Dual-Constrained Diffusion Model for Nonlinear Microstructure Inverse Design with Direct Evaluation

Jecheon Yu [1*], Jinun Byun [1*], Hyeonbin Moon [1], Junhyeong Lee [1], Seunghwa Ryu [1‡]

[1] Department of Mechanical Engineering, Korea Advanced Institute of Science and Technology (KAIST), Daejeon, Republic of Korea

* Jecheon Yu and Jinun Byun contributed equally to this work.

Abstract: Diffusion models have emerged as powerful tools for inverse design, yet their stochastic nature typically requires surrogate models or additional numerical simulations to verify target-condition satisfaction. To address this bottleneck, we propose a dual-constrained diffusion framework that jointly generates structural designs and solution fields under both physical and conditional constraints. The framework combines a dual-constrained diffusion model (DCDM) for training-stage constraint enforcement with dual-guided diffusion sampling (DGS) for sampling-stage guidance, enabling direct evaluation without additional numerical simulations. We demonstrate the framework through the inverse design of hyperelastic microstructures conditioned on prescribed strain energy under uniaxial extension. Compared with a conventional diffusion model, DCDM with DGS reduces physics residual errors by more than an order of magnitude. In addition, 1,000 generated candidates can be screened directly in approximately 0.27 seconds. Importantly, because the dual constraints are incorporated during training, a single DCDM supports target-aligned unguided generation and can be combined with DGS for guided generation. For high-density in-distribution targets, unguided generation produces structurally diverse candidates with reliable direct evaluation at one-fifth the sampling time of guided generation. When higher target fidelity is desired or out-of-distribution targets are considered, DGS can be activated on the same trained model to increase the number of valid candidates and maintain reliable target satisfaction beyond the training distribution. The proposed framework provides a general and practical approach to inverse design for systems governed by physical laws, requiring only the definition of physics and condition residuals.

Keywords: diffusion model; inverse design; physics-informed diffusion model; physics-guided sampling; hyperelastic microstructures

1. Introduction

The macroscopic mechanical behavior of materials, including stiffness, strength, and energy-related responses, is closely linked to the structure of their underlying microstructures [1–3]. In particular, the geometric configuration and spatial distribution of material properties within a microstructure largely determine these mechanical responses. Understanding the relationship between microstructure and mechanical responses is essential for designing microstructures with prescribed mechanical properties. To this end, microstructure design has traditionally relied on trial-and-error or optimization-based approaches. However, these approaches remain challenging due to the high-dimensional design space, non-convexity, and the nonlinear relationships between design variables and material responses [4].

In recent years, machine learning (ML)-based approaches have emerged as promising alternatives for addressing these challenges in inverse design problems [5]. ML models can effectively capture the complex and nonlinear relationships between microstructures and mechanical responses [6]. These models can be employed as surrogate models for computationally expensive numerical simulations and coupled with meta-heuristic optimization methods to search for microstructures with desired mechanical properties [7]. Meanwhile, generative models such as variational autoencoders (VAEs) [8] and generative adversarial networks (GANs) [9] have been adopted for inverse design, but their practical applicability remains limited due to issues such as sample smoothness and training instability [10]. More recently, diffusion models [11,12] have emerged as powerful generative models capable of learning complex data distributions through a progressive denoising process. These models have been successfully applied to inverse design problems [4,13–15]. However, because diffusion models generate samples probabilistically, the generated designs do not always strictly satisfy the target conditions. As a result, additional validation or filtering is typically required using surrogate models [4] or computationally expensive numerical

simulations [14,16]. This reliance on external validation creates a critical bottleneck in generative inverse design, particularly when large numbers of candidates must be screened. Bastek and Kochmann [13] partially addressed this limitation by jointly generating displacement and stress fields for evaluation. However, because mechanical consistency is not explicitly enforced, discrepancies with high-fidelity numerical simulations may still remain. Incorporating physical constraints into diffusion models therefore provides a promising direction to improve the physical reliability of generated designs and to reduce the reliance on external validation.

Inspired by physics-informed neural networks (PINN) [17], recent studies have attempted to incorporate physics-based information into machine learning models, including generative models [18,19]. These efforts have been extended to diffusion models to enhance the physical consistency of generated samples. For example, Shu et al. [20] employed partial differential equation (PDE) residuals as physics-informed conditioning information to guide the denoising process for high-fidelity flow field reconstruction. Bastek et al. [21] incorporated PDE constraints into diffusion model training by augmenting the training objective with a physics-based residual loss. Alternatively, physical constraints can be incorporated during the sampling process by introducing gradient-based guidance during the diffusion sampling process [22,23]. CoCoGen [24] builds a conditional diffusion model and enforces physical consistency by appending residual-minimization steps to the reverse diffusion solver, directly correcting the intermediate diffusion samples during sampling. To solve inverse problems under partial observation, DiffusionPDE [25] guides the sampling process using gradients derived from observation losses and PDE losses evaluated on the denoised estimate. To the best of our knowledge, existing diffusion models generally enforce physical constraints either during training or sampling, whereas a unified framework that enforces both physical and conditional constraints across both stages and clarifies their complementary roles in inverse

design remains largely unexplored.

In this work, we propose a dual-constrained diffusion framework that directly addresses this validation bottleneck by jointly generating design and solution fields under physical and conditional constraints. The framework integrates a dual-constrained diffusion model (DCDM) for training-stage constraint enforcement with dual-guided diffusion sampling (DGS) for sampling-stage enforcement. By ensuring physical consistency of the generated joint fields, the proposed framework enables direct evaluation of mechanical responses from generated samples without surrogate models or additional numerical simulations for a controlled benchmark, Mechanical MNIST [26] dataset consisting of microstructures with hyperelastic discretized blocks. As a result, direct sample-based evaluation becomes both reliable and computationally efficient, with 1,000 candidates evaluated in 0.27 seconds while maintaining close agreement with numerical simulations. Under high-density in-distribution conditions, unguided DCDM enables faster and more diverse generation with reliable direct evaluation, whereas DGS improves target satisfaction. Under out-of-distribution (OOD) target conditions, reliable inverse design requires DGS to be applied to a model trained with physical constraint, highlighting that sampling-stage guidance alone is insufficient when the learned generative distribution is not physically consistent. The overall workflow of the proposed framework is summarized in **Fig. 1,** illustrating dual-constrained generation of physically consistent joint fields and direct evaluation of the response without additional numerical simulations.

The main contributions of this study are as follows:

(1) A generative inverse design framework is established for direct and reliable evaluation of design candidates without surrogate models or additional numerical simulations, by jointly generating design and solution fields under physical consistency.

(2) A dual-constrained formulation is developed by incorporating physical and conditional residuals into DCDM training and DGS sampling, enabling constraint enforcement throughout the generation process.

(3) The complementary roles of training- and sampling-stage constraint enforcement are demonstrated: training-stage dual constraints enable physically consistent, target-aligned unguided generation with greater structural diversity under in-distribution conditions, whereas adding sampling-stage dual guidance to the same trained DCDM increases candidate validity and enables reliable out-of-distribution performance.

The remainder of this paper is organized as follows. **Section 2** introduces the proposed framework and describes how physical and conditional constraints are incorporated into both the training and sampling stages. **Section 3** presents the problem setting, including the inverse design formulation, the data generation procedure, and the weak-form formulations of the physics and condition residuals. **Section 4** presents the inverse design results and evaluates the performance of the proposed framework. **Section 5** concludes the paper with a summary of the main findings.

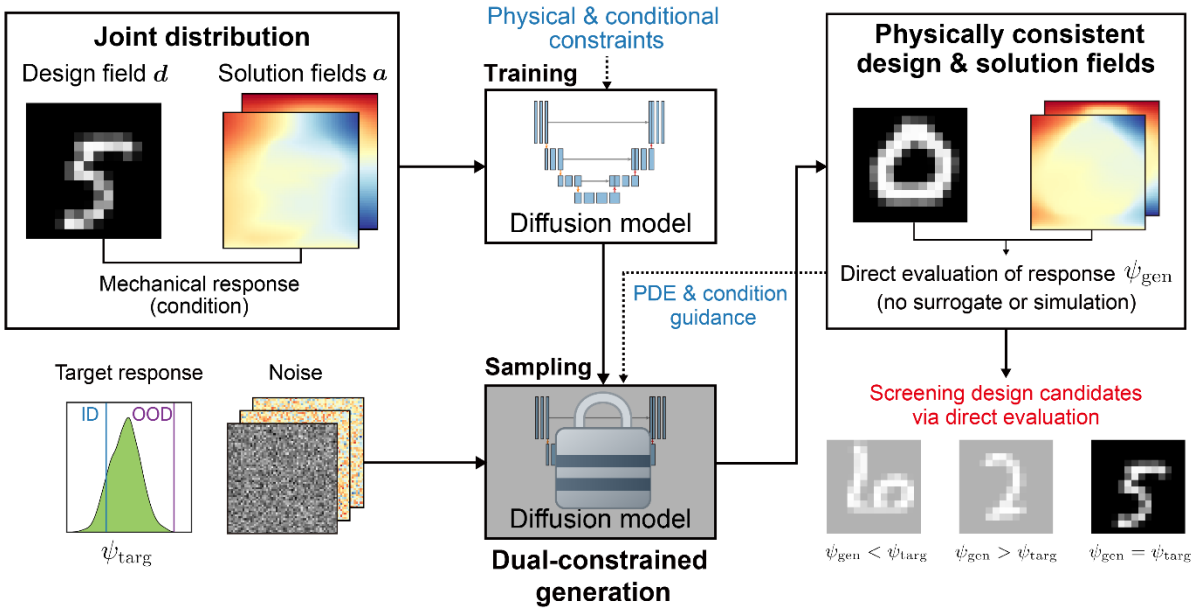

Figure 1. Overview of the proposed dual-constrained diffusion framework for inverse design. DCDM is trained to learn the joint distribution of the design field $\boldsymbol{d}$ and the corresponding solution fields $\boldsymbol{a}$ under physical and conditional constraints, while DGS additionally enforces these constraints during sampling. Dual-constrained generation enables direct evaluation of the mechanical response from generated samples without surrogate models or additional numerical simulations, thereby allowing efficient screening of design candidates. The framework is assessed under both in-distribution and out-of-distribution target conditions.

2. Methods

2.1. Theoretical background

2.1.1. Denoising diffusion models

The denoising diffusion model has recently emerged as one of the most prominent generative models [11,12,22]. It is trained to generate samples from the target data distribution $q(\boldsymbol{x}_0)$ by learning to progressively remove noise from samples drawn from a simple prior $p(\boldsymbol{x}_T)$, typically a standard Gaussian $\mathcal{N}(\boldsymbol{0}, \boldsymbol{I})$. A diffusion model consists of two processes: a forward diffusion process and a reverse denoising process. In the forward process, Gaussian noise is progressively injected into a clean data sample $\boldsymbol{x}_0 \sim q(\boldsymbol{x}_0)$ according to a predefined variance schedule $\{\beta_t\}_{t=1}^{T}$, expressed as

$$q(\boldsymbol{x}_{1:T}|\boldsymbol{x}_0) = \prod_{t=1}^{T} q(\boldsymbol{x}_t|\boldsymbol{x}_{t-1}), \quad q(\boldsymbol{x}_t|\boldsymbol{x}_{t-1}) = \mathcal{N}\left(\boldsymbol{x}_t; \sqrt{1-\beta_t}\boldsymbol{x}_{t-1}, \beta_t \boldsymbol{I}\right), \tag{1}$$

where $\beta_t \in (0,1)$ is the noise variance at timestep t. Defining $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{i=1}^{t} \alpha_i$, the forward process can be written as

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

In the reverse process, the clean signal $\boldsymbol{x}_0$ is progressively reconstructed from the noisy sample $\boldsymbol{x}_t$. The reverse process, parameterized by a neural network with $\boldsymbol{\theta}$, is given by

$$p(\boldsymbol{x}_{0:T}) = p(\boldsymbol{x}_T)\prod_{t=1}^{T} p_{\boldsymbol{\theta}}(\boldsymbol{x}_{t-1}|\boldsymbol{x}_t), \qquad p_\theta(\boldsymbol{x}_{t-1}|\boldsymbol{x}_t) = \mathcal{N}\big(\boldsymbol{x}_{t-1}; \boldsymbol{\mu}_{\boldsymbol{\theta}}(\boldsymbol{x}_t, t), \boldsymbol{\Sigma}(\boldsymbol{x}_t, t)\big), \tag{3}$$

where $\boldsymbol{\mu}_{\boldsymbol{\theta}}(\boldsymbol{x}_t, t)$ is trained to approximate the mean of the true posterior distribution $q(\boldsymbol{x}_{t-1}|\boldsymbol{x}_t)$ while the posterior covariance is fixed to $\boldsymbol{\Sigma}(\boldsymbol{x}_t, t) = \frac{1-\bar{\alpha}_{t-1}}{1-\bar{\alpha}_t}\beta_t \boldsymbol{I} = \Sigma_t \boldsymbol{I}$. Instead of directly modeling $\boldsymbol{\mu}_\theta(\boldsymbol{x}_t, t)$, alternative parameterizations can be obtained using the reparameterization trick together with Bayes' theorem [8,12]:

$$\mu_{\boldsymbol{\theta}}(\boldsymbol{x}_t, t) = \frac{1}{\sqrt{\alpha_t}}\left(\boldsymbol{x}_t - \frac{\beta_t}{\sqrt{1-\bar{\alpha}_t}}\boldsymbol{\epsilon}_{\boldsymbol{\theta}}(\boldsymbol{x}_t, t)\right) = \frac{\sqrt{\alpha_t}(1-\bar{\alpha}_{t-1})}{1-\bar{\alpha}_t}\boldsymbol{x}_t + \frac{\sqrt{\bar{\alpha}_{t-1}}\beta_t}{1-\bar{\alpha}_t}\hat{\boldsymbol{x}}_{0,\theta}(\boldsymbol{x}_t, t), \tag{4}$$

where $\boldsymbol{\epsilon}_\theta$ is the model prediction of injected noise, and $\hat{\boldsymbol{x}}_{0,\theta}$ is the estimate of clean data. Since $\boldsymbol{x}_t$ is available during training, estimating $\boldsymbol{\mu}_{\boldsymbol{\theta}}$ is equivalent to predicting either $\boldsymbol{\epsilon}_\theta$ or $\hat{\boldsymbol{x}}_{0,\theta}$ as these quantities are linearly related in **Eq. (2)**. In contrast to the widely adopted $\boldsymbol{\epsilon}$-prediction [12], we adopt an $\boldsymbol{x}_0$-prediction formulation and consider the following training objective:

$$L(\theta) = \mathbb{E}_{t,\boldsymbol{x}_0,\boldsymbol{\epsilon}}\left[\lambda_t \left\|\boldsymbol{x}_0 - \hat{\boldsymbol{x}}_{0,\theta}(\boldsymbol{x}_t, t)\right\|^2\right], \tag{5}$$

where λ_t follows Min-SNR-5 weighting strategy [27], improving training stability by balancing contributions from different noise levels. For notation simplicity, the parametric dependence on θ in $\hat{\boldsymbol{x}}_{0,\theta}$ is omitted in the following.

After training, a new sample $\boldsymbol{x}_0$ is generated by iteratively applying the learned reverse diffusion process starting from Gaussian noise $\boldsymbol{x}_T \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I})$.

$$\boldsymbol{x}_{t-1} = \frac{\sqrt{\alpha_t}(1-\bar{\alpha}_{t-1})}{1-\bar{\alpha}_t}\boldsymbol{x}_t + \frac{\sqrt{\bar{\alpha}_{t-1}}\beta_t}{1-\bar{\alpha}_t}\hat{\boldsymbol{x}}_0(\boldsymbol{x}_t, t) + \sqrt{\Sigma_t}\boldsymbol{z}, \qquad \boldsymbol{z} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I}) \tag{6}$$

2.1.2. Residual formulation for physical consistency

Physical laws governing a system can be formulated as PDEs with associated boundary conditions defined over a physical domain $\Omega \subset \mathbb{R}^d$, expressed as:

$$\begin{cases} \boldsymbol{\mathcal{F}}[\boldsymbol{d}(\boldsymbol{\xi}), \boldsymbol{a}(\boldsymbol{\xi})] = \mathbf{0} \ \text{ in } \Omega \\ \boldsymbol{\mathcal{B}}[\boldsymbol{a}(\boldsymbol{\xi})] = \mathbf{0} \ \text{ on } \partial\Omega \end{cases}, \tag{7}$$

where $\boldsymbol{\mathcal{F}}[\cdot]$ and $\boldsymbol{\mathcal{B}}[\cdot]$ denote the differential operators and the boundary operator, respectively, $\partial\Omega$ is the boundary of the domain Ω, and $\boldsymbol{\xi} \in \Omega$ is the spatial coordinate. $\boldsymbol{d}$ is the spatially distributed design field that parameterizes the governing PDEs. $\boldsymbol{a}$ is the corresponding solution field that satisfies the PDEs parameterized by $\boldsymbol{d}$ in Ω and associated boundary conditions on $\partial\Omega$.

In this work, we aim to train a diffusion model to learn the joint distribution of data $\boldsymbol{x} := (\boldsymbol{d}, \boldsymbol{a})$ such that the generated samples satisfy **Eq. (7)**. The spatial coordinate $\boldsymbol{\xi}$ is discretized on an $n \times n$ grid, so that $\boldsymbol{d}(\boldsymbol{\xi})$ and $\boldsymbol{a}(\boldsymbol{\xi})$ are discretized accordingly. The physics residual $\boldsymbol{\mathcal{R}}_{\text{phys}}(\boldsymbol{x})$, defined for the discretized fields $\boldsymbol{x}$, measures the extent to which the governing PDEs and boundary conditions are satisfied:

$$\boldsymbol{\mathcal{R}}_{\text{phys}}(\boldsymbol{x}) \coloneqq \begin{bmatrix} \boldsymbol{\mathcal{F}}[\boldsymbol{x}] \\ \boldsymbol{\mathcal{B}}[\boldsymbol{x}] \end{bmatrix} \tag{8}$$

Conventional diffusion models are trained on discretized data $\boldsymbol{x}$ obtained from numerical simulations that satisfy **Eq. (7)**. The learned joint distribution implicitly reflects the underlying physical consistency. However, since no explicit physical constraint is incorporated into the training objective, physical consistency is not directly enforced during generation. To

explicitly incorporate physical constraints into the generative framework, prior studies have proposed two main strategies: (i) integrating such constraints into the training objective [21] and (ii) guiding diffusion sampling to enforce physical consistency [20,24,25]. The former strategy is first introduced through the physics-informed diffusion model in **Section 2.1.3** and is further extended in **Section 2.2.1** by incorporating conditional constraints while the sampling-based strategy in **Section 2.2.2** applies guidance to enforce these constraints during generation.

2.1.3. Physics-informed diffusion models

A physics-informed diffusion model (PIDM) [21] integrates constraints derived from governing physical equations directly into the diffusion model training process, thereby enforcing generated samples to satisfy the physical laws. To incorporate such physical constraints into the training objective while preserving the probabilistic formulation of generative models, PIDM introduces the physics residual as a virtual observable [28], $\hat{\boldsymbol{r}} = 0$. The virtual observable is assumed to be sampled from the following Gaussian distribution centered at $\boldsymbol{\mathcal{R}}_{\text{phys}}(\boldsymbol{x}_0)$:

$$q_{\mathcal{R}_{\text{phys}}}(\hat{\boldsymbol{r}}|\boldsymbol{x}_0) = \mathcal{N}\left(\hat{\boldsymbol{r}}; \boldsymbol{\mathcal{R}}_{\text{phys}}(\boldsymbol{x}_0), \sigma^2 \boldsymbol{I}\right), \tag{9}$$

where σ^2 is the variance controlling the strength of the imposed physical constraint, with $\boldsymbol{\mathcal{R}}_{\text{phys}}(\boldsymbol{x}_0) = 0$ being enforced more strongly as σ^2 decreases. The physics residual is evaluated with respect to the clean sample $\boldsymbol{x}_0$. The virtual likelihood $p_\theta(\hat{\boldsymbol{r}})$ can be obtained by marginalizing over the clean sample $\boldsymbol{x}_0$ as follows:

$$p_\theta(\hat{\boldsymbol{r}}) = \int p_\theta(\hat{\boldsymbol{r}}, \boldsymbol{x}_0)\, d\boldsymbol{x}_0 = \int q_{\mathcal{R}_{\text{phys}}}(\hat{\boldsymbol{r}}|\boldsymbol{x}_0) p_\theta(\boldsymbol{x}_0) d\boldsymbol{x}_0 = \mathbb{E}_{\mathbf{x}_0 \sim p_\theta(\mathbf{x}_0)}\left[q_{\mathcal{R}_{\text{phys}}}(\hat{\boldsymbol{r}}|\boldsymbol{x}_0)\right] \tag{10}$$

The training objective of PIDM in **Eq. (11)** includes the virtual likelihood term $p_\theta(\hat{\boldsymbol{r}})$ as well as the data likelihood term $p_\theta(\boldsymbol{x}_0)$ in purely data-driven diffusion models. This guides the generative distribution toward a physically consistent subspace supported by the training data.

$$\arg\max_{\theta} \mathbb{E}_{\boldsymbol{x}_0 \sim q(\boldsymbol{x}_0)}[\log p_\theta(\boldsymbol{x}_0)] + \mathbb{E}_{\mathbf{x}_0 \sim p_\theta(\mathbf{x}_0)}\left[\log q_{\mathcal{R}_{\text{phys}}}(\hat{\boldsymbol{r}} = \mathbf{0}|\boldsymbol{x}_0)\right] \tag{11}$$

In the reverse process of the denoising diffusion algorithm, the clean sample $\boldsymbol{x}_0$ is not directly accessible and only the noisy sample $\boldsymbol{x}_t$ is available, on which the physics residual is not meaningfully defined. The denoising MSE objective in **Eq. (5)** implies that the diffusion model output $\hat{\boldsymbol{x}}_0$ is an estimate of $\mathbb{E}[\boldsymbol{x}_0|\boldsymbol{x}_t]$ and this estimate is used to evaluate the physics residual [11]. Accordingly, **Eq. (11)** can be rewritten as:

$$\arg\max_{\theta} \mathbb{E}_{\boldsymbol{x}_0 \sim q(\boldsymbol{x}_0)}[\log p_\theta(\boldsymbol{x}_0)] + \mathbb{E}_{\boldsymbol{x}_{1:T} \sim p_\theta(\boldsymbol{x}_{1:T})}\left[\log q_{\mathcal{R}_{\text{phys}}}(\hat{\boldsymbol{r}} = \mathbf{0}|\hat{\boldsymbol{x}}_0(\boldsymbol{x}_{1:T})\right]. \tag{12}$$

