

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 GUIDED DIFFUSION BY OPTIMIZED LOSS FUNCTIONS ON RELAXED PARAMETERS FOR INVERSE MATERIAL DESIGN

Anonymous authors

Paper under double-blind review

ABSTRACT

Inverse design problems are common in engineering and materials science. The forward direction, i.e., computing output quantities from design parameters, typically requires running a numerical simulation, such as a FEM, as an intermediate step, which is often an optimization problem by itself. In many scenarios, several design parameters can lead to the same or similar output values. For such cases, multi-modal probabilistic approaches are advantageous to obtain diverse solutions. Additional difficulties arise if the design problem is constrained. We propose a novel inverse design method based on diffusion models. The model learns a prior over possible approximate designs in a relaxed parameter space. Parameters are sampled using guided diffusion for which we leverage implicit differentiation of the simulation to evaluate the loss function. A design sample is obtained by backprojecting the sampled parameters. We develop our approach for a composite material design problem where the forward process is modeled as a linear FEM problem. We evaluate with the objective of finding designs that match a specified bulk modulus. We demonstrate that our method can propose diverse designs within 1% relative error margin from medium to high target bulk moduli in 2D and 3D settings. We also demonstrate that the material density of generated samples can be minimized simultaneously by using a multi-objective loss function.

1 INTRODUCTION

Guided diffusion has recently been demonstrated as a promising approach for synthesizing solutions for complex inverse design problems such as in engineering (Bastek & Kochmann, 2023) or computational biology (Vignac et al., 2023). The diffusion model can be trained to generate samples from the prior distribution of possible solutions. The reverse diffusion process is typically guided towards solutions that satisfy specific design criteria by gradients of loss functions (Chung et al., 2023; Song et al., 2023).

In this paper, we propose a novel approach for inverse design by guided diffusion. We consider problems in which the loss function requires to solve an inner optimization problem such as finite element methods (FEM) in material science. Instead of learning surrogate models of the design measures, our novel approach uses implicit differentiation of a simulation for loss guidance of the denoising process. By reformulating the original parameter space into a grid of continuous variables, e.g. pixels or voxels, the diffusion model can learn a prior over a relaxed parameter space and the loss can be differentiated for the relaxed parameters. The sampled parameters are backprojected into the original design space. We demonstrate our approach for a material design problem which infers properties of composite materials.

We consider the design of composite materials whose microstructure is inhomogeneous due to the presence of spherical particles. To make numerical simulations feasible, the complex microstructure is represented by a smaller sample called a representative volume element (RVE). At the macroscale, this RVE is treated as if it were homogeneous, a concept known as homogenization. We consider isotropic, linear elastic properties for both matrix (the base or filling phase) and spherical included particles, which allows to compute the bulk modulus from the stress response (Huet, 1990) by solving a linear FEM problem. The design space consists of the choice of base materials for matrix and

054 particles as well as the particle volume fraction and radius. Base materials exist as discrete instances
 055 and the particle volume fraction and radius are not easily differentiable due to the implicit constraint of
 056 integer numbers of particles. We relax the design problem by allowing arbitrary material properties in
 057 each element of the discretized microstructure, resulting in a pixel (2D) or voxel (3D) representation.
 058 We train a diffusion model on a dataset of plausible microstructures and then use it in conjunction
 059 with gradients computed by a FEM solver on the relaxed problem to optimize an objective function,
 060 while simultaneously staying on the manifold of plausible microstructures. Importantly, the diffusion
 061 model is independent of the objective function and can be reused for other targets, as long as we
 062 can compute the respective gradients. We evaluate our approach for 2D and 3D design problems for
 063 simulated materials where the goal is to achieve a prescribed macroscopic bulk modulus. Our results
 064 demonstrate that our approach generates diverse material samples within 1% relative error margin
 065 from medium to high target bulk moduli.

066 In summary, we contribute the following: (1) We propose a novel approach for guided diffusion with
 067 optimized loss functions. (2) We develop and evaluate our approach for a material design problem.
 068 We propose to relax its parameter space to enable loss-guided diffusion. The diffusion model acts
 069 as a prior over approximate designs. (3) Our results demonstrate that our approach can propose
 070 materials that diversely and closely achieve a wide range of medium to high target bulk moduli. **We**
 071 **also demonstrate that material density can be minimized simultaneously by a multi-objective loss.**

072 2 RELATED WORK

073 **Optimal Experimental Design.** Optimal Experimental Design methods Franceschini & Macchietto
 074 (2008) such as Bayesian optimization (BO (Foster et al., 2019; Frazier & Wang, 2016)) can be used to
 075 search for feasible design parameters for inverse design problems. Trials are conducted sequentially
 076 at promising parameters according to measures like information gain. BO uses the outcomes to
 077 update a regression function of the target value of an objective function. However, the regressor
 078 typically needs to be trained for each specific target value. Our diffusion model-based approach
 079 learns a prior over feasible designs in a relaxed space of the partially discrete design parameters. It
 080 generates diverse samples which are guided at test-time to achieve the design objective in a zero-shot
 081 manner.

082 **Guided Diffusion.** Several methods have been proposed that use diffusion models as priors and guide
 083 the reverse diffusion process by additional constraints. Diffusion Posterior Sampling (DPS, (Chung
 084 et al., 2023)) adds a Gaussian measurement to perform posterior inference. The diffusion is guided
 085 by the derivative of the squared residual between measurement and its expected value given the
 086 denoised state. The works in (Yu et al., 2023; Song et al., 2023) generalize this concept by energy-
 087 and loss-guided diffusion, respectively. In our approach, determining the expected measurement
 088 or loss requires solving an inner optimization problem. Some methods add proximal optimization
 089 steps to the reverse diffusion process (Song et al., 2022; Chung et al., 2024). Universal Guidance
 090 Diffusion (Bansal et al., 2023) proposes to learn a classifier function which is used to guide the
 091 diffusion process. In (Ye et al., 2024), several guidance approaches are unified in a single formulation.
 092 However, the above approaches do not consider inner optimization problems like our approach.

093 **Inverse Material Design.** Previous work on the design of microstructures in the context of homoge-
 094 nization has primarily relied on evolutionary algorithms, such as Genetic Algorithms (GA) (Zohdi,
 095 2003), or on deep reinforcement learning (DRL) approaches (Würz & Weißenfels, 2025). GAs have
 096 proven effective in tackling even non-convex optimization problems, like the case presented here.
 097 However, they are highly sensitive to initial conditions and often converge to local optima. Moreover,
 098 they struggle with handling complex and continuous state representations. DRL-based methods, on
 099 the other hand, explore the design space by trial-and-error to identify microstructures that exhibit
 100 desired effective material properties. **If the inverse design problem is directly differentiable for the**
 101 **parameters, gradient-based optimization can be employed (Xue et al., 2023). However, the design**
 102 **problem considered in this work has discrete parameters and is non-differentiable.**

103 Diffusion approaches have been applied for related problems. In metamaterial design, the aim is
 104 to optimize small-scale structures which have a single material either added or removed at each
 105 location and give rise to specific macroscopic properties. The approach in (Bastek & Kochmann,
 106 2023) uses classifier-free guided diffusion to infer metamaterial shapes represented as 2D Gaussian
 107 random field. In a recent preprint (Liu et al., 2025), signed distance functions are instead used to

model the metamaterial. Differently to our training-free method, the diffusion model needs to be trained on conditionals. Similar to our method, the preprint (Yang et al., 2024) proposes to use loss-guided diffusion for inverse design of metamaterials, but using a learned regressor of material properties. These works do not incorporate an FEM solver for guidance like our approach. In the preprint (Zampini et al., 2025), constraint projection is proposed for guided diffusion for metamaterial design. While the method uses FEM to determine the stress-strain curve, differently to our method, the solver is differentiated numerically using Monte-Carlo estimates for the guidance.

3 PRELIMINARIES

3.1 DENOISING DIFFUSION PROBABILISTIC MODELS

In recent years, denoising diffusion probabilistic models (DDPM, (Ho et al., 2020)) have become a popular tool for image generation and inverse design. DDPM models a forward process that iteratively diffuses a data sample \mathbf{x}_0 in a latent variable \mathbf{x}_t in successive time steps t . The forward diffusion process is modelled by a Gaussian distribution $q(\mathbf{x}_t | \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t; \sqrt{\bar{\alpha}_t} \mathbf{x}_0, (1 - \bar{\alpha}_t) \mathbf{I})$ conditional on the data sample, where $\bar{\alpha}_t := \prod_{s=1}^t \alpha_s$, $\alpha_t = 1 - \beta_t$, and β_t are parameters of the variance schedule. We can write $\mathbf{x}_t(\mathbf{x}_0, \epsilon) = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon$ for $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ in terms of \mathbf{x}_0 and $\epsilon = \epsilon(\mathbf{x}_t, \mathbf{x}_0)$. For $T \rightarrow \infty$, the latent variable tends to $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Ho et al. (2020) showed that this diffusion process can be reverted by iteratively sampling a latent variable for the previous time step, where a neural network with parameters θ predicts the noise $\epsilon_\theta(\mathbf{x}_t, t)$. It is trained on the forward diffusion noise $\epsilon(\mathbf{x}_t, \mathbf{x}_0)$. Song et al. (2020) further showed that this reverse process can be generalized, allowing an arbitrary number of sampling steps N , where then each $i \in \{0, N - 1\}$ corresponds to a certain timestep t from a subset of the original $T - 1$ to 0 timesteps. Note that $\epsilon_\theta(\mathbf{x}_t, t)$ is closely related to the score of $p_t(\mathbf{x}_t) = \int q(\mathbf{x}_t | \mathbf{x}_0) p(\mathbf{x}_0) d\mathbf{x}_0$ and it holds $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t) \approx -\frac{1}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(\mathbf{x}_t, t)$ (Luo, 2022).

