

Random-Walk Microstructures for Differentiable Topology Optimization

Samuel Silverman
Boston University
USA
sssilver@bu.edu

Keith A. Brown
Boston University
USA
brownka@bu.edu

Dylan Balter
Boston University
USA
dbalter@bu.edu

Emily Whiting
Boston University
USA
whiting@bu.edu

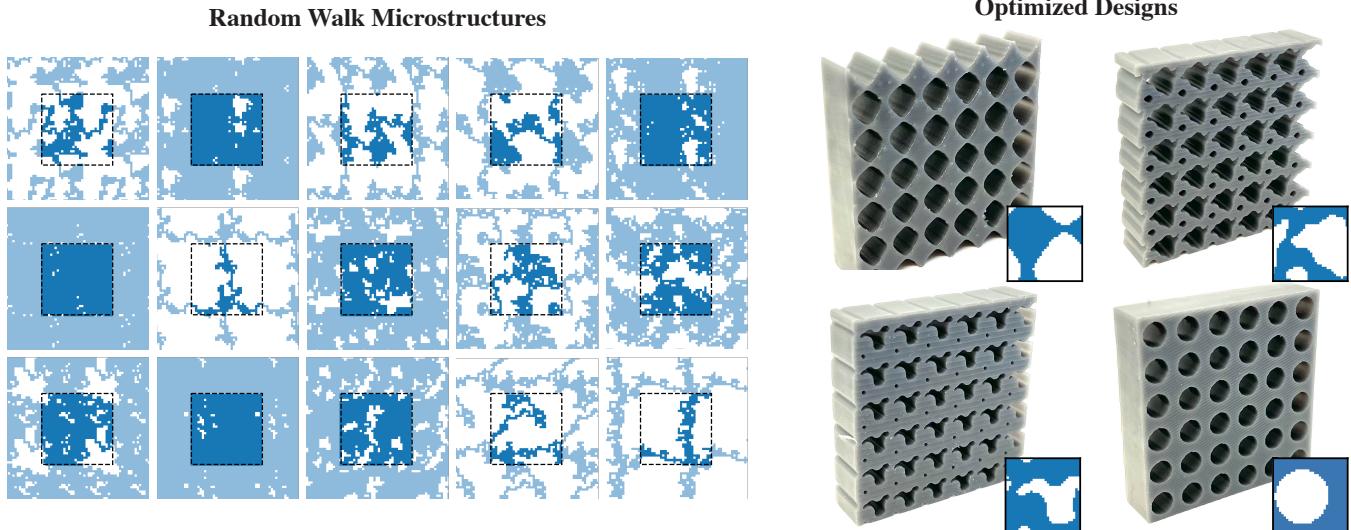


Figure 1: We create a dataset of random walk microstructures (left) to train a surrogate model that maps microstructure geometry to effective stiffness. This surrogate enables efficient inverse design via gradient-based topology optimization, producing printable architectures with targeted directional stiffness properties (right).

Abstract

This paper presents a differentiable pipeline for topology optimization of high-resolution mechanical metamaterials on grid domains, enabling complete geometric freedom within a fixed-resolution design space. Our method begins with a microstructure generation procedure based on random walks, which avoids hand-crafted parameterizations and populates the design space without strong geometric priors, yielding a diverse set of mechanically meaningful microstructures. We train a convolutional neural network to predict homogenized stiffness tensors from these microstructures,

enabling a fast and differentiable approximation of mechanical behavior without the need for finite element solves. By plugging this surrogate into a topology optimization loop, we can backpropagate through mechanical objectives and discover high-resolution, fabricable designs across a wide range of densities and target behaviors. We demonstrate our pipeline's inverse design capabilities, producing microstructures with both isotropic and anisotropic stiffness, and validate our predictions through mechanical testing.

CCS Concepts

- Applied computing → Computer-aided design; • Computer systems organization → Neural networks; • Computing methodologies → Physical simulation.

Keywords

microstructures, random walks, inverse design, neural networks, homogenization



This work is licensed under a Creative Commons Attribution 4.0 International License.
SCF '25, Cambridge, MA, USA
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2034-5/2025/11
<https://doi.org/10.1145/3745778.3766645>

ACM Reference Format:

Samuel Silverman, Dylan Balter, Keith A. Brown, and Emily Whiting. 2025. Random-Walk Microstructures for Differentiable Topology Optimization. In *ACM Symposium on Computational Fabrication (SCF '25), November 20–21, 2025, Cambridge, MA, USA*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3745778.3766645>

1 Introduction

Microstructures offer a powerful means of tailoring mechanical properties at the macroscale. By carefully designing the internal geometry of a periodic unit cell, it is possible to achieve effective stiffness [Surjadi et al. 2025], compliance [Huang et al. 2024], and anisotropy [Kumar et al. 2020] that would be impossible with homogeneous bulk materials. Such functionally engineered materials, commonly referred to as mechanical metamaterials, derive their properties not just from the base material itself but from the intricate spatial arrangement of voids and solids. Mechanical metamaterials are widely used in lightweight structural design, soft robotics, mechanical filtering, and energy absorption systems.

Creating these geometries with desired mechanical properties is a challenging inverse problem, and optimal designs often feature non-intuitive features. To handle this complexity, topology optimization (TO) [Bendsøe and Sigmund 1999] has emerged as a powerful tool, using gradients of mechanical objectives to evolve the geometry. However, traditional TO pipelines are computationally expensive. Each design iteration requires solving a partial differential equation (PDE), often via finite element methods (FEM), making large-scale or real-time exploration inefficient.

A promising remedy is to replace the forward FEM solve with a differentiable surrogate, typically a neural network that maps a candidate design to its homogenized mechanical properties and whose gradients can be back-propagated through the TO loop. Existing surrogates [Kumar et al. 2020; Li et al. 2023; Van 't Sant et al. 2023], however, are commonly trained on specific, parametric shape families, which limits their generalization to arbitrary designs. To support reliable learning across broader design spaces, it is essential to curate a large dataset that spans a wide range of physically plausible and fabricable microstructures without introducing strong geometric biases.

In this work, we introduce a random-walk growth procedure that produces a large, bias-free set of binary microstructures spanning the full density spectrum and a wide range of mechanical responses. Leveraging this dataset, we train a periodic convolutional neural network (CNN) surrogate that can be directly inserted into gradient-based topology-optimization loops on regular grids, enabling truly free-form, high-resolution design beyond the limitations of shape-parametric surrogates.

2 Related Work

2.1 Metamaterial Design

Mechanical metamaterials derive their effective properties from geometry rather than their underlying material composition. To explore and exploit this structure–property relationship, researchers have developed a wide variety of parameterized microstructure

families. These parameterizations serve two key roles: enabling expressive control over geometry and facilitating efficient exploration of mechanical responses.

Truss-based parameterizations offer a lightweight, physically intuitive scaffold for designing microstructures with tunable elasticity. Researchers combine truss networks with connectivity graphs to span a diverse range of Poisson's ratios and moduli [Panetta et al. 2015; Tozoni et al. 2020]. Huang et al. [2024] explore shock-absorbing lattices generated by modifying 2D truss connectivity. Bastek et al. [2022] construct trusses via composable topologies and affine deformations to enable shear coupling. Researchers have combined truss and woven networks to emulate dual-polymer systems [Surjadi et al. 2025] and introduced programmable digital weaves with sliding yarn-like connections [Li et al. 2022].

Another prominent strategy is to represent microstructures through tiling patterns. Isohedral tilings form a compact yet expressive family, with geometric degrees of freedom that modulate the shape of the unit cell. Schumacher et al. [2018] and Li et al. [2023] both use isohedral tilings to generate diverse families with varying elastic responses. Deng et al. [2022] use a set of control points to perturb a base unit tile. Martínez et al. [2019] introduce tilings via generalized Voronoi diagrams, with Efremov et al. [2021] extending the process to 3D.