However, evaluating the physics residual using the estimate $\hat{\boldsymbol{x}}_0$ can degrade the effectiveness of physical enforcement, since the physics residual on the posterior mean, $\boldsymbol{\mathcal{R}}_{\text{phys}}(\mathbb{E}[\boldsymbol{x}_0|\boldsymbol{x}_t])$, is not equivalent to the expected residual, $\mathbb{E}\left[\boldsymbol{\mathcal{R}}_{\text{phys}}(\boldsymbol{x}_0)|\boldsymbol{x}_t\right]$, a discrepancy commonly referred to as the Jensen gap [21,29,30]. As the timestep t increases, the conflict between the data and residual loss terms in the objective becomes more pronounced. To mitigate this, the variance of residual likelihood is increased with t, thereby progressively relaxing the physical constraint as $t \to T$. Accordingly, $q_{\mathcal{R}_{\text{phys}}}\left(\hat{\boldsymbol{r}}\middle|\, \hat{\boldsymbol{x}}_0(\boldsymbol{x}_t, t)\right)$ is modeled as a Gaussian distribution whose variance is given by a scaled version of the fixed variance schedule of the denoising process as:

$$q_{\mathcal{R}_{\text{phys}}}\left(\hat{\boldsymbol{r}}|\hat{\boldsymbol{x}}_0(\boldsymbol{x}_{1:T})\right) = \prod_{t=1}^{T} q_{\mathcal{R}_{\text{phys}}}\left(\hat{\boldsymbol{r}}|\hat{\boldsymbol{x}}_0(\boldsymbol{x}_t, t)\right)$$
$$\text{, where } \; q_{\mathcal{R}_{\text{phys}}}\left(\hat{\boldsymbol{r}}\middle|\, \hat{\boldsymbol{x}}_0(\boldsymbol{x}_t, t)\right) = \mathcal{N}(\hat{\boldsymbol{r}}; \boldsymbol{\mathcal{R}}_{\text{phys}}(\hat{\boldsymbol{x}}_0), \Sigma_t / c_{\text{phys}} \boldsymbol{I}), \tag{13}$$

where $\Sigma_t = \frac{1-\bar{\alpha}_{t-1}}{1-\bar{\alpha}_t}\beta_t$ is the fixed posterior variance of reverse diffusion process, and the hyperparameter c_{phys} serves as a variance scale factor that controls the strength of penalization for deviations from $\boldsymbol{\mathcal{R}}_{\text{phys}}(\hat{\boldsymbol{x}}_0) = 0$. Since sampling from $\boldsymbol{x}_{1:T} \sim p_\theta(\boldsymbol{x}_{1:T})$ in **Eq. (12)** is intractable during training, it is simplified by sampling from the forward diffusion distribution

$q(\boldsymbol{x}_{1:T})$, leading to the training objective of PIDM in **Eq. (14)**. A detailed derivation of this formulation can be found in [21].

$$L_{\text{PIDM}}(\boldsymbol{\theta}) = \mathbb{E}_{t\sim[1,T],\mathbf{x}_{0:T}\sim q(\mathbf{x}_{0:T})}\left[\lambda_t\|\boldsymbol{x}_0 - \hat{\boldsymbol{x}}_0(\boldsymbol{x}_t,t)\|^2 + \frac{c_{\text{phys}}}{2\Sigma_t}\left\|\boldsymbol{\mathcal{R}}_{\text{phys}}\left(\hat{\boldsymbol{x}}_0(\boldsymbol{x}_t,t)\right)\right\|^2\right] \quad (14)$$

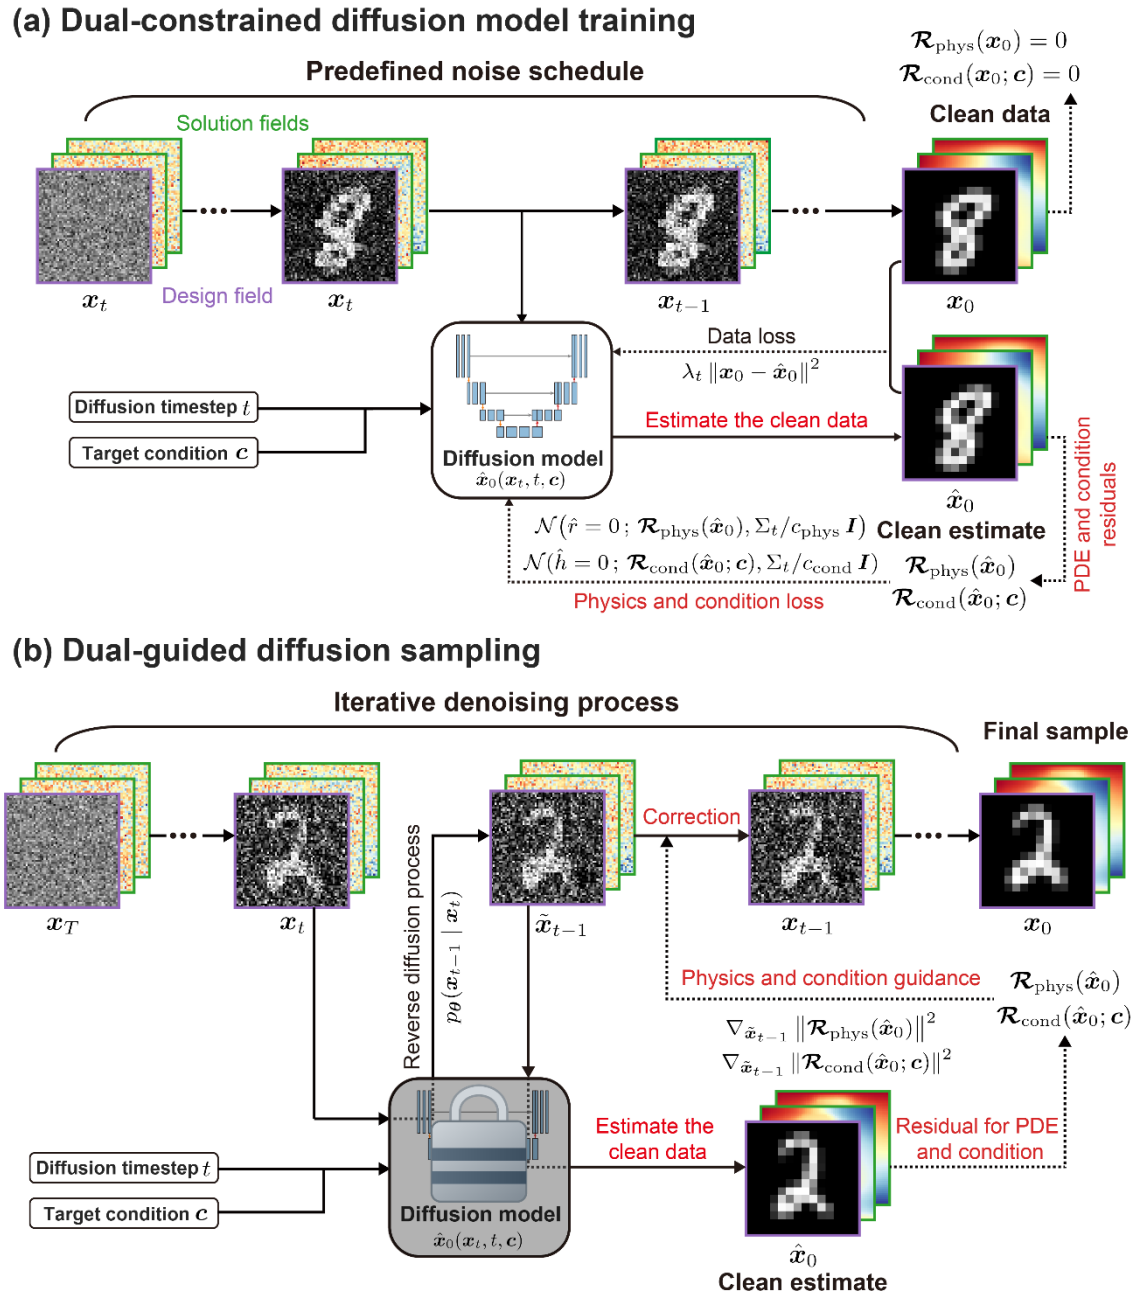

Figure 2. (a) Training of the dual-constrained diffusion model. The diffusion model predicts the clean sample from noisy inputs $\boldsymbol{x}_t$. PDE and condition residuals evaluated on the clean estimate $\hat{\boldsymbol{x}}_0$ are incorporated through residual likelihoods, enforcing physical consistency and condition satisfaction. (b) Dual-guided diffusion sampling. During sampling, residual gradients computed from the estimated clean sample $\hat{\boldsymbol{x}}_0$ are used to guide the reverse diffusion process toward samples that satisfy both the governing physics and the prescribed target conditions.

2.2. Proposed dual-constrained diffusion framework

The proposed dual-constrained diffusion framework consists of two components: a dual-constrained diffusion model (DCDM), which enforces physical and conditional constraints during training, and dual-guided diffusion sampling (DGS), which additionally enforces these constraints during sampling. These components enable the generation of samples that are not only physically consistent but also aligned with prescribed target conditions. The overall framework is summarized in **Fig. 2**.

2.2.1. Dual-constrained training: integration of physical and conditional constraints

Although PIDM can generate physically consistent samples in an unconditional manner, it is insufficient for inverse design problems requiring prescribed target properties. To enable generation of physically consistent samples conditioned on prescribed target conditions, we extend the PIDM to a conditional diffusion model by incorporating the target conditions $\boldsymbol{c}$ as an additional input vector. The corresponding training objective is given by

$$L_{\mathrm{cPIDM}}(\boldsymbol{\theta}) = \mathbb{E}_{t\sim[1,T],\boldsymbol{x}_{0:T}\sim q(\boldsymbol{x}_{0:T})}\left[\lambda_t\|\boldsymbol{x}_0 - \hat{\boldsymbol{x}}_0(\boldsymbol{x}_t,t,\boldsymbol{c})\|^2 + \frac{C_{\mathrm{phys}}}{2\Sigma_t}\left\|\boldsymbol{\mathcal{R}}_{\mathrm{phys}}\left(\hat{\boldsymbol{x}}_0(\boldsymbol{x}_t,t,\boldsymbol{c})\right)\right\|^2\right]. \quad (15)$$

By minimizing the objective in **Eq. (15),** the conditional reverse transition $p_{\boldsymbol{\theta}}(\boldsymbol{x}_{t-1}|\boldsymbol{x}_t,\boldsymbol{c})$ is learned, enabling conditional generation of samples given $\boldsymbol{c}$.

Furthermore, since the target properties can be directly computed from the generated design and solution fields, we introduce a conditional constraint that is incorporated into the training objective in a probabilistic manner analogous to the physical constraint, improving consistency with the prescribed target properties during generation. The condition residual $\boldsymbol{\mathcal{R}}_{\mathrm{cond}}(\boldsymbol{x}_0;\boldsymbol{c})$ that measures the extent to which a generated sample satisfies the target conditions $\boldsymbol{c}$ is defined by

$$\boldsymbol{\mathcal{R}}_{\mathrm{cond}}(\boldsymbol{x}_0;\boldsymbol{c}) := \boldsymbol{\mathcal{H}}(\boldsymbol{x}_0) - \boldsymbol{c}, \quad (16)$$

where $\boldsymbol{\mathcal{H}}(\cdot)$ denotes an operator that computes the properties from the clean sample $\boldsymbol{x}_0$. To integrate the conditional constraint into the training objective, the condition residual is

introduced as a virtual observable, $\widehat{\boldsymbol{h}} = 0$, assumed to be sampled from a Gaussian distribution centered at $\boldsymbol{\mathcal{R}}_{\text{cond}}(\boldsymbol{x}_0; \boldsymbol{c})$:

$$q_{\boldsymbol{\mathcal{R}}_{\text{cond}}}(\widehat{\boldsymbol{h}}|\boldsymbol{x}_0, \boldsymbol{c}) = \mathcal{N}(\widehat{\boldsymbol{h}}; \boldsymbol{\mathcal{R}}_{\text{cond}}(\boldsymbol{x}_0; \boldsymbol{c}), \sigma^2 \boldsymbol{I}) \tag{17}$$

The virtual likelihood $p_\theta(\widehat{\boldsymbol{h}}|\boldsymbol{c})$ is obtained by marginalizing over the $\boldsymbol{x}_0$ as follows:

$$p_\theta(\widehat{\boldsymbol{h}}|\boldsymbol{c}) = \int p_\theta(\widehat{\boldsymbol{h}}, \boldsymbol{x}_0|\boldsymbol{c})\, d\mathbf{x}_0 = \int q_{\mathcal{R}_{\text{cond}}}(\widehat{\boldsymbol{h}}|\boldsymbol{x}_0, \boldsymbol{c}) p_\theta(\boldsymbol{x}_0|\boldsymbol{c}) d\boldsymbol{x}_0 = \mathbb{E}_{\mathbf{x}_0 \sim p_\theta(\boldsymbol{x}_0|\boldsymbol{c})}\left[q_{\mathcal{R}_{\text{cond}}}(\widehat{\boldsymbol{h}}|\boldsymbol{x}_0, \boldsymbol{c})\right] \tag{18}$$

Analogously to $q_{\mathcal{R}_{\text{phys}}}(\hat{\boldsymbol{r}}|\boldsymbol{x}_0)$, $q_{\mathcal{R}_{\text{cond}}}\left(\widehat{\boldsymbol{h}}\middle|\widehat{\boldsymbol{x}}_0(\boldsymbol{x}_{1:T})\right)$ is modeled as a Gaussian distribution whose variance follows a scaled version of the fixed schedule of the denoising process:

$$q_{\mathcal{R}_{\text{cond}}}(\widehat{\boldsymbol{h}}|\widehat{\boldsymbol{x}}_0(\boldsymbol{x}_{1:T}), \boldsymbol{c}) = \prod_{t=1}^{T} q_{\mathcal{R}_{\text{cond}}}(\widehat{\boldsymbol{h}}|\widehat{\boldsymbol{x}}_0(\boldsymbol{x}_t, t, \boldsymbol{c}), \boldsymbol{c})$$

$$\text{, where } \; q_{\mathcal{R}_{\text{cond}}}\left(\widehat{\boldsymbol{h}}\middle|\widehat{\boldsymbol{x}}_0(\boldsymbol{x}_t, t, \boldsymbol{c})\right) = \mathcal{N}(\widehat{\boldsymbol{h}}; \boldsymbol{\mathcal{R}}_{\text{cond}}(\widehat{\boldsymbol{x}}_0(\boldsymbol{x}_t, t, \boldsymbol{c}); \boldsymbol{c}), \Sigma_t / c_{\text{cond}} \boldsymbol{I}), \tag{19}$$

where Σ_t is the fixed posterior variance of reverse diffusion process, and the hyperparameter c_{cond} is a variance scale factor that controls the strength of penalization for deviations from $\boldsymbol{\mathcal{R}}_{\text{cond}}(\widehat{\boldsymbol{x}}_0; \boldsymbol{c}) = 0$.

By incorporating the maximization of the condition virtual residual log-likelihood into the training objective in **Eq. (15)**, the proposed objective in this work is given by

$$L_{\text{DCDM}}(\boldsymbol{\theta}) = \mathbb{E}_{t\sim[1,T], \boldsymbol{x}_{0:T}\sim q(\boldsymbol{x}_{0:T})}\Bigg[\underbrace{\lambda_t \|\boldsymbol{x}_0 - \widehat{\boldsymbol{x}}_0(\boldsymbol{x}_t, t, \boldsymbol{c})\|^2}_{\text{data loss}} + \underbrace{\frac{c_{\text{phys}}}{2\Sigma_t} \left\|\boldsymbol{\mathcal{R}}_{\text{phys}}\left(\widehat{\boldsymbol{x}}_0(\boldsymbol{x}_t, t, \boldsymbol{c})\right)\right\|^2}_{\text{physics loss}}$$
$$+ \underbrace{\frac{c_{\text{cond}}}{2\Sigma_t} \|\boldsymbol{\mathcal{R}}_{\text{cond}}(\widehat{\boldsymbol{x}}_0(\boldsymbol{x}_t, t, \boldsymbol{c}); \boldsymbol{c})\|^2}_{\text{condition loss}}\Bigg]. \tag{20}$$

Enforcing both physical consistency and condition satisfaction during training enables the diffusion model to learn a conditional generative process that produces physically valid designs satisfying the prescribed target conditions. The overall training procedure of DCDM employed in this work is summarized in **Algorithm 1**.

Algorithm 1 Dual-constrained training procedure of DCDM

1: **Require**: Target condition $\mathbf{c}$, Variance scale factors c_{phys}, c_{cond}, Variance schedule $\{\Sigma_t\}_{t=1}^{T}$, Data loss weights $\{\lambda_t\}_{t=1}^{T}$, Diffusion steps T, Physics residual $\boldsymbol{\mathcal{R}}_{\text{phys}}(\cdot)$, Condition residual $\boldsymbol{\mathcal{R}}_{\text{cond}}(\cdot, \mathbf{c})$

2: **Repeat**

3: Sample $(\boldsymbol{x}_0, \boldsymbol{c})$ from dataset $\left\{\boldsymbol{x}_0^{(i)}, \boldsymbol{c}^{(i)}\right\}_{i=1}^{N}$

4: $t \sim \text{Uniform}\{1, \dots, T\}$

5: $\boldsymbol{\epsilon} \sim \boldsymbol{\mathcal{N}}(\mathbf{0}, \mathbf{I})$

6: $\boldsymbol{x}_t \leftarrow \sqrt{\overline{\alpha}_t}\boldsymbol{x}_0 + \sqrt{1-\overline{\alpha}_t}\boldsymbol{\epsilon}$

7: $\widehat{\boldsymbol{x}}_0 \leftarrow \widehat{\boldsymbol{x}}_{0,\boldsymbol{\theta}}(\boldsymbol{x}_t, t, \boldsymbol{c})$

8: Update $\boldsymbol{\theta}$ by taking a gradient descent step on

$$\nabla_{\boldsymbol{\theta}}\left[\lambda_t\|\boldsymbol{x}_0 - \widehat{\boldsymbol{x}}_0\|^2 + \frac{c_{\text{phys}}}{2\Sigma_t}\left\|\boldsymbol{\mathcal{R}}_{\text{phys}}(\widehat{\boldsymbol{x}}_0)\right\|^2 + \frac{c_{\text{cond}}}{2\Sigma_t}\|\boldsymbol{\mathcal{R}}_{\text{cond}}(\widehat{\boldsymbol{x}}_0; \boldsymbol{c})\|^2\right]$$

9: **Until** converged

2.2.2. Dual-guided diffusion sampling for physical and conditional constraints

Despite incorporating both physical and conditional constraints during training, the DCDM may exhibit degraded performance when target conditions are underrepresented in the training data. To alleviate this issue, guidance mechanisms are applied during sampling. Classifier-free guidance (CFG) [31], which linearly combines conditional and unconditional model outputs, is commonly used to strengthen conditioning in diffusion models. Although successful in engineering inverse design applications [13,16], CFG-based approaches typically generate only the design field. In contrast, our framework jointly generates both design and solution fields, where the linear combination in CFG does not guarantee preservation of physical consistency. To overcome this limitation, DiffusionPDE [25] extends guided diffusion methods [22] by incorporating physics-based residual gradients into the reverse process.

Building upon this approach, we further incorporate an additional guidance term to promote condition satisfaction along with physical consistency during sampling. This dual-guided update is applied after each reverse diffusion step and is given by

$$\boldsymbol{x}_t = \widetilde{\boldsymbol{x}}_t - \zeta_{\text{phys}}\nabla_{\widetilde{\boldsymbol{x}}_t}\left\|\boldsymbol{\mathcal{R}}_{\text{phys}}\left(\widehat{\boldsymbol{x}}_0(\widetilde{\boldsymbol{x}}_t, t, \boldsymbol{c})\right)\right\|^2 - \zeta_{\text{cond}}\nabla_{\widetilde{\boldsymbol{x}}_t}\|\boldsymbol{\mathcal{R}}_{\text{cond}}(\widehat{\boldsymbol{x}}_0(\widetilde{\boldsymbol{x}}_t, t, \boldsymbol{c}); \mathbf{c})\|^2, \quad (21)$$

where $\tilde{\boldsymbol{x}}_t$ denotes the intermediate sample obtained by updating $\boldsymbol{x}_{t+1}$ through the learned reverse diffusion process in **Eq. (6).** ζ_{phys} and ζ_{cond} are hyperparameters controlling the guidance strengths for the physics and condition residuals, balancing physical consistency and condition satisfaction during sampling. The overall procedure of dual-guided diffusion sampling (DGS) employed in this study is summarized in **Algorithm 2**.

Algorithm 2 Dual-guided diffusion sampling procedure of DCDM

1: **Require**: Trained model $\hat{\boldsymbol{x}}_{0,\boldsymbol{\theta}}(\boldsymbol{x}_t, t, c)$, Target condition $\mathbf{c}$, Guidance weights $\zeta_{\text{phys}}, \zeta_{\text{cond}}$
Variance schedule $\{\Sigma_t\}_{t=1}^{T}$, Diffusion steps T, Physics residual $\boldsymbol{\mathcal{R}}_{\text{phys}}(\cdot)$,
Condition residual $\boldsymbol{\mathcal{R}}_{\text{cond}}(\cdot, \mathbf{c})$,
2: Sample $\boldsymbol{x}_T \sim \mathcal{N}(0, \boldsymbol{I})$
3: **for** $t = T$ to 1 **do**
4: $\quad \mathbf{z} \sim \boldsymbol{\mathcal{N}}(\mathbf{0}, \mathbf{I})$
5: $\quad \tilde{\boldsymbol{x}}_{t-1} \leftarrow \frac{\sqrt{\overline{\alpha}_t}(1-\overline{\alpha}_{t-1})}{1-\overline{\alpha}_t}\, \boldsymbol{x}_t + \frac{\sqrt{\overline{\alpha}_{t-1}}\beta_t}{1-\overline{\alpha}_t}\, \hat{\boldsymbol{x}}_{0,\boldsymbol{\theta}}(\boldsymbol{x}_t, t, c) + \sqrt{\Sigma_t}\mathbf{z}$
6: $\quad \boldsymbol{x}_{t-1} \leftarrow \tilde{\boldsymbol{x}}_{t-1} - \zeta_{\text{phys}}\nabla_{\tilde{\boldsymbol{x}}_{t-1}} \left\| \boldsymbol{\mathcal{R}}_{\text{phys}}\left(\hat{\boldsymbol{x}}_{0,\boldsymbol{\theta}}(\tilde{\boldsymbol{x}}_{t-1}, t-1, \boldsymbol{c})\right) \right\|^2 - \zeta_{\text{cond}}\nabla_{\tilde{\boldsymbol{x}}_{t-1}} \left\| \boldsymbol{\mathcal{R}}_{\text{cond}}\left(\hat{\boldsymbol{x}}_{0,\boldsymbol{\theta}}(\tilde{\boldsymbol{x}}_{t-1}, t-1, \boldsymbol{c}); \mathbf{c}\right) \right\|^2$
7: Return $\boldsymbol{x}_0$

3. Problem settings

3.1. Problem formulation

To evaluate the proposed framework, we consider the inverse design problem of identifying microstructures that achieve a prescribed strain energy under 50% uniaxial extension. The nonlinear mechanical behavior of the microstructure is modeled using a compressible Neo-Hookean hyperelastic material model, and the corresponding strain energy functional W is given by:

$$W(\boldsymbol{F}) = \frac{1}{2}\mu[\boldsymbol{F}:\boldsymbol{F} - 3 - 2\ln(\det \boldsymbol{F})] + \frac{1}{2}\lambda\left[\frac{1}{2}((\det \boldsymbol{F})^2 - 1) - \ln(\det \boldsymbol{F})\right], \tag{22}$$

where $\boldsymbol{F}$ is the deformation gradient and μ, λ are the Lamé parameters, which are related to Young's modulus E and Poisson's ratio ν in Eq. **(23)**.

$$\mu = \frac{E}{2(1+\nu)} \quad \text{and} \quad \lambda = \frac{E\nu}{(1+\nu)(1-2\nu)} \tag{23}$$

Here, the design field $\boldsymbol{d}$ is defined as a spatial distribution of Young's modulus over a 14×14 regular cell-based design domain, while Poisson's ratio is fixed at 0.3 throughout the domain. The solution field $\boldsymbol{a}$ is defined over a 57×57 regular grid-based domain and corresponds to the displacement field obtained from a 50% uniaxial extension test, as shown in **Fig. 3a**. The diffusion model is formulated to jointly represent the Young's modulus field $\boldsymbol{d}$ (design field) and the displacement field $\boldsymbol{a} = (\boldsymbol{u}_1, \boldsymbol{u}_2)$ (solution field), conditioned on the prescribed target strain energy ψ_{targ} (condition $\boldsymbol{c}$). This formulation enables inverse design of hyperelastic microstructures consistent with the prescribed condition. **Table 1** summarizes the key symbols used in the problem formulation and the denoising diffusion framework.