There exist equivalent formulations for the noise ϵ , such as the “predicted clean sample” $\hat{\mathbf{x}}_0 = \frac{1}{\sqrt{\bar{\alpha}_t}} \mathbf{x}_t - \sqrt{1 - \bar{\alpha}_t} \epsilon$. An alternative is the velocity parameterization proposed in (Salimans & Ho, 2022): $\mathbf{v} = \sqrt{\bar{\alpha}_t} \epsilon - \sqrt{1 - \bar{\alpha}_t} \hat{\mathbf{x}}_0$. All of the three quantities can be used as targets for the neural network. During denoising, they can be converted from one to another as shown. We train our network on the \mathbf{v} formulation. For the further derivation, we use $\hat{\mathbf{x}}_0$.

We apply DDIM (Song et al., 2020) to reduce the number of iterations in the reverse diffusion process. A sample from the previous DDIM distribution \mathbf{x}_{i-1} can be obtained as

$$\mathbf{x}_{i-1} = \frac{\sqrt{\bar{\alpha}_i}(1 - \bar{\alpha}_{i-1})}{1 - \bar{\alpha}_i} \mathbf{x}_i + \frac{\sqrt{\bar{\alpha}_{i-1}} \tilde{\beta}_i}{1 - \bar{\alpha}_i} \hat{\mathbf{x}}_0 + \sigma_i \mathbf{z} \quad (1)$$

where $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, $\tilde{\alpha}_i = \bar{\alpha}_i / \bar{\alpha}_{i-1}$, $\tilde{\beta}_i = 1 - \tilde{\alpha}_i$ and $\sigma_i = \sqrt{(1 - \bar{\alpha}_{i-1}) / (1 - \bar{\alpha}_i) \tilde{\beta}_i}$, and i is the DDIM iteration.

3.2 LOSS-GUIDED DIFFUSION

Using Bayes’ rule, diffusion posterior sampling (DPS (Chung et al., 2023)) combines the prior learned by the diffusion model with likelihoods of additional measurements \mathbf{y} , $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t | \mathbf{y}) = \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t) + \nabla_{\mathbf{x}_t} \log p_t(\mathbf{y} | \mathbf{x}_t)$. DPS proposes to approximate the intractable likelihood $p_t(\mathbf{y} | \mathbf{x}_t)$ with $p_t(\mathbf{y} | \hat{\mathbf{x}}_0)$. Loss-guided diffusion (Song et al., 2023) extends this approach to arbitrary loss functions $\ell_{\mathbf{y}}(\hat{\mathbf{x}}_0)$ by choosing $p(\mathbf{y} | \hat{\mathbf{x}}_0) = \frac{1}{Z} \exp(-\ell_{\mathbf{y}}(\hat{\mathbf{x}}_0))$, where Z is the partition function. This allows for approximating

$$\nabla_{\mathbf{x}_t} \log p_t(\mathbf{y} | \mathbf{x}_t) \approx \nabla_{\mathbf{x}_t} \log p_t(\mathbf{y} | \hat{\mathbf{x}}_0) = -\nabla_{\mathbf{x}_t} \ell_{\mathbf{y}}(\hat{\mathbf{x}}_0). \quad (2)$$

To implement the guidance, we follow DPS and add $-\rho_D \nabla_{\mathbf{x}_t} \ell_{\mathbf{y}}(\hat{\mathbf{x}}_0)$ to \mathbf{x}_{i-1} in the denoising step, where ρ_D is a scaling parameter.

162 4 GUIDED DIFFUSION WITH OPTIMIZED LOSS FUNCTIONS
163164 4.1 OPTIMIZED LOSS FUNCTIONS FOR INVERSE DESIGN
165

166 We consider discretized inverse design problems, with parameters $\theta \in \Theta$ for which the quality of the
167 design is evaluated using an objective function $J(\theta, \mathbf{u})$. J depends on the design parameters θ and
168 also on the solution $\mathbf{u}(\theta) \in \mathbb{R}^N$ of a forward optimization problem that depends on the design θ . For
169 example, one can calculate the inner displacements in a microstructure with FEM given the design
170 parameters. The solution \mathbf{u} needs to satisfy some constraint function $c(\theta, \mathbf{u}) = 0$. The goal is to
171 minimize the objective function $J(\theta, \mathbf{u})$, i.e., $\min_{\theta \in \Theta, \mathbf{u} \in \mathbb{R}^N} J(\theta, \mathbf{u})$ s.t. $c(\theta, \mathbf{u}) = 0$. Note that
172 these constraints are not the constraints on the designs. For many problems, J is not differentiable
173 w.r.t. to θ . For example, one needs to generate a valid discretized microstructure $\mathbf{m} \in \mathcal{M} \subset \mathbb{R}^K$
174 from θ for FEM simulation. We relax the parameter space into another continuous parameter space
175 $\mathcal{X} = \mathbb{R}^K$, i.e., $\mathcal{M} \subset \mathcal{X}$, which is potentially of much higher dimensionality than the original Θ .
176 Additionally, for each $\theta \in \Theta$, there exists a $\mathbf{x} \in \mathcal{X}$ with $\mathbf{x} = \mathbf{x}(\theta)$. We further assume that J can be
177 expressed in \mathbf{x} as $J(\mathbf{x}, \mathbf{u}(\mathbf{x}))$ and that J is differentiable w.r.t. to \mathbf{u} and \mathbf{x} . An example of \mathbf{m} , as used
178 in our experiments, is a 2D or 3D discretization into finite elements of a microstructure consisting
179 of a matrix and particles. The \mathbf{x} are parameter grids with arbitrary material properties of individual
180 pixels or voxels at each element.

181 In the relaxed parameter space, the implicit function theorem can be used to determine the total
182 differential of J w.r.t. \mathbf{x} (Xue et al., 2023), i.e.

$$184 \frac{dJ}{d\mathbf{x}} = -\frac{\partial J}{\partial \mathbf{u}} \left(\frac{\partial c}{\partial \mathbf{u}} \right)^{-1} \frac{\partial c}{\partial \mathbf{x}} + \frac{\partial J}{\partial \mathbf{x}} \quad \text{by using} \quad \frac{d\mathbf{u}}{d\mathbf{x}} = -\left(\frac{\partial c}{\partial \mathbf{u}} \right)^{-1} \frac{\partial c}{\partial \mathbf{x}}. \quad (3)$$

187 Note that the presented formulation might be impractical for computation and one can instead
188 use an adjoint formulation (Strang, 2007). Also note that optimizing the objective function using
189 gradient descent is not sufficient due to the high-dimensional ambiguous parameter space and missing
190 constraints of the original design space (e.g., discrete number and spherical particles, discrete set of
191 materials) which requires suitable means for regularization.

192 4.2 REGULARIZATION BY GUIDED DIFFUSION
193

195 Instead of solving the inverse design problem in the original design space, we propose to find possible
196 solutions in the relaxed parameter space. To approximate the constraints of the design problem, we
197 regularize gradient-based optimization in the relaxed space using a diffusion model of valid relaxed
198 parameters as prior. We first train an unconditional diffusion model on a training set of plausible
199 designs which are in \mathcal{M} . We then use loss guidance to infer relaxed parameters that remain close to
200 the training data manifold and minimize the loss function $\ell_y(\hat{\mathbf{x}}_0) = J(\mathbf{y}, \hat{\mathbf{y}}(\hat{\mathbf{x}}_0, \mathbf{u}))$. The loss can
201 be chosen, for instance, as the squared error between the expected measurement $\hat{\mathbf{y}}(\hat{\mathbf{x}}_0, \mathbf{u})$ and its
202 target value \mathbf{y} , i.e., $J(\mathbf{y}, \hat{\mathbf{y}}(\hat{\mathbf{x}}_0, \mathbf{u})) = \|\mathbf{y} - \hat{\mathbf{y}}(\hat{\mathbf{x}}_0)\|_2^2$. The expected measurement is calculated from
203 parameters $\hat{\mathbf{x}}_0$ and the solution \mathbf{u} for the constraint function c . For example, we can choose \mathbf{y} as a
204 target bulk modulus of a material, and $\hat{\mathbf{y}}(\hat{\mathbf{x}}_0, \mathbf{u})$ as the bulk modulus of the material generated by
205 guided diffusion.

206 5 MATRIX-PARTICLE INVERSE MATERIAL DESIGN BY LOSS-GUIDED
207 DIFFUSION
208

209 We consider an inverse material design problem in which circular respectively spherical particles of a
210 specific material, radius, and volume fraction are mixed into a matrix made of another material. Our
211 goal is to infer the parameters of the particles and the matrix materials in a microstructure of specific
212 square 2D or cubic 3D size.

214 Following the idea of homogenization, the macroscopic, averaged physical properties are calculated
215 on the microstructure (Temizer & Zohdi, 2007). This requires a specific set of possible loadings (i.e.,
displacements at the boundary) on the microstructure to fulfill energetic requirements (Hill, 1972).