Alternatively, researchers have explored foam-like microstructures using either procedural generation [Martínez et al. 2016, 2018, 2017] or through phasor-noise-based textures [Tricard et al. 2020]. Implicit methods, such as spinodoids [Kumar et al. 2020], trade explicit geometry for fields: a few shape parameters modulate complex topologies via level sets or random fields, offering both smoothness and diversity. Makatura et al. [2023] introduce a unified procedural graph framework that generalizes many existing parameterizations, including beams, shells, and triply periodic surfaces, into a compact, composable system for metamaterial design. Microstructures defined on voxel grids [Schumacher et al. 2015] give fine-grained control and a natural fit for topology optimization, breaking free from geometric parameterization. In this work, we adopt a grid-based design representation, which offers the highest geometric flexibility within a fixed resolution.

2.2 Homogenization

Directly simulating a fine-scale microstructure is computationally intractable. To address this, homogenization provides a principled approach to approximating the behavior of a periodic material by computing its effective mechanical properties. We adopt a classic homogenization formulation based on the method of Andreassen and Andreasen [2014] to obtain effective stiffness behavior for our designs. This technique remains the dominant approach for linear elastic regimes. Other approaches extend homogenization to more complex scenarios; Kharevych et al. [2009] introduce a mesh coarsening technique that aggregates fine-scale elastic heterogeneity into effective coarse-scale elements.

When the constituent material is nonlinear, traditional linear homogenization is insufficient. Recent methods address this with nonlinear homogenization frameworks. Li et al. [2023] run a homogenization procedure at various directions and magnitudes, using nonlinear material properties to construct an effective nonlinear

behavior. Huang et al. [2024] develop a nonlinear homogenization simulation accounting for self-contact between microstructural features during large deformations. Sperl et al. [2020] built a nonlinear homogenization procedure for woven and knitted fabrics. Zhang et al. [2023a] created an adaptive interpolant for homogenization at finite strains.

2.3 Inverse Design

Traditional approaches to inverse design define a differentiable objective function over the microstructure geometry and use the homogenization simulation’s gradients to optimize the design iteratively. These gradients are often computed via adjoint methods instead of finite differences, which enables efficient backpropagation through expensive simulations. Panetta et al. [2015] perform shape optimization of periodic truss networks by differentiating the homogenization procedure itself. Schumacher et al. [2015] use topology optimization where the material layout is relaxed to continuous densities and the effective properties are optimized via simulation gradients. Huang et al. [2024] develop a contact-aware nonlinear simulation pipeline and apply gradient-based optimization. In contrast to gradient-based methods, Martínez et al. [2016] adopt a sampling-based approach, fitting a low-order polynomial to a dataset of designs, bypassing the need for differentiable simulation altogether.

More recently, machine-learning-based surrogate models have emerged as a promising alternative to traditional simulation-driven pipelines. These models approximate the homogenization map directly, predicting the effective mechanical properties of a given microstructure design. The key advantage is that once trained, surrogates replace FEM solves, and their differentiability makes them ideal for integration into inverse design loops. Rao and Liu [2020] train a neural network to predict homogenized stiffness tensors from voxelized inputs, although they do not utilize it for design purposes. We develop a similar model for use in downstream gradient-based topology optimization. Li et al. [2023] train a network to predict the strain energy density field for a given design, enabling gradients of the predicted energy to yield effective stiffness tensors. Peng et al. [2022] train a CNN that predicts the displacement fields of parallelepiped microstructures used for homogenization calculations. Li et al. [2024] propose a data-driven discovery pipeline that uses a CNN surrogate, allowing for optimization over trade-offs between competing objectives. Tandem learning strategies [Kumar et al. 2020; Silverman et al. 2025; Van ’t Sant et al. 2023], where one network maps design parameters to effective properties while a second learns the reverse mapping, enable inverse design without requiring a differentiable simulation.

3 Method

We develop a fast, differentiable surrogate model that maps microstructure designs to their elastic properties. To train this model, we generate a dataset of valid periodic microstructures paired with ground-truth stiffness tensors computed via finite element homogenization (Sections 3.1 and 3.2). To improve generalization, we apply symmetry-based data augmentation, exploiting the periodic and geometric invariances of the domain (Section 3.3). Finally, we present the architecture of our surrogate model (Section 3.4).

3.1 Random Walk Microstructures

3.1.1 Microstructure Growth. We require a large and diverse dataset of valid microstructure designs to train a model that generalizes across a broad range of physical behaviors. In particular, we seek structures that are:

- (1) **Physically meaningful:** Solid material should form a single connected component, forming one continuous material.
- (2) **Expressive:** The dataset should encompass a range of material fractions and spatial configurations.
- (3) **Unbiased:** The designs should not be restricted to a narrow parametric family nor confined to the output of a specific optimization algorithm.

One strategy for generating grid-based microstructures is to use parametric shape families, such as lattices, trusses, or tiled primitives, that can be procedurally varied to produce a collection of designs. While effectively constraining the space to fabricable and connected structures, such parameterizations inherently bias the dataset toward specific geometric motifs and may limit generalization. At the opposite end, purely random sampling methods (e.g., assigning each pixel as solid or void independently via Bernoulli draws) are unbiased. However, even with post-selection to enforce a target volume fraction, these randomly generated patterns typically lack global connectivity and are poorly suited for mechanical tasks.

To address these limitations, we introduce a simple and effective *random walk growth* procedure (Figure 2), which produces connected, diverse microstructures while offering direct control over the total amount of material.

Given a grid of resolution $l_x \times l_y$, we construct each binary microstructure as follows:

- (1) Select a target number of filled cells $m \in [l_x + l_y - 1, l_x l_y]$. The lower bound ($l_x + l_y - 1$) corresponds to the minimum number of steps required to span the domain in a valid connected path. The upper bound ($l_x l_y$) yields a fully solid design.
- (2) Randomly choose a grid cell to initialize the walk.
- (3) Perform a random walk over the grid. At each step, the walk moves to a random neighboring cell (up, down, left, or right) and marks it as solid if it has not already been visited. We allow the walk to wrap around grid edges, which does not violate the material connectivity requirements when tiling the microstructure. The process stops once m distinct cells are filled.

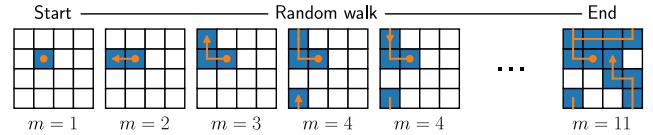


Figure 2: A microstructure is generated by performing a random walk over a 4×4 grid, growing until a target number of filled cells (11) is reached. Because each new cell is adjacent to the existing cluster, counting wraparound adjacency, this procedure guarantees *local connectivity* in the periodically tiled design.

3.1.2 Microstructure Verification. The random walk procedure ensures local connectivity as solid pixels form a continuous path, assuming seamless tiling across boundaries. However, the procedure does not guarantee global connectivity of the resulting periodic structure. For example, a walk that reaches only the horizontal edges may appear connected under local wrapping yet still produce disconnected components when the unit cell is tiled in the plane.

To address this, we introduce a global connectivity *verification procedure* (Figure 3) that checks for strong connectivity across the tiled designs. The steps are as follows:

- (1) Tile the candidate design into a 2×2 array. This ensures that all periodic boundary transitions are represented.
- (2) Choose a random solid cell within the tiled domain as the starting point.
- (3) Perform a flood fill to explore all reachable solid cells. The fill wraps at the domain boundaries, mimicking the effect of infinite periodic tiling, and continues until no further material can be reached.

A design passes the verification check if the flood fill reaches all solid cells in the tiled domain. These structures are considered strongly connected and are retained for the dataset.

