tuned nonlinear material properties, *Computer Methods in Applied Mechanics and Engineering* 413 (2023) 116126.
40. J.-H. Bastek, D. M. Kochmann, Inverse design of nonlinear mechanical metamaterials via video denoising diffusion models, *Nature Machine Intelligence* 5 (12) (2023) 1466–1475.
41. E. Azqadan, H. Jahed, A. Arami, Predictive microstructure image generation using denoising diffusion probabilistic models, *Acta Materialia* 261 (2023) 119406.
42. X. Lyu, X. Ren, Microstructure reconstruction of 2d/3d random materials via diffusion-based deep generative models, *Scientific Reports* 14 (1) (2024) 5041.
43. K.-H. Lee, G. J. Yun, Microstructure reconstruction using diffusion-based generative models, *Mechanics of Advanced Materials and Structures* 31 (18) (2024) 4443–4461.

44. H. Wang, Z. Du, F. Feng, Z. Kang, S. Tang, X. Guo, Diffmat: Data-driven inverse design of energy-absorbing metamaterials using diffusion model, *Computer Methods in Applied Mechanics and Engineering* 432 (2024) 117440.
45. R. W. Ogden, Large deformation isotropic elasticity—on the correlation of theory and experiment for incompressible rubberlike solids, *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences* 326 (1567) (1972) 565–584.
46. M. Fernández, F. Fritzen, O. Weeger, Material modeling for parametric, anisotropic finite strain hyperelasticity based on machine learning with application in optimization of metamaterials, *International Journal for Numerical Methods in Engineering* 123 (2) (2022) 577–609.
47. O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, in: *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III* 18, Springer, 2015, pp. 234–241.
48. O. Oktay, J. Schlemper, L. L. Folgoc, M. Lee, M. Heinrich, K. Misawa, K. Mori, S. McDonagh, N. Y. Hammerla, B. Kainz, et al., Attention u-net: Learning where to look for the pancreas, *arXiv preprint arXiv:1804.03999* (2018).
49. Hibbitt, Karlsson, Sorensen, *ABAQUS: Theory manual, Vol. 2*, Hibbitt, Karlsson & Sorensen, 1997.
50. P. Wang, Implementation of imagen, google's text-to-image neural network that beats dall-e2, Pytorch. GitHub <https://github.com/lucidrains/imagen-pytorch> (2022).

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.