To determine macroscopic properties, either the stress or the strain must be constant on average over the microstructure and independent of the distribution and composition of the materials. In this study, the constant strain approach is applied. In addition, we impose displacements linearly to all points of the surface to fulfill (Hill, 1972). In the case of isotropy, the material can be described using two parameters, such as Young's modulus E and Poisson ratio ν (Zohdi & Wriggers, 2008). Additionally, we consider the density ρ of the material. This leads to parameters (E_m, ν_m, ρ_m) for the matrix, in which particles of common radius r_p with material properties (E_p, ν_p, ρ_p) are mixed. The total volume fraction of particles is denoted by f_p which is implemented by an integer number of particles in our generated microstructures. Our design parameters are therefore $\theta = (E_m, E_p, \nu_m, \nu_p, \rho_m, \rho_p, r_p, f_p)$ and we are mainly interested in achieving a specific macroscopic bulk modulus K . The averaged, homogenized K is calculated from the average of the stress over the entire volume and the previously defined strain. $K = \frac{\text{tr}(\sigma)}{3 \text{tr} \varepsilon}$ where ε denotes the prescribed strain and $\langle \sigma \rangle$ the stress averaged over the microstructure. Boundary conditions are applied as $\bar{u}(\mathbf{q}) = \varepsilon \mathbf{q}$ where $\bar{u}(\mathbf{q})$ denotes the displacement at the surface point of the microstructure with position vector \mathbf{q} (Zohdi & Wriggers, 2008).

5.1 DISCRETE IMPLEMENTATION

The finite element method (FEM) is used to calculate the stress distribution within the microstructure (Zohdi & Wriggers, 2008) for calculating its bulk modulus. For this purpose, the volume is subdivided into equally sized elements. The adjacent connection points of the elements are referred to as nodes. The linear displacement is imposed on all boundary nodes on the surface. A discretized microstructure $\mathbf{m} \in \mathcal{M}$ consists of such an element grid where each element is assigned to either matrix or particle material. Particles are circles resp. spheres and equally sized. We additionally assume that there is an integer number of particles and they do not intersect the boundaries. Consequently, one can calculate design parameters θ from \mathbf{m} , but not all θ can directly be represented as a single microstructure. To determine reliable macroscopic measures, several such microstructures need to be sampled and their results averaged to assess the homogenization. Solutions for the FEM can be obtained by solving a linear system $\mathbf{A}\mathbf{u} = \mathbf{b}$. Note that $\frac{\partial c}{\partial \mathbf{u}} = \mathbf{A}$ and since \mathbf{A} is symmetric for this problem, the left term of Equation (3) can be obtained as $\mathbf{p} = \frac{\partial J}{\partial \mathbf{u}} \mathbf{A}^{-1}$ by solving the system $\mathbf{A}^\top \mathbf{p}^\top = \mathbf{A} \mathbf{p}^\top = (\frac{\partial J}{\partial \mathbf{u}})^\top$. Here, \mathbf{A} depends on the microstructure \mathbf{m} which itself depends on θ .

Our first considered objective function is measuring the squared error of the predicted material's K to a specific prescribed K^* , i.e., $J_1(K, K^*) = (K - K^*)^2$. Our second considered objective function is finding a prescribed K^* while also minimizing the average density, i.e., $J_2(K, K^*, \rho_m, \rho_p, f_p) = (K - K^*)^2 + \lambda ((1 - f_p) \rho_m + f_p \rho_p)$

5.2 GUIDED DIFFUSION OF MATERIAL PARAMETERS

While these objective functions are differentiable w.r.t. the material properties $(E_m, E_p, \nu_m, \nu_p, \rho_m, \rho_p)$, the spatial configuration and the radius and number of particles for a volume fraction are highly ambiguous. Moreover, optimizing the discrete number of particles in valid spatial configurations is challenging. Also, the possible set of materials is constrained to a specific discrete set of usable known materials (e.g., specific types of rubber, metals). Without proper regularization, this optimization problem is ill-posed. We use diffusion models to learn a prior over valid relaxed material parameters. We use it in guided diffusion to perform the regularized optimization.

To parametrize materials in a continuous form which is suited for diffusion models, we represent materials as finite element discretizations $\mathbf{x} \in \mathcal{X}$ similar to \mathcal{M} , but each element e can have arbitrary material properties (E^e, ν^e, ρ^e) . The relaxed material parameters which we optimize by guided diffusion are thus the element materials in the finite element grid. This representation is 2D-image- or 3D-grid-like and can be embedded into a latent representation using a common U-Net (Ronneberger et al., 2015) architecture.

We train the diffusion model with samples that are in \mathcal{M} to learn a prior over valid microstructures in the relaxed parameter space. To generate an example, we randomly sample the properties of matrix and particle from a table of possible materials. We also sample volume fraction f_p and particle radius r_p randomly, determine the number of particles by rounding the quotient of f_p and the area πr_p^2 resp. volume $\frac{4}{3}r_p^3$ and create a microstructure \mathbf{m} with randomized, non-overlapping particle positions.

270
271

5.3 BACKPROJECTION

272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293

Given an objective function, e.g., to achieve a bulk modulus, our guided diffusion yields parameters \mathbf{x} , for which we now need to find the most accurate match $\hat{\theta} \in \Theta$. Ideally, we expect the resulting sample to have a structure where particles can be distinguished from the matrix, all particles have the same radius and the materials for all particles and for all matrix pixels, respectively, are identical. In principle, a domain expert can analyze the result and determine validity and estimate a matching $\hat{\theta}$. However, we describe an automatic approach that we apply in our experiments. First, we fit a 2-component Gaussian mixture model on the vectorized parameter image or grid on the 3-channel material data E, ν, ρ . For this model, we prescribe spherical covariances for simplicity. We expect two peaks, if the composite material consists of different materials for particles and matrix, or one, if there are no particles or they have the same material. The means are the estimates for the materials $\hat{E}_1, \hat{E}_2, \hat{\nu}_1, \hat{\nu}_2, \hat{\rho}_1, \hat{\rho}_2$. Then, we try to distinguish elements into matrix or particle material. Each element is assigned to the closest of both detected materials. We test both possibilities of assigning a material to the particles and identify circles (2D) or spheres (3D) and their median radius based on skeletonization (Lee et al., 1994). We choose the assignment with lowest variance in detected radii. This yields r_p as mean of matched radii and we obtain the volume fraction f_p by counting the number of particle and matrix material assignments to elements. Note that this does not necessarily result in an integer number of particles. Details of our algorithm can be found in the supplementary material section B. In the list of available materials, we search for the nearest neighbor of our predicted materials for matrix and particle. We measure distance in a normalized space $[-1, 1]^3$ as metric. The nearest existing materials together with r_p and f_p form our prediction $\hat{\theta}$ in the original parameter space Θ . To approximate the expected value of the bulk modulus for such $\hat{\theta}$, we sample several microstructures with discrete number of particles and average the results.

294

6 EXPERIMENTS

295
296
297

6.1 EXPERIMENT SETUP

298
299
300
301
302
303
304
305
306
307
308

For our experiments, we selected properties (E, ν, ρ) of 500 materials according to the online database MatWeb¹ with properties $0 < E \leq 500$, $0 < \nu < 0.5$, $0 < \rho < 10$. Since more than 8000 materials fulfill these properties and their distribution density varies greatly (e.g. there exist hundreds of entries for steel with similar properties), we slice the three property dimensions into 10 equidistant segments each and obtain 1000 chunks. We query the database for each of the individual chunks and subsample retrieved materials per chunk so that the total number does not exceed 500, a limitation required by MatWeb’s terms of use. To this end, we compute a maximum number of materials that any chunk can contain and subsample accordingly. 168 of the chunks are nonempty. We show the distribution of the materials in the supplementary material section A. We normalize each dimension between $[-1, 1]$, which constitutes the space in which our model is trained and our distances for nearest neighbor-matching, as well as variances for material matching are reported.

309
310
311
312
313
314
315
316
317
318

To obtain a balanced training set, we first sample one of the nonempty chunks and then uniformly sample one of the available materials inside that chunk for both matrix and particles. We then sample volume fraction and particle radius: For 2D, we sample v_f uniformly in $[0.05, 0.5]$ and $2 \cdot r_p$ (diameter) uniformly in $[0.15, 0.4]$ where the unit refers to the relative size to the microstructure. For 3D, the ranges for v_f are $[0.05, 0.45]$ and for $2 \cdot r_p$ (diameter) $[0.15, 0.35]$. Note that there are limits to those quantities, since circles or spheres need to be packed tightly, which becomes hard or even impossible with random sampling for higher volume fractions. We then sample the particle positions (see Section 5) and obtain a discretized microstructure for diffusion and FEM. For simpler computation, we use a 3D FEM solver also for our 2D case, where we build a thin 3D plane with a thickness of one element in one dimension. This represents plane stress conditions (the structure can move in this dimension, but the stress is zero). For our training dataset, we sample 10,000 examples.

319
320
321
322
323

For our analysis, we need to choose a set of target values for K^* . To this end, we create a new dataset of 10k samples similar to the training set and determine the 1 and 99-percentiles of K values on the individual examples, which constitute an interval from which we select 5 uniformly spaced values, including start and end. See the supplementary material section A for a histogram of K values. We

¹<https://www.matweb.com/>

324 tune guidance parameters on the start, end and midpoint values (for 2D these are 4.8, 168.5, 332.2)
 325 and additionally report results on the 25% and 75% positions in the interval (86.6, 250.4).
 326

327 6.2 DENOISING, TRAINING, AND EVALUATION METRICS

329 **Model.** For implementing our model and diffusion process, we use the Diffusers library (von Platen
 330 et al., 2022). The architecture is based on UNet (Ronneberger et al., 2015) using either 2D or 3D
 331 convolutions. The model employs ResNet layers and two downsampling steps, each halving the input
 332 dimensions. Before the respective upsampling stage, two more ResNet layers with self attention
 333 layers are employed. We detail the model architecture in the supplementary material section C.