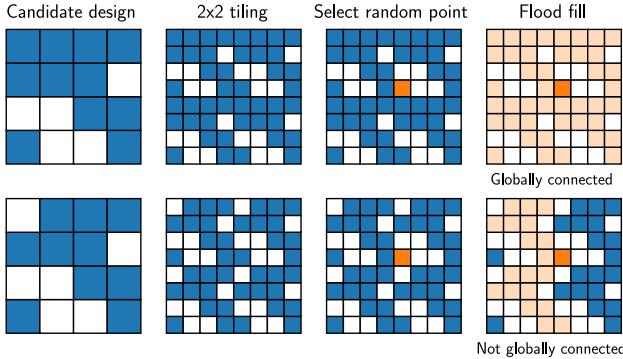


Figure 3: A valid microstructure exhibits *global connectivity*, remaining strongly connected after 2×2 tiling and flood fill (top). Removing a single pixel yields a design that maintains local connectivity but fails the verification procedure (bottom).

Using this two-stage pipeline, growth followed by verification, we generate a dataset of 100,000 distinct microstructure designs on a 32×32 grid (Figure 1). If a valid sample cannot be found within 1000 attempts for a given target number of filled cells m , we resample m . We discard duplicate designs using a hash of each binary pattern. The process is trivially parallelizable and supports efficient large-scale sampling. Figure 4a shows our dataset covering a wide range of effective stiffness values, spanning four orders of magnitude in relative E_x across the density range. The number of valid microstructures drops off sharply near the theoretical minimum density as these structures leave little room for random deviations, while the rest of the distribution is approximately uniform.

3.2 Homogenization

Simulating the complete heterogeneous microstructure for every load case is prohibitively expensive. Homogenization offers a practical alternative by replacing the complex, spatially varying elastic behavior with an equivalent, uniform approximation that captures the material's average response.

In a heterogeneous linear-elastic material, the local stiffness tensor $C(x)$ relates the infinitesimal strain

$$\varepsilon(x) = \frac{1}{2}(\nabla u(x) + \nabla u(x)^\top)$$

to the stress $\sigma(x)$ via

$$\sigma(x) = C(x) : \varepsilon(x),$$

where $u(x)$ is the displacement field. In 2D Voigt notation, this reads

$$\begin{bmatrix} \sigma_{11} \\ \sigma_{22} \\ \sigma_{12} \end{bmatrix} = \begin{bmatrix} C_{11} & C_{12} & C_{13} \\ C_{12} & C_{22} & C_{23} \\ C_{13} & C_{23} & C_{33} \end{bmatrix} \begin{bmatrix} \varepsilon_{11} \\ \varepsilon_{22} \\ 2\varepsilon_{12} \end{bmatrix}.$$

C_{11} and C_{22} correspond to axial stiffness in the x - and y -directions, respectively. C_{33} represents shear stiffness, and C_{12} , C_{13} , and C_{23} represent coupling terms.

$C(x)$ exhibits sharp discontinuities between solid and void regions in the microstructure domain, leading to highly heterogeneous stress-strain relationships. Homogenization aims to approximate such materials with a constant, effective stiffness tensor \bar{C} , such that the average stress under any macroscopic strain $\bar{\varepsilon}$ satisfies

$$\langle \sigma \rangle = \bar{C} \bar{\varepsilon},$$

where $\langle \cdot \rangle$ denotes the volume average over the unit cell Ω .

We adopt the standard periodic-cell homogenization approach following Andreassen and Andreassen [2014], involving three main steps:

- (1) Prescribe three independent macroscopic strain fields $\bar{\varepsilon}^{(i)}$ (in Voigt form)
- $\bar{\varepsilon}^{(1)} = [1, 0, 0]^\top, \quad \bar{\varepsilon}^{(2)} = [0, 1, 0]^\top, \quad \bar{\varepsilon}^{(3)} = [0, 0, 1]^\top,$

corresponding to uniaxial strain in x , uniaxial strain in y , and pure shear strain, respectively.

- (2) For each macroscopic strain, solve the boundary value problem for the corresponding fluctuation displacement field $\tilde{u}^{(i)}(x)$:

$$\nabla \cdot (C(x)(\bar{\varepsilon}^{(i)} + \tilde{\varepsilon}^{(i)}(x))) = 0 \text{ in } \Omega, \\ \tilde{u}^{(i)}(x) \text{ is periodic on } \partial\Omega.$$

Here, $\tilde{\varepsilon}^{(i)}(x)$ is the strain of the fluctuation field. The fluctuation field accounts for local deviations from the imposed macroscopic strain due to material heterogeneity. While periodic boundary conditions eliminates all rigid-body rotations and non-uniform translations, the fluctuation field remains ambiguous up to a global constant shift. We therefore fix all displacement components at a single node to anchor the system and ensure uniqueness.

- (3) Compute the homogenized stiffness tensor

$$\bar{C}_{ij} = \frac{1}{|\Omega|} \int_{\Omega} (\bar{\varepsilon}^{(i)} + \tilde{\varepsilon}^{(i)}(x))^\top C(x)(\bar{\varepsilon}^{(j)} + \tilde{\varepsilon}^{(j)}(x)) dV,$$

where $|\Omega|$ is the unit cell area.

We use eight-node quadrilateral elements for each grid cell, avoiding the need for meshing. Solid pixels are modeled as linear isotropic materials under plane strain conditions, with a Young's modulus $E = 2100$ MPa and Poisson's ratio $\nu = 0.39$, corresponding to Formlabs Tough 2000 resin [Decker et al. 2023; Formlabs 2025]. Void pixels are assigned a stiffness value 10^{-10} times the stiffness of solid pixels to maintain numerical stability during solves. Figure 4b shows the Ashby plot of the dataset's homogenized properties.

We enforce periodic boundary conditions by coupling degrees of freedom on opposing edges. The global stiffness matrix is assembled using Eigen in C++, and each linear system is solved using an LDL^\top factorization. On our hardware, the solve time per microstructure is approximately 2.2 seconds. Incorporating optimized GPU solvers [Zhang et al. 2023b] could further improve simulation times.

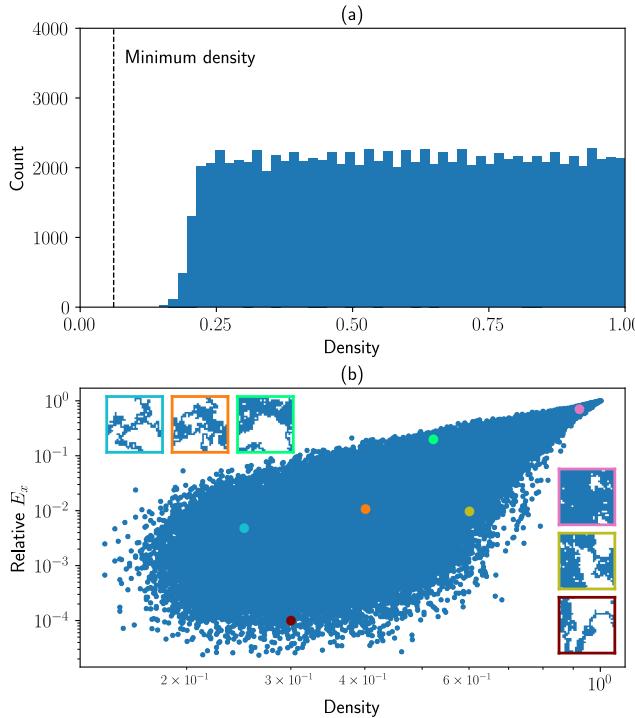


Figure 4: (a) Distribution of volume fractions across dataset samples. While coverage is approximately uniform for most densities, the count drops off sharply near the theoretical minimum since such designs are hard to achieve stochastically. (b) Ashby plot of relative x -directional Young's modulus versus density showing that our dataset spans over four orders of magnitude in stiffness across the full density range. Color-framed insets highlight representative microstructures at different density-stiffness regimes.

3.3 Data Augmentation

We augment our dataset by exploiting the inherent symmetries of periodic, grid-based microstructures (Figure 5). These symmetries fall into two categories:

- **Invariances**, where a transformation leaves the homogenized stiffness tensor \bar{C} unchanged.
- **Equivariances**, where \bar{C} transforms in a predictable, structure-preserving way.

By applying these transformations on the fly during training, we expose the model to a wider variety of patterns and orientations without requiring additional FEM solves.

