

Article

OPERA: A Unified Framework for AI-Assisted Polymer Metamaterial Design Through Operator Learning, Physics Embedding, and Normalizing-Flow Inverse Architecture

Koffi Enakoutsa ^{1,*}  and Ivan Giorgio ² 

¹ Department of Mathematics, University of California, Los Angeles, CA 90095, USA

² Department of Civil, Construction-Architectural and Environmental Engineering (DICEAA), University of L'Aquila, 67100 L'Aquila, Italy

* Correspondence: koffi@math.ucla.edu

Abstract

Additive manufacturing has opened an extraordinary design space for polymer metamaterials, enabling microstructures whose macroscopic mechanical behavior is governed largely by geometry rather than by chemical composition. A principled design framework must solve two coupled problems: a *forward* problem (given a microstructure, predict effective properties) and an *inverse* problem (given target properties, generate a microstructure). Convolutional neural networks (CNNs) solve the forward problem accurately, but the inverse problem remains more challenging for three reasons reported in the literature: (i) many surrogates predict only a scalar proxy rather than the full second-order elastic tensor; (ii) fixed or randomly initialized inverse decoders create a distribution-shift gap between surrogate predictions and physical re-evaluation; and (iii) dataset bias toward near-solid configurations limits exploration of low-density and anisotropic designs. We present a unified framework, the *Operator-Physics-Enhanced Reverse Architecture* (OPERA), that addresses all three issues. First, the forward surrogate predicts the complete 3×3 plane-stress stiffness tensor C^{eff} in Voigt notation, with an analytical layer enforcing $C_{ij} = C_{ji}$ and positive definiteness by construction, achieving $R^2 > 0.99$ on the directional moduli and density and $R^2 > 0.88$ on the off-diagonal coupling term C_{16} and the effective Poisson ratio. Second, a normalizing-flow decoder $\mathcal{F}_\phi$, jointly trained with the forward surrogate, keeps inverse design on the training manifold and reduces the surrogate–PDE re-evaluation gap from more than 30% to below 6% on held-out targets. Third, a five-family dataset with uniform coverage of $\rho \in [0.10, 0.95]$ is augmented through an expected-improvement active-learning loop. We embed minimum-feature-size, connectivity, and print-direction constraints into the optimization through differentiable regularization and report agreement of $R^2 = 0.987$ between predictions and tensile measurements on ten FDM-printed specimens. The framework is demonstrated on five problems (auxetic, extreme anisotropy, isotropic low-density, chiral, and hierarchical), with an average target error of 6.8%. The results are framed relative to a reproduced scalar-proxy baseline; we provide an explicit statistical uncertainty analysis, a baseline-reproduction protocol, and a discussion of the method's assumptions and numerical enforcement.

Keywords: polymer metamaterials; additive manufacturing; machine learning; homogenization; normalizing flows; inverse design; auxetic; chiral; active learning; elastic tensor; operator learning



Academic Editor: H. N. Cheng

Received: 15 May 2026

Revised: 8 June 2026

Accepted: 13 June 2026

Published: 15 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article

distributed under the terms and

conditions of the [Creative Commons](#)

[Attribution \(CC BY\)](#) license.

1. Introduction

1.1. Polymer Metamaterials and the Design Challenge

Polymer metamaterials—architected materials whose macroscopic mechanical behavior is dominated by internal geometry rather than chemical composition—have become one of the most consequential classes of designer materials enabled by additive manufacturing [1–6]. Fused deposition modeling (FDM), stereolithography, and multi-jet fusion can produce unit-cell geometries with feature sizes below 200 μm and spatial periods below 5 mm at centimeter-scale part dimensions [7–13].

This capability has opened access to a broad range of mechanical phenomena. These include negative-Poisson-ratio (auxetic) behavior arising from re-entrant or chiral cells [14–17]; direct and indirect Poynting effects in pantographic architectures [18–21]; high deformability and strong energy absorption associated with second- and higher-gradient energies [22–28]; extreme elastic anisotropy through unidirectional banding [29]; programmable buckling [30]; and bandgap-based acoustic filtering [31].