The Mechanical MNIST [26] dataset was used as a controlled benchmark because its reduced and structured design space enables efficient and transparent evaluation of the proposed framework while retaining the essential characteristics of the inverse design problem. This dataset consists of discretized two-dimensional microstructures with a hyperelastic constitutive model. The inverse design problem is formulated to generate microstructures that achieve a prescribed target property defined as the total absorbed strain energy under uniaxial extension. Furthermore, the PDE residual corresponding to the governing equations is formulated in weak form and incorporated into the training objective and the sampling process of the diffusion model, promoting both physical consistency and condition satisfaction in the generated samples.

Table 1. Summary of symbols used in the problem formulation and the denoising diffusion process

Symbol	Description
$\boldsymbol{x}$	Joint state modeled by the diffusion process, defined as $\boldsymbol{x} \coloneqq (\boldsymbol{d}, \boldsymbol{a})$, where $\boldsymbol{d}$ is the design field and $\boldsymbol{a}$ is the solution field.
$\boldsymbol{d}$	Design field; in this problem, the spatially distributed Young's modulus field.
$\boldsymbol{a}$	Solution field; in this problem, the displacement field $(\boldsymbol{u}_1, \boldsymbol{u}_2)$.
$\boldsymbol{c}$	Condition variable representing the prescribed target total strain energy of the microstructure.
ψ_{targ}	Target total strain energy used as the condition $\boldsymbol{c}$.
ψ_{gen}	Total strain energy computed from the generated design and displacement fields $(\boldsymbol{d}, \boldsymbol{a})$.
ψ_{fem}	Total strain energy obtained from finite element analysis using the generated design field $\boldsymbol{d}$.
$\boldsymbol{\mathcal{R}}_{\text{phys}}(\boldsymbol{x})$	Physics residual defined by the governing equilibrium equations and boundary conditions.
$\mathcal{R}_{\text{cond}}(\boldsymbol{x}; \boldsymbol{c})$	Condition residual defined as the discrepancy between ψ_{gen} and the target condition $\boldsymbol{c}$.
$\mathcal{R}_{\text{phys}}^{\text{MAE}}(\boldsymbol{x})$	Mean absolute error of the physics residual.

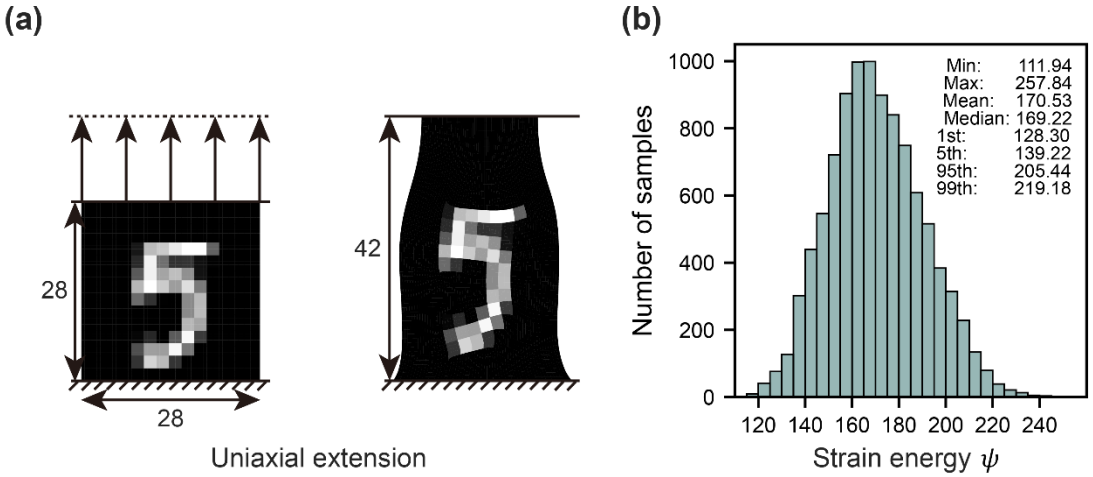

Figure 3. (a) Schematic illustration of the loading and boundary conditions for the MNIST-based microstructure subjected to 50% uniaxial extension. (b) Distribution of the total strain energy ψ of the microstructure across the full dataset of 10,000 microstructures.

3.2. Data generation

The grayscale value of each pixel $\beta \in [0,255]$ is converted into Young's modulus E as

$$E = \frac{\beta}{255.0}(100.0 - 1.0) + 1.0 \tag{24}$$

While retaining the material property mapping rule of Mechanical MNIST, we adjust the design resolution. Specifically, the original 28×28 MNIST images are downsampled to 14×14 using 2×2 average pooling to reduce the design dimensionality. The resulting

14×14 images are then mapped to Young's modulus values for each cell using **Eq. (24)**, yielding a continuous design field with $E \in [1,100]$.

To implement the 50% uniaxial extension test, the bottom boundary of the domain is fixed, while a vertical displacement corresponding to 50% of the total height is applied at the top boundary. For the finite element analysis, the domain is discretized on a 57×57 regular nodal grid, from which the displacement fields $\boldsymbol{u}_1$ and $\boldsymbol{u}_2$ and the total strain energy are computed. In our framework, the diffusion model uses a U-Net [32] backbone, which requires all input fields to share the same spatial resolution when concatenated along the channel dimension. Since the design field is originally defined on a 14×14 cell-based grid, it is upsampled to 56×56 by uniformly replicating each cell into a 4×4 block with identical values. The Young's modulus field is then padded to obtain the final 57×57 resolution, ensuring spatial alignment with the nodal displacement fields. This padding is introduced solely to match the input resolution required by the network and does not contribute to the computation of the physics or condition losses.

The Young's modulus field is linearly normalized to the range $[-1, 1]$ to construct the dataset for the design field $\boldsymbol{d}$. The displacement fields $\boldsymbol{u}_1$ and $\boldsymbol{u}_2$ are independently normalized to the same range using their respective global minimum and maximum values and then combined to form the dataset for the solution field $\boldsymbol{a}$. The total strain energy is normalized to [-1,1] using the physical range [50, 300], and the normalized value is used as the target condition ψ_{targ} in the dataset. The dataset consists of 10,000 samples, which are randomly split into 9,000 training samples and 1,000 validation samples. The distribution of ψ_{targ} in the full dataset and its summary statistics are shown in **Fig. 3b**.

3.3. Enforcing physical consistency and conditioning

In this section, the physics residual $\boldsymbol{\mathcal{R}}_{\mathrm{phys}}(\boldsymbol{x})$ and the condition residual $\boldsymbol{\mathcal{R}}_{\mathrm{cond}}(\boldsymbol{x};\boldsymbol{c})$ are defined to measure physical consistency and the satisfaction of the target condition in the inverse problem considered in this study.

Specifically, we consider a static equilibrium problem governed by the following mechanical equations:

$$\nabla_{\boldsymbol{X}} \cdot \boldsymbol{P}(\boldsymbol{X}) = 0 \text{ in } \Omega \tag{25a}$$

$$\boldsymbol{u}(\boldsymbol{X}) = \overline{\boldsymbol{u}} \text{ on X} \in \partial\Omega_{\mathrm{u}} \tag{25b}$$

$$\boldsymbol{P}(\boldsymbol{X}) \cdot \boldsymbol{n}(\boldsymbol{X}) = \overline{\boldsymbol{t}} \text{ on } \boldsymbol{X} \in \partial\Omega_t, \tag{25c}$$

where $\Omega \subset \mathbb{R}^2$ is the reference domain and its boundary $\partial\Omega$ is partitioned into $\partial\Omega_{\mathrm{u}}$ and $\partial\Omega_t$, on which Dirichlet and Neumann boundary conditions are applied, prescribing the displacement $\overline{\boldsymbol{u}}$ and the traction $\overline{\boldsymbol{t}}$, respectively. The boundary satisfies $\partial\Omega = \partial\Omega_t \cup \partial\Omega_u$ and $\partial\Omega_t \cap \partial\Omega_u = \emptyset$. Here, $\nabla_{\boldsymbol{X}} \cdot$ is the divergence operator in the reference coordinates, $\boldsymbol{P}$ is the first Piola-Kirchhoff stress tensor, $\boldsymbol{F}$ is the deformation gradient, $\boldsymbol{u}$ is the displacement vector, and $\boldsymbol{n}$ is the unit normal vector defined in the reference configuration. Neglecting body forces, the weak form of the linear momentum balance equation can be written as:

$$\int_{\Omega} \boldsymbol{P} : \nabla_X \boldsymbol{v} dV - \int_{\partial\Omega_{\mathrm{t}}} \overline{\boldsymbol{t}} \cdot \boldsymbol{v} dS = 0, \qquad \forall \text{ admissible } \boldsymbol{v} \tag{26}$$

where $\boldsymbol{v}$ is an admissible test function that vanishes on the Dirichlet boundary $\partial\Omega_{\mathrm{u}}$.

After discretizing the reference domain into finite elements, the set of all nodal degrees of freedom is denoted by $\mathcal{D} = \{(a,i): a = 1, \ldots, n_n \,; i = 1,2\}$. The set of degrees of freedom is partitioned into the fixed degrees of freedom $\mathcal{D}^{\mathrm{fix}}$ and the free degrees of freedom $\mathcal{D}^{\mathrm{free}}$ such that $\mathcal{D} = \mathcal{D}^{\mathrm{fix}} \cup \mathcal{D}^{\mathrm{free}}$ and $\mathcal{D}^{\mathrm{fix}} \cap \mathcal{D}^{\mathrm{free}} = \emptyset$. In the Bubnov–Galerkin method, the displacement field $\boldsymbol{u}$ and the test function $\boldsymbol{v}$ are approximated using the same shape functions

N^a, leading to

$$\boldsymbol{u}(\boldsymbol{X}) = \sum_{a=1}^{n_n} N^a(\boldsymbol{X})\, \boldsymbol{u}^a \tag{27}$$

$$\boldsymbol{v}(\boldsymbol{X}) = \sum_{a=1}^{n_n} N^a(\boldsymbol{X})\, \boldsymbol{v}^a\,, \qquad \text{with } v_i^a = 0 \text{ if } (a,i) \in \mathcal{D}^{\text{fix}} \tag{28}$$

The deformation gradient is expressed as

$$\boldsymbol{F}(\boldsymbol{X}) = \boldsymbol{I} + \sum_{a=1}^{n_n} \boldsymbol{u}^a \otimes \nabla_{\boldsymbol{X}} N^a(\boldsymbol{X}), \tag{29}$$

where $\nabla_{\boldsymbol{X}}$ denotes the gradient operator with respect to the reference coordinates.

Substituting **Eq. (28)** into the weak form in **Eq. (26)** reduces to

$$\sum_{a=1}^{n_n} \boldsymbol{v}^a \cdot \left[\int_\Omega \boldsymbol{P}\ \nabla_{\boldsymbol{X}} N^a dV - \int_{\partial\Omega_t} \bar{\boldsymbol{t}}\, N^a dS \right] = 0. \tag{30}$$

Since **Eq. (30)** must hold for all admissible test functions, the residual for the free degrees of freedom is given by **Eq. (31),** with the constitutive relation $\boldsymbol{P} = \partial W / \partial \boldsymbol{F}$. For the fixed degrees of freedom prescribed by Dirichlet boundary conditions, the residual is defined as the displacement mismatch in **Eq. (32)**.

$$\mathcal{R}_{\text{free}}^{a,i} = \int_\Omega \frac{\partial W}{\partial F_{ij}}\, \nabla_j N^a dV - \int_{\partial\Omega_{\text{t}}} \bar{t}_i\, N^a dS, \qquad \forall\, (a,i) \in D^{\text{free}} \tag{31}$$

$$\mathcal{R}_{\text{fix}}^{a,i} = u_i^a - \bar{u}_i(\boldsymbol{X}^a), \qquad \forall\, (a,i) \in D^{\text{fix}} \tag{32}$$

For the present inverse design problem, the physics residual for a generated sample $\boldsymbol{x}$ during denoising is formulated as

$$\boldsymbol{\mathcal{R}}_{\text{phys}}(\boldsymbol{x}) = \begin{bmatrix} \boldsymbol{\mathcal{R}}_{\text{free}}(\boldsymbol{x}) \\ \lambda_{\text{fix}} \boldsymbol{\mathcal{R}}_{\text{fix}}(\boldsymbol{x}) \end{bmatrix}, \tag{33}$$

where $\boldsymbol{\mathcal{R}}_{\text{free}}(\boldsymbol{x}) = \left[\mathcal{R}_{\text{free}}^{a,i}(\boldsymbol{x})\right]_{(a,i)\in D^{\text{free}}}$ and $\boldsymbol{\mathcal{R}}_{\text{fix}}(\boldsymbol{x}) = \left[\mathcal{R}_{\text{fix}}^{a,i}(\boldsymbol{x})\right]_{(a,i)\in D^{\text{fix}}}$ denote the equilibrium residuals at the free degrees of freedom and the displacement constraint residuals

associated with the Dirichlet boundary conditions, respectively. $\lambda_{\text{fix}} > 0$ is a weighting hyperparameter that balances the relative contributions of the two residual components. The prescribed Dirichlet boundary conditions are not hard-enforced through the network architecture or post-processing, but their satisfaction is encouraged through $\boldsymbol{\mathcal{R}}_{\text{fix}}(\boldsymbol{x})$ as a PINN-like soft penalty in the training objective. For each generated sample, the physics residual error is measured using the mean absolute error (MAE) as

$$\mathcal{R}_{\text{phys}}^{\text{MAE}}(\boldsymbol{x}) = \frac{1}{|D^{\text{free}}| + |D^{\text{fix}}|}\left(\sum_{(a,i)\in D^{\text{free}}} \left|\mathcal{R}_{\text{free}}^{a,i}(\boldsymbol{x})\right| + \sum_{(a,i)\in D^{\text{fix}}} \lambda_{\text{fix}} \left|\mathcal{R}_{\text{fix}}^{a,i}(\boldsymbol{x})\right| \right) \tag{34}$$

, where $\left|D^{\text{free}}\right|$ and $\left|D^{\text{fix}}\right|$ denote the number of free and fixed degrees of freedom, respectively.

The condition residual quantifies the difference between the target condition c and the total strain energy $\psi_{\text{gen}}(\boldsymbol{x})$ computed from the generated sample within the finite element formulation, given by

$$\mathcal{R}_{\text{cond}}(\boldsymbol{x}; c) = \psi_{\text{gen}}(\boldsymbol{x}) - c, \quad \psi_{\text{gen}}(\boldsymbol{x}) = \sum_{e=1}^{n_{el}} \int_{\Omega_e} W(\boldsymbol{F})\, dV, \tag{35}$$

where Ω_e denotes the reference domain of element e. For a generated sample $\boldsymbol{x} = (\boldsymbol{d}, \boldsymbol{a})$, the deformation gradient is computed from **Eq. (29),** and the strain energy density W is evaluated using the material parameters from the design field $\boldsymbol{d}$. The total strain energy $\psi_{\text{gen}}(\boldsymbol{x})$ is then obtained by integrating over all elements. Unlike conventional generative inverse design approaches that rely on separate surrogate models [4] or additional finite element analysis [14,16] for post hoc validation, the proposed framework enables the rapid screening of design candidates via direct evaluation of condition satisfaction from the jointly generated design and solution fields.

4. Results and discussion

4.1. Training of diffusion models for microstructure inverse design

In this section, diffusion models were trained to jointly generate the Young's modulus field and the corresponding displacement field, conditioned on a prescribed target strain energy ψ_{targ}. Different training strategies were investigated to assess the effects of incorporating physics loss and condition loss on physical consistency and condition satisfaction. Detailed descriptions of the network architecture, implementation, and hyperparameters are provided in **Appendix A**. Sensitivity analyses of the training settings and key hyperparameters are presented in **Appendix B**.

Three models sharing the same network architecture were compared, differing only in their objective functions: (i) Standard DM, trained using only the data loss; (ii) PIDM, incorporating both data and physics losses; and (iii) DCDM, incorporating data, physics, and condition losses. DGS can be independently applied to each of these trained models at the sampling stage, and the corresponding constraint settings for all compared models are summarized in **Table 2**. Model training and validation were conducted using four different random train–validation splits, each consisting of 9,000 training samples and 1,000 validation samples. The distributions of the total strain energy of the microstructures in the training and validation subsets of each split are shown in **Fig. A1**.

Table 2. Summary of training- and sampling-stage constraints in the compared diffusion models.

Model	Training-stage constraints	Sampling-stage constraints
Standard DM	None	None
PIDM [21]	Physics	None
DCDM	Physics + condition	None
Standard DM+DGS	None	Physics + condition
PIDM+DGS	Physics	Physics + condition
DCDM+DGS	Physics + condition	Physics + condition

The Standard DM was trained for a total of 600,000 iterations using only the data loss. For PIDM and DCDM, the physics loss was not introduced at the beginning of training. During

the early stage of training, generated samples may not sufficiently represent the underlying data distribution, and enforcing the physical constraints prematurely can result in non-physical deformations and numerical instability. To mitigate this issue, both PIDM and DCDM were first pretrained for 300,000 iterations using only the data loss. The physics and condition losses were then introduced, and training continued for an additional 300,000 iterations based on their respective objective functions. All experiments, including both training and sampling, were conducted on a single NVIDIA RTX 5090 GPU.

Figures 4a and 4b present the evolution of the data loss, the physics residual error $\mathcal{R}_{\text{phys}}^{\text{MAE}}(\hat{\boldsymbol{x}}_0)$, and the condition residual error $|\mathcal{R}_{\text{cond}}(\hat{\boldsymbol{x}}_0; c)|$ during training and validation. The condition residual error was evaluated based on the total strain energy normalized to the range $[-1, 1]$. For PIDM and DCDM, the training data loss exhibited a temporary increase at 300,000 iterations, coinciding with the introduction of the physics and condition losses. As training progressed, the data losses of PIDM and DCDM converged to approximately twice that of Standard DM, while their physics residual errors remained approximately one order of magnitude lower. Notably, although PIDM does not explicitly include a condition loss term, it achieved a lower condition residual error than Standard DM. This indicates that improved physical consistency can indirectly enhance agreement between the generated strain energy ψ_{gen} and the target condition. DCDM further reduced the condition residual error while maintaining a physics residual error comparable to that of PIDM.

Table 3 separately reports the equilibrium and Dirichlet BC residuals evaluated on the validation data, providing a more detailed assessment of the physical consistency of the generated solution fields. Compared with Standard DM, DCDM and PIDM substantially reduced both residual components, with only marginal increases in validation data loss. In particular, DCDM maintained equilibrium and Dirichlet BC residuals comparable to those of PIDM while reducing the condition residual to less than half of that achieved by PIDM. These

trends were consistently observed across the four random train–validation splits, indicating limited sensitivity to data partitioning.

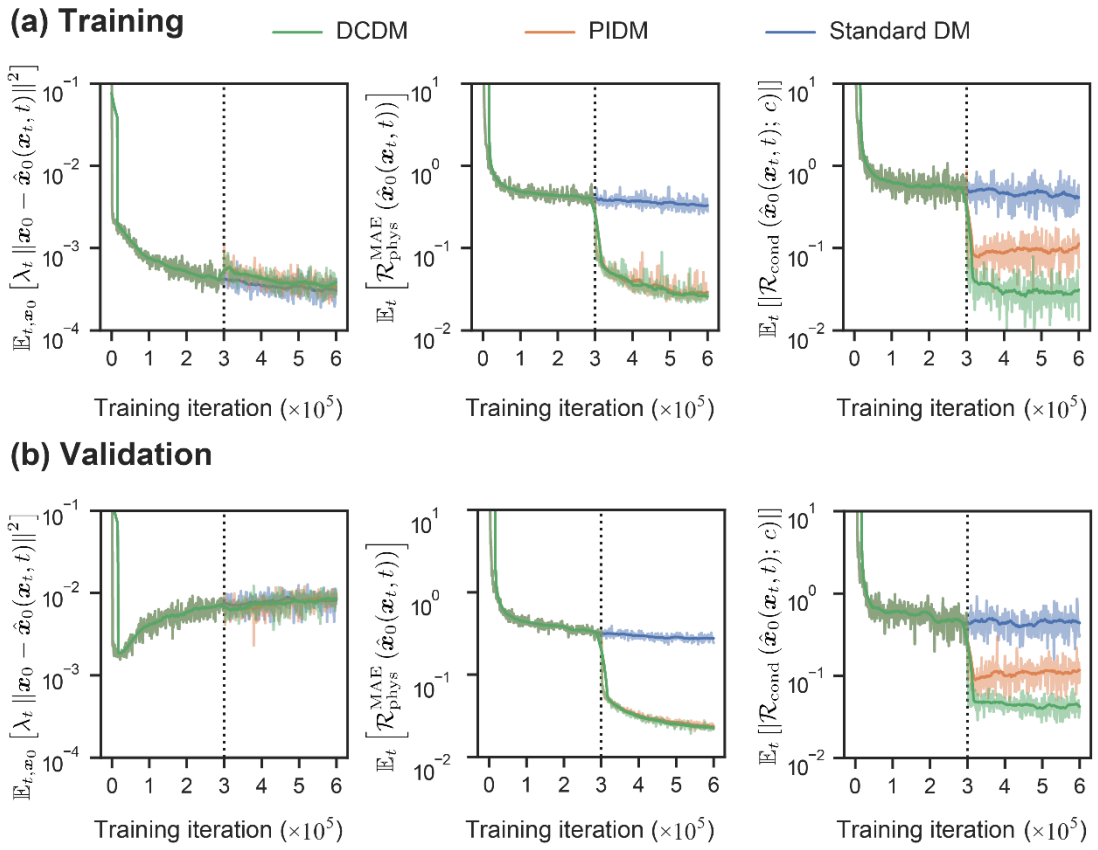

Figure 4. (a) Training and (b) validation histories of the data loss, physics residual error, and condition residual error for Standard DM, PIDM, and DCDM. The vertical dashed line at 3×10^5 iterations marks the introduction of the physics and condition losses. The curves were smoothed using a moving average with a window size of 30,000 iterations.

Table 3. Validation metrics for DCDM, Standard DM, and PIDM. Each model was independently trained and evaluated on four different random 9:1 train–validation splits of the same dataset using the hyperparameters listed in **Table A2**. Values are reported as the mean ± standard deviation over the four splits.