3.3.1 Translation Invariance. Because each microstructure is defined as a periodic unit cell, any integer grid shift results in the same effective properties. Let $\rho \in \{0, 1\}^{32 \times 32}$ denote the binary material layout. For any offset $a, b \in \{0, \dots, 31\}$, the translated design

$$\rho'_{ij} = \rho(i+a) \bmod 32, (j+b) \bmod 32$$

satisfies $\bar{C}' = \bar{C}$. This translation gives 1024 invariances per sample.

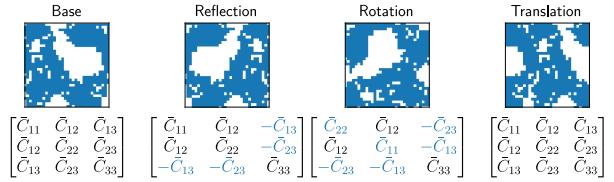


Figure 5: A base microstructure and its transformed variants via reflection, rotation, and translation. Each transformation corresponds to a specific change (or invariance) in the homogenized stiffness tensor \bar{C} . Reflections and 90° rotations induce signed permutations of the tensor components, while periodic translations leave \bar{C} unchanged.

3.3.2 Rotation and Reflection Equivariance. Unlike translations, which leave the homogenized tensor \bar{C} unchanged, rotations and reflections modify the orientation of material features, causing \bar{C} to transform predictably.

We restrict our symmetry operations to rotations by 90° intervals and reflections about principal axes. These transformations form the 2D dihedral group D_4 , which includes:

- rotations a^k by $90^\circ \times k$ (for $k = 1, 2, 3$),
- and reflections b, ba, ba^2, ba^3 across vertical, diagonal, and horizontal axes.

We use these transformations because they are grid-aligned and can be applied exactly to binary images without interpolation or aliasing.

Given a transformation matrix $R \in \mathbb{R}^{2 \times 2}$, the full fourth-order stiffness tensor transforms as:

$$C'_{mnop} = R_{im} R_{jn} R_{ko} R_{lp} C_{ijkl}.$$

This expression induces a signed permutation of entries in the homogenized Voigt-form tensor:

- Reflection about the y -axis (b):

$$\bar{C}_{13} \mapsto -\bar{C}_{13},$$

$$\bar{C}_{23} \mapsto -\bar{C}_{23}.$$

- Counterclockwise rotation by 90° (a):

$$\bar{C}_{11} \leftrightarrow \bar{C}_{22},$$

$$\bar{C}_{13} \mapsto -\bar{C}_{23},$$

$$\bar{C}_{23} \mapsto -\bar{C}_{13}.$$

Combining 1024 periodic translations with eight rotation/reflection variants gives 8192 symmetry-derived versions of every microstructure. Thus, our dataset expands to over 8×10^8 samples. We implement this augmentation stochastically during training: a random symmetry transformation is applied to both the design and its label at each batch. This strategy exposes the network to diverse inputs while preserving physical and geometric consistency.

3.4 Surrogate Neural Network

To eliminate the computational bottleneck of finite element homogenization, we train a data-driven surrogate model to predict effective material properties from binary microstructures. Specifically, we learn a mapping $f : \{0, 1\}^{32 \times 32} \rightarrow \mathbb{R}^6$ that predicts the six independent components of the homogenized stiffness tensor:

$$f(\rho) \approx \bar{c} = [\bar{C}_{11}, \bar{C}_{22}, \bar{C}_{33}, \bar{C}_{12}, \bar{C}_{13}, \bar{C}_{23}]^\top.$$

To improve numerical conditioning during training, we apply a transformation to each component of \bar{c} . For the diagonal terms \bar{C}_{ii} , which are always non-negative, we apply a standard logarithmic transform $\log(1 + \bar{C}_{ii})$. For the off-diagonal terms \bar{C}_{ij} , which may be negative, we use a sign-preserving log transform $\log(1 + |\bar{C}_{ij}|) \cdot \text{sign}(\bar{C}_{ij})$. Each transformed component is then standardized to have a mean of zero and a variance of one using statistics computed over the training set.

Figure 6 illustrates the structure of our CNN, comprising two components: a convolutional feature extractor and a multilayer perceptron (MLP) regressor.

In the feature extractor, the binary input is first extended with 6 pixels of periodic padding on each side. This padding ensures that the first convolutional layer correctly accounts for the periodic boundary conditions of the microstructure. The padded input then passes through six convolutional blocks with Swish activations [Ramaiah et al. 2017]. Every second block applies stride-2 downsampling, reducing spatial resolution while increasing channel capacity. The final feature map is flattened and passed to an MLP consisting of three hidden layers (each of width 256) with Swish activations. The last linear layer outputs the six predicted stiffness components.

4 Results

4.1 Training

We train the surrogate network using a mean squared error (MSE) loss between predicted and simulated homogenized stiffness tensors:

$$\mathcal{L} = \frac{1}{6n} \sum_{i=1}^n \sum_{j=1}^6 (\bar{c}_j^{(i)} - f(\rho^{(i)})_j)^2.$$

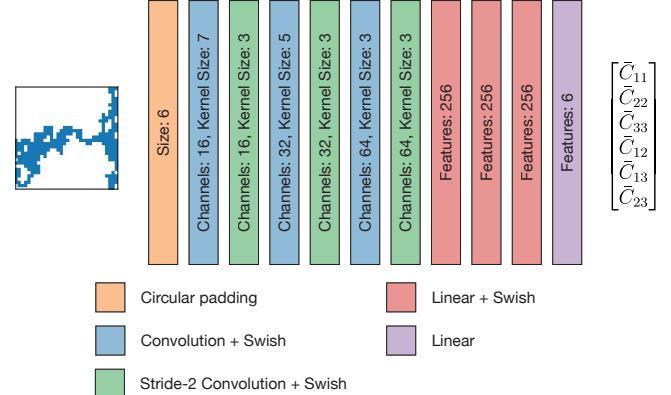


Figure 6: The architecture of our surrogate convolutional neural network (CNN). A binary 32×32 microstructure is first periodically padded (orange), then passed through a series of convolutional layers with Swish activations (blue), with downsampling every second block (blue). The resulting feature map is flattened and fed into a multilayer perceptron with three hidden layers (red), culminating in a linear output layer (purple) that predicts the six independent components of the homogenized stiffness tensor.

While prior work has explored incorporating physics-based loss terms for hyperelastic materials, such as strain energy density [Han et al. 2024; Li et al. 2023], we found that such additions offered no meaningful improvement in accuracy or generalization for our linear-elastic dataset. As a result, we use the simplest possible formulation: direct regression via MSE.

We train the network using the Adam optimizer [Kingma and Ba 2017] with a learning rate of $1e^{-4}$, a batch size of 64, and early stopping based on the validation loss over 1000 epochs. We split the dataset into 90% training, 5% validation, and 5% test sets. Data augmentations (Section 3.3) are applied randomly per batch during training. Our implementation is in PyTorch [Paszke et al. 2019], with training completed in approximately 1.5 hours on a MacBook M1, and inference requiring approximately 15 ms per sample.

To contextualize prediction errors, we report RMSE values as a percentage of the base material modulus (2100 MPa). On the held-out test set, the network achieves an RMSE of 38.8 MPa, or a relative error of 1.8%, and exhibits a high correlation with the simulated results. Figure 7 shows parity plots comparing CNN predictions and FEM-computed values for each independent stiffness component. The model achieves excellent accuracy across all terms, with R^2 values above 0.99 for the major components (\bar{C}_{11} , \bar{C}_{22} , \bar{C}_{33} , and \bar{C}_{12}) and slightly lower but still strong agreement on the off-diagonal shear coupling terms (\bar{C}_{13} and \bar{C}_{23}).

To assess the surrogate’s ability to generalize beyond the training distribution, we evaluate its performance on four structured microstructures that differ substantially from the random walk-based dataset (Figure 8): A truncated square grid, an H-shaped cutout, a perforated layout with circular voids, and a fractal pattern. These designs reflect the structured geometries commonly encountered in topology-optimized or analytically derived materials.