A common thread is that these properties are highly sensitive to geometry: a small change in wall thickness, node connectivity, or pore aspect ratio can convert an auxetic structure into a conventional positive-Poisson-ratio foam [32]. This sensitivity makes physics-based computational design both essential and computationally demanding.

1.2. Homogenization and Its Computational Cost

The standard framework linking microstructure to effective properties is periodic homogenization [33]. For a periodic unit cell Y with indicator $\chi(\mathbf{x}) \in \{0, 1\}$, the homogenized stiffness satisfies

$$C_{ijkl}^{\text{eff}} E_{kl} = \frac{1}{|Y|} \int_Y C_{ijkl}(\mathbf{x}) (E_{kl} + \varepsilon_{kl}(w(\mathbf{x}))) \, d\mathbf{x}, \quad (1)$$

where w solves a family of cell problems and $C(\mathbf{x}) = \chi(\mathbf{x})C_0$. Each FEH evaluation requires solving three boundary-value problems on the unit cell, at a cost of seconds to minutes. For inverse design, which evaluates hundreds to thousands of candidates, this is prohibitive.

1.3. Machine Learning Surrogates: Promise and a Precise Gap

CNNs can learn accurate surrogates for the forward map $\mathcal{H} : \chi \mapsto C^{\text{eff}}$, reducing per-evaluation cost from seconds to milliseconds [34,35]. Image-based architectures capture the relation between microstructural geometry and emergent properties, and deep generative inverse-design frameworks exploit latent representations to accelerate exploration of large design spaces [36–40].

However, high accuracy on the training distribution does not guarantee physically admissible inverse designs. Three limitations recur:

1. **Scalar proxy instead of the full tensor.** Several surrogates train on scalar conductivity analogies or reduced descriptors and predict only a subset of properties [41,42]; tensorial couplings (shear, anisotropic interactions, effective Poisson effects) then remain unresolved, preventing reliable design of chiral or auxetic architectures [43].
2. **Poorly constrained latent representations.** When the decoder is weakly constrained by physics or manufacturability, optimization leaves the admissible manifold and yields checkerboard or disconnected geometries; FEH re-evaluation then deviates substantially from the surrogate prediction [36,44].
3. **Dataset bias and limited diversity.** Many datasets concentrate on near-solid microstructures because they are easier to generate and homogenize, so the surrogate generalizes poorly to sparse, strongly anisotropic, or topologically complex configurations [38,39].

Table 1 situates the present framework against representative recent approaches along the dimensions that matter for inverse design.

Table 1. Systematic comparison of representative machine learning approaches to metamaterial property prediction and inverse design. “Full tensor” denotes prediction of all independent components of C^{eff} ; “Mfg.” indicates embedded manufacturing constraints; “Exp.” indicates experimental validation.

Reference	Output	Inverse Mechanism	Diversity	Mfg.	Exp.
Guo et al. [34]	Scalar/props	Review	—	—	—
Liu et al. [35]	Scalar	Optimization	Low	No	No
Zheng et al. [43]	Props	DNN generator	Med.	No	Yes
Ha et al. [36]	Curve	Generative + surrogate	Med.	Partial	Yes
Lee et al. [38]	Review	Review	—	—	—
Hao et al. [40]	Tensor (GCNN)	DB/mapping	Med.	No	No
Brzin & Brojan [44]	Beam shape	GAN	Low	No	No
This work	Full 3×3	Norm. flow + AL	High (5)	Yes	Yes (10)

Operator-learning approaches to homogenization have advanced rapidly: finite operator-learning techniques map elastic-property fields to mechanical response [45], and neural-operator homogenization has been applied to triply periodic minimal surfaces and other complex geometries [46]. Bayesian-optimization-driven inverse design has also matured [47]. The present framework combines these threads—operator-style full-tensor prediction, a generative inverse decoder, active learning, and manufacturing constraints—within a single pipeline.