	Data loss	Equilibrium residual	Dirichlet BC residual	Condition residual
	$\|\|\boldsymbol{x}_0 - \hat{\boldsymbol{x}}_0\|\|_2^2$	$\|\|\boldsymbol{\mathcal{R}}_{\text{free}}(\hat{\boldsymbol{x}}_0)\|\|_1 / \|\mathcal{D}_{\text{free}}\|$	$\|\|\boldsymbol{\mathcal{R}}_{\text{fix}}(\hat{\boldsymbol{x}}_0)\|\|_1 / \|\mathcal{D}_{\text{fix}}\|$	$\|\mathcal{R}_{\text{cond}}(\hat{\boldsymbol{x}}_0; c)\|$
DCDM	$(8.98 \pm 2.17) \times 10^{-3}$	$(1.84 \pm 0.07) \times 10^{-2}$	$(8.06 \pm 0.97) \times 10^{-3}$	$(4.68 \pm 1.01) \times 10^{-2}$
Standard DM	$(8.76 \pm 0.84) \times 10^{-3}$	$(2.72 \pm 0.29) \times 10^{-1}$	$(2.74 \pm 0.48) \times 10^{-2}$	$(4.70 \pm 1.16) \times 10^{-1}$
PIDM	$(9.10 \pm 0.73) \times 10^{-3}$	$(1.92 \pm 0.12) \times 10^{-2}$	$(7.33 \pm 0.53) \times 10^{-3}$	$(1.06 \pm 0.09) \times 10^{-1}$

4.2. Inverse design of hyperelastic microstructures using the proposed DCDM+DGS framework

In this section, we evaluate whether the proposed DCDM+DGS framework enables reliable direct evaluation of generated candidates without additional numerical simulations for the hyperelastic inverse design problem, in comparison with Standard DM and PIDM introduced in **Section 4.1**. While the baseline models do not incorporate conditional constraints during training and employ unguided reverse diffusion without sampling-stage constraint enforcement, DCDM+DGS integrates a trained DCDM with DGS, enabling both physical and conditional constraints to be enforced across training and sampling. The key objective of this comparison is to assess whether the generated samples can be reliably evaluated through direct sample-based evaluation, while also examining their ability to satisfy the prescribed target conditions. To this end, we first examine condition satisfaction and physical consistency using quantities directly computed from the generated samples, namely the total strain energy ψ_{gen} and the physics residual error $\mathcal{R}_{\mathrm{phys}}^{\mathrm{MAE}}(\boldsymbol{x}_0)$ in **Section 4.2.1**. We then evaluate the reliability of this sample-based assessment by comparing ψ_{gen} with finite element method (FEM)-computed strain energy ψ_{fem} in **Section 4.2.2**. Detailed quantitative results for all evaluated target conditions are provided in **Tables S1–S48** in **the Supplementary Materials**.

The target conditions were categorized into three groups based on the target strain energy distribution presented in **Fig. 3b**: (i) high-density in-distribution (ID) conditions within the 5^{th}–95^{th} percentile range of the data distribution, (ii) low-density ID conditions within the observed data range but outside the 5^{th}–95^{th} percentile range, and (iii) out-of-distribution (OOD) conditions outside the observed data range. In this section, the high-density ID targets were set to $\psi_{\mathrm{targ}} = 140$, 160, and 200, while the low-density ID targets were set to $\psi_{\mathrm{targ}} = 120$ and 250. Target conditions outside the observed data range, including $\psi_{\mathrm{targ}} = 110$ and 260, were classified as OOD conditions.

4.2.1. Evaluation of condition satisfaction and physical consistency

For each of the three models, 1,000 hyperelastic microstructures were generated under prescribed target strain energies ψ_{targ}. The generated samples were then directly evaluated in terms of the strain energy ψ_{gen} and the physics residual error $\mathcal{R}_{\mathrm{phys}}^{\mathrm{MAE}}(\boldsymbol{x}_0)$.

For the high-density ID conditions, conditional generation was performed at $\psi_{\mathrm{targ}} =$ 140, 160, and 200, with 1,000 samples generated for each target. The distributions of ψ_{gen} and $\mathcal{R}_{\mathrm{phys}}^{\mathrm{MAE}}(\boldsymbol{x}_0)$ for representative condition $\psi_{\mathrm{targ}} = 160$ are shown in **Fig. 5a**. Across all high-density ID targets, DCDM+DGS produced ψ_{gen} distributions most closely concentrated around the prescribed targets while maintaining low physics residual errors. For DCDM+DGS, the relative errors between the mean values of ψ_{gen} and the prescribed targets were 0.14%, 0.13%, and 0.06%, while the corresponding mean values of $\mathcal{R}_{\mathrm{phys}}^{\mathrm{MAE}}(\boldsymbol{x}_0)$ were 0.0087, 0.0096, 0.0131. In contrast, Standard DM showed the weakest performance in both condition satisfaction and physical consistency, with relative errors of about 5–6% and physics residual errors that were approximately one order of magnitude higher than those of DCDM+DGS. PIDM showed improved physical consistency compared with Standard DM, achieving physics residual errors of a similar order of magnitude to those of DCDM+DGS, but its condition satisfaction remained less accurate, with relative errors of about 1–2%.

For the low-density ID conditions, conditional generation was conducted at $\psi_{\mathrm{targ}} =$ 120 and 250, while the OOD target conditions were set to $\psi_{\mathrm{targ}} = 110$ and 260. The distributions of ψ_{gen} and $\mathcal{R}_{\mathrm{phys}}^{\mathrm{MAE}}(\boldsymbol{x}_0)$ for the representative OOD condition $\psi_{\mathrm{targ}} = 260$ are shown in **Fig. 5b**. Under these more challenging conditions, DCDM+DGS consistently maintained both accurate condition satisfaction and low physics residuals across the low-density ID and OOD conditions. For DCDM+DGS, the relative errors of the mean ψ_{gen} values

were 0.30%, 0.28%, −0.28%, and −0.33% for $\psi_{\mathrm{targ}} = 110$, 120, 250, and 260, respectively. The corresponding mean values of $\mathcal{R}_{\mathrm{phys}}^{\mathrm{MAE}}(\boldsymbol{x}_0)$ were 0.0095, 0.0095, 0.0249, and 0.0295. In contrast, Standard DM showed the least reliable performance, with relative errors increasing to about 5–16% with respect to the prescribed targets and physics residual errors remaining more than one order of magnitude larger than those of DCDM+DGS. PIDM showed better overall physical consistency than Standard DM, but still exhibited larger deviations from the prescribed targets and higher physics residual errors than DCDM+DGS, with relative errors of about 2–4% with respect to the prescribed targets.

In summary, DCDM+DGS achieved the most reliable overall performance among the three models by maintaining both low physics residual errors and accurate condition satisfaction. Both DCDM+DGS and PIDM exhibited substantially lower physics residual errors than Standard DM across all target conditions. However, while PIDM achieved low physics residuals, it showed limited accuracy in satisfying the target condition, whereas DCDM+DGS consistently improved condition satisfaction. This difference was relatively small under high-density ID conditions but generally became more pronounced under low-density ID and OOD conditions. Standard DM and PIDM, neither of which incorporated conditional constraints during training, exhibited more limited condition satisfaction than DCDM+DGS across the evaluated target conditions. Notably, although PIDM did not explicitly enforce the target condition, it still showed better condition satisfaction than Standard DM, suggesting that physics-informed training can indirectly facilitate conditional generation. The relationship between low physics residual errors and the agreement between ψ_{gen} and ψ_{fem} obtained from FEM simulations is further examined in **Section 4.2.2**.

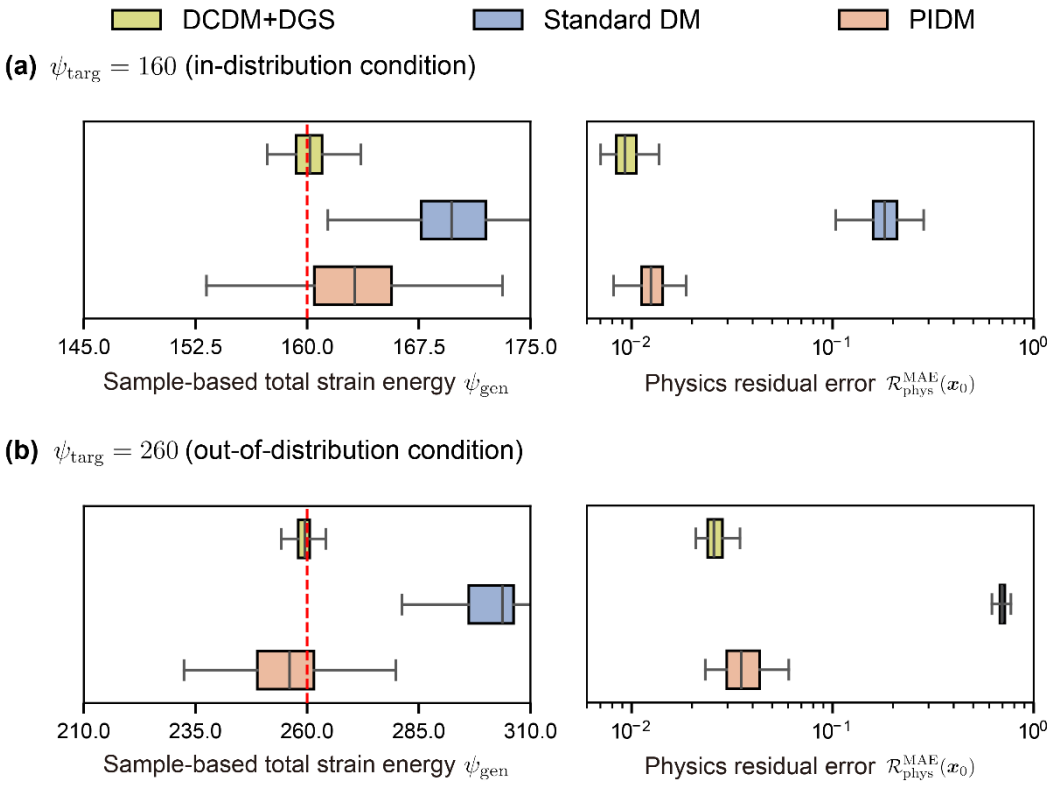

Figure 5. Distributions of the sample-based total strain energy ψ_{gen} (left) and the physics residual error $\mathcal{R}^{\mathrm{MAE}}_{\mathrm{phys}}(\boldsymbol{x}_0)$ (right) for hyperelastic microstructures generated by DCDM+DGS, Standard DM, and PIDM under representative target conditions. (a) High-density ID condition with ψ_{targ} =160. (b) OOD condition with ψ_{targ} =260. For each model, 1,000 samples were generated per target condition. Red dashed lines indicate the prescribed target strain energy ψ_{targ}.

4.2.2. FEM-based validation of generated hyperelastic microstructures

In this section, FEM-based validation was performed for samples generated under prescribed ψ_{targ}. For each generated sample, finite element analysis was conducted using only the design field (Young's modulus field) under the prescribed boundary and loading conditions, and the corresponding ψ_{fem} was obtained. The agreement between the sample-based evaluation and FEM validation was quantified using the relative error $|\psi_{\mathrm{gen}} - \psi_{\mathrm{fem}}|/\psi_{\mathrm{fem}}$. The validation was conducted for the same ID and OOD conditions considered in **Section 4.2.1**, using 100 generated samples for each target condition.

For the high-density ID conditions, FEM-based validation was performed at $\psi_{\mathrm{targ}} =$ 140, 160, and 200. **Figure 6a** presents the distributions of ψ_{fem} and $|\psi_{\mathrm{gen}} - \psi_{\mathrm{fem}}|/\psi_{\mathrm{fem}}$ for

the representative condition $\psi_{\text{targ}} = 160$. Across all high-density ID targets, DCDM+DGS produced FEM-validated strain energy distributions most closely aligned with the prescribed targets while maintaining low relative errors between ψ_{gen} and ψ_{fem}. For DCDM+DGS, the relative errors between the mean values of ψ_{fem} and the prescribed targets were 0.16%, 0.04%, 0.01%, while the corresponding mean relative errors $|\psi_{\text{gen}} - \psi_{\text{fem}}|/\psi_{\text{fem}}$ were 0.13 %, 0.13%, and 0.15%. In contrast, Standard DM showed the weakest performance, with target errors of about 1–3% and relative errors between ψ_{gen} and ψ_{fem} of about 3–4%. PIDM showed better agreement than Standard DM, but its relative errors between ψ_{gen} and ψ_{fem} were approximately twice those of DCDM+DGS, and its FEM-validated strain energies deviated from the prescribed targets by about 0.5–1.5%.

For the low-density ID conditions $\psi_{\text{targ}} = 120$ and 250 and OOD conditions $\psi_{\text{targ}} = 110$ and 260, FEM-based validation was carried out. **Figure 6b** illustrates the distributions of ψ_{fem} and $|\psi_{\text{gen}} - \psi_{\text{fem}}|/\psi_{\text{fem}}$ for the representative OOD condition $\psi_{\text{targ}} = 260$. Under these more challenging conditions, DCDM+DGS showed the most reliable overall performance among the three models. The relative errors between the mean values of ψ_{fem} and the prescribed targets were −0.05%, 0.03%, −0.68%, and −1.25% for $\psi_{\text{targ}} = 110$, 120, 250, and 260, respectively. The corresponding mean relative errors $|\psi_{\text{gen}} - \psi_{\text{fem}}|/\psi_{\text{fem}}$ were 0.44%, 0.27%, 0.51%, and 0.86%, demonstrating close agreement between sample-based evaluation and FEM validation across the low-density ID and OOD conditions. In contrast, Standard DM exhibited the least reliable performance under these conditions. Although the relative errors between the mean values of ψ_{fem} and the prescribed targets remained within about 0–3%, the relative errors between ψ_{gen} and ψ_{fem} were 2.37%, 2.97%, 11.81%, and 16.41%, with substantially larger discrepancies observed at the high-energy targets. PIDM

showed improved overall agreement compared with Standard DM, with relative errors between ψ_{gen} and ψ_{fem} remaining below 1% across all four targets. However, its FEM-validated strain energies deviated from the prescribed targets by approximately 2–3%, indicating less accurate target satisfaction than DCDM+DGS.

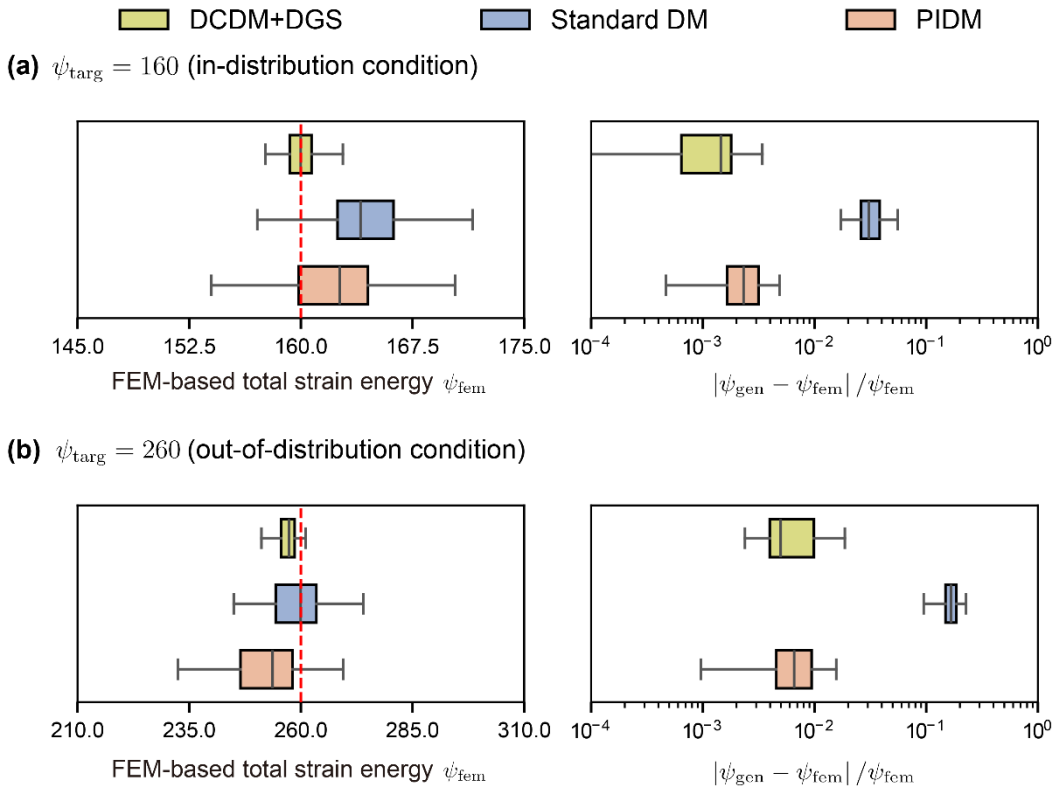

Figure 6. Distributions of FEM-validated total strain energy ψ_{fem} (left) and the relative error $|\psi_{\text{gen}} - \psi_{\text{fem}}|/\psi_{\text{fem}}$ (right) for hyperelastic microstructures generated by DCDM+DGS, Standard DM, and PIDM under representative target conditions. (a) High-density ID condition with ψ_{targ} =160. (b) OOD condition with ψ_{targ} =260. For each model, 100 samples were generated per target condition. Red dashed lines indicate the prescribed target strain energy ψ_{targ}.

To further illustrate the differences among the models, representative samples corresponding to the median value of ψ_{gen} and high-error samples at the 99th percentile of the relative direct-evaluation error were selected for $\psi_{\text{targ}} = 160$ and 260, as shown in **Figs. 7 and A3**. The generated Young's modulus fields, the generated and FEM-computed displacement fields, and the corresponding free-DOF equilibrium residual fields are visualized, enabling

direct comparison of physical consistency and agreement with FEM-based evaluation, as well as the accuracy of target condition satisfaction.

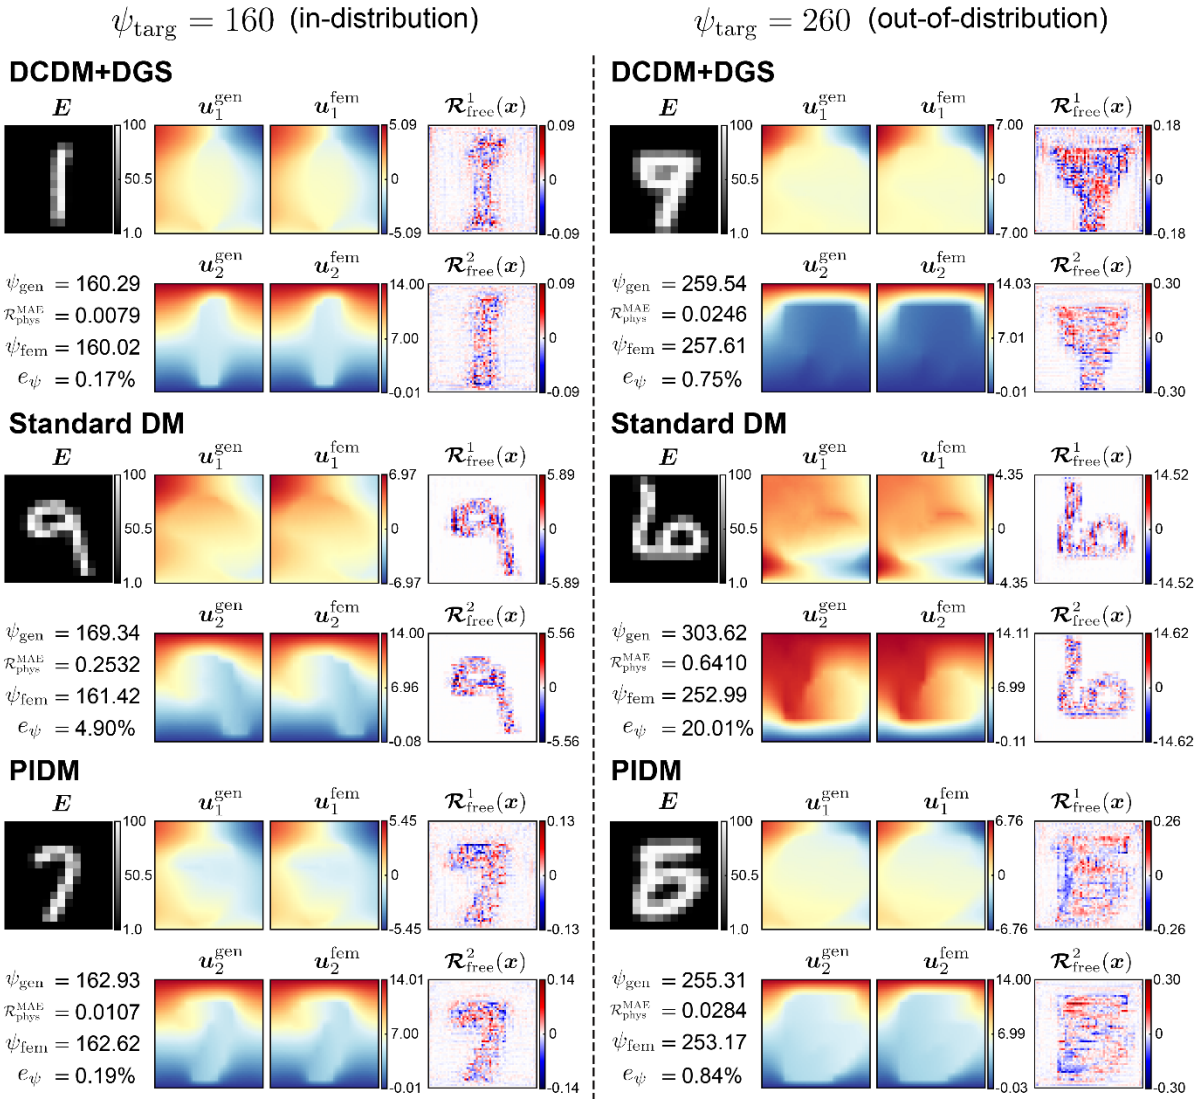

Figure 7. Representative samples corresponding to the median value of ψ_{gen} for DCDM+DGS, Standard DM, and PIDM under the high-density ID condition $\psi_{\text{targ}} = 160$ (left) and the OOD condition $\psi_{\text{targ}} = 260$ (right). For each sample, the generated Young's modulus field $\boldsymbol{E}$ and the generated displacement fields $\boldsymbol{u}_1^{\text{gen}}, \boldsymbol{u}_2^{\text{gen}}$, and the corresponding FEM solutions $\boldsymbol{u}_1^{\text{fem}}, \boldsymbol{u}_2^{\text{fem}}$ are shown. In addition, the free-DOF equilibrium residual fields $\boldsymbol{\mathcal{R}}_{\text{free}}^1(\boldsymbol{x})$ and $\boldsymbol{\mathcal{R}}_{\text{free}}^2(\boldsymbol{x})$ are presented to quantify the degree of physics violation in each displacement component. The error $e_\psi = |\psi_{\text{gen}} - \psi_{\text{fem}}|/\psi_{\text{fem}}$ measures the discrepancy between the sample-based and FEM-based evaluation.

In addition to the target conditions discussed above, **Fig. 8** extends the evaluation of DCDM+DGS to a broader range of $\psi_{\text{targ}} = 100$–300, including more distant extrapolation conditions. The two relative errors represent relative target error, $|\psi_{\text{gen}} - \psi_{\text{targ}}|/\psi_{\text{targ}}$ and the

direct evaluation error, $|\psi_{\text{gen}} - \psi_{\text{fem}}|/\psi_{\text{fem}}$, respectively. A sample was considered valid when both relative errors were within 2%, and n_{valid} denotes the number of such samples. From $\psi_{\text{targ}} = 100$ to 270, at least 84 samples satisfied both criteria, and both mean relative errors remained below approximately 1.1%. For the more distant high-energy OOD targets, the performance gradually deteriorated. At $\psi_{\text{targ}} = 300$, 41 samples remained valid, and both mean errors were approximately 2.4%.