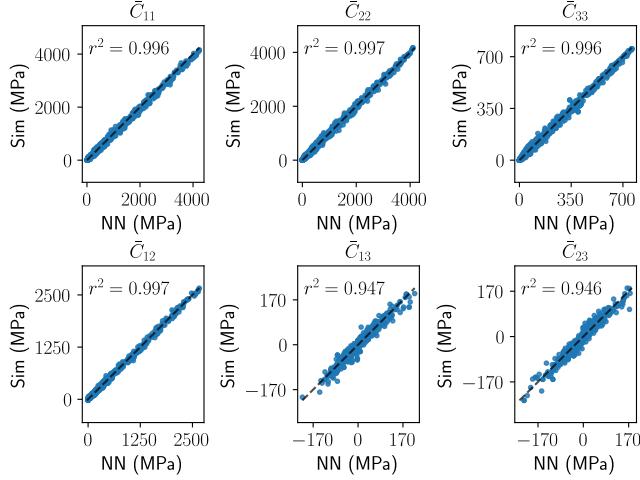


Figure 7: Parity plots comparing neural network predictions and FEM-computed stiffness tensor components from the test set.

The truncated square, perforated, and H-shaped designs exhibit relative errors of 2.4%, 2.3%, and 4.2%, respectively, comparable to or slightly above the test set average. The fractal pattern yields a larger error of 16.5%, likely due to its sharp corners and thin features, which are underrepresented in the training distribution. Despite this, all designs show strong agreement with FEM predictions, demonstrating the surrogate’s ability to generalize meaningfully beyond the random-walk distribution.

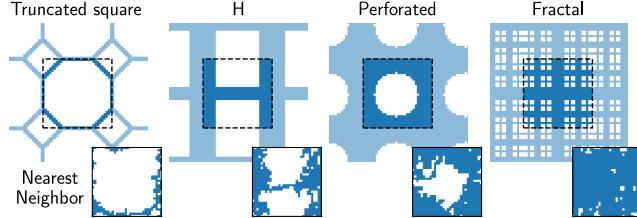


Figure 8: Four structured microstructures (top) used to assess the surrogate model’s ability to generalize beyond the training distribution. Each is paired with its nearest neighbor from the training set based on Hamming distance (bottom). The dashed box indicates the base unit cell before tiling.

4.2 Topology Optimization

Our TO framework seeks to identify a binary material layout $\rho_{ij} \in \{0, 1\}$ that minimizes a mechanical performance objective $\mathcal{J}(\rho)$ while adhering to physical constraints.

Direct optimization over binary designs leads to a highly non-convex combinatorial problem. A common workaround is to relax the design domain to continuous densities $\rho_{ij} \in [0, 1]$ and then apply filtering, penalization, or projection to steer the solution toward binary values. While effective, gradient-based optimization

methods introduce nontrivial heuristics and require access to *simulation derivatives*, i.e., the sensitivity of the effective stiffness tensor to changes in the density field:

$$\frac{\partial \mathcal{J}}{\partial \rho} = \frac{\partial \mathcal{J}}{\partial \bar{C}} \frac{\partial \bar{C}}{\partial \rho}$$

Computing $\frac{\partial \bar{C}}{\partial \rho}$ typically requires differentiating through the FEM pipeline using adjoint methods, which can be complex to implement and expensive to evaluate.

Our surrogate model enables efficient and straightforward gradient-based optimization (Figure 9). Because the network f is fully differentiable, the sensitivity $\frac{\partial f}{\partial \rho} \approx \frac{\partial \bar{C}}{\partial \rho}$ is obtained automatically via backpropagation. This allows us to define objectives in terms of the predicted stiffness tensor and optimize directly using standard deep learning optimizers.

Although the surrogate is trained on binary microstructures, we optimize over a continuous density field. To reconcile this, we use a straight-through estimator (STE):

- During the forward pass, each pixel is thresholded at 0.5 to produce a binary image (i.e. $\rho_{ij}^{\text{binary}} = \mathbb{1}[\rho_{ij}^{\text{cont}} \geq 0.5]$).
- During the backward pass, gradients are propagated through the continuous field as if the thresholding were the identity function.

To ensure valid pixel values, we apply a sigmoid activation to constrain the density field to the range $[0, 1]$. The STE lets us maintain continuous optimization’s expressive power and stability while producing valid binary layouts compatible with our surrogate’s training distribution.

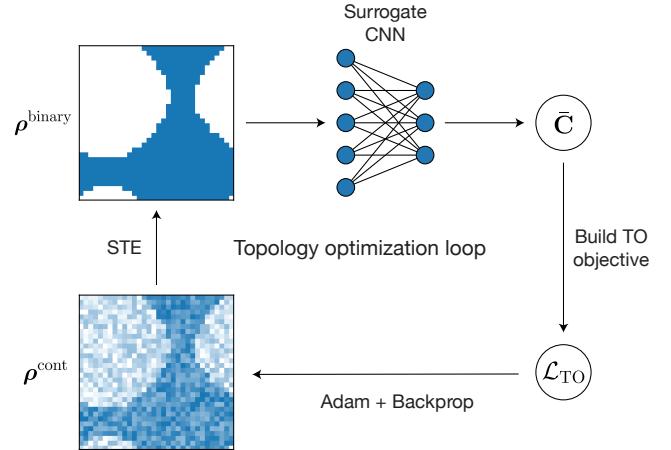


Figure 9: In our topology optimization (TO) loop, a continuous density field ρ^{cont} is passed through a straight-through estimator (STE) to produce a binary microstructure ρ^{binary} . The binary layout is evaluated by the surrogate CNN, which predicts the homogenized stiffness tensor \bar{C} . We assemble the total TO loss \mathcal{L}_{TO} using \bar{C} and regularizers. Adam uses the gradients from backpropagation to update ρ^{cont} , closing the loop.

4.2.1 Objective Functions. We evaluate the performance of our surrogate in TO using two objectives. The first objective promotes overall isotropic stiffness by maximizing the directional moduli in both x - and y -directions:

$$\mathcal{J}_{\text{Iso}}(\rho^{\text{binary}}) = \bar{S}_{11} + \bar{S}_{22},$$

where $\bar{S} = \bar{C}^{-1}$ is the homogenized compliance tensor reconstructed from the surrogate output $\bar{C} = f(\rho^{\text{binary}})$. Here, $\bar{S}_{11} = 1/\bar{E}_x$ and $\bar{S}_{22} = 1/\bar{E}_y$ represent the inverse of the Young's moduli along the x and y directions, respectively. Although \mathcal{J}_{Iso} does not explicitly enforce $\bar{E}_x = \bar{E}_y$, maximizing the sum of both directional stiffnesses naturally drives the design toward layouts where neither direction is sacrificed, encouraging a near-isotropic response.

The second objective encourages directional stiffness anisotropy. We define a soft constraint that penalizes deviations from a target ratio $E_x/E_y = 3$:

$$\mathcal{J}_{\text{Ortho}}(\rho^{\text{binary}}) = \frac{1}{2} \left(\frac{\bar{S}_{22}}{\bar{S}_{11}} - 3 \right)^2.$$

This objective encourages designs that are three times stiffer along the x -axis than the y -axis without requiring explicit directional filtering or symmetry constraints.

4.2.2 Density Penalty. To constrain material usage during optimization, we add a quadratic penalty that encourages the design to match a target volume fraction $d \in [0, 1]$. The density regularization term is:

$$\mathcal{R}_d(\rho^{\text{binary}}) = \lambda_d \left(\frac{1}{32^2} \sum_{i=1}^{32} \sum_{j=1}^{32} \rho_{ij}^{\text{binary}} - d \right)^2,$$

where $\lambda_d \in \mathbb{R}$ is a user-defined weight that controls the strength of the penalty for deviating from the target density.