1.4. A Unified Framework: OPERA

We refer to the proposed framework as OPERA (*Operator-Physics-Enhanced Reverse Architecture*). It addresses each limitation with a targeted component: (1) a physics-embedded CNN predicting the full 3×3 Voigt tensor, with an analytical symmetry layer and a positive-definiteness projection (Section 4); (2) a normalizing-flow decoder jointly trained with the forward surrogate, so inverse design operates on the training manifold, reducing the surrogate–PDE gap from >30% to <6%; (3) a five-family dataset with uniform density coverage $\rho \in [0.10, 0.95]$, augmented by EI active learning; and (4) differentiable FDM constraints with experimental validation. Throughout, we frame the results relative to a reproduced scalar-proxy baseline (Section 7) and avoid claims of absolute novelty unless supported by that baseline or the cited literature.

These four components are assembled into a single pipeline, as summarized in Figure 1. Panel (a) displays the five microstructure families as binary 32×32 indicator images, illustrating the geometric diversity—from the near-solid elliptic void to the open chiral and banded lattices—that the dataset must span. Panel (b) traces the two data paths: a forward path in which the physics-embedded CNN $\hat{\mathcal{H}}_\theta$ maps a microstructure to its full elastic tensor, and an inverse path in which the normalizing-flow decoder $\mathcal{F}_\phi$ maps a target property back to a manufacturable design, with an active-learning loop closing the two. Panel (c) makes the dataset-bias problem concrete: the reproduced original dataset (grey) clusters tightly around the isotropic line $E_x^{\text{eff}} = E_y^{\text{eff}}$, whereas the present dataset (teal) populates the strongly anisotropic corners of the property space that inverse design must be able to reach. The remainder of this paper develops each component in turn, beginning with the homogenization operator that all of them approximate.




Figure 1. Framework overview. (a) Five microstructure families (binary 32×32 indicator images). (b) Architecture: forward CNN $\hat{\mathcal{H}}_\theta$ predicts the full tensor; flow decoder $\mathcal{F}_\phi$ solves the inverse problem; an active-learning loop refines both. (c) Property-space coverage: each point is an FEH-computed pair $(E_x^{\text{eff}}, E_y^{\text{eff}})$ normalized by the base-polymer modulus E_0 ; the reproduced original dataset (grey) concentrates near isotropy, whereas the present dataset (teal) covers strongly anisotropic configurations.

2. Theoretical Framework

2.1. Homogenization as a Nonlinear Operator

The base solid phase is linear elastic and isotropic with Young's modulus $E_0 = 3.0$ GPa and Poisson's ratio $\nu_0 = 0.38$ (PLA); the void phase is a compliant solid with $E_{\text{void}} = 10^{-4}E_0$ to keep the elasticity problem well posed. Periodic boundary conditions are imposed on the unit cell, and three independent macroscopic strain states recover all components of $\mathbb{C}^{\text{eff}}$. A plane-stress assumption is used (Section 2.2).

Definition 1 (Homogenization operator). Let $Y \subset \mathbb{R}^2$ be a periodic unit cell and $\mathcal{A} = \{\chi \in L^\infty(Y; \{0,1\}) \mid \chi \text{ satisfies connectivity}\}$. The homogenization operator is $\mathcal{H} : \mathcal{A} \rightarrow \mathbb{S}_{3 \times 3}^+$, $\chi \mapsto \mathbb{C}^{\text{eff}}$, where $\mathbb{S}_{3 \times 3}^+$ is the cone of symmetric positive-definite 3×3 matrices (the Voigt representation of the plane-stress effective stiffness).

$\mathcal{H}$ is nonlinear and nonlocal. We learn a surrogate $\hat{\mathcal{H}}_\theta \approx \mathcal{H}$ (forward) and a parametric pseudo-inverse $\mathcal{F}_\phi : \mathbb{S}_{3 \times 3}^+ \rightarrow \mathcal{A}$ (inverse).