334 **Denoising.** We use DDIM with $\eta = 1$ and by default $N = 100$ denoising time steps. We use a
 335 linear β -schedule between $\beta_0 = 10^{-5}$ and $\beta_T = 10^{-2}$ and employ β -rescaling as described by (Lin
 336 et al., 2024) which ensures that no information of the clean sample is left at $t = T$. We sample
 337 time steps that correspond to their “trailing” strategy. After denoising, we clip the sample to lie in
 338 the boundaries $[-1, 1]$. We use DPS with constant scaling parameter $\rho_D = 1$ for guidance in our
 339 experiments. **We scale the gradients from the FEM solver by 0.5 for E and by 0.02 for ν which we**
 340 **determined in a grid search. Please refer to the supp. mat. section D for more details on the choice of**
 341 **guidance parameters.**

342 **Training.** For training, we use a cosine learning schedule with peak learning rate 10^{-3} and 5,000
 343 warmup steps. We use the AdamW optimizer (Loshchilov & Hutter, 2019) with PyTorch default
 344 parameters $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 10^{-8}$, $\lambda = 10^{-2}$. We clip the gradient norm at 1. We train
 345 our models for 100,000 steps with a batch size of 128.

346 **Evaluation metrics.** For assessing how well our approach finds parameters that achieve a desired K^* ,
 347 we first perform backprojection (cf. Section 5.3) on the generated samples, which yields a set of valid
 348 materials and particle parameters. To estimate the bulk modulus corresponding to these extracted
 349 parameters, we generate $n = 10$ microstructures with random spatial distribution of particles, similar
 350 to our dataset generation. Since we require an integer number of discrete particles for the sampled
 351 microstructure, we sample this number to accurately represent the predicted volume fraction (e.g., if
 352 the required number of particles computed by volume fraction and diameter is 3.2, we use 3 circles
 353 with probability 0.8 and 4 with 0.2). The bulk modulus is computed again via FEM and averaged over
 354 the n samples, yielding the quantity K_θ . **For the qualitative experiments, we additionally compute**
 355 **the bulk modulus of the particular generated sample and report it as K_s .** We consider the relative
 356 error $\epsilon_r = |K_\theta - K^*|/K^*$ as primary metric. However, since we propose a probabilistic method that
 357 generates random samples by diffusion, it is interesting to know how many samples fall within some
 358 error margin. Therefore, for a set of generated samples with same target K^* , we compute how many
 359 of the samples have $\epsilon_r < 1\%$ and $\epsilon_r < 5\%$ and denote this fraction by $frac$. For small sample sets,
 360 this quantity has a high variance, which can be reduced by obtaining more samples. We additionally
 361 assess the diversity of generated samples. Due to the high variability in material sample density,
 362 we assess the diversity by the number of unique material chunks in which the model generated
 363 samples (considering the nearest actually available materials as done by the backprojection). For a
 364 set of generated samples, the metric cov counts these chunks and computes the fraction of it by the
 365 nonempty chunks. Note that this metric is highly dependent on the number of samples.

366 6.3 RESULTS


367 We first evaluate our approach on models trained on 10k samples for a 2D problem discretized into
 368 64×64 elements. In Table 1a, we assess the quality of 200 samples obtained for the J_1 objective
 369 function for each target K^* (averaged over 3 model training seeds). We observe that more samples
 370 satisfy the error margins at medium K^* values. For all targets, at least one sample can be found in the
 371 $\epsilon_r < 1\%$ margin. We note that the samples exhibit a diverse coverage of material chunks, as, e.g., for
 372 $K^* = 168.5$, the average 100 samples that satisfy $\epsilon_r < 5\%$ cover 77 of the 168 material chunks on
 373 average. We show the best samples from a single trained model in terms of ϵ_r and for higher quantiles
 374 of ϵ_r for some K^* in Figure 1. One can see that, e.g., for $K^* = 168.5$, many diverse designs are
 375 generated, as can be seen by the varying shapes and colors. Even for higher error quantiles of that
 376 target, one can observe that the generated sample still matches the target well (K_s similar to K^*).
 377 For the worst sample (100% quantile), the result after the backprojection (K_θ) deviates stronger. One
 378 possibility is that the model generated materials that are not available or the generated structure could

378
379 Table 1: Evaluation of 200 samples each at different target K^* , guiding for J_1 , averaged over 3 model
380 training seeds. K^* marked with \dagger are part of the targets used for guidance parameter optimization.

(a) Evaluation on 2D problem							(b) Evaluation on 3D problem						
	K^* :	4.8 \dagger	86.6	168.5 \dagger	250.4	332.2 \dagger		K^* :	0.9 \dagger	83.3	165.6 \dagger	248.0	330.3 \dagger
$\epsilon_r < 1\%$	frac	0.003	0.052	0.110	0.043	0.033	$\epsilon_r < 1\%$	frac	0.000	0.037	0.100	0.075	0.100
	cov	0.008	0.097	0.171	0.038	0.036		cov	0.000	0.077	0.157	0.093	0.022
$\epsilon_r < 5\%$	frac	0.033	0.278	0.500	0.195	0.107	$\epsilon_r < 5\%$	frac	0.005	0.220	0.455	0.387	0.388
	cov	0.056	0.321	0.458	0.127	0.075		cov	0.012	0.292	0.411	0.210	0.048

389
390 Table 2: Evaluation of alternative inverse design approaches on the 2D problem. Metrics are computed
391 over 200 samples each at different target K^* of J_1 , averaged over 3 seeds.

(a) Bayesian optimization							(b) Conditional diffusion model						
	K^* :	4.8	86.6	168.5	250.4	332.2		K^* :	4.8	86.6	168.5	250.4	332.2
$\epsilon_r < 1\%$	frac	0.002	0.103	0.090	0.110	0.148	$\epsilon_r < 1\%$	frac	0.008	0.058	0.130	0.120	0.063
	cov	0.004	0.099	0.149	0.097	0.111		cov	0.018	0.121	0.214	0.129	0.083
$\epsilon_r < 5\%$	frac	0.007	0.482	0.408	0.580	0.608	$\epsilon_r < 5\%$	frac	0.040	0.335	0.570	0.517	0.287
	cov	0.016	0.266	0.322	0.248	0.248		cov	0.077	0.427	0.456	0.302	0.264



400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419 Figure 1: Inverse 2D material designs. Generated samples for selected bulk moduli K^* , ordered by
420 relative error quantile. Best on the left, worst on the right. Labels show $K_s / K_\theta / \epsilon_r$. The values
421 (E, ν, ρ) in the normalized coordinate space are encoded as (r, g, b) values of the image. Our
422 model is able to propose diverse and plausible designs close to the target bulk moduli.

424
425
426 not be well matched by the backprojection. For the higher target bulk moduli, the model proposes
427 several plausible designs, while implausible designs can be observed for the worst samples and the
428 50% quantile for $K^* = 332.2$. Results for further K^* can be found in the supp. mat. section F.

429
430 **Ablations** In Table 3a, we compare the performance of our default models with models that are
431 trained on only 1k samples. We see that the latter perform almost equally well, hinting that fewer
432 training samples could be used than in our default setting. We also compare our default denoising time




Figure 2: Inverse 3D material designs. Generated samples for selected bulk moduli K^* , ordered by relative error quantile. Best on the left, worst on the right. Labels show $K_s / K_\theta / \epsilon_r$. The values (E, ν, ρ) in the normalized coordinate space are encoded as (r, g, b) values of the image. Our model is able to propose diverse and plausible designs close to the target bulk moduli.

Table 3: Performance for guidance by J_1, J_2 with ablations. Mean metrics over all targets K^* for 200 samples each and over 3 model training seeds.

(a) Performance for guidance by J_1 , ablations.

	$\epsilon_r < 1\%$		$\epsilon_r < 5\%$	
	frac	cov	frac	cov
ds 10k	0.049	0.073	0.226	0.213
ds 1k	0.047	0.069	0.213	0.180
$N = 50$	0.043	0.068	0.200	0.185
$N = 200$	0.055	0.086	0.245	0.224
unguided	0.002	0.004	0.019	0.037
project	0.066	0.069	0.247	0.148

(b) Performance for guidance by J_2 (also minimize density). Additionally we show the average density of samples that fall in the respective error margin.

	$\epsilon_r < 1\%$		$\epsilon_r < 5\%$			
	ρ_{avg}	frac	cov	ρ_{avg}	frac	cov
$\lambda = 0$	4.546	0.049	0.073	4.263	0.226	0.213
$\lambda = 10^{-4}$	2.994	0.051	0.078	2.904	0.236	0.207
$\lambda = 10^{-3}$	2.338	0.008	0.014	2.200	0.052	0.040
$\lambda = 10^{-2}$	1.713	0.006	0.007	1.714	0.031	0.017

steps of 100 against 50 and 200 and observe that more timesteps yield slightly better performance, both in frac and cov. However, the differences are relatively small, so that a denoising with $N = 50$ timesteps (meaning also only 49 gradient computations) could be a viable speedup, since the time depends linearly on N . For example, obtaining 200 samples with $N = 100$ took approximately 1.5h on a cluster node using 16 CPU cores and an A40 GPU. [Further details on runtime are found in the supp. mat. section F](#). We also provide a baseline that uses no guidance at all which demonstrates that guidance strongly improves sample adherence to the objective. In addition, we experiment with a variant that uses $N = 200$ timesteps, but alternates between a guidance step and a step that performs part of the backprojection on \hat{x}_0 , replacing material values with the nearest neighbors of actually available materials. This encourages the diffusion model to generate designs that use existing materials and uses a similar amount of gradient computations. We see that indeed the frac metrics improve, but at the cost of reduced diversity as measured by cov.