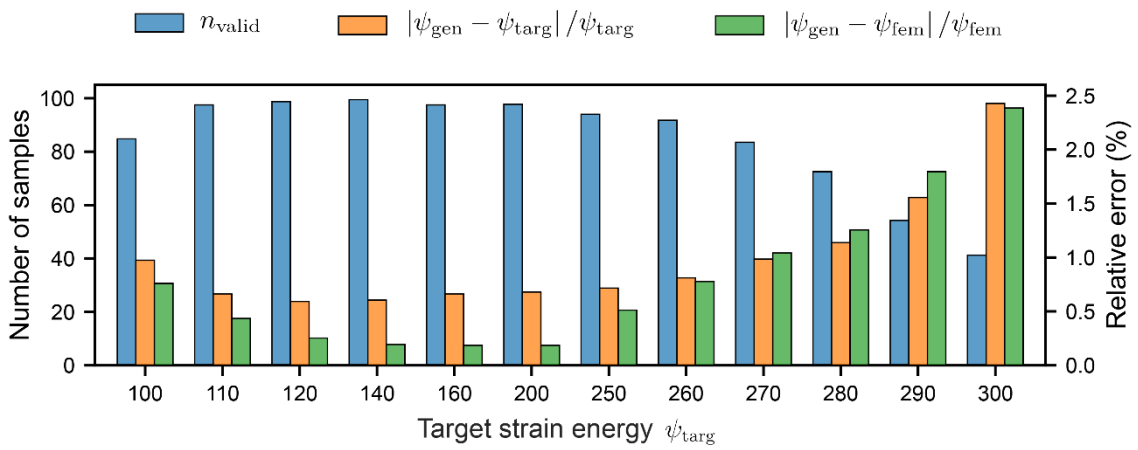

Figure 8. Number of valid samples n_{valid} (left axis) and mean relative target errors and direct evaluation errors (right axis) for DCDM+DGS over the target condition range $\psi_{\text{targ}} = 100$–300. A sample was considered valid when both $|\psi_{\text{gen}} - \psi_{\text{targ}}|/\psi_{\text{targ}}$ and $|\psi_{\text{gen}} - \psi_{\text{fem}}|/\psi_{\text{fem}}$ were within 2%. The relative error values were averaged over 100 generated samples at each target condition.

In summary, DCDM+DGS demonstrates the strongest performance in FEM-based validation by achieving the closest match to the prescribed targets and strong agreement between sample-based and FEM-based evaluation. Across nearly all target conditions, the discrepancy between the FEM-based evaluation and the prescribed target is smaller for DCDM+DGS than for the other two models. This improvement can be attributed to applying DGS, which simultaneously incorporates physics and condition constraints during sampling, to the more physics-consistent generative distribution learned by DCDM, thereby guiding generated samples more effectively toward the prescribed target conditions while maintaining physical consistency, rather than relying solely on implicit conditional generation. In addition,

DCDM+DGS yields substantially smaller differences between direct sample-based evaluation and FEM-based evaluation than Standard DM across all target conditions, while yielding differences comparable to or slightly smaller than those of PIDM. This trend is consistent with the physics residual results in **Section 4.2.1**, supporting that the low residual errors of DCDM+DGS translate into close agreement between direct evaluation and FEM validation.

By achieving both accurate condition satisfaction and strong agreement between direct sample-based evaluation and FEM validation, DCDM+DGS enables efficient and accurate screening of inverse design candidates without relying on additional surrogate models or iterative FEM-based validation. For example, FEM-based validation required approximately 221 s to evaluate 1,000 samples, whereas direct sample-based evaluation required about 0.27 s for the same number of samples. As shown in **Table 4,** using a batch size of 100 over 10 repeated sampling trials, DCDM+DGS required 14.75 ± 0.05 seconds to generate 100 candidates, approximately five times the sampling time required by Standard DM and PIDM. This increase arises because DGS requires physics and condition residual evaluations and the corresponding gradient computations at every reverse-diffusion timestep. Taken together with its strong agreement with FEM validation, these results demonstrate that DCDM+DGS makes direct evaluation a practical and reliable basis for large-scale screening of condition-aligned inverse design candidates. This advantage is expected to become even more significant for problems with stronger nonlinearities or larger numbers of degrees of freedom.

Table 4. Wall-clock sampling time for generating 100 candidates using DCDM, Standard DM, and PIDM with and without DGS under the same batch size, evaluated over 10 repeated sampling trials.

Model	Without DGS	With DGS
DCDM	3.01 ± 0.04 s	14.75 ± 0.05 s
Standard DM	3.01 ± 0.04 s	14.90 ± 0.09 s
PIDM	2.98 ± 0.04 s	14.85 ± 0.09 s

For Standard DM, the error relative to the prescribed target ψ_{targ} was smaller for ψ_{fem} than for ψ_{gen}, and this mismatch became especially large under high-energy OOD conditions. This indicates that the model learned the conditional distribution of the design field to a certain extent. However, when the generated displacement field was used together with the Young's modulus field to compute ψ_{gen}, the error relative to ψ_{targ} increased substantially, reflecting insufficient physical consistency of the generated fields due to the absence of the physics loss during training. These results highlight the limited reliability of sample-based assessment in such regimes and suggest that physically consistent joint generation is essential for accurate and rapid screening of inverse design candidates.

These observations indicate that explicitly incorporating both physics and condition constraints into the diffusion framework is important for generating target-aligned design candidates and for filtering them accurately and efficiently through direct evaluation. **Section 4.3** therefore further investigates this point through an ablation study on the roles of training-stage and sampling-stage constraint enforcement.

4.3. Ablation study on dual-constrained training and dual-guided diffusion sampling

As discussed in **Section 4.2**, explicitly incorporating both physics and condition constraints is important for generating target-aligned candidates while preserving physical consistency. To clarify the respective roles of training-stage and sampling-stage constraint enforcement, this section compares DCDM+DGS with the additional three models employing different constraint-injection strategies. DCDM+DGS incorporates both physics and condition

constraints during training and sampling, whereas DCDM incorporates both constraints only during training. Standard DM+DGS incorporates both constraints only during sampling, while PIDM+DGS incorporates the physical constraint during training and both physics and condition constraints during sampling. Using ψ_{targ} as the prescribed condition, **Section 4.3.1** evaluates the generated samples in terms of ψ_{gen} and $\mathcal{R}_{\mathrm{phys}}^{\mathrm{MAE}}(\boldsymbol{x}_0)$, and **Section 4.3.2** validates the generated samples using FEM by examining the discrepancy between ψ_{gen} and ψ_{fem}. Detailed quantitative results for all evaluated target conditions are provided in **Tables S1–S48** in **the Supplementary Materials**.

4.3.1. Evaluation of condition satisfaction and physical consistency

For each model, 1,000 hyperelastic microstructure samples were generated under each target condition and evaluated in terms of the total strain energy ψ_{gen} and the physics residual error $\mathcal{R}_{\mathrm{phys}}^{\mathrm{MAE}}(\boldsymbol{x}_0)$. The high-density ID conditions considered in this section were $\psi_{\mathrm{targ}} = 140$, 160, and 200, while the OOD conditions were set to $\psi_{\mathrm{targ}} = 100$, 110, 260, and 270. The OOD analysis considered both low- and high-energy extrapolation targets to evaluate the model performance on both sides of the observed data range.

Samples were generated under the high-density ID conditions $\psi_{\mathrm{targ}} = 140$, 160, and 200, and the distributions of ψ_{gen} and $\mathcal{R}_{\mathrm{phys}}^{\mathrm{MAE}}(\boldsymbol{x}_0)$ for the representative condition $\psi_{\mathrm{targ}} = 160$ are presented in **Fig. 9a**. DCDM maintained the relative errors between the mean values of ψ_{gen} and the prescribed targets below 0.5% across all high-density ID conditions, with physics residual errors comparable to those of DCDM+DGS. Standard DM+DGS exhibited relative target errors of 0.1–2.1%, while its physics residual errors remained approximately one order of magnitude larger than those of DCDM+DGS. PIDM+DGS showed relative errors below 0.2% and its physics residual errors remained comparable in magnitude to those of DCDM+DGS. These results indicate that, under high-density ID conditions, DCDM can achieve accurate

condition satisfaction and physical consistency without sampling-stage constraint enforcement, whereas applying sampling-stage constraint enforcement to a model trained without physical constraints is insufficient to achieve comparable physical consistency.

For the OOD conditions $\psi_{\mathrm{targ}} = 100$, 110, 260, and 270, samples were generated at each target condition, while the distributions of ψ_{gen} and $\mathcal{R}_{\mathrm{phys}}^{\mathrm{MAE}}(\boldsymbol{x}_0)$ for the representative condition $\psi_{\mathrm{targ}} = 260$ are presented in **Fig. 9b**. DCDM exhibited mean relative errors of approximately 5.84%, 0.60%, 3.57%, and 4.21% at $\psi_{\mathrm{targ}} = 100$, 110, 260, and 270, respectively, with physics residual errors generally higher than those of DCDM+DGS. Standard DM+DGS showed mean relative errors of approximately 1–3%, and physics residual errors approximately one order of magnitude larger than those of DCDM+DGS under the high-energy OOD conditions. PIDM+DGS achieved slightly lower physics residual errors than DCDM+DGS, whereas DCDM+DGS showed slightly better condition satisfaction. Nevertheless, both methods performed at comparable levels for both metrics. These results indicate that, under OOD conditions, applying DGS to a diffusion model trained with physical constraints is effective for achieving both accurate condition satisfaction and physical consistency.

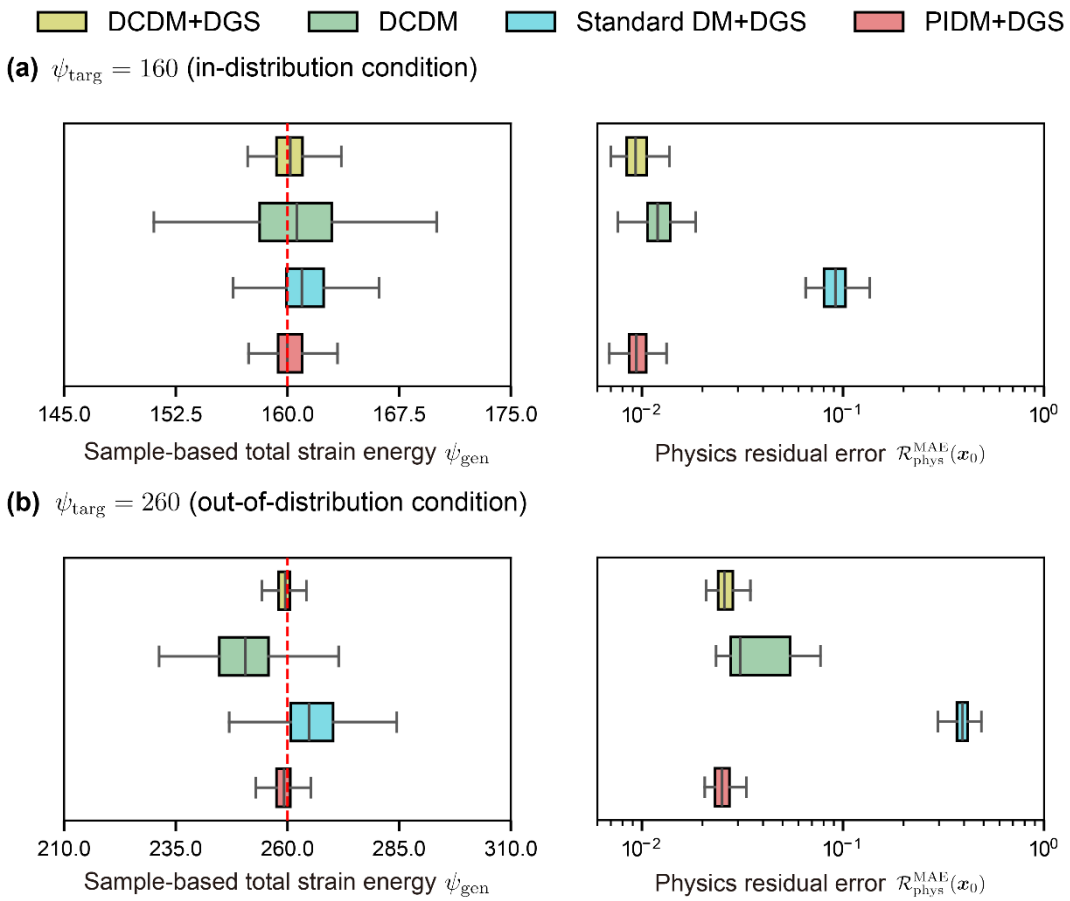

Figure 9. Distributions of the sample-based total strain energy ψ_{gen} (left) and the physics residual error $\mathcal{R}^{\text{MAE}}_{\text{phys}}(\boldsymbol{x}_0)$ (right) for hyperelastic microstructures generated by DCDM+DGS, DCDM, Standard DM+DGS, and PIDM+DGS under representative target conditions. (a) High-density ID condition with ψ_{targ} =160. (b) OOD condition with ψ_{targ} =260. For each model, 1,000 samples were generated per target condition. Red dashed lines indicate the prescribed target strain energy ψ_{targ}.

In summary, these results show that the effectiveness of DGS critically depends on the degree of physical consistency embedded in the generative distribution learned during training. Although Standard DM+DGS generated samples with relatively accurate condition satisfaction, their large physics residual errors limited the physical reliability of the generated solutions. In contrast, PIDM+DGS and DCDM+DGS achieved both accurate condition satisfaction and low physics residual errors, indicating that DGS is effective when applied to a generative distribution shaped by physical constraints during training. Under high-density ID conditions, training-stage physics and condition constraints enable DCDM to achieve strong physical consistency and reasonably accurate condition satisfaction without sampling-stage guidance,

while requiring approximately one-fifth of the sampling time of DCDM+DGS, as reported in **Table 4**. However, under high-energy OOD conditions, the performance of DCDM degrades, as the model tends to generate samples biased toward regions well represented in the training data, leading to systematic deviations from the prescribed targets. This behavior reflects the limited coverage of high-energy conditions in the training data, which prevents the model from adequately learning the corresponding conditional distribution. In such cases, sampling-stage guidance becomes necessary, as it steers the generation process toward the prescribed OOD targets by adapting structural patterns observed near the boundary of the training distribution, thereby improving both condition satisfaction and physical consistency, as illustrated in **Figs. A5 and A6.**

4.3.2. FEM-based validation and diversity analysis of generated hyperelastic microstructures

In this section, FEM-based validation was performed for samples generated by the three additional models under prescribed ψ_{targ}. Following the same evaluation protocol as in **Section 4.2.2**, ψ_{fem} was computed for the generated samples under both ID and OOD conditions, and the discrepancy between ψ_{gen} and ψ_{fem} was evaluated using the relative error, with 100 generated samples for each target condition. In addition, the diversity of samples generated by all six models was analyzed using pairwise L1 distance and structural similarity index measure (SSIM) [33], as described in **Appendix C**.

For each of the high-density ID conditions $\psi_{\mathrm{targ}} = 140$, 160, and 200, 100 samples were generated. **Figure 10a** illustrates the distributions of ψ_{fem} and $\left|\psi_{\mathrm{gen}} - \psi_{\mathrm{fem}}\right|/\psi_{\mathrm{fem}}$ for the representative condition $\psi_{\mathrm{targ}} = 160$. DCDM maintained strong agreement with the prescribed targets across all high-density ID conditions, with relative deviations of the mean ψ_{fem} from ψ_{targ} of 0.61%, 0.38%, −0.48%, respectively.

Its mean relative errors between ψ_{gen} and ψ_{fem} also remained very small at 0.1–0.2%, comparable to those of DCDM+DGS. Standard DM+DGS showed larger deviations from the prescribed targets, with relative deviations of the mean of ψ_{fem} from ψ_{targ} of 0.39%, −1.49%, and −2.89%. Its mean relative errors between ψ_{gen} and ψ_{fem} were also substantially higher at 2.0–3.1%, indicating that sample-based evaluation remained less reliable for filtering design candidates. PIDM+DGS yielded relative deviations of 0.04%, −0.11%, and −0.16% between the mean ψ_{fem} and ψ_{targ}, respectively, while the mean relative errors between ψ_{gen} and ψ_{fem} remained at 0.08–0.12%. Overall, its FEM-based performance was comparable to that of DCDM+DGS.

Table 5 summarizes the number and diversity of valid samples at $\psi_{\mathrm{targ}} = 160$. A generated sample was considered valid only when it simultaneously satisfied the condition satisfaction criterion, $|\psi_{\mathrm{gen}} - \psi_{\mathrm{targ}}|/\psi_{\mathrm{targ}} < \epsilon$, and the direct evaluation criterion, $|\psi_{\mathrm{gen}} - \psi_{\mathrm{fem}}|/\psi_{\mathrm{fem}} < \epsilon$. For the training dataset, samples satisfying $|\psi - \psi_{\mathrm{targ}}|/\psi_{\mathrm{targ}} < \epsilon$ were selected, and diversity was evaluated for these samples and the valid generated samples. Representative valid samples generated by DCDM and DCDM+DGS at $\psi_{\mathrm{targ}} = 160$ and 200 are presented in **Figs. A7 and A8**, respectively, to illustrate the structural diversity of the generated microstructures. At both $\epsilon = 1\%$ and 2%, all models incorporating training- or sampling-stage constraints exhibited lower diversity than the corresponding training-data subset. Across both tolerance levels, DCDM produced approximately 10 more valid samples than PIDM on average owing to the additional condition constraint imposed during training, while maintaining comparable diversity in terms of pairwise L1 distance and mean SSIM. Even under the stricter $\epsilon = 1\%$ criterion, DCDM+DGS and PIDM+DGS produced an average of 79.8 and 75.3 valid samples, respectively, but exhibited lower diversity than their unguided

counterparts, DCDM and PIDM. This indicates that DGS improves condition satisfaction and direct evaluation accuracy at the expense of some diversity.

Table 5. Diversity analysis and number of valid samples at $\psi_{\mathrm{targ}} = 160$. Valid generated samples were defined as those simultaneously satisfying the relative error criteria $|\psi_{\mathrm{gen}} - \psi_{\mathrm{targ}}|/\psi_{\mathrm{targ}} < \epsilon$ and $|\psi_{\mathrm{gen}} - \psi_{\mathrm{fem}}|/\psi_{\mathrm{fem}} < \epsilon$. Training samples were selected using $|\psi - \psi_{\mathrm{targ}}|/\psi_{\mathrm{targ}} < \epsilon$. Values are reported as the mean ± standard deviation over four train–validation splits.

		$\psi_{\mathrm{targ}} = 160$		
		n_{valid}	$d_{\mathrm{L1}}(\uparrow)$	$\overline{\mathrm{SSIM}}$ $(\downarrow)$
$\epsilon = 1\%$	Training dataset	565.8 ± 3.3	0.265 ± 0.001	0.482 ± 0.002
	DCDM	31.5 ± 3.4	0.242 ± 0.008	0.521 ± 0.017
	Standard DM	0	–	–
	PIDM	21.3 ± 2.4	0.231 ± 0.004	0.533 ± 0.016
	DCDM+DGS	79.8 ± 8.8	0.181 ± 0.004	0.619 ± 0.006
	Standard DM+DGS†	5.3 ± 6.1	0.042 ± 0.015	0.880 ± 0.054
	PIDM+DGS	75.3 ± 5.9	0.173 ± 0.011	0.634 ± 0.018
$\epsilon = 2\%$	Training dataset	1147.3 ± 7.8	0.264 ± 0.001	0.484 ± 0.002
	DCDM	53.3 ± 5.9	0.235 ± 0.008	0.531 ± 0.014
	Standard DM	0	–	–
	PIDM	44.3 ± 1.0	0.233 ± 0.006	0.530 ± 0.015
	DCDM+DGS	97.5 ± 1.7	0.182 ± 0.009	0.617 ± 0.014
	Standard DM+DGS	46.3 ± 20.3	0.072 ± 0.017	0.775 ± 0.057
	PIDM+DGS	96.8 ± 1.9	0.175 ± 0.009	0.633 ± 0.013

† For Standard DM+DGS, pairwise metrics were computed using three splits because one split contained no valid samples.

For the OOD targets $\psi_{\mathrm{targ}} = 100$, 110, 260, and 270, 100 samples were generated for each target condition and evaluated using FEM. **Figure 10b** presents the distributions of ψ_{fem} and $|\psi_{\mathrm{gen}} - \psi_{\mathrm{fem}}|/\psi_{\mathrm{fem}}$ for the representative condition $\psi_{\mathrm{targ}} = 260$. DCDM exhibited relatively small discrepancies between ψ_{gen} and ψ_{fem}, with mean relative errors of 0.3–1.8%. However, the relative deviations of the mean ψ_{fem} from ψ_{targ} were −7.39%, 0.40%, −5.10%, and −5.82% for $\psi_{\mathrm{targ}} = 100$, 110, 260, and 270, respectively. By contrast, Standard DM+DGS showed mean relative errors between ψ_{gen} and ψ_{fem} of 1.3–22.6%. Its mean ψ_{fem} deviated from ψ_{targ} by −0.60%, −3.94%, −11.99%, and −14.61%, respectively. This deterioration at the high-energy OOD targets may result from applying condition guidance to a less physically consistent generative distribution, such that improved alignment of ψ_{gen} with ψ_{targ} does not translate into accurate FEM-validated target satisfaction. PIDM+DGS yielded relative deviations of 0.37%, 0.04%, −0.93%, and −1.71% between the mean ψ_{fem} and ψ_{targ},

respectively, with mean relative errors between ψ_{gen} and ψ_{fem} of 0.4–0.9%. Overall, PIDM+DGS showed FEM-validated target satisfaction comparable to that of DCDM+DGS, while achieving only marginally better direct evaluation accuracy.

Table 6 shows the number and diversity of valid samples at $\psi_{\mathrm{targ}} = 260$. At both tolerance levels, only one training sample satisfied the selection criterion, precluding pairwise diversity analysis of the training dataset. Representative valid samples generated by DCDM+DGS at $\psi_{\mathrm{targ}} = 110$ and 260 are presented in **Figs. A5 and A6**, respectively, to illustrate the structural diversity under low- and high-energy OOD conditions. DCDM+DGS and PIDM+DGS produced substantially more valid samples than their unguided counterparts, with mean numbers of valid samples of 63.3 and 65.5 under the $\epsilon = 1\%$ criterion and 91.8 and 96.0 under the $\epsilon = 2\%$ criterion, respectively. Because DCDM and PIDM yielded relatively small valid sets, their diversity estimates should be interpreted with caution, whereas Standard DM+DGS produced no valid samples at either tolerance level. Of the two guided models, DCDM+DGS consistently exhibited slightly higher diversity at both tolerance levels while maintaining a comparable number of valid samples.

Table 6. Diversity analysis and number of valid samples at $\psi_{\mathrm{targ}} = 260$. Valid generated samples were defined as those simultaneously satisfying the relative error criteria $|\psi_{\mathrm{gen}} - \psi_{\mathrm{targ}}|/\psi_{\mathrm{targ}} < \epsilon$ and $|\psi_{\mathrm{gen}} - \psi_{\mathrm{fem}}|/\psi_{\mathrm{fem}} < \epsilon$. Training samples were selected using $|\psi - \psi_{\mathrm{targ}}|/\psi_{\mathrm{targ}} < \epsilon$. Values are reported as the mean ± standard deviation over four train–validation splits.