2.2. Full Plane-Stress Elastic Tensor

Under plane stress,

$$\begin{bmatrix} \sigma_{11}^{\text{eff}} \\ \sigma_{22}^{\text{eff}} \\ \sigma_{12}^{\text{eff}} \end{bmatrix} = \begin{bmatrix} C_{11} & C_{12} & C_{16} \\ C_{12} & C_{22} & C_{26} \\ C_{16} & C_{26} & C_{66} \end{bmatrix} \begin{bmatrix} E_{11} \\ E_{22} \\ 2E_{12} \end{bmatrix}. \quad (2)$$

For microstructures without special symmetry, all six independent components are nonzero. The coupling terms C_{16} and C_{26} vanish only for orthotropic or isotropic microstructures and signal chirality; they are not recoverable by a scalar proxy (the full symmetry-class classification is given in Appendix B).

Remark 1 (Interpretation of C_{16}). A nonzero C_{16} couples normal stress σ_{11} to shear strain $2E_{12}$: the mechanical footprint of handedness. Designing for $C_{16} \neq 0$ requires the full tensor model.

2.3. Physics-Embedded Neural Operator

Definition 2 (Physics-embedded CNN). The forward surrogate $\hat{\mathcal{H}}_\theta$ consists of:

- (i) A convolutional encoder $f_\theta : \{0, 1\}^{32 \times 32} \rightarrow \mathbb{R}^{512}$;
- (ii) A fully connected head $\mathbb{R}^{512} \rightarrow \mathbb{R}^6$;
- (iii) A symmetry layer $\hat{C}_{ij} \leftarrow (\hat{C}_{ij} + \hat{C}_{ji})/2$; and
- (iv) A positive-definite projection (Section 4).

The symmetry and positivity constraints are analytical rather than learned, so they hold exactly at inference regardless of the network weights. This property is formalized next.

Theorem 1 (Symmetry and positivity preservation). $\hat{\mathcal{H}}_\theta$ produces $\hat{C} \in \mathcal{S}_{3 \times 3}^+$ for all $\chi \in \{0, 1\}^{32 \times 32}$, by construction of the symmetry layer and PD projection.

Proof. The symmetry layer makes $\hat{C}_{ij} = \hat{C}_{ji}$. The PD projection (Equation 7) shifts all eigenvalues to $\geq \varepsilon > 0$. $\square$

The surrogate is trained by minimizing the forward loss:

$$\mathcal{L}_{\text{fwd}}(\theta) = \frac{1}{N} \sum_{n=1}^N \|\hat{\mathcal{H}}_\theta(\chi_n) - C_n^{\text{eff}}\|_F^2 + \lambda \|\hat{\rho}_n - \rho_n\|^2, \quad \lambda = 0.5, \quad (3)$$

where the second term weights an auxiliary density prediction.

The resulting forward accuracy is reported in Figure 2. The heat maps in panels (b,c) confirm that the model reproduces the qualitative structure of C^{eff} , including the negative off-diagonal term $C_{12} < 0$, which is the tensorial signature of auxetic behavior (panel b), and its absence in an isotropic reference (panel c). Panel (d) shows that the effective Poisson ratio spans $\nu^{\text{eff}} \in [-0.5, +0.4]$ across the five families, with the auxetic family occupying the negative range that scalar proxies cannot represent. The contrast is quantified in panel (e): the reproduced scalar proxy attains $R^2 > 0.99$ on the directional moduli but $R^2 = 0$ on C_{12} and C_{16} , whereas the present model reaches $R^2 = 0.912$ and 0.887 on those coupling terms, with the corresponding parity plot for C_{12} shown in panel (f).

The architecture that produces these predictions, together with its training behavior, is shown in Figure 3. Panel (a) shows the schematic—four convolutional blocks, global average pooling, the fully connected readout, and the two analytical output layers—whereas the feature maps in panels (b–d) illustrate the progression from local geometry in the early layers to the global, stiffness-relevant patterns consumed by the readout head. Panel (e) shows that training converges stably for both the present model and the reproduced scalar baseline; the difference is not the rate of convergence but the richness of the output. Panel (f) makes the value of the enforcement layer concrete: without it, the predicted tensors exhibit symmetry violations $|C_{ij} - C_{ji}|$ of up to a few percent, whereas with enforcement, they are reduced to machine precision. Because this constraint is built into the forward pass rather than learned, the guarantee is independent of how well the network is trained. The ability to predict the full tensor accurately is precisely what makes inverse design of chiral and auxetic targets possible, which we take up next.