Multiple objectives. We also evaluate our method on a different objective function, J_2 , which incorporates minimizing the average density of the sample as additional objective. Results in Table 3b for varying factors of the density penalty term show that indeed with stronger penalty, the average densities are reduced further. This also leads to lower frac and cov, which is expected since the set of

486 acceptable samples is reduced. This experiment demonstrates how our approach can be used to adapt
 487 the objective function in a zero-shot way without retraining the model.
 488

489 **Comparison to alternative methods.** We compare our approach to established methods for inverse
 490 material design methods on the 2D problem. Details can be found in the supp. mat. section E.
 491 Firstly, we compare to a Bayesian optimization approach which models a surrogate cost function
 492 using Gaussian processes. We perform an optimization over 1000 steps for the same targets K^* for
 493 the objective function J_1 as used to evaluate our method. Results are presented in table Table 2a. For
 494 $K^* = 4.8$ and 168.5 , our approach outperforms dedicated Bayesian optimizations in both metrics and
 495 error margins, for $K^* = 86.6$ it is similar or better in the cov metric. For higher target bulk moduli, a
 496 higher fraction and more diverse samples can be found with the Bayesian optimization approach. We
 497 note, however, that for each of the target bulk moduli (or generally, different objectives), the whole
 498 optimization process, including exploration of the space and fitting the surrogate function, has to be
 499 performed from scratch. To obtain a sample, sampling and minimization of the surrogate function
 500 has to be performed. Our method in contrast can be applied directly to new objectives and only uses
 FEM gradients of the relaxed problem.

501 We also compare our approach to a setting where the diffusion model is trained on a dataset with
 502 annotated bulk moduli K as conditional input. This allows to perform sampling with classifier-free
 503 guidance (Ho & Salimans, 2022) as for example employed by Yang et al. (2024). Results in Table 2b
 504 show that in this setting, higher metrics can be achieved compared to our approach in almost all
 505 cases. Our approach exhibits similar performance for targets until $K^* = 168.5$ and has a larger
 506 performance gap for higher targets. We note, however, that conditional diffusion models require an
 507 explicitly provided label to condition on. A loss function like J_2 which penalizes density without an
 508 explicit density target cannot be implemented in this setting without further guidance mechanisms.
 509 Also, a modification of the type of conditional is not possible without retraining, while our approach
 510 is more general and can be adapted to various objective functions in a zero-shot way.
 511

512 **3D problem.** We also train models on 10k samples of a 32x32x32 discretized 3D problem and
 513 perform guided sampling with the same parameters as in 2D for J_1 . Results for different target K^*
 514 are shown in Table 1b. One can observe that performance is worse compared to 2D up until the
 515 midpoint $K^* = 165.6$, but much better in terms of frac for higher targets. Except for $K^* = 0.9$,
 516 always several samples are found in $\epsilon_r < 1\%$. We note that this case is especially difficult, since
 517 the respective absolute error margin is only 0.009 and that samples could be found in $\epsilon_r < 5\%$. We
 518 provide visualizations of 3D samples generated by a single trained model in Figure 2.
 519

520 7 CONCLUSION

521 In this paper, we develop a novel approach for loss-guided diffusion in which the loss is evaluated
 522 by solving an inner optimization problem, and evaluate our method for an inverse material design
 523 problem which requires solving a linear FEM to assess the bulk modulus of composite material
 524 microstructures. The microstructure consists of spherical particles in a matrix. Our approach operates
 525 on a relaxed reparametrization of the original parameter space which allows for denoising 2D and 3D
 526 grid representations of the microstructures. The diffusion model acts as prior and approximates the
 527 constraints of the design problem. Approximate design samples are projected back into the original
 528 design space. Our approach can directly leverage physics-based simulation to determine the loss
 529 function and does not require training a surrogate model for the loss. We evaluate our method using a
 530 dataset of real material properties and demonstrate that our approach finds diverse samples within a
 531 relative error of 1% from medium to high target bulk moduli in 2D and 3D settings. **Our approach can**
 532 **optimize multiple objectives, which we show by also minimizing the density of generated samples in**
 533 **addition to matching a specified bulk modulus.** We anticipate that our approach will inspire future
 534 applications of guided diffusion with optimization-based loss functions for inverse design in various
 535 application areas. In this work, we only considered linear FEMs as inner optimization problems.
 536 Future work could also evaluate and extend the method for non-linear optimization problems if
 537 implicit differentiation is possible. **Another interesting direction of future research is to investigate**
 538 **implementing parallel solvers on GPU or incremental solvers which can reuse solutions of the inner**
 539 **optimization problem from previous diffusion iterations to improve runtime efficiency, especially for**
 larger-scale or non-linear problems. From the material design perspective, investigating non-linear
 material properties and anisotropic materials is an interesting avenue of future research.

540 REPRODUCIBILITY STATEMENT
541

542 The section Experiments and the supplementary material provide details on model architecture,
543 hyperparameters, and the algorithm for reproducing the experiments. Additionally, we plan to make
544 the source code for our method publicly available after acceptance of the paper. The datasets cannot
545 be made publicly available due to license terms. However, we describe how one can create a similar
546 dataset.

547

548 REFERENCES
549

- 550 Jason Ansel, Edward Yang, Horace He, Natalia Gimelshein, Animesh Jain, Michael Voznesensky,
551 Bin Bao, Peter Bell, David Berard, Evgeni Burovski, Geeta Chauhan, Anjali Chourdia, Will
552 Constable, Alban Desmaison, Zachary DeVito, Elias Ellison, Will Feng, Jiong Gong, Michael
553 Gschwind, Brian Hirsh, Sherlock Huang, Kshitij Kalmabarkar, Laurent Kirsch, Michael Lazos,
554 Mario Lezcano, Yanbo Liang, Jason Liang, Yinghai Lu, CK Luk, Bert Maher, Yunjie Pan, Christian
555 Puhrsch, Matthias Reso, Mark Saroufim, Marcos Yukio Siraichi, Helen Suk, Michael Suo, Phil
556 Tillet, Eikan Wang, Xiaodong Wang, William Wen, Shunting Zhang, Xu Zhao, Keren Zhou,
557 Richard Zou, Ajit Mathews, Gregory Chanan, Peng Wu, and Soumith Chintala. PyTorch 2: Faster
558 Machine Learning Through Dynamic Python Bytecode Transformation and Graph Compilation.
559 In *29th ACM International Conference on Architectural Support for Programming Languages and
Operating Systems, Volume 2 (ASPLOS '24)*, 2024. doi: 10.1145/3620665.3640366.
- 560 Arpit Bansal, Hong-Min Chu, Avi Schwarzschild, Roni Sengupta, Micah Goldblum, Jonas Geiping,
561 and Tom Goldstein. Universal Guidance for Diffusion Models. In *Proc. of the International
Conference on Learning Representations (ICLR)*, 2023.
- 562 Jan-Hendrik Bastek and Dennis M. Kochmann. Inverse design of nonlinear mechanical metamaterials
563 via video denoising diffusion models. *Nat. Mac. Intell.*, 5(12), 2023.
- 564 Hyungjin Chung, Jeongsol Kim, Michael Thompson McCann, Marc Louis Klasky, and Jong Chul
565 Ye. Diffusion posterior sampling for general noisy inverse problems. In *Proc. of the International
Conference on Learning Representations (ICLR)*, 2023.
- 566 Hyungjin Chung, Suhyeon Lee, and Jong Chul Ye. Decomposed diffusion sampler for accelerating
567 large-scale inverse problems. In *Proc. of the International Conference on Learning Representations
(ICLR)*, 2024.
- 568 Adam Foster, Martin Jankowiak, Elias Bingham, Paul Horsfall, Yee Whye Teh, Thomas Rainforth, and
569 Noah Goodman. Variational bayesian optimal experimental design. In H. Wallach, H. Larochelle,
570 A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information
Processing Systems*, 2019.
- 571 Gaia Franceschini and Sandro Macchietto. Model-based design of experiments for parameter
572 precision: State of the art. *Chemical Engineering Science*, 63(19):4846–4872, 2008.
- 573 Peter I. Frazier and Jialei Wang. *Bayesian Optimization for Materials Design*, pp. 45–75. Springer
574 International Publishing, Cham, 2016.
- 575 Rodney Hill. On constitutive macro-variables for heterogeneous solids at finite strain. *Proceedings
of the Royal Society of London. A. Mathematical and Physical Sciences*, 326(1565), 1972.
- 576 Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*,
577 2022.
- 578 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising Diffusion Probabilistic Models. *Advances in
579 Neural Information Processing systems (NeurIPS)*, 33, 2020.
- 580 Christian Huet. Application of variational concepts to size effects in elastic heterogeneous bodies.
581 *Journal of the Mechanics and Physics of Solids*, 38(6), 1990.
- 582 Ta-Chih Lee, Rangasami L. Kashyap, and Chong-Nam Chu. Building skeleton models via 3-D
583 medial surface/axis thinning algorithms. *CVGIP Graph. Model. Image Process.*, 56(6), 1994.