		$\psi_{\mathrm{targ}} = 260$		
		n_{valid}	$d_{\mathrm{L1}}(\uparrow)$	$\overline{\mathrm{SSIM}}$ $(\downarrow)$
$\epsilon = 1\%$	Training dataset	1	–	–
	DCDM	10.8 ± 5.2	0.254 ± 0.067	0.498 ± 0.160
	Standard DM	0	–	–
	PIDM	12.3 ± 5.6	0.289 ± 0.079	0.509 ± 0.105
	DCDM+DGS	63.3 ± 9.0	0.327 ± 0.029	0.412 ± 0.026
	Standard DM+DGS	0	–	–
	PIDM+DGS	65.5 ± 6.9	0.298 ± 0.008	0.454 ± 0.017
$\epsilon = 2\%$	Training dataset	1	–	–
	DCDM	22.3 ± 6.0	0.319 ± 0.021	0.397 ± 0.029
	Standard DM	0	–	–
	PIDM	28.3 ± 9.9	0.335 ± 0.046	0.416 ± 0.058
	DCDM+DGS	91.8 ± 9.3	0.340 ± 0.014	0.403 ± 0.015
	Standard DM+DGS	0	–	–
	PIDM+DGS	96.0 ± 2.4	0.304 ± 0.018	0.450 ± 0.036

To illustrate the effects of training-stage and sampling-stage constraint enforcement for the models compared in the ablation study, representative samples corresponding to the median value of ψ_{gen} and high-error samples at the 99th percentile of the relative direct-evaluation error were selected for $\psi_{\text{targ}} = 160$ and 260, as shown in **Figs. 11 and A4, respectively**. The generated Young's modulus fields, the generated and FEM-computed displacement fields, and the corresponding free-DOF equilibrium residual fields are visualized, enabling direct comparison of physical consistency, agreement with FEM-based evaluation, and the effectiveness of constraint enforcement.

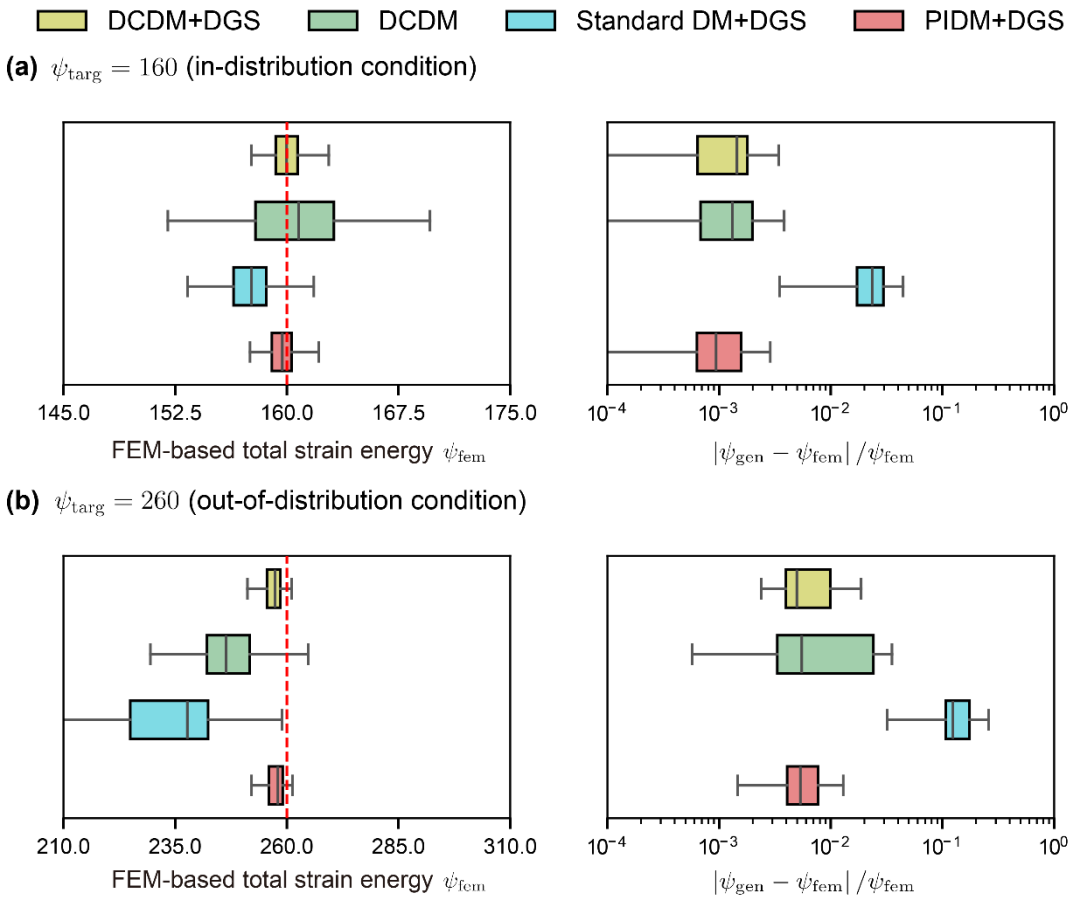

Figure 10. Distributions of FEM-validated total strain energy ψ_{fem} (left) and the relative error $|\psi_{\text{gen}} - \psi_{\text{fem}}|/\psi_{\text{fem}}$ (right) for hyperelastic microstructures generated by DCDM+DGS, DCDM, Standard DM+DGS, and PIDM+DGS under representative target conditions. (a) High-density ID condition with ψ_{targ} =160. (b) OOD condition with ψ_{targ} =260. For each model, 100 samples were generated per target condition. Red dashed lines indicate the prescribed target strain energy ψ_{targ}.

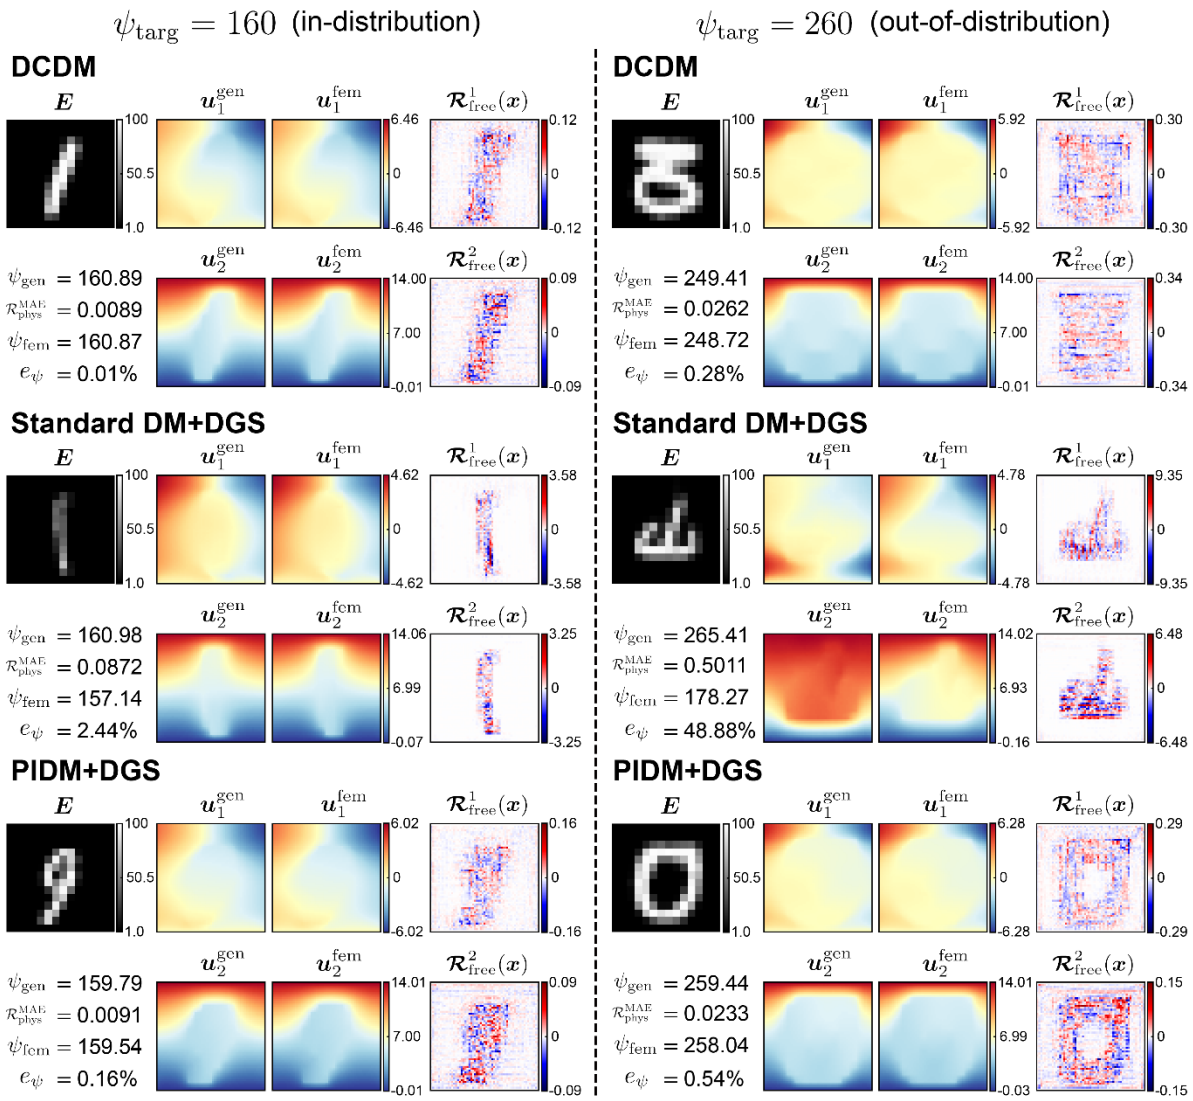

Figure 11. Representative samples corresponding to the median value of ψ_{gen} for DCDM, Standard DM+DGS, and PIDM+DGS under the high-density ID condition $\psi_{\text{targ}} = 160$ (left) and the OOD condition $\psi_{\text{targ}} = 260$ (right). For each sample, the generated Young's modulus field $\boldsymbol{E}$ and the generated displacement fields $\boldsymbol{u}_1^{\text{gen}}, \boldsymbol{u}_2^{\text{gen}}$, and the corresponding FEM solutions $\boldsymbol{u}_1^{\text{fem}}, \boldsymbol{u}_2^{\text{fem}}$ are shown. In addition, the free-DOF equilibrium residual fields $\boldsymbol{\mathcal{R}}_{\text{free}}^1(\boldsymbol{x})$ and $\boldsymbol{\mathcal{R}}_{\text{free}}^2(\boldsymbol{x})$ are presented to quantify the degree of physics violation in each displacement component. The error $e_\psi = |\psi_{\text{gen}} - \psi_{\text{fem}}|/\psi_{\text{fem}}$ measures the discrepancy between the sample-based and FEM-based evaluation.

In summary, under high-density ID conditions, DCDM+DGS and PIDM+DGS generated the largest numbers of valid candidates, whereas DCDM remained a practical unguided alternative by providing reliable target satisfaction and direct evaluation together with greater structural diversity, as shown in **Fig. 12**. Moreover, DCDM required only

approximately one-fifth of the sampling time of DCDM+DGS and PIDM+DGS. The comparison between DCDM and PIDM further showed that the training-stage condition constraint enabled DCDM to achieve better alignment of both ψ_{gen} and ψ_{fem} with ψ_{targ} while maintaining comparable structural diversity, making DCDM the more favorable unguided model under high-density ID conditions. Under OOD conditions, however, training-stage constraints alone maintained relatively good agreement between sample-based and FEM-based evaluations but were insufficient to ensure accurate target satisfaction. Although DGS can be applied to any trained diffusion model, its effectiveness strongly depends on the underlying learned distribution; applying DGS to Standard DM did not provide reliable FEM performance because the model had not learned a physically consistent generative distribution. Reliable OOD generation therefore required DGS to be combined with a physically constrained model, as in DCDM+DGS and PIDM+DGS. These two guided models showed broadly comparable FEM-validated target satisfaction and direct evaluation accuracy. **Figure 12** shows that they also generated comparable numbers of valid samples across most target conditions, except at the low-energy OOD target $\psi_{\mathrm{targ}} = 100$, where DCDM+DGS generated substantially more valid samples. DCDM+DGS exhibited greater structural diversity across the evaluated targets up to $\psi_{\mathrm{targ}} = 270$, whereas PIDM+DGS showed greater diversity at the highest-energy OOD targets.

Overall, these results demonstrate the complementary roles of joint physics and condition enforcement during training and additional constraint guidance during sampling within the DCDM framework. The same trained DCDM can support efficient and diverse unguided generation under ID conditions, while its combination with DGS increases the number of valid samples under ID conditions and enables reliable target satisfaction under OOD conditions, thereby providing a unified and flexible approach for inverse design.

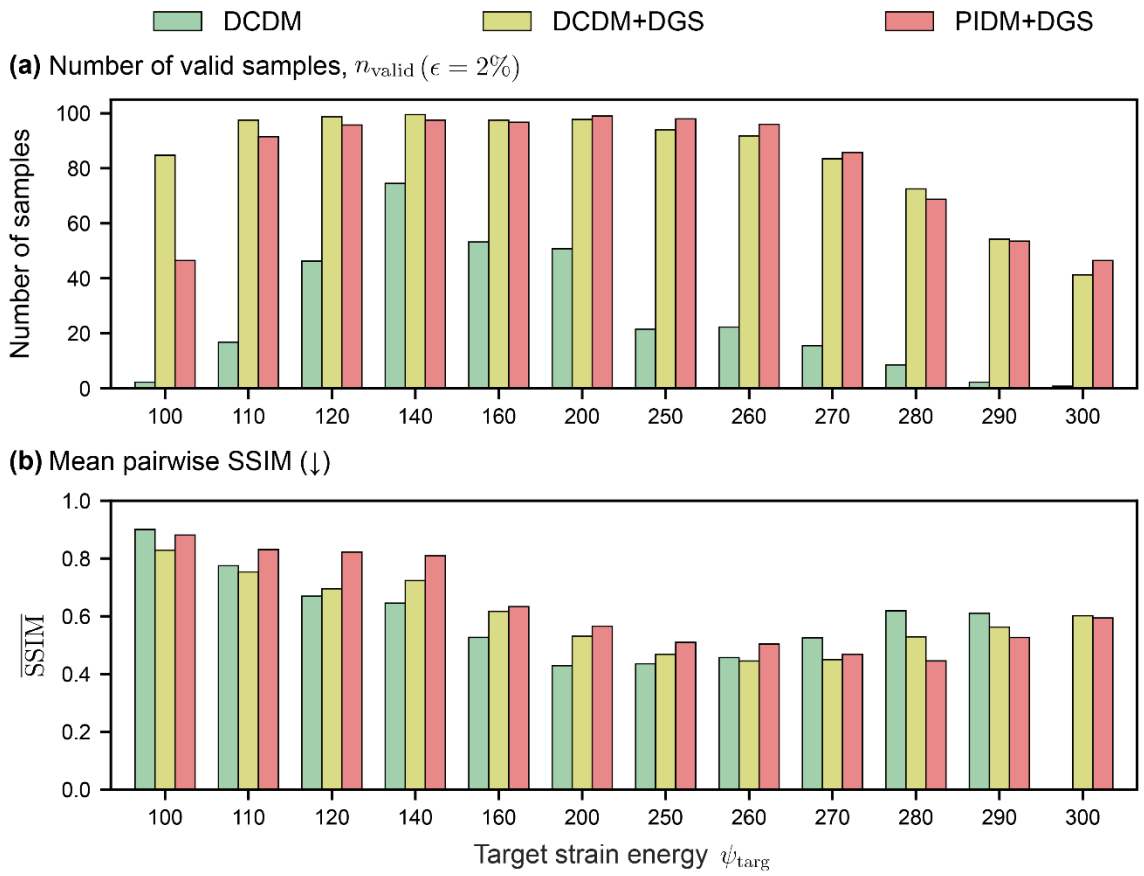

Figure 12. Comparison of the number and structural diversity of valid samples generated by DCDM, DCDM+DGS, and PIDM+DGS across the evaluated target strain energies. (a) Mean number of valid samples, n_{valid}, under the $\epsilon = 2\%$ criterion. (b) Mean pairwise SSIM of the valid samples, where lower values indicate greater structural diversity.

5. Conclusion

This study proposed a dual-constrained diffusion framework that jointly generates design and solution fields by enforcing physical and conditional constraints during training and employing dual-guided diffusion sampling. By integrating these constraints into the generative process, the proposed framework enables direct evaluation of generated candidates without relying on surrogate models or additional numerical simulations. The effectiveness of the framework was demonstrated through the inverse design of hyperelastic microstructures, where the generated samples showed strong agreement between sample-based evaluation and FEM validation. Under high-density in-distribution conditions, activating DGS on the trained DCDM increased the number of candidates jointly satisfying the target condition and FEM validation relative to unguided generation, while preserving reliable direct evaluation. DCDM remained a practical unguided alternative, providing comparable direct-evaluation reliability, greater structural diversity, and requiring only approximately one-fifth of the sampling time. Under out-of-distribution conditions, reliable target satisfaction and FEM-validated performance required DGS to be combined with a model trained under physical constraints. Thus, the same trained DCDM provides a flexible choice between efficient and diverse unguided generation and guided generation with higher validity under in-distribution conditions, while enabling reliable guided generation under out-of-distribution conditions, thereby providing a unified framework for inverse design. The present study is limited to a structured grid setting and a scalar condition, demonstrated on a hyperelastic benchmark problem. Future work will extend the framework to more complex physical systems, richer condition spaces, and unstructured discretizations to broaden its applicability to practical engineering design problems.

Declaration of Competing Interest: The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

CRediT authorship contribution statement:

Jecheon Yu: Conceptualization, Methodology, Formal analysis, Investigation, Validation, Visualization, Data curation, Writing – original draft. **Jinun Byun:** Methodology, Investigation, Validation, Data curation, Writing – review & editing. **Hyeonbin Moon:** Methodology, Investigation, Validation, Writing – review & editing. **Junhyeong Lee**: Methodology, Investigation, Writing – review & editing. **Seunghwa Ryu:** Conceptualization, Supervision, Project administration, Funding acquisition, Writing – review & editing.

Declaration of generative AI and AI-assisted technologies in the writing process:

During the preparation of this work, the authors used OpenAI's ChatGPT to improve the language and readability of the text. The authors reviewed and edited the output as needed and take full responsibility for the content of the published article.

Data Availability Statement: Data will be made available on request.

Acknowledgements: This work was supported by the InnoCORE program (N10250154) of the Ministry of Science and ICT, the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. RS-2025–16070951 and RS-2023–00222166).

Appendix A. Model architecture and training details

We adopt a conditional denoising diffusion architecture with a U-Net [32] backbone to jointly generate the microstructure design field and the corresponding displacement fields. Each sample is represented as a 57×57 tensor with three channels $(\boldsymbol{u}_1, \boldsymbol{u}_2, \boldsymbol{E})$. The target strain energy ψ_{targ} is provided as a scalar condition and mapped through an MLP into the time-embedding space, where it is additively combined with the sinusoidal time embedding to enable conditional denoising.

The network follows a four-level encoder–decoder structure. At each resolution level, two ResNet blocks and an attention module are employed. During downsampling, the spatial resolution decreases as $57 \times 57 \rightarrow 29 \times 29 \rightarrow 15 \times 15 \rightarrow 8 \times 8$, while the latent channel dimension increases as $32 \rightarrow 64 \rightarrow 128 \rightarrow 256$. The attention module [34,35] uses a head dimension of 32 with 8 attention heads. Detailed architectural hyperparameters are summarized in **Table A1.**

The diffusion process is trained with 100 timesteps using a cosine β_t schedule [22]. Optimization is performed with Adam [36] (learning rate 10^{-4}), and an exponential moving average (EMA) with decay 0.99 is applied after 1,000 iterations to stabilize training. The variance scale factors for the physics residual and condition residual $(c_{\mathrm{phys}}, c_{\mathrm{cond}})$, as well as hyperparameters controlling the guidance strengths $(\zeta_{\mathrm{phys}}, \zeta_{\mathrm{cond}})$, are listed in **Table A2**.

Table A1. Denoising diffusion architecture hyperparameters

Hyperparameter	Value
ResNet blocks per down- and upsampling pass	2
ResNet block normalization	Group Normalization
ResNet block activation function	SiLU
Attention block normalization	Layer Normalization
Feature map resolutions (downsampling pass)	$57 \times 57 \rightarrow 29 \times 29 \rightarrow 15 \times 15 \rightarrow 8 \times 8$
Latent dimension	$32 \rightarrow 64 \rightarrow 128 \rightarrow 256$
Attention head dimension	32
Number of attention heads	8

Table A2. Denoising diffusion process and training hyperparameters

Hyperparameter	**Value**
Number of diffusion timesteps	100
β_t-scheduler	Cosine schedule
Batch size	32
Training Iterations	600k
Learning rate	10^{-4}
Optimization algorithm	Adam
Exponential Moving Average (EMA) start iteration	1,000
EMA decay	0.99
Dropout	None
c_{phys}	10^{-3}
c_{cond}	10^{-4}
λ_{fix}	$\sqrt{0.5}$
$(\zeta_{\text{phys}}, \zeta_{\text{cond}})$	Standard DM+DGS: (0.15, 0.05) PIDM+DGS: (0.15, 0.2) DCDM+DGS: (0.15, 0.2)

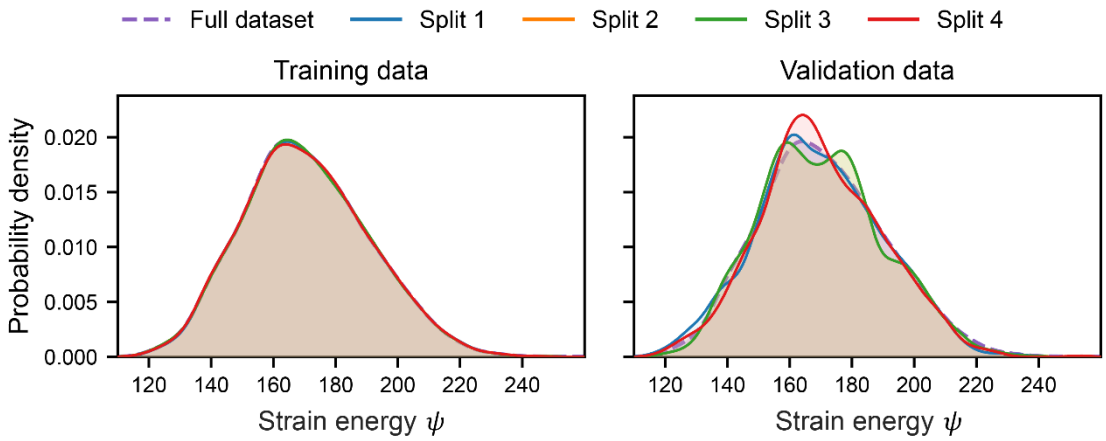

Figure A1. Probability density distributions of the total strain energy ψ for the training and validation subsets from each of four different random 9:1 train–validation splits of the full dataset of 10,000 microstructures. The distribution of the full dataset is shown for reference.