4.2.3 Smoothness Penalty. Although the surrogate is trained only on valid, globally connected microstructures, the optimization process may still yield disconnected or noisy artifacts when starting from continuous densities. To suppress such patterns and promote cleaner designs, we introduce a smoothness regularizer inspired by discrete Laplacian energy. Let κ denote the 3×3 Laplacian kernel over the Moore neighborhood:

$$\kappa = \begin{bmatrix} 1 & 1 & 1 \\ 1 & -8 & 1 \\ 1 & 1 & 1 \end{bmatrix}.$$

We apply this kernel to the continuous density field ρ^{cont} using periodic padding and define the smoothness penalty as:

$$\mathcal{R}_s(\rho^{\text{cont}}) = \frac{\lambda_s}{32^2} \sum_{i=1}^{32} \sum_{j=1}^{32} (\kappa * \rho^{\text{cont}})_{ij}^2,$$

where $\lambda_s \in \mathbb{R}$ is a user-defined weight that controls the strength of smoothing in the design.

4.2.4 Implementation. The final loss is defined as the sum of the objective function, the density penalty, and the smoothness penalty:

$$\mathcal{L}_{\text{TO}} = \mathcal{J}(\rho^{\text{binary}}) + \mathcal{R}_d(\rho^{\text{binary}}) + \mathcal{R}_s(\rho^{\text{cont}}).$$

We use the Adam optimizer to minimize this loss for each objective across target densities $d \in \{0.5, 0.6, 0.7, 0.8\}$ (Figure 10). The

continuous density field is initialized to $\rho^{\text{cont}} = 0$, which becomes a uniform density of 0.5 after applying the sigmoid activation.

For the isotropic stiffness objective \mathcal{J}_{Iso} , we run the optimization for 10000 epochs. The regularization weights are linearly annealed over time: $\lambda_d : 0.01 \rightarrow 1000$ and $\lambda_s : 1 \rightarrow 1000$. For the orthotropic stiffness objective $\mathcal{J}_{\text{Ortho}}$, we run the optimization for 1000 epochs with more aggressive smoothness regularization early on: $\lambda_d : 0.01 \rightarrow 1000$ and $\lambda_s : 100 \rightarrow 1000$.

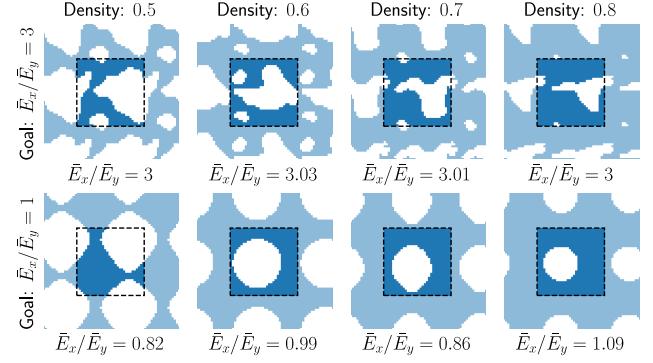


Figure 10: Optimized microstructures from the surrogate-based TO pipeline at varying target densities (0.5, 0.6, 0.7, and 0.8). The relative directional Young's moduli are reported below each structure. Designs that targeting anisotropy $E_x/E_y = 3$ are shown in the top row while designs that maximize isotropic stiffness $E_x/E_y = 1$ are shown in the bottom row.

4.3 Fabrication

We fabricate all eight optimized microstructures from the topology optimization results. To prepare each design for printing, we developed a post-processing pipeline that converts the binary unit cell into a watertight, printable mesh (Figure 11).

Starting with the binary design, we tile the unit cell in the plane (6 repetitions in both x - and y -directions) to produce a periodic structure. To mitigate aliasing and jagged edges caused by the discrete grid, we apply Gaussian smoothing to a $2 \times$ upsampled version of the tiled image. We then extract the 0.5 level set of the smoothed field, yielding a clean, closed boundary curve. This 2D shape is triangulated and extruded to a fixed thickness, producing a solid 3D mesh suitable for fabrication.

All samples are printed using a Formlabs Form 3+ stereolithography (SLA) printer with Tough 2000 resin (Figure 12). Each fabricated sample has final dimensions of $50\text{mm} \times 50\text{mm} \times 12.5\text{mm}$. To obtain experimental measurements of the directional Young's moduli, we perform uniaxial compression testing using an Instron 5965 universal testing machine in both axial directions (Figure 13). The resulting stress-strain data is used to estimate a Young's moduli ratio, which we compare against predictions from both the surrogate and full FEM simulation.

We observe good agreement across all three sources, confirming that our optimized designs not only meet their anisotropy targets in simulation but also translate accurately to physical behavior. The largest deviation occurs for the orthotropic design at density

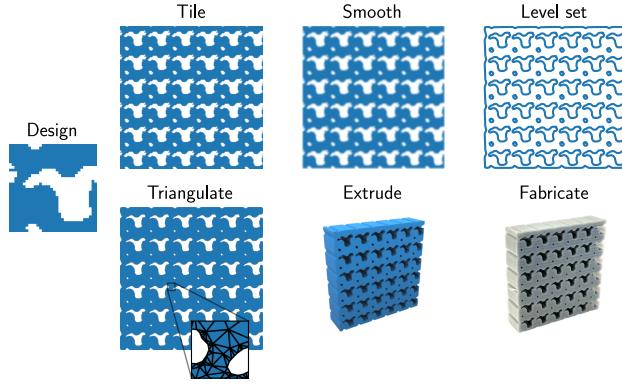


Figure 11: Starting from a binary design, the structure is tiled, smoothed, converted to a level-set boundary, triangulated, and extruded to produce a printable model. The final sample is fabricated in Tough 2000 resin using stereolithography (SLA) printing.

$d = 0.5$, where the printed structure is extremely sparse, amplifying fabrication and testing noise. Additionally, we note that simulation results at lower densities (e.g., $d = 0.5$ and $d = 0.7$) underestimate the target anisotropy, suggesting that high anisotropy is more challenging to realize in sparse structures with limited topological expressivity.

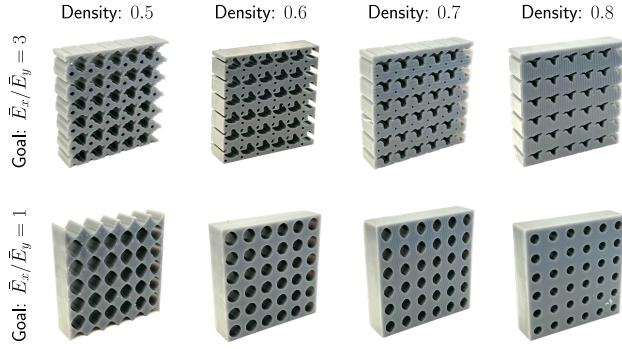


Figure 12: To perform experimental measurements on the directional Young's moduli, we fabricate each design from Figure 10 following the procedure outlined in Figure 11.

4.4 Ablation Study: Optimizer Choice

We obtain our main results with Adam, but in principle, any optimizer can drive the TO loop. To test robustness, we repeated the TO run at target density $d = 0.6$ using stochastic gradient descent (SGD), RMSprop [Graves 2014], and Adadelta [Zeiler 2012] with identical learning-rate schedules (Figure 14).