- 594 Shanchuan Lin, Bingchen Liu, Jiashi Li, and Xiao Yang. Common diffusion noise schedules and
 595 sample steps are flawed. In *Proc. of the IEEE/CVF winter conference on applications of computer*
 596 *vision*, 2024.
- 597 Qibang Liu, Seid Koric, Diab Abueidda, Hadi Meidani, and Philippe Geubelle. Towards signed
 598 distance function based metamaterial design: Neural operator transformer for forward prediction
 599 and diffusion model for inverse design. *arXiv preprint arXiv:2504.01195*, 2025.
- 601 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *Proc. of the Interna-*
 602 *tional Conference on Learning Representations (ICLR)*, 2019.
- 603 Calvin Luo. Understanding diffusion models: A unified perspective. *arXiv preprint arXiv:2208.11970*,
 604 2022.
- 606 Fernando Nogueira. Bayesian Optimization: Open source constrained global optimization
 607 tool for Python, 2014. URL [https://github.com/bayesian-optimization/](https://github.com/bayesian-optimization/BayesianOptimization)
 608 BayesianOptimization.
- 610 Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical
 611 image segmentation. In *Proc. of the International Conference on Medical Image Computing and*
 612 *Computer-Assisted Intervention (MICCAI)*, volume 9351 of *Lecture Notes in Computer Science*,
 613 2015.
- 614 Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. In
 615 *Proc. of the International Conference on Learning Representations (ICLR)*, 2022.
- 616 Yifei Shen, Xinyang Jiang, Yifan Yang, Yezhen Wang, Dongqi Han, and Dongsheng Li. Understanding
 617 and improving training-free loss-based diffusion guidance. *Advances in Neural Information*
 618 *Processing Systems*, 37:108974–109002, 2024.
- 620 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising Diffusion Implicit Models. In *Proc. of*
 621 *the International Conference on Learning Representations (ICLR)*, 2020.
- 622 Jiaming Song, Qinsheng Zhang, Hongxu Yin, Morteza Mardani, Ming-Yu Liu, Jan Kautz, Yongxin
 623 Chen, and Arash Vahdat. Loss-Guided Diffusion Models for Plug-and-Play Controllable Genera-
 624 tion. In *Proc. of the International Conference on Machine Learning (ICML)*, 2023.
- 626 Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. Solving inverse problems in medical
 627 imaging with score-based generative models. In *Proc. of the International Conference on Learning*
 628 *Representations (ICLR)*, 2022.
- 630 Gilbert Strang. *Computational Science and Engineering*. Wellesley-Cambridge Press, Philadelphia,
 631 PA, 2007. doi: 10.1137/1.9780961408817. URL [https://epubs.siam.org/doi/abs/](https://epubs.siam.org/doi/abs/10.1137/1.9780961408817)
 632 10.1137/1.9780961408817.
- 633 Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh
 634 Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high
 635 frequency functions in low dimensional domains. In *Advances in Neural Information Processing*
 636 *Systems (NeurIPS)*, volume 33, 2020.
- 637 I Temizer and TI Zohdi. A numerical method for homogenization in non-linear elasticity. *Computa-*
 638 *tional Mechanics*, 40, 2007.
- 639 Clément Vignac, Nagham Osman, Laura Toni, and Pascal Frossard. MiDi: Mixed graph and 3D
 640 denoising diffusion for molecule generation. In *European Conference on Machine Learning and*
 641 *Knowledge Discovery in Databases: Research Track (ECML PKDD)*, volume 14170 of *Lecture*
 642 *Notes in Computer Science*, 2023.
- 643 Patrick von Platen, Suraj Patil, Anton Lozhkov, Pedro Cuenca, Nathan Lambert, Kashif Rasul,
 644 Mishig Davaadorj, Dhruv Nair, Sayak Paul, William Berman, Yiyi Xu, Steven Liu, and Thomas
 645 Wolf. Diffusers: State-of-the-art diffusion models. [https://github.com/huggingface/](https://github.com/huggingface/diffusers)
 646 diffusers, 2022.

- 648 Valentin Würz and Christian Weißenfels. Inverse material design using deep reinforcement learning
 649 and homogenization. *Computer Methods in Applied Mechanics and Engineering*, 435, 2025.
 650
- 651 Tianju Xue, Shuheng Liao, Zhengtao Gan, Chanwook Park, Xiaoyu Xie, Wing Kam Liu, and
 652 Jian Cao. JAX-FEM: A differentiable GPU-accelerated 3D finite element solver for automatic
 653 inverse design and mechanistic data science. *Computer Physics Communications*, 291, 2023. doi:
 654 10.1016/j.cpc.2023.108802.
- 655 Yanyan Yang, Lili Wang, Xiaoya Zhai, Kai Chen, Wenming Wu, Yunkai Zhao, Ligang Liu, and Xiao-
 656 Ming Fu. Guided Diffusion for Fast Inverse Design of Density-based Mechanical Metamaterials.
 657 *arXiv preprint arXiv:2401.13570*, 2024.
- 658 Haotian Ye, Haowei Lin, Jiaqi Han, Minkai Xu, Sheng Liu, Yitao Liang, Jianzhu Ma, James Y. Zou,
 659 and Stefano Ermon. TFG: unified training-free guidance for diffusion models. In *Advances in
 660 Neural Information Processing Systems (NeurIPS)*, 2024.
- 661 Jiwen Yu, Yinhua Wang, Chen Zhao, Bernard Ghanem, and Jian Zhang. FreeDoM: Training-free
 662 energy-guided conditional diffusion model. In *Proc. of the IEEE/CVF International Conference
 663 on Computer Vision (ICCV)*, 2023.
- 664 Stefano Zampini, Jacob K Christopher, Luca Oneto, Davide Anguita, and Ferdinando
 665 Fioretto. Training-free constrained generation with stable diffusion models. *arXiv preprint
 666 arXiv:2502.05625*, 2025.
- 667 Tarek I Zohdi and Peter Wriggers. *An introduction to computational micromechanics*. Springer
 668 Science & Business Media, 2008.
- 669 TI Zohdi. Constrained inverse formulations in random material design. *Computer Methods in Applied
 670 Mechanics and Engineering*, 192(28-30), 2003.
- 671
 672
 673
 674
 675
 676
 677
 678
 679
 680
 681
 682
 683
 684
 685
 686
 687
 688
 689
 690
 691
 692
 693
 694
 695
 696
 697
 698
 699
 700
 701

Supplementary Material

A MATERIAL LIST AND DATASET DETAILS




Figure 3: Visualization of base materials used. Color represents the index in the 168 non-empty chunks.

Material list For our experiments, we selected properties (E, ν, ρ) of 500 materials from the online database MatWeb². The original data is subject to copyright and terms of use of MatWeb. Due to license terms, our derived datasets and models cannot be made publicly available. Figure 3 shows the distribution of the base materials we used. The value ranges are $E \in [0.0055, 462]$, $\nu \in [0.032, 0.499]$, $\rho \in [0.032, 9.99]$.


Dataset sampling When sampling an example for our datasets, we proceed in the following way: For both materials of the matrix and the particles, we first uniformly sample a non-empty chunk (out of 168) and then uniformly sample a base material that is contained in that chunk. The same distribution is used for matrix and particle material. After sampling the volume fraction v_f and particle radius r_p (see main paper), the number of particles is determined by dividing the volume fraction by the area (2D) or volume (3D) of a circle resp. sphere of that radius and rounding the result. This can result in a slightly different volume fraction than initially sampled. Note that we consider a unit square resp. cube.

To determine non-overlapping sphere positions (likewise for circles), we first randomly sample positions for all spheres so that boundaries are not intersected. We then compute the distances between all spheres and return if there are no intersections. If there are intersections, we determine the directions between all sphere centers and add small random vectors to better resolve penetrations. For each pairwise intersection between spheres, the delta in the direction that would resolve this intersection is added times 1.5 to a total delta per sphere position. Then, these deltas are applied at once for all spheres and the intersection check is done again. We stop this iterative resolving after 10,000 unsuccessful updates and re-sample initial positions in that case.

After obtaining an example microstructure, we can determine its bulk modulus K with a FEM solver. Note that this is, however, not necessary to train the diffusion model.

Histogram of K in datasets As described in section Section 6.1, we create datasets similar to our training datasets for the purpose of finding suitable target bulk moduli K^* . We show histograms of

²<https://www.matweb.com/>



(a) K histogram (50 bins) of 2D problem (64×64), otherwise unused seed. Cut at $K = 450$, maximum 784.5.

(b) K histogram (50 bins) of 3D problem ($32 \times 32 \times 32$), otherwise unused seed. Cut at $K = 450$, maximum 795.4.

Figure 4: Histograms of bulk modulus K of individual examples in datasets. Guidance targets are chosen uniformly spaced between the 1 and 99 percentile.




Figure 5: Relation between quotient of Elastic modulus of matrix and particle (y-axis) and quotient of bulk modulus K and lower (b^-) / upper (b^+) Voigt-Reuss bounds (Zohdi & Wriggers, 2008). Shown are 500 samples generated similarly to our 3D training dataset. The lower bound fraction is clipped at 2 for better display.

the bulk moduli of those datasets in Figure 4. We choose the range of target bulk moduli K^* for evaluation as the 1% and 99% quantiles from these datasets.

We validate the resulting bulk modulus of samples for the 3D problem by computing the Voigt-Reuss bounds (Zohdi & Wriggers, 2008) from the material of matrix and particle for 500 samples generated similarly to our 3D training dataset and show the results in Figure 5. Out of the 500 samples, 33 were outside the bounds, but only with minor deviation, as evident from the plot. Due to the allocation of materials to the entire element, sawtooth-shaped transitions occur. We hypothesize that slightly excessive stresses may occur at the edges that can lead to values slightly above the bounds.