Appendix B. Sensitivity analysis

B.1. Sensitivity analysis of the introduction schedule for the physics and condition residual losses

To assess the sensitivity of the final model performance to the introduction schedule for the physics and condition residual losses, we compared the abrupt-switching schedule used in this study with an alternative linear-ramp schedule. Under the abrupt-switching schedule, the variance scale factors c_{phys} and c_{cond} were set to zero for the first 300,000 iterations and then changed directly to their prescribed values. This abrupt activation of the residual losses

changed the training objective and caused the sudden changes in the data, physics, and condition loss curves observed at approximately 300,000 iterations in **Fig. 4a**. Under the linear-ramp schedule, the applicable variance scale factors were increased linearly from zero to their prescribed values between 300,000 and 350,000 iterations and were kept fixed thereafter. All other training settings were kept unchanged. As summarized in **Table A3**, the abrupt-switching and linear-ramp schedules yielded comparable validation metrics for both models. The differences between the two schedules were generally within the split-to-split variability, indicating that the final model performance is not materially sensitive to the introduction schedule for the physics and condition residual losses.

Table A3. Validation metrics under abrupt-switching and linearly ramped residual loss schedules. Each model was trained four times with different random 9:1 train-validation splits. Values are reported as the mean ± standard deviation.

	Data loss	Equilibrium residual	Dirichlet BC residual	Condition residual
	$\|\boldsymbol{x}_0 - \hat{\boldsymbol{x}}_0\|_2^2$	$\|\boldsymbol{\mathcal{R}}_{\text{free}}(\hat{\boldsymbol{x}}_0)\|_1/\|\mathcal{D}_{\text{free}}\|$	$\|\boldsymbol{\mathcal{R}}_{\text{fix}}(\hat{\boldsymbol{x}}_0)\|_1/\|\mathcal{D}_{\text{fix}}\|$	$\|\mathcal{R}_{\text{cond}}(\hat{\boldsymbol{x}}_0; c)\|$
DCDM (abrupt-switching)	$(8.98 \pm 2.17)\times 10^{-3}$	$(1.84 \pm 0.07)\times 10^{-2}$	$(8.06 \pm 0.97)\times 10^{-3}$	$(4.68 \pm 1.01)\times 10^{-2}$
DCDM (linear ramp)	$(8.98 \pm 1.62)\times 10^{-3}$	$(1.85 \pm 0.08)\times 10^{-2}$	$(8.29 \pm 1.19)\times 10^{-3}$	$(4.28 \pm 0.51)\times 10^{-2}$
PIDM (abrupt-switching)	$(9.10 \pm 0.73)\times 10^{-3}$	$(1.92 \pm 0.12)\times 10^{-2}$	$(7.33 \pm 0.53)\times 10^{-3}$	$(1.06 \pm 0.09)\times 10^{-1}$
PIDM (linear ramp)	$(9.50 \pm 0.70)\times 10^{-3}$	$(1.96 \pm 0.10)\times 10^{-2}$	$(8.11 \pm 0.98)\times 10^{-3}$	$(1.11 \pm 0.09)\times 10^{-1}$

B.2. Sensitivity analysis of training-stage hyperparameters

To examine the sensitivity of model performance to the training-stage hyperparameters, a one-at-a-time sensitivity analysis was conducted. For PIDM, increasing λ_{fix} within the stable range improved the satisfaction of the Dirichlet boundary conditions, while increasing c_{phys} reduced both the equilibrium and Dirichlet BC residuals. However, excessively large values of either parameter resulted in unstable training. Based on these results, $\lambda_{\text{fix}} = \sqrt{0.5}$ and $c_{\text{phys}} = 10^{-3}$ were fixed for DCDM, and c_{cond} was then varied to isolate the effect of the condition residual term. Increasing c_{cond} improved target-condition satisfaction but increased the

Dirichlet BC residual, whereas the equilibrium residual remained nearly unchanged. Accordingly, $c_{\mathrm{cond}} = 10^{-4}$ was selected to balance condition satisfaction and Dirichlet BC satisfaction. The detailed results are summarized in **Tables A4–A6.**

Table A4. Sensitivity to the weighting hyperparameter λ_{fix}, which balances the equilibrium and Dirichlet BC residuals. Validation metrics are reported with fixed $c_{\mathrm{phys}} = 0.001$ and $c_{\mathrm{cond}} = 0$. Each model was trained four times with different random 9:1 train-validation splits. The PIDM baseline corresponds to $(\lambda_{\mathrm{fix}}^2, c_{\mathrm{phys}}, c_{\mathrm{cond}}) = (0.50, 0.001, 0)$.

$(\boldsymbol{\lambda_{\mathrm{fix}}^2}, c_{\mathrm{phys}}, c_{\mathrm{cond}})$	Data loss	Equilibrium residual	Dirichlet BC residual	Condition residual
	$\|\boldsymbol{x}_0 - \widehat{\boldsymbol{x}}_0\|_2^2$	$\|\boldsymbol{\mathcal{R}}_{\mathrm{free}}(\widehat{\boldsymbol{x}}_0)\|_1/\|\mathcal{D}_{\mathrm{free}}\|$	$\|\boldsymbol{\mathcal{R}}_{\mathrm{fix}}(\widehat{\boldsymbol{x}}_0)\|_1/\|\mathcal{D}_{\mathrm{fix}}\|$	$\|\mathcal{R}_{\mathrm{cond}}(\widehat{\boldsymbol{x}}_0; c)\|$
(**1.00**, 0.001, 0)	(diverged)	(diverged)	(diverged)	(diverged)
(**0.50**, 0.001, 0)	$(9.10 \pm 0.73)\times 10^{-3}$	$(1.92 \pm 0.12)\times 10^{-2}$	$(7.33 \pm 0.53)\times 10^{-3}$	$(1.06 \pm 0.09)\times 10^{-1}$
(**0.25**, 0.001, 0)	$(7.27 \pm 1.79)\times 10^{-3}$	$(2.04 \pm 0.06)\times 10^{-2}$	$(9.89 \pm 1.02)\times 10^{-3}$	$(1.52 \pm 0.17)\times 10^{-1}$

Table A5. Sensitivity to the weighting hyperparameter c_{phys}, which controls the relative contribution of the physics residual term to the data loss. Validation metrics are reported with fixed $\lambda_{\mathrm{fix}}^2 = 0.50$ and $c_{\mathrm{cond}} = 0$. Each model was trained four times with different random 9:1 train-validation splits. The PIDM baseline corresponds to $(\lambda_{\mathrm{fix}}^2, c_{\mathrm{phys}}, c_{\mathrm{cond}}) = (0.50, 0.001, 0)$.

$(\lambda_{\mathrm{fix}}^2, \boldsymbol{c_{\mathrm{phys}}}, c_{\mathrm{cond}})$	Data loss	Equilibrium residual	Dirichlet BC residual	Condition residual
	$\|\boldsymbol{x}_0 - \widehat{\boldsymbol{x}}_0\|_2^2$	$\|\boldsymbol{\mathcal{R}}_{\mathrm{free}}(\widehat{\boldsymbol{x}}_0)\|_1/\|\mathcal{D}_{\mathrm{free}}\|$	$\|\boldsymbol{\mathcal{R}}_{\mathrm{fix}}(\widehat{\boldsymbol{x}}_0)\|_1/\|\mathcal{D}_{\mathrm{fix}}\|$	$\|\mathcal{R}_{\mathrm{cond}}(\widehat{\boldsymbol{x}}_0; c)\|$
(0.50, **0.005**, 0)	(diverged)	(diverged)	(diverged)	(diverged)
(0.50, **0.001**, 0)	$(9.10 \pm 0.73)\times 10^{-3}$	$(1.92 \pm 0.12)\times 10^{-2}$	$(7.33 \pm 0.53)\times 10^{-3}$	$(1.06 \pm 0.09)\times 10^{-1}$
(0.50, **0.0002**, 0)	$(1.01 \pm 0.13)\times 10^{-2}$	$(2.68 \pm 0.24)\times 10^{-2}$	$(1.13 \pm 1.04)\times 10^{-2}$	$(1.57 \pm 0.50)\times 10^{-1}$

Table A6. Sensitivity to the weighting hyperparameter c_{cond}, which controls the relative contribution of the condition residual term in the overall training objective. Validation metrics are reported with fixed $\lambda_{\mathrm{fix}}^2 = 0.50$ and $c_{\mathrm{phys}} = 0.001$. Each model was trained four times with different random 9:1 train-validation splits. The DCDM baseline corresponds to $(\lambda_{\mathrm{fix}}^2, c_{\mathrm{phys}}, c_{\mathrm{cond}}) = (0.50, 0.001, 0.0001)$.

$(\lambda_{\mathrm{fix}}^2, c_{\mathrm{phys}}, \boldsymbol{c_{\mathrm{cond}}})$	Data loss	Equilibrium residual	Dirichlet BC residual	Condition residual
	$\|\boldsymbol{x}_0 - \widehat{\boldsymbol{x}}_0\|_2^2$	$\|\boldsymbol{\mathcal{R}}_{\mathrm{free}}(\widehat{\boldsymbol{x}}_0)\|_1/\|\mathcal{D}_{\mathrm{free}}\|$	$\|\boldsymbol{\mathcal{R}}_{\mathrm{fix}}(\widehat{\boldsymbol{x}}_0)\|_1/\|\mathcal{D}_{\mathrm{fix}}\|$	$\|\mathcal{R}_{\mathrm{cond}}(\widehat{\boldsymbol{x}}_0; c)\|$
(0.50, 0.001, **0.001**)	$(6.84 \pm 1.05)\times 10^{-3}$	$(1.83 \pm 0.08)\times 10^{-2}$	$(1.19 \pm 0.20)\times 10^{-2}$	$(3.36 \pm 0.61)\times 10^{-2}$
(0.50, 0.001, **0.0001**)	$(8.98 \pm 2.17)\times 10^{-3}$	$(1.84 \pm 0.07)\times 10^{-2}$	$(8.06 \pm 0.97)\times 10^{-3}$	$(4.68 \pm 1.01)\times 10^{-2}$
(0.50, 0.001, **0.00001**)	$(8.72 \pm 2.16)\times 10^{-3}$	$(1.83 \pm 0.06)\times 10^{-2}$	$(7.39 \pm 1.41)\times 10^{-3}$	$(7.31 \pm 1.71)\times 10^{-2}$

B.3. Sensitivity analysis of sampling-stage hyperparameters

For the sampling-stage hyperparameters, a full grid-search sensitivity analysis was conducted over $\zeta_{\mathrm{phys}}, \zeta_{\mathrm{cond}} \in [0, 0.5]$ with a step size of 0.05, resulting in 121 parameter

combinations for each model. The grid-search space included the single-guidance limiting cases of physics-only guidance $(\zeta_{\text{phys}} > 0, \zeta_{\text{cond}} = 0)$ and condition-only guidance $(\zeta_{\text{phys}} = 0, \zeta_{\text{cond}} > 0)$, in addition to dual-guidance settings with both coefficients nonzero. The Pareto fronts were identified using the mean physics and condition residuals of the generated samples. Model-specific guidance settings were then selected based on the Pareto trade-offs at $\psi_{\text{targ}} = 260$, and their Pareto optimality was additionally confirmed at $\psi_{\text{targ}} = 160$. The selected settings were $(\zeta_{\text{phys}}, \zeta_{\text{cond}}) = (0.15, 0.05)$ for Standard DM+DGS and $(0.15, 0.20)$ for both PIDM+DGS and DCDM+DGS. Standard DM+DGS favored a smaller condition-guidance weight, whereas PIDM+DGS and DCDM+DGS favored a larger condition-guidance weight within the tested range. **Figure A2** presents the grid-search results at $\psi_{\text{targ}} = 260$.

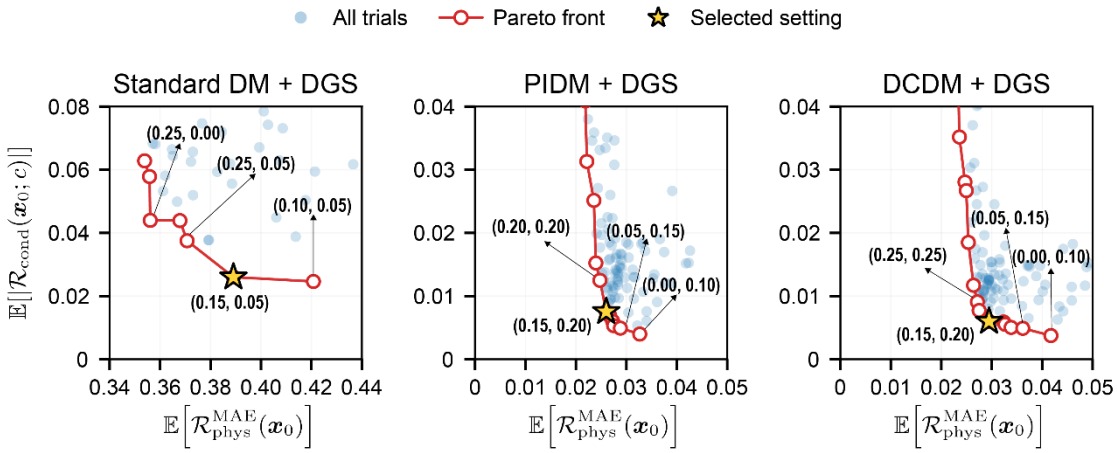

Figure A2. Grid-search sensitivity analysis of the DGS guidance parameters $(\zeta_{\text{phys}}, \zeta_{\text{cond}})$ at $\psi_{\text{targ}} = 260$. The red curves indicate the Pareto fronts, and the stars denote the selected settings. Numbers in parentheses denote $(\zeta_{\text{phys}}, \zeta_{\text{cond}})$.

B.4. Sensitivity analysis of the number of diffusion timesteps

In this study, we used 100 diffusion timesteps, following the default setting adopted in [21]. To investigate the sensitivity of the proposed framework to the number of diffusion timesteps, DCDM, Standard DM, and PIDM were additionally trained using $T = 1{,}000$, while all other training settings and hyperparameters were kept identical to those reported in **Tables**

A1 and A2. The corresponding validation metrics are presented in **Table S49**. For both DCDM and PIDM, increasing the number of diffusion timesteps from 100 to 1000 resulted in a slight increase in data loss but substantially reduced the equilibrium, Dirichlet boundary-condition, and condition residuals, while increasing the wall-clock sampling time by nearly one order of magnitude, as reported in **Table A7**. The corresponding numbers of valid samples and sample-diversity results are reported in **Table A8**. For $\psi_{\mathrm{targ}} = 160$, the number of valid samples generated by unguided PIDM decreased substantially. The diversity of DCDM, DCDM+DGS, and PIDM+DGS also decreased, with the largest reduction observed for PIDM+DGS. For $\psi_{\mathrm{targ}} = 260$, the numbers of valid samples generated by unguided DCDM and PIDM decreased. In contrast, DCDM+DGS and PIDM+DGS maintained similar valid-sample counts but showed reduced diversity, with a larger reduction for PIDM+DGS.

Table A7. Wall-clock sampling time for generating 100 candidates with Standard DM, PIDM, and DCDM, with and without DGS, using $T = 1{,}000$ diffusion steps. All measurements were conducted using the same batch size over 10 repeated sampling trials.

Model	Without DGS	With DGS
DCDM	29.44 ± 0.28 s	151.15 ± 1.51 s
Standard DM	29.69 ± 0.51 s	150.21 ± 0.83 s
PIDM	29.41 ± 0.08 s	150.21 ± 0.87 s

Table A8. Diversity analysis and number of valid samples using $T = 1{,}000$ diffusion steps at $\psi_{\mathrm{targ}} = 160$ and 260. Valid generated samples were defined as those simultaneously satisfying the relative error criteria $|\psi_{\mathrm{gen}} - \psi_{\mathrm{targ}}|/\psi_{\mathrm{targ}} < \epsilon$ and $|\psi_{\mathrm{gen}} - \psi_{\mathrm{fem}}|/\psi_{\mathrm{fem}} < \epsilon$. Training samples were selected using $|\psi - \psi_{\mathrm{targ}}|/\psi_{\mathrm{targ}} < \epsilon$.

		$\psi_{\mathrm{targ}} = 160$			$\psi_{\mathrm{targ}} = 260$		
		n_{valid}	$d_{\mathrm{L1}}(\uparrow)$	$\overline{\mathrm{SSIM}}$ $(\downarrow)$	n_{valid}	$d_{\mathrm{L1}}(\uparrow)$	$\overline{\mathrm{SSIM}}$ $(\downarrow)$
	Training dataset	1141	0.265	0.483	1	–	–
	DCDM	58	0.218	0.551	7	0.349	0.358
	Standard DM	0	–	–	0	–	–
$\epsilon = 2\%$	PIDM	22	0.228	0.524	20	0.282	0.421
	DCDM+DGS	100	0.157	0.661	88	0.251	0.507
	Standard DM+DGS	3	0.056	0.801	0	–	–
	PIDM+DGS	93	0.067	0.822	96	0.144	0.682

Appendix C. Diversity metrics

Pairwise L1 distance and SSIM [33] were used to quantify the structural diversity among the n_{valid} valid samples. Here, n_{valid} denotes the number of generated samples satisfying both $|\psi_{\text{gen}} - \psi_{\text{targ}}|/\psi_{\text{targ}} < \epsilon$ and $|\psi_{\text{gen}} - \psi_{\text{fem}}|/\psi_{\text{fem}} < \epsilon$. The diversity analysis was performed on the Young's modulus fields $\mathbf{d}$, which were normalized to the range $[-1, 1]$. The pairwise L1 distance in **Eq. (36)** quantifies the average pixel-wise difference between Young's modulus fields over all valid sample pairs. Here, H and W denote the height and width of the Young's modulus field, respectively. The pairwise SSIM in **Eqs. (37) and (38)** quantifies the average structural similarity between Young's modulus fields over all valid sample pairs. In **Eq. (38)**, x and y denote corresponding local windows of two Young's modulus fields; $\mu_x$and μ_y are the local means, σ_x^2 and σ_y^2 are the local variances, and σ_{xy} is the local covariance. The stabilization constants are defined as $C_1 = (K_1 L)^2$ and $C_2 = (K_2 L)^2$, where $L = 2$ is the data range of the normalized Young's modulus fields. SSIM was calculated using a 7 × 7 uniform sliding window with K_1 = 0.01 and K_2 = 0.03.

$$d_{\text{L1}} = \frac{2}{n_{\text{valid}}(n_{\text{valid}} - 1)} \sum_{1 \le i < j \le n_{\text{valid}}} \frac{1}{HW} \left\| \mathbf{d}_i - \mathbf{d}_j \right\|_1 \tag{36}$$

$$\overline{\text{SSIM}} = \frac{2}{n_{\text{valid}}(n_{\text{valid}} - 1)} \sum_{1 \le i < j \le n_{\text{valid}}} \text{SSIM}(\mathbf{d}_i, \mathbf{d}_j) \tag{37}$$

$$\text{SSIM}(x, y) = \frac{(2\mu_x\mu_y + C_1)(2\sigma_{xy} + C_2)}{(\mu_x^2 + \mu_y^2 + C_1)(\sigma_x^2 + \sigma_y^2 + C_2)} \tag{38}$$

Appendix D. Additional examples of generated microstructures and FEM-based validation

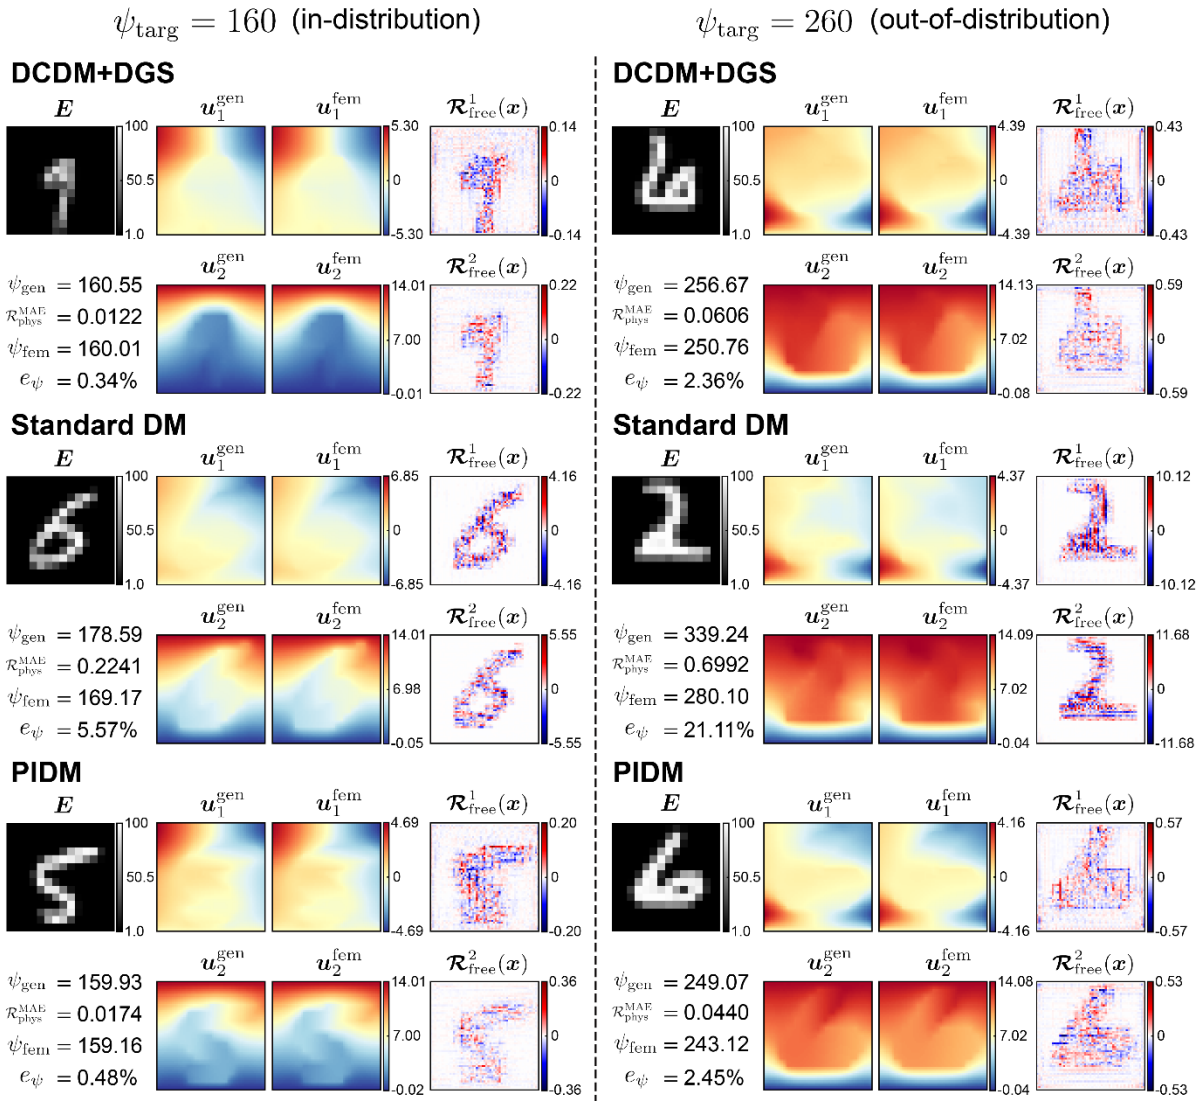

Figure A3. Representative high-error samples at the 99th percentile of the relative direct-evaluation error $e_\psi = |\psi_{\text{gen}} - \psi_{\text{fem}}|/\psi_{\text{fem}}$ for DCDM+DGS, Standard DM, and PIDM under the high-density ID condition ψ_{targ} = 160 (left) and the OOD condition ψ_{targ} = 260 (right). For each sample, the generated Young's modulus field $\boldsymbol{E}$ and the generated displacement fields $\boldsymbol{u}_1^{\text{gen}}, \boldsymbol{u}_2^{\text{gen}}$, and the corresponding FEM solutions $\boldsymbol{u}_1^{\text{fem}}, \boldsymbol{u}_2^{\text{fem}}$ are shown. In addition, the free-DOF equilibrium residual fields $\boldsymbol{\mathcal{R}}_{\text{free}}^1(\boldsymbol{x})$ and $\boldsymbol{\mathcal{R}}_{\text{free}}^2(\boldsymbol{x})$ are presented to quantify the degree of physics violation in each displacement component.