Adam and SGD both converge to fabricable designs with comparable final losses. SGD reaches a slightly lower value but requires approximately 5× more iterations. Adadelta minimizes the loss most aggressively yet produces a porous layout full of sub-pixel holes, illustrating that loss reduction alone is an insufficient quality

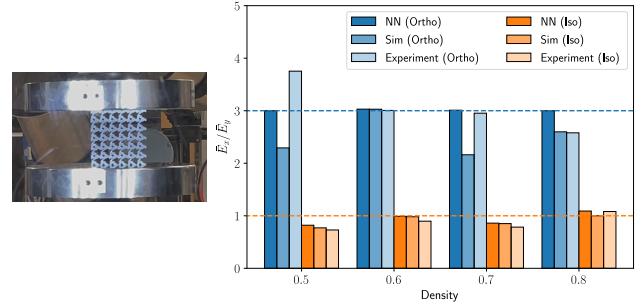


Figure 13: Printed sample undergoing uniaxial compression in an Instron 5965 mechanical testing machine (left). We compare the directional stiffness ratio \bar{E}_x/\bar{E}_y across densities for orthotropic and isotropic design targets shown in Figure 12 (right). Predictions from the surrogate, FEM simulation, and experimental measurements are plotted. Horizontal lines depict the target stiffness ratios for each objective.

metric. RMSprop stagnates early and yields the highest loss and the noisiest geometry. Therefore, we keep Adam as our default, as it strikes a good compromise between speed, objective value, and visual and physical plausibility. However, the framework is not overly sensitive to the choice of optimizer between Adam and SGD.

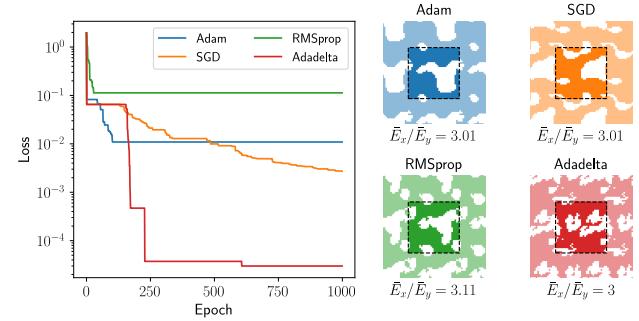


Figure 14: We conduct an ablation study comparing the Adam, stochastic gradient descent (SGD), RMSprop, and Adadelta optimizers used in the surrogate-based TO loop for the anisotropic objective at a target density of 0.6 over 1000 epochs. Adam and SGD yield clean, fabricable structures with similar directional performance, whereas Adadelta and RMSprop produce porous, impractical geometries.

5 Limitations and Future Work

We have introduced a data-driven procedure for training a surrogate model that predicts homogenized stiffness properties directly from binary microstructure designs. Unlike prior approaches that either rely on fixed parametric families, our surrogate enables differentiable topology optimization over a nonparametric space of binary inputs. To train the model, we generate a large dataset of microstructures using a random-walk growth procedure, ensuring

structural diversity and periodic connectivity. We evaluate performance through a suite of optimization tasks that target directional stiffness ratios under various volume constraints. The designs exhibit clear anisotropic trends, and physical fabrication followed by mechanical testing confirms strong alignment between surrogate predictions, FEM simulation, and experimental behavior. However, there are also several potential improvements to the framework we present.

Extension to 3D. While our random walk generator naturally extends to 3D grids, scaling the full pipeline becomes challenging due to the increased cost of finite element simulations. As design resolution and physical fidelity grow, the number of degrees of freedom and independent loading conditions required by homogenization increases sharply. Efficient surrogate training in such settings may require multi-resolution modeling or physics-aware dimensionality reduction.

Hyperparameter Tuning. Our TO procedure depends on manually tuned schedules for regularization terms such as smoothness and density penalties. Although surrogate-based optimization enables rapid iteration, the tuning process remains heuristic and problem-specific. Future work could automate this process using reinforcement learning (RL) or differentiable hyperparameter optimization.

Fabrication Constraints. Although our designs successfully print and function as intended, the pipeline does not explicitly account for printability constraints, such as minimum feature size or printer resolution limits. Incorporating fabrication-aware constraints could improve real-world reliability. Additionally, extending this framework to 3D would require addressing constraints such as overhang and supports.

Global Optimization. While gradient-based methods like Adam are efficient, they are inherently local and prone to getting trapped in suboptimal minima, particularly for objectives involving strong anisotropy or extreme sparsity. Global optimization techniques, such as Bayesian optimization, RL, or genetic algorithms, offer a promising alternative by enabling a broader exploration of the design space. This is especially promising, given recent progress in scalable, high-dimensional Bayesian optimization methods [Hvarfner et al. 2024].

Nonlinear and Multiphysics Behavior. Our surrogate and homogenization pipeline is tailored to linear elasticity, but many real-world applications involve nonlinear, rate-dependent, or coupled multi-physics responses, such as hyperelastic deformation or electromechanical effects. Extending the method to these regimes would require more expressive training data and potentially new surrogate models that respect underlying physical constraints. Careful adaptation of data augmentation and boundary conditions would also be necessary.

Acknowledgments

The authors would like to thank Ruichen Liu for fabricating the optimized microstructure designs and Alec Ewe for running the

compression tests on the fabricated samples. This work was supported by a Focused Research Program from the Rafik B. Hariri Institute for Computing and Computational Science & Engineering at Boston University and by the National Science Foundation (DMR-2323728).

References

- Erik Andreassen and Casper Schouboe Andreasen. 2014. How to determine composite material properties using numerical homogenization. *Computational Materials Science* 83 (2014), 488–495. doi:10.1016/j.commatsci.2013.09.006
- Jan-Hendrik Bastek, Siddhant Kumar, Bastian Telgen, Raphaël N. Glaesener, and Dennis M. Kochmann. 2022. Inverting the structure–property map of truss metamaterials by deep learning. *Proceedings of the National Academy of Sciences* 119, 1 (2022), e2111505119. doi:10.1073/pnas.2111505119 arXiv:<https://www.pnas.org/doi/pdf/10.1073/pnas.2111505119>
- M. P. Bendsøe and O. Sigmund. 1999. Material interpolation schemes in topology optimization. *Archive of Applied Mechanics* 69, 9 (01 Nov 1999), 635–654. doi:10.1007/s004190050248
- Thierry Decker, Slawomir Kedziora, and Elvin Museyibov. 2023. Comparison of Strength Properties of Common Powder Bed Fusion and Stereolithography Materials. *Key Engineering Materials* 969 (2023), 11–20. doi:10.4028/p-Sgu1sR
- Bolei Deng, Ahmad Zareei, Xiaoxiao Ding, James C. Weaver, Chris H. Rycroft, and Katia Bertoldi. 2022. Inverse Design of Mechanical Metamaterials with Target Nonlinear Response via a Neural Accelerated Evolution Strategy. *Advanced Materials* 34, 41 (2022), 2206238. doi:10.1002/adma.202206238 arXiv:<https://advanced.onlinelibrary.wiley.com/doi/pdf/10.1002/adma.202206238>
- Semyon Efremov, Jonàs Martínez, and Sylvain Lefebvre. 2021. 3D Periodic Cellular Materials with Tailored Symmetry and Implicit Grading. *Computer-Aided Design* 140 (2021), 103086. doi:10.1016/j.cad.2021.103086
- Formlabs. 2025. *Tough 2000 Resin*. Technical Data Sheet FLTO2001. Somerville, MA.
- Alex Graves. 2014. Generating Sequences With Recurrent Neural Networks. arXiv:1308.0850 [cs.NE] <https://arxiv.org/abs/1308.0850>
- Xingjian Han, Yu Jiang, Weiming Wang, Guoxin Fang, Simeon Gill, Zhiqiang Zhang, Shengfa Wang, Jun Saito, Deepak Kumar, Zhongxuan Luo, Emily Whiting, and Charlie C.L. Wang. 2024. Motion-Driven Neural Optimizer for Prophylactic Braces Made by Distributed Microstructures. In *SIGGRAPH Asia 2024 Conference Papers* (Tokyo, Japan) (SA '24). Association for Computing Machinery, New York, NY, USA, Article 57, 11 pages. doi:10.1145/3680528.3687661
- Zizhou Huang, Daniele Panozzo, and Denis Zorin. 2024. Optimized shock-protecting microstructures. *ACM Trans. Graph.* 43, 6, Article 181 (Nov. 2024), 21 pages. doi:10.1145/3687765
- Carl Hvarfner, Erik Orm Hellsten, and Luigi Nardi. 2024. Vanilla Bayesian Optimization Performs Great in High Dimensions. In *Proceedings of the 41st International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 235)*. Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (Eds.). PMLR, 20793–20817. <https://proceedings.mlr.press/v235/hvarfner24a.html>
- Lily Kharlevych, Patrick Mullen, Houman Owhadi, and Mathieu Desbrun. 2009. Numerical coarsening of inhomogeneous elastic materials. *ACM Trans. Graph.* 28, 3, Article 51 (July 2009), 8 pages. doi:10.1145/153126.1531357
- Diederik P. Kingma and Jimmy Ba. 2017. Adam: A Method for Stochastic Optimization. arXiv:1412.6980 [cs.LG] <https://arxiv.org/abs/1412.6980>
- Siddhant Kumar, Stephanie Tan, Li Zheng, and Dennis M. Kochmann. 2020. Inverse-designed spinodoid metamaterials. *npj Computational Materials* 6, 1 (05 Jun 2020), 73. doi:10.1038/s41524-020-0341-6
- Beichen Li, Bolei Deng, Wan Shou, Tae-Hyun Oh, Yuanming Hu, Yiyu Luo, Liang Shi, and Wojciech Matusik. 2024. Computational discovery of microstructured composites with optimal stiffness-toughness trade-offs. *Science Advances* 10, 5 (2024), eadk4284. doi:10.1126/sciadv.adk4284 arXiv:<https://www.science.org/doi/pdf/10.1126/sciadv.adk4284>
- Yue Li, Stelian Coros, and Bernhard Thomaszewski. 2023. Neural Metamaterial Networks for Nonlinear Material Design. *ACM Trans. Graph.* 42, 6, Article 186 (Dec. 2023), 13 pages. doi:10.1145/3618325
- Yue Li, Juan Montes, Bernhard Thomaszewski, and Stelian Coros. 2022. Programmable Digital Weaves. *IEEE Robotics and Automation Letters* 7, 2 (2022), 2891–2896. doi:10.1109/LRA.2022.3145948
- Liane Makatura, Bohan Wang, Yi-Lu Chen, Bolei Deng, Chris Wojtan, Bernd Bickel, and Wojciech Matusik. 2023. Procedural Metamaterials: A Unified Procedural Graph for Metamaterial Design. *ACM Trans. Graph.* 42, 5, Article 168 (July 2023), 19 pages. doi:10.1145/3605389
- Jonàs Martínez, Jérémie Dumas, and Sylvain Lefebvre. 2016. Procedural voronoi foams for additive manufacturing. *ACM Trans. Graph.* 35, 4, Article 44 (July 2016), 12 pages. doi:10.1145/2897824.2925922
- Jonàs Martínez, Samuel Hornus, Haichuan Song, and Sylvain Lefebvre. 2018. Polyhedral voronoi diagrams for additive manufacturing. *ACM Trans. Graph.* 37, 4, Article 129