B BACKPROJECTION DETAILS AND EVALUATION

In the following, we detail how we identify particles in the microstructure using the GMMs fitted to the material channels. For each element, we obtain a binary label indicating which component of the GMMs it belongs to with highest probability. This allows to represent the microstructure as a binary mask. The next step requires determining which label corresponds to the matrix and which to particles. We try to match circles resp. spheres for both cases of treating either the one- or zero-labels as foreground pixels, where the foreground is the candidate for the particles. To this end, we obtain the 2D resp. 3D skeleton by skeletonization (Lee et al., 1994) and their minimum

	K^*	4.8	86.6	168.5	250.4	332.2
V_m	0.007	0.001	0.001	0.040	0.077	
d_m	0.097	0.112	0.131	0.143	0.236	
frac $\epsilon_r < 1\% \text{ sample}$	0.020	0.490	0.557	0.183	0.473	
frac $\epsilon_r < 5\% \text{ sample}$	0.110	0.995	0.998	0.530	0.865	
frac $\epsilon_r < 1\% \text{ no closest material}$	0.018	0.478	0.528	0.175	0.213	
frac $\epsilon_r < 5\% \text{ no closest material}$	0.115	0.997	1.000	0.383	0.535	
frac $\epsilon_r < 1\%$	0.003	0.052	0.110	0.043	0.033	
frac $\epsilon_r < 5\%$	0.033	0.278	0.500	0.195	0.107	

Table 4: Further analysis of generated samples. Guidance for J_1 on the 2D problem with our main settings, 200 generated samples, averaged over 3 model training seeds. Results in the last two rows are identical to Table 1a.

distance to background pixels. We use some domain knowledge and filter out points with too small distances to the boundary of the segment or skip the whole case if there are too large distances. We also skip the case if more than half of the edge pixels resp. face voxels are positives, which should not be possible since we consider problems where particles do not intersect with the boundary. **For all remaining foreground pixels i in the skeleton and their minimum distance to background d_i , we check if there is another foreground pixel j in the skeleton within d_i that itself has $d_j < d_i$.** In this case, we remove j . Afterwards, we are left with likely centers of circles resp. spheres and can use their distances as radius. If both cases have not been skipped, we compute the variance of remaining distances and choose the case with lower variance.

Additional metrics To assess the quality of (potentially unconditionally) generated samples themselves, without comparing to an external objective, we employ the following metrics: Firstly, the sum of the variances of material fitting V_m : During the backprojection, we fit a 2-component GMM to the 3-channel material data. We sum up the fitted variances of both components to form our metric V_m . In the ideal case, this variance is 0, meaning that only 1 or 2 values ever occur as material values. This measures how well a model can enforce consistency of the materials in a sample.

Secondly, we use the nearest neighbor distance to existing material d_m : In the backprojection, once the material parameters are fitted by the GMMs and correspondences are established, we look up the nearest neighbors of the two materials in our material list. The sum of the distances to these neighbors in the normalized space constitutes our metric d_m . Note that this part is especially challenging, since the model needs to learn which parameter vectors (E, ν, ρ) are plausible. Compare Figure 3 for the allowed resp. plausible base materials.

Evaluation of backprojection In Table 4 we provide further metrics for the experiment shown in Table 1a. One can see that the material fitting variance V_m is only slightly increased from the unguided case (see Table 5) for low to medium target bulk moduli and more strongly increased for higher targets. A similar observation holds for the distance to existing materials d_m .

In the following, we inspect the frac metric at several stages of the backprojection. The case "sample" refers to metrics computed over the generated sample directly (K_s in the main paper). The setting "no closest material" extracts parameters from the sample by fitting the material GMM and performing skeletonization. It evaluates these parameters with averaging over random spatial distributions of particles, identically to the main metrics. However it uses the extracted material parameters as-is and does look up the closest existing material. One can see that the metrics for "sample" and "no closest material" are similar for most targets, with a bigger difference for higher targets. This suggests that for high targets, more samples are generated whose spatial arrangement is not representative of the resulting bulk modulus. However, the "no closest material" metrics are still much higher than the final metrics after the material lookup. This indicates that the model can generate a high fraction of samples that are close to the target bulk modulus, but doing so interpolates materials, so that the used materials are not close to available ones.

864 To evaluate the backprojection, we also run it on our 10k 2D training dataset and report various
 865 metrics: For distance to closest material d_m , the mean, 99-th percentile and maximum are 1.3e-5,
 866 7.8e-14, 0.13. Only for one sample in the dataset, d_m was larger than 1e-5, which indicates a fail
 867 in recovering the correct material. For the absolute error between predicted and actual volume
 868 fraction, the mean, 99-th percentile and maximum are 1.3e-3, 0 and 0.5, respectively. Only in 42
 869 out of the 10000 samples, the volume fraction was determined with non-zero error. For the absolute
 870 error between predicted and actual radius of the particles, the mean, 99-th percentile and maximum
 871 are 6.6e-3, 1.5e-2, 0.2. For reference, for dataset generation, we sample radii uniformly between
 872 0.075 and 0.175. These results show that the backprojection can produce incorrect results on the
 873 clean training data, but does so only in very few cases. Also, the circle radius, which is difficult to
 874 determine, is estimated with low error.

875 C MODEL ARCHITECTURE DETAILS

876 We implement our approach in Pytorch (Ansel et al., 2024). Our diffusion model implementation is
 877 build on the `Unet2DModel` from the `Diffusers` library (von Platen et al., 2022). In the following,
 878 we detail the model architecture. If a setting is not specified, the default from the `Diffusers` library is
 879 assumed. First, the 3-channel input (normalized E, ν, ρ) is embedded by a convolutional layer with
 880 kernel size 1 into 8 channels. The current timestep, varying between 0 and 999, is embedded to 16
 881 channels with a Gaussian Fourier embedding (Tancik et al., 2020). It is then processed by a 2-layer
 882 MLP with `SILU` activation function to 32 channels.

883 The two down-blocks have output sizes 32 and 64. They process the input with two ResNet layers each,
 884 where their output sizes are equal to the whole block’s output size. The ResNet layers use a kernel
 885 size of 3 for each existing spatial dimension and employ a group normalization with constant number
 886 of groups 8. They use `swish` as nonlinearity. First, the group normalization and the nonlinearity
 887 is applied, followed by the first convolution. The timestep embedding is linearly projected to the
 888 output size and added to the hidden states, after which a second group normalization is applied. This
 889 is followed again by the nonlinearity and the second convolutional layer. Both convolutional layers
 890 use the same output size. Finally, this processed result is added to the input (residual connection).
 891 To achieve this, the input is first mapped by a convolution with kernel size 1 to match the output
 892 size. After both ResNet layers, the output is downsampled with an average pooling and kernel size
 893 and stride two for existing spatial dimensions. Due to the two down-blocks used, this results in a
 894 reduction of factor 4 for the spatial dimensions in the middle of the model.

895 After the down-blocks, the data is processed by a mid-block which features three ResNet layers with
 896 an attention layer between each. The hidden sizes are 128 and 128 and the last layer maps back to 64,
 897 as in the input to the mid-block. The attention layers use 16 heads and a 3-dimensional positional
 898 encoding³.

899 The up-blocks are build similarly and symmetrically to the down-blocks, only that they also take the
 900 output of the respective down-block as additional input (“skip connection”). The current result is
 901 concatenated with the output of the respective down-block and fed as input to each ResNet Layer
 902 (meaning, a ResNet layer by itself, which again consists of two convolutional layers). Before the input
 903 is passed to the respective ResNet layers, it is upsampled by 2 for each existing spatial dimension
 904 with nearest mode of `torch.nn.functional.interpolate`.

905 As last step in the model, the output of original spatial size and embedded channel size 8 is processed
 906 by a convolutional layer with kernel size 1 to project back to the original 3 data channels.

907 D HYPERPARAMETERS

912 In this section, we detail our model training and metrics which we used to optimize hyperparameters.
 913 For a definition of additional metrics used here, refer to Section B.

915 **Dataset size for unconditional generation** We train our model as described in the main paper for
 916 different dataset sizes of the 2D 64x64 problem. After training, we obtain 1000 samples with 1000

917 ³<https://github.com/tatp22/multidim-positional-encoding>

918
 919 Table 5: Comparison of different training dataset sizes for the 2D problem. 3 models with different
 920 seeds are trained for each row and 1000 samples generated each, averaged. Generation is unguided
 921 with 1000 diffusion steps.
 922

ds size	$V_m \downarrow$	$d_m \downarrow$	cov \uparrow
1k	0.0007	0.1074	0.9742
2.5k	0.0007	0.1137	0.9802
5k	0.0013	0.1162	0.9841
10k	0.0011	0.1164	0.9802

923 diffusion steps each and compute the mean metrics. These results are shown in Table 5. We see that
 924 all metrics are relatively similar between the considered dataset sizes, with a consistent tendency of
 925 slightly lower d_m for smaller sizes.
 926

927 **Training and diffusion hyperparameters** We obtained the specified hyperparameters by an
 928 empirical search. Initial experiments were conducted on models trained with a single seed on a
 929 32×32 problem and then parameter choices were refined on models trained with three seeds on a
 930 64×64 problem. From each trained model, 1000 samples were generated (without guidance) and
 931 evaluated according to the metrics specified above. Starting from default values provided by the
 932 framework, we iteratively searched by varying likely related parameters (e.g. β_0, β_T together with
 933 the type of β -schedule) and checking whether it improved upon our previous results. If there was no
 934 considerable improvement, we kept the previous parameters. We tried out constant and polynomial
 935 (exponent 0.5) learning rate schedules with different learning rates. For β -schedules, we tried out
 936 squared cosine and sigmoid schedules with varying β_0 and β_T . We also found the β -rescaling to be
 937 quite important for sample diversity according to our cov metric. As prediction targets, we compared
 938 predicting the noise ϵ , the clean sample x_0 and the velocity v and found that ϵ -prediction performed
 939 the worst and v prediction the best. We found larger batch sizes to perform similar to the baseline of
 940 128. Regarding training step sizes, we found diminishing improvements increasing the training steps
 941 between 50k, 100k and 200k steps.
 942