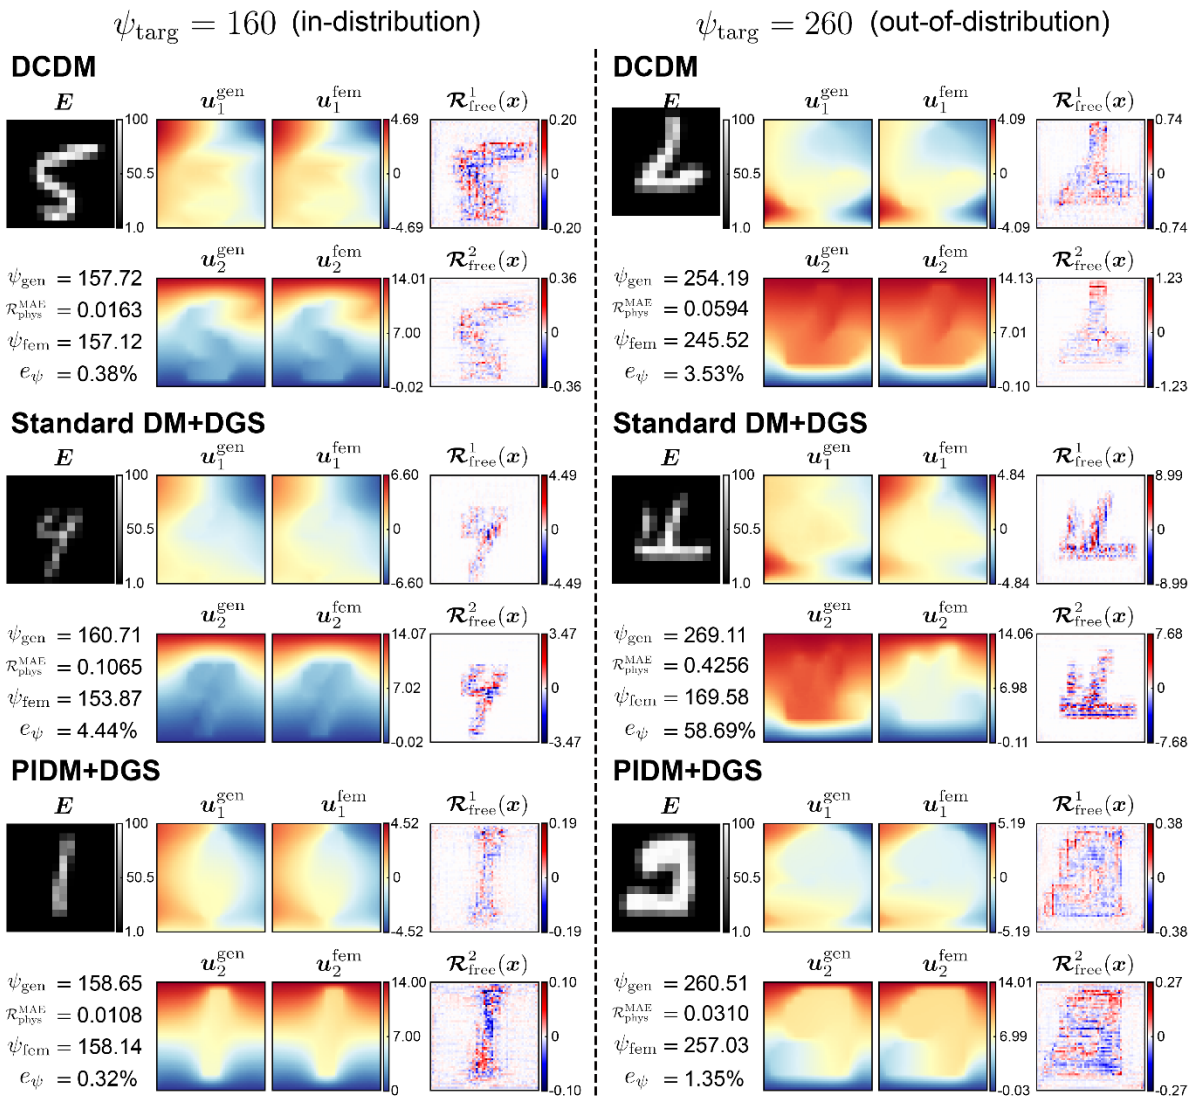

Figure A4. Representative high-error samples at the 99th percentile of the relative direct-evaluation error $e_\psi = |\psi_{\text{gen}} - \psi_{\text{fem}}|/\psi_{\text{fem}}$ for DCDM, Standard DM+DGS, and PIDM+DGS under the high-density ID condition $\psi_{\text{targ}} = 160$ (left) and the OOD condition $\psi_{\text{targ}} = 260$ (right). For each sample, the generated Young's modulus field $\boldsymbol{E}$ and the generated displacement fields $\boldsymbol{u}_1^{\text{gen}}, \boldsymbol{u}_2^{\text{gen}}$, and the corresponding FEM solutions $\boldsymbol{u}_1^{\text{fem}}, \boldsymbol{u}_2^{\text{fem}}$ are shown. In addition, the free-DOF equilibrium residual fields $\boldsymbol{\mathcal{R}}_{\text{free}}^1(\boldsymbol{x})$ and $\boldsymbol{\mathcal{R}}_{\text{free}}^2(\boldsymbol{x})$ are presented to quantify the degree of physics violation in each displacement component.

Appendix E. Diversity of microstructures generated by DCDM and DCDM+DGS

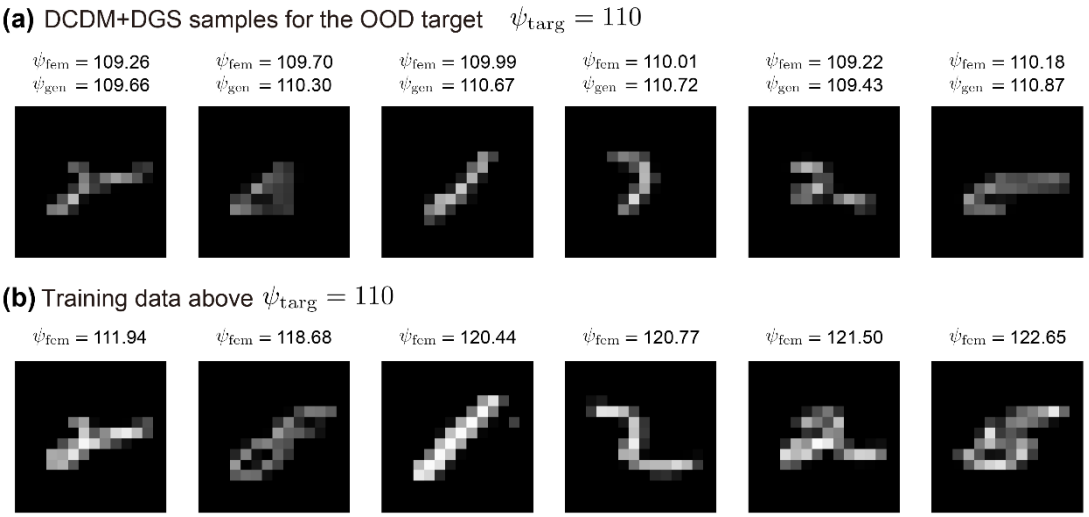

Figure A5. Comparison of generated and training samples near the low-energy OOD target $\psi_{\text{targ}} = 110$. **(a)** Selected samples generated by DCDM+DGS, together with their directly evaluated strain energies ψ_{gen} and corresponding FEM results ψ_{fem}. **(b)** Morphologically similar training samples selected from the low-energy end of the training distribution.

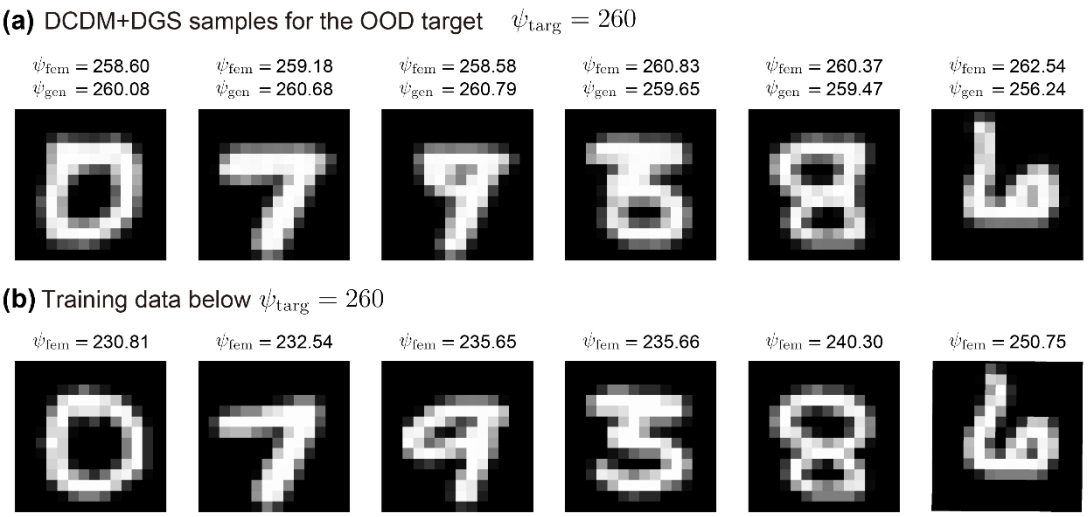

Figure A6. Comparison of generated and training samples near the high-energy OOD target $\psi_{\text{targ}} = 260$. **(a)** Selected samples generated by DCDM+DGS, together with their directly evaluated strain energies ψ_{gen} and corresponding FEM results ψ_{fem}. **(b)** Morphologically similar training samples selected from the high-energy end of the training distribution.

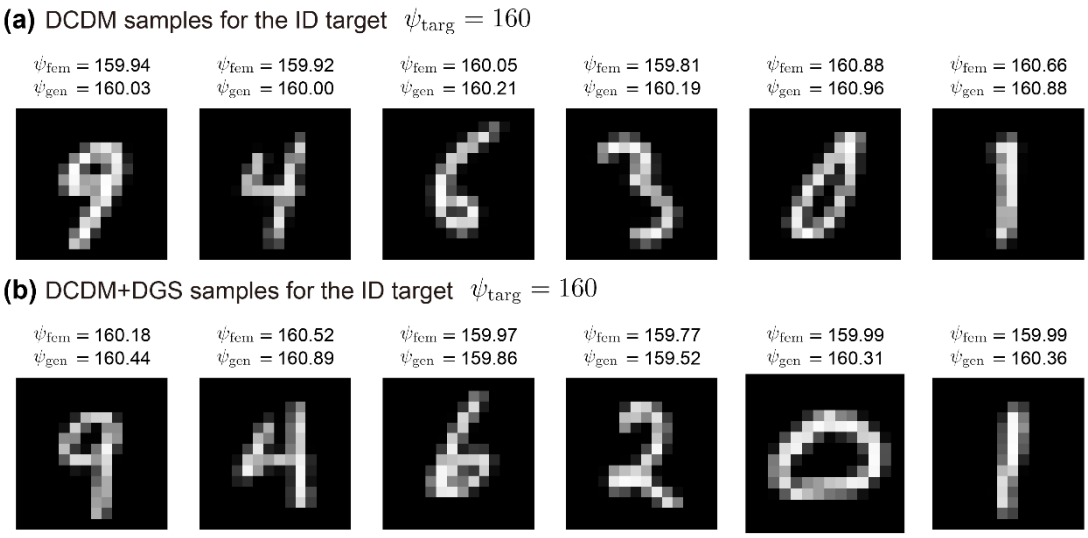

Figure A7. Representative valid samples generated by (a) DCDM and (b) DCDM+DGS at the high-density ID target ψ_{targ} = 160.

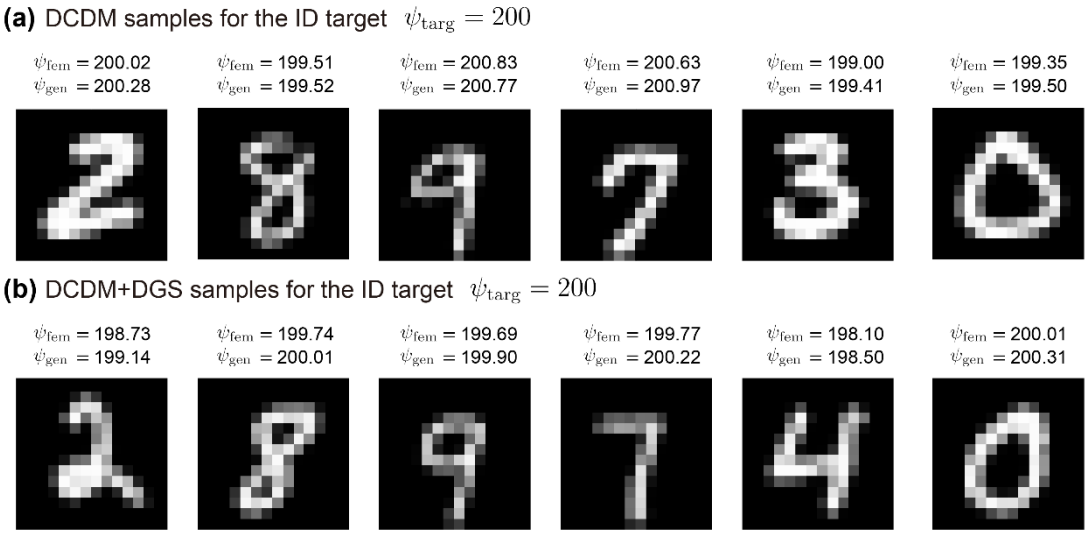

Figure A8. Representative valid samples generated by (a) DCDM and (b) DCDM+DGS at the high-density ID target ψ_{targ} = 200.

REFERENCE

[1] S. Kumar, S. Tan, L. Zheng, D.M. Kochmann, Inverse-designed spinodoid metamaterials, NPJ Comput. Mater. 6 (2020) 73-. https://doi.org/10.1038/s41524-020-0341-6.

[2] U.G.K. Wegst, H. Bai, E. Saiz, A.P. Tomsia, R.O. Ritchie, Bioinspired structural

materials, Nature Materials 2014 14:1 14 (2014) 23–36. https://doi.org/10.1038/nmat4089.

[3] C.S. Ha, D. Yao, Z. Xu, C. Liu, H. Liu, D. Elkins, M. Kile, V. Deshpande, Z. Kong, M. Bauchy, X. (Rayne) Zheng, Rapid inverse design of metamaterials based on prescribed mechanical behavior through machine learning, Nature Communications 2023 14:1 14 (2023) 5765-. https://doi.org/10.1038/s41467-023-40854-1.

[4] N.N. Vlassis, W.C. Sun, Denoising diffusion algorithm for inverse design of microstructures with fine-tuned nonlinear material properties, Comput. Methods Appl. Mech. Eng. 413 (2023) 116126. https://doi.org/10.1016/j.cma.2023.116126.

[5] J. Lee, D. Park, M. Lee, H. Lee, K. Park, I. Lee, S. Ryu, Machine learning-based inverse design methods considering data characteristics and design space size in materials design and manufacturing: a review, Mater. Horiz. 10 (2023) 5436–5456. https://doi.org/10.1039/d3mh00039g.

[6] D. Park, J. Lee, H. Lee, G.X. Gu, S. Ryu, Deep generative spatiotemporal learning for integrating fracture mechanics in composite materials: inverse design, discovery, and optimization, Mater. Horiz. 11 (2024) 3048–3065. https://doi.org/10.1039/d4mh00337c.

[7] B. Deng, A. Zareei, X. Ding, J.C. Weaver, C.H. Rycroft, K. Bertoldi, Inverse Design of Mechanical Metamaterials with Target Nonlinear Response via a Neural Accelerated Evolution Strategy, Advanced Materials 34 (2022) 2206238. https://doi.org/10.1002/adma.202206238.

[8] D.P. Kingma, M. Welling, Auto-Encoding Variational Bayes, 2nd International Conference on Learning Representations, ICLR 2014 - Conference Track Proceedings (2022). http://arxiv.org/abs/1312.6114.

[9] I.J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, Generative Adversarial Nets, Adv. Neural Inf. Process. Syst. 27 (2014). http://www.github.com/goodfeli/adversarial.

[10] L. Regenwetter, A.H. Nobari, F. Ahmed, Deep Generative Models in Engineering Design: A Review, Journal of Mechanical Design 144 (2022). https://doi.org/10.1115/1.4053859.

[11] Y. Song, J. Sohl-Dickstein, D.P. Kingma, A. Kumar, S. Ermon, B. Poole, Score-Based Generative Modeling through Stochastic Differential Equations, ICLR 2021 - 9th International Conference on Learning Representations (2021). http://arxiv.org/abs/2011.13456.

[12] J. Ho, A. Jain, P. Abbeel, Denoising Diffusion Probabilistic Models, Adv. Neural Inf. Process. Syst. 33 (2020) 6840–6851.

[13] J.H. Bastek, D.M. Kochmann, Inverse design of nonlinear mechanical metamaterials via video denoising diffusion models, Nat. Mach. Intell. 5 (2023) 1466–1475. https://doi.org/10.1038/s42256-023-00762-x.

[14] G. Li, D. Liu, Inverse design of composite materials with desired mechanical behaviors

based on conditional diffusion model, Eng. Struct. 345 (2025) 121417. https://doi.org/10.1016/j.engstruct.2025.121417.

[15] Y. Zhang, T. Long, H. Zhang, Stable diffusion for the inverse design of microstructures, (2024). https://arxiv.org/pdf/2409.19133.

[16] S. Deng, Q. Zhang, Z. Cai, Generative method for aerodynamic optimization based on classifier-free guided denoising diffusion probabilistic model, Structural and Multidisciplinary Optimization 2025 68:8 68 (2025) 165-. https://doi.org/10.1007/s00158-025-04111-x.

[17] M. Raissi, P. Perdikaris, G.E. Karniadakis, Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, J. Comput. Phys. 378 (2019) 686–707. https://doi.org/10.1016/j.jcp.2018.10.045.

[18] K. Ciftci, K. Hackl, A physics-informed GAN framework based on model-free data-driven computational mechanics, Comput. Methods Appl. Mech. Eng. 424 (2024) 116907. https://doi.org/10.1016/j.cma.2024.116907.

[19] W. Zhong, H. Meidani, PI-VAE: Physics-Informed Variational Auto-Encoder for stochastic differential equations, Comput. Methods Appl. Mech. Eng. 403 (2023) 115664. https://doi.org/10.1016/j.cma.2022.115664.

[20] D. Shu, Z. Li, A. Barati Farimani, A physics-informed diffusion model for high-fidelity flow field reconstruction, J. Comput. Phys. 478 (2023) 111972. https://doi.org/10.1016/j.jcp.2023.111972.

[21] J.-H. Bastek, W. Sun, D.M. Kochmann, Physics-Informed Diffusion Models, 13th International Conference on Learning Representations, ICLR 2025 (2025) 82766–82791. http://arxiv.org/abs/2403.14404.

[22] P. Dhariwal, A. Nichol, Diffusion Models Beat GANs on Image Synthesis, Adv. Neural Inf. Process. Syst. 11 (2021) 8780–8794. http://arxiv.org/abs/2105.05233.

[23] H. Chung, J. Kim, M.T. Mccann, M.L. Klasky, J.C. Ye, Diffusion Posterior Sampling for General Noisy Inverse Problems, 11th International Conference on Learning Representations, ICLR 2023 (2024). http://arxiv.org/abs/2209.14687.

[24] C. Jacobsen, Y. Zhuang, K. Duraisamy, COCOGEN: PHYSICALLY CONSISTENT AND CONDITIONED SCORE-BASED GENERATIVE MODELS FOR FORWARD AND INVERSE PROBLEMS, SIAM Journal on Scientific Computing 47 (2025) C399–C425. https://doi.org/10.1137/24M1636071.

[25] J. Huang, G. Yang, Z. Wang, J.J. Park, DiffusionPDE: Generative PDE-Solving under Partial Observation, Adv. Neural Inf. Process. Syst. 37 (2024) 130291–130323. https://doi.org/10.52202/079017-4140.

[26] E. Lejeune, Mechanical MNIST: A benchmark dataset for mechanical metamodels, Extreme Mech. Lett. 36 (2020) 100659. https://doi.org/10.1016/j.eml.2020.100659.

[27] T. Hang, S. Gu, C. Li, J. Bao, D. Chen, H. Hu, X. Geng, B. Guo, Efficient Diffusion

Training via Min-SNR Weighting Strategy, Proceedings of the IEEE International Conference on Computer Vision (2024) 7407–7417. http://arxiv.org/abs/2303.09556.

[28] M. Rixner, P.S. Koutsourelakis, A probabilistic generative model for semi-supervised training of coarse-grained surrogates and enforcing physical constraints through virtual observables, J. Comput. Phys. 434 (2021) 110218. https://doi.org/10.1016/J.JCP.2021.110218.

[29] X. Gao, M. Sitharam, A.E. Roitberg, Bounds on the Jensen Gap, and Implications for Mean-Concentrated Distributions, Australian Journal of Mathematical Analysis and Applications 16 (2020) 1–16. http://arxiv.org/abs/1712.05267.

[30] Y. Zhang, D. Zou, Physics-Informed Distillation of Diffusion Models for PDE-Constrained Generation, (2025). http://arxiv.org/abs/2505.22391.

[31] J. Ho, T. Salimans, Classifier-Free Diffusion Guidance, (2022). http://arxiv.org/abs/2207.12598.

[32] O. Ronneberger, P. Fischer, T. Brox, U-Net: Convolutional Networks for Biomedical Image Segmentation, Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) 9351 (2015) 234–241. https://doi.org/10.1007/978-3-319-24574-4_28.

[33] Z. Wang, A.C. Bovik, H.R. Sheikh, E.P. Simoncelli, Image quality assessment: From error visibility to structural similarity, IEEE Transactions on Image Processing 13 (2004) 600–612. https://doi.org/10.1109/TIP.2003.819861.

[34] A. Vaswani, G. Brain, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is All you Need, Adv. Neural Inf. Process. Syst. 30 (2017).

[35] A. Katharopoulos, A. Vyas, N. Pappas, F. Fleuret, Transformers are RNNs: Fast Autoregressive Transformers with Linear Attention, (2020) 5156–5165. https://proceedings.mlr.press/v119/katharopoulos20a.html.

[36] D.P. Kingma, J.L. Ba, Adam: A Method for Stochastic Optimization, 3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings (2014). https://arxiv.org/pdf/1412.6980.