- (July 2018), 15 pages. doi:10.1145/3197517.3201343
- Jonàs Martínez, Mélina Skouras, Christian Schumacher, Samuel Hornus, Sylvain Lefebvre, and Bernhard Thomaszewski. 2019. Star-shaped metrics for mechanical metamaterial design. *ACM Trans. Graph.* 38, 4, Article 82 (July 2019), 13 pages. doi:10.1145/3306346.3322989
- Jonàs Martínez, Haichuan Song, Jérémie Dumas, and Sylvain Lefebvre. 2017. Orthotropic k-nearest foams for additive manufacturing. *ACM Trans. Graph.* 36, 4, Article 121 (July 2017), 12 pages. doi:10.1145/3072959.3073638
- Julian Panetta, Qingnan Zhou, Luigi Malomo, Nico Pietroni, Paolo Cignoni, and Denis Zorin. 2015. Elastic textures for additive fabrication. *ACM Trans. Graph.* 34, 4, Article 135 (July 2015), 12 pages. doi:10.1145/2766937
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. *PyTorch: an imperative style, high-performance deep learning library*. Curran Associates Inc., Red Hook, NY, USA.
- Hao Peng, An Liu, Jingcheng Huang, Lingxin Cao, Jikai Liu, and Lin Lu. 2022. PH-Net: Parallellepipiped microstructure homogenization via 3D Convolutional Neural Networks. *Additive Manufacturing* 60 (2022), 103237. doi:10.1016/j.addma.2022.103237
- Prajit Ramachandran, Barret Zoph, and Quoc V. Le. 2017. Searching for Activation Functions. arXiv:1710.05941 [cs.NE] <https://arxiv.org/abs/1710.05941>
- Chengping Rao and Yang Liu. 2020. Three-dimensional convolutional neural network (3D-CNN) for heterogeneous material homogenization. *Computational Materials Science* 184 (2020), 109850. doi:10.1016/j.commatsci.2020.109850
- Christian Schumacher, Bernd Bickel, Jan Rys, Steve Marschner, Chiara Daraio, and Markus Gross. 2015. Microstructures to control elasticity in 3D printing. *ACM Trans. Graph.* 34, 4, Article 136 (July 2015), 13 pages. doi:10.1145/2766926
- Christian Schumacher, Steve Marschner, Markus Gross, and Bernhard Thomaszewski. 2018. Mechanical characterization of structured sheet materials. *ACM Trans. Graph.* 37, 4, Article 148 (July 2018), 15 pages. doi:10.1145/3197517.3201278
- Samuel Silverman, Kelsey L. Snapp, Keith A. Brown, and Emily Whiting. 2025. Data-Driven Nonlinear Deformation Design of 3D-Printable Shells. *3D Printing and Additive Manufacturing* (2025). doi:10.1089/3dp.2024.0175 arXiv:<https://doi.org/10.1089/3dp.2024.0175>
- Georg Sperl, Rahul Narain, and Chris Wojtan. 2020. Homogenized yarn-level cloth. *ACM Trans. Graph.* 39, 4, Article 48 (Aug. 2020), 16 pages. doi:10.1145/3386569.3392412
- James Utama Surjadi, Bastien F. G. Aymon, Molly Carton, and Carlos M. Portela. 2025. Double-network-inspired mechanical metamaterials. *Nature Materials* 24, 6 (01 Jun 2025), 945–954. doi:10.1038/s41563-025-02219-5
- Davi Colli Tozoni, Jérémie Dumas, Zhongshi Jiang, Julian Panetta, Daniele Panozzo, and Denis Zorin. 2020. A low-parametric rhombic microstructure family for irregular lattices. *ACM Trans. Graph.* 39, 4, Article 101 (Aug. 2020), 20 pages. doi:10.1145/3386569.3392451
- Thibault Tricard, Vincent Tavernier, Cédric Zanni, Jonàs Martínez, Pierre-Alexandre Hugron, Fabrice Neyret, and Sylvain Lefebvre. 2020. Freely orientable microstructures for designing deformable 3D prints. *ACM Trans. Graph.* 39, 6, Article 211 (Nov. 2020), 16 pages. doi:10.1145/3414685.3417790
- Sikko Van 't Sant, Prakash Thakolkaran, Jonàs Martínez, and Siddhant Kumar. 2023. Inverse-designed growth-based cellular metamaterials. *Mechanics of Materials* 182 (2023), 104668. doi:10.1016/j.mechmat.2023.104668
- Matthew D. Zeiler. 2012. ADADELTA: An Adaptive Learning Rate Method. arXiv:1212.5701 [cs.LG] <https://arxiv.org/abs/1212.5701>
- Di Zhang, Xiaoya Zhai, Ligang Liu, and Xiao-Ming Fu. 2023b. An optimized, easy-to-use, open-source GPU solver for large-scale inverse homogenization problems. *Structural and Multidisciplinary Optimization* 66, 9 (09 Sep 2023), 207. doi:10.1007/s00158-023-03657-y
- Zhan Zhang, Christopher Brandt, Jean Jouve, Yue Wang, Tian Chen, Mark Pauly, and Julian Panetta. 2023a. Computational Design of Flexible Planar Microstructures. *ACM Trans. Graph.* 42, 6, Article 185 (Dec. 2023), 16 pages. doi:10.1145/3618396