943 **Guidance parameters** We use the cov of $\epsilon_r < 5\%$ metric for tuning of the guidance parameters,
 944 averaged over the previously introduced subset of targets K^* . Guidance with DPS leaves the following
 945 parameters: The constant factor of the DPS gradient ρ_D (not to be confused with a density ρ), the
 946 number of denoising time steps N and the scaling factors of the gradients of E and ν . We first
 947 calibrated the scaling factors of the gradients with a DPS guidance only and DPS factor $\rho_D = 1$ in
 948 several hand-crafted trials. We then tried out several values of the ρ_D factor and found that 1 still
 949 performs best. We provide a comparison with $N = 50$ and 200 guidance steps in the main paper. In
 950 Figure 6, we show the effect of the two gradient scaling hyperparameters on the objective J_1 .
 951

952 We experimented with other means of scaling the gradients, for example normalizing by the distance
 953 to goal as proposed by Chung et al. (2023), clipping individual components or clipping the magnitude
 954 of the gradients. We also tried scaling by the predicted noise as proposed by Shen et al. (2024).
 955 However, we found that none of these methods provided better results than the individual scaling
 956 factors found as explained before.
 957

958 E DETAILS OF ALTERNATIVE METHOD EXPERIMENTS

959 **Bayesian optimization** We implement this method with the framework of Nogueira (2014). A
 960 Gaussian process is fit to evaluated data points and their objective value. The next data point to
 961 evaluate is determined via an upper confidence bound (UCB) formulation. Concretely, the term
 962 $\mu + \kappa \sigma$ is sought to be maximized. We conducted several hand-crafted experiments and found $\kappa = 1$
 963 to yield best results in terms of our cov metric in the $\epsilon_r < 5\%$ interval. Notably, finding the next data
 964 point to evaluate requires solving (even if not to optimality) an optimization problem, which takes
 965 significant time. This time is increased with the number of obtained data points. For our experiments,
 966 we found that between 750 and 1000 iterations (albeit more data points are added; see below) it takes
 967 roughly 10 seconds to suggest a new data point.
 968




Figure 6: Effect of gradient scaling hyperparameters on J_1 . Computed over 200 samples, averaged over 3 target bulk moduli. Black dots show evaluated data points, the red cross indicates the value used in our main experiments.

We represent the parameter space Θ continuously as subset of \mathbb{R}^8 . Since only several discrete material parameters are possible, we project the data point suggested by the acquisition function to closest existing materials, similarly to our backprojection. We then evaluate this parameter with random sampling of microstructures, identically to the evaluation of our method (compare Section 6.2) and assign the resulting value to both the suggested and actually evaluated data point. Note that the former is required to reduce the variance at the suggested point so that it will not be suggested again. Formally, the continuous space of material properties of matrix and particles can be represented as a Voronoi diagram, where the given set of points are the available materials. Each suggested point in a Voronoi region will be projected to the closest existing material and therefore has identical cost value. We experimented with sampling the Voronoi region after the first point in it is evaluated to add multiple data points to the Gaussian process, but found that the increased computation time for suggestion outweighs the potentially redundant evaluations. To obtain a parameter for evaluation, a value is first randomly sampled from the space and used as initial value for the maximization of μ (without an exploration term).

Conditionally trained diffusion models Firstly, we compute the bulk modulus K on all examples in the training set. To embed the bulk modulus, we first map the interval of occurring values to the interval $[0, 1]$. We then embed it similarly to the diffusion timestep and after processing with a dedicated MLP, both embeddings are added before being input to the convolutional layers. Following Ho & Salimans (2022), we trained models with a probability to replace the conditional input by a null embedding $p \in \{0, 0.1, 0.2\}$. For the models with non-zero probability, we tried out several values for the guidance scale w and found that the combination of $p = 0.1, w = 0.5$ performed best (using the scale formulation w as in the cited paper). In the main paper, we report results obtained with that parameter combination. Apart from that, we use identical settings for the diffusion model and sampling as for our approach (but not using any gradient guidance).

F ADDITIONAL RESULTS

Qualitative results of remaining targets We show the visualizations of remaining targets K^* in Figure 7 and Figure 8.

Runtime analysis We investigate execution times of our approach further in Table 6. Results are reported for a cluster node using 16 CPU cores and an A40 GPU. Three instances of the FEM solver run in parallel on CPU and each instance uses multiple threads. The dominating factor of the pipeline is the solution of the FEM problem and the computation of FEM gradients in the analytical solver ("FEM solve & diff").




Figure 7: Inverse 2D material designs. Generated samples for selected bulk moduli K^* , ordered by relative error quantile. Best on the left, worst on the right. Labels show $K_s / K_\theta / \epsilon_r$. The values (E, ν, ρ) in the normalized coordinate space are encoded as (r, g, b) values of the image.




Figure 8: Inverse 3D material designs. Generated samples for selected bulk moduli K^* , ordered by relative error quantile. Best on the left, worst on the right. Labels show $K_s / K_\theta / \epsilon_r$. The values (E, ν, ρ) in the normalized coordinate space are encoded as (r, g, b) values of the image.

Standard deviations of tables We provide the standard deviations of Table 1 and Table 3 in Table 7 and Table 9, respectively. One can see that the variance over model training seeds is relatively low, as evident from Table 7. This shows that the method is robust to training variability of the model without the need to adjust guidance parameters. The standard deviations in Table 9 are much larger and likely stem from the varying performance over material targets, which can be seen in Table 1. Regarding alternative methods in Table 8, one can see that the conditional diffusion model exhibits low variance over trained model seeds, similar to our guidance results. The Bayesian optimization approach exhibits highest variance over initial random states in most cases.

setting	batch size	total time	FEM solve & diff. (fraction)
2D 32x32	50	387 s	385 s (99.4%)
2D 64x64	50	1241 s	1236 s (99.5%)

Table 6: Runtime analysis of our approach. Results are reported for the generation of one batch of the specified size for J_1 , averaged over 5 target bulk moduli, 200 samples total. Guided diffusion with 100 steps (99 gradient computations per sample). The FEM gradient computation clearly dominates the runtime.

1080
 1081
 1082
 1083
1084 Table 7: Standard deviations of Table 1 (main results in 2D and 3D for different targets for J_1),
 1085 **computed over 3 model training seeds**

(a) 2D problem							(b) 3D problem						
	K^* :	4.8 [†]	86.6	168.5 [†]	250.4	332.2 [†]		K^* :	0.9 [†]	83.3	165.6 [†]	248.0	330.3 [†]
$\epsilon_r < 1\%$	frac	0.003	0.023	0.044	0.003	0.010	$\epsilon_r < 1\%$	frac	0.000	0.012	0.031	0.013	0.067
	cov	0.007	0.044	0.055	0.009	0.006		cov	0.000	0.036	0.065	0.018	0.007
$\epsilon_r < 5\%$	frac	0.016	0.019	0.013	0.010	0.028	$\epsilon_r < 5\%$	frac	0.005	0.058	0.049	0.028	0.151
	cov	0.018	0.012	0.021	0.018	0.034		cov	0.012	0.055	0.016	0.033	0.026

1094
 1095
 1096
 1097
 1098
 1099
 1100
1101 Table 8: Standard deviations of Table 2 (alternative methods), computed over 3 seeds
 1102

(a) Bayesian optimization						(b) Conditional diffusion model							
	K^* :	4.8	86.6	168.5	250.4	332.2		K^* :	4.8	86.6	168.5	250.4	332.2
$\epsilon_r < 1\%$	frac	0.003	0.019	0.044	0.088	0.060	$\epsilon_r < 1\%$	frac	0.003	0.010	0.030	0.035	0.026
	cov	0.007	0.019	0.058	0.061	0.035		cov	0.010	0.012	0.042	0.034	0.021
$\epsilon_r < 5\%$	frac	0.008	0.162	0.143	0.283	0.137	$\epsilon_r < 5\%$	frac	0.005	0.035	0.036	0.035	0.003
	cov	0.018	0.043	0.118	0.102	0.051		cov	0.006	0.023	0.033	0.003	0.018

1111
 1112
 1113
 1114
 1115
 1116
 1117
1118 Table 9: Standard deviations of Table 3 (ablations and guidance with J_2), computed over 3 model
 1119 **training seeds and all 5 target bulk moduli**

(a) Guidance by J_1 , ablations					(b) Guidance by J_2 (also minimize density)						
	$\epsilon_r < 1\%$	$\epsilon_r < 5\%$	$\epsilon_r < 1\%$	$\epsilon_r < 5\%$	ρ_{avg}	frac	cov	ρ_{avg}	frac	cov	
ds 10k	0.041	0.070	0.169	0.167	$\lambda = 0$	1.782	0.041	0.070	1.878	0.169	0.167
ds 1k	0.043	0.061	0.170	0.149	$\lambda = 10^{-4}$	1.120	0.039	0.067	1.115	0.179	0.167
$N = 50$	0.037	0.067	0.190	0.164	$\lambda = 10^{-3}$	1.057	0.009	0.016	1.142	0.050	0.036
$N = 200$	0.037	0.067	0.167	0.157	$\lambda = 10^{-2}$	0.845	0.012	0.015	0.672	0.050	0.026
unguided	0.003	0.007	0.021	0.040							
project	0.066	0.238	0.07	0.133							