Figure 2. Full elastic-tensor prediction. (a) Example auxetic microstructure. (b,c) Heat maps of C^{eff} (FEH-computed, normalized by E_0 ; color scale from low (dark blue) to high (yellow) stiffness) for an auxetic ($C_{12} < 0$) and an isotropic reference. (d) Kernel-density estimate of ν^{eff} (dimensionless) per family, FEH over the full dataset. (e) Held-out R^2 : reproduced scalar proxy vs. full-tensor model. (f) Parity for C_{12}^{eff} (predicted vs. FEH reference, normalized), $R^2 = 0.912$, $N_{\text{test}} = 2000$.




Figure 3. Physics-embedded CNN architecture. (a) The forward surrogate: four convolutional blocks, global average pooling, a fully connected readout, and the analytical symmetry-enforcement and positive-definite projection layers. (b–d) Feature maps at layers 1–3 (blue-to-yellow colormap indicating low-to-high activation magnitude), showing the progression from local geometry to global stiffness-relevant features. (e) Training convergence (MSE loss, E_0 -normalized) for the present model (teal) and the reproduced scalar model (grey); both converge, but the present model predicts a richer output. (f) Symmetry violation $|C_{ij} - C_{ji}|$ (normalized): without the enforcement layer (grey), violations reach a few percent; with enforcement (green), they are reduced to machine precision.

3. Normalizing-Flow Decoder for Inverse Design

3.1. The Distribution-Shift Failure Mode

The CNN trained on $\mathcal{D}_{\text{train}}$ achieves low error on $\mathcal{D}_{\text{test}}$. A fixed/random decoder optimizes in a pre-image unrelated to the training distribution, so χ_{orig}^* satisfies $d(\chi_{\text{orig}}^*, \mathcal{D}_{\text{train}}) \gg 0$, and the surrogate prediction has a large error relative to the FEH ground truth.

3.2. Normalizing-Flow Architecture and Binary Microstructures

Definition 3 (Normalizing-flow decoder). $\mathcal{F}_\phi : \mathcal{Z} \rightarrow \mathcal{A}$ is a composition of K affine coupling layers. Let $z = [z_a, z_b]$, with $z'_a = z_a$ and $z'_b = z_b \odot \exp(s_\phi(z_a)) + t_\phi(z_a)$, where s_ϕ and t_ϕ are CNNs. After $K = 8$ layers, transposed convolutions and a sigmoid produce a continuous relaxation $\hat{\chi} \in [0, 1]^{32 \times 32}$.

Continuous relaxation and differentiability.

The flow operates on $\hat{\chi} \in [0, 1]^{32 \times 32}$; a binary indicator is needed only at the final manufacturing/evaluation step. During latent optimization, the hard threshold $\chi^* = (\hat{\chi} > 0.5)$ is replaced with the smooth Heaviside projection (Equation (8)) with continuation $\beta_H \in \{1, 2, 4, 8, 16, 32\}$. Because this projection is differentiable for finite β_H , gradients flow through it; the hard threshold is applied only once, after convergence, so it never enters the gradient computation. This avoids the non-differentiability of direct binary thresholding.

Physical admissibility after projection.

Because $\mathcal{F}_\phi$ is trained by maximum likelihood on physically admissible (connected, manufacturable) training microstructures, its output lies near the training manifold for any z . The regularizer $\mathcal{R}$ (Section 6) penalizes disconnected components and sub-resolution features, and after projection, the candidate is re-evaluated by FEH; if the tolerance is exceeded, the active-learning loop adds the sample and re-trains. Admissibility is thus enforced by construction, by regularization, and by verification.

The triangular Jacobian enables exact log-likelihood computation, and the flow is trained by

$$\mathcal{L}_{\text{flow}}(\phi) = -\frac{1}{N} \sum_{n=1}^N \log p_\phi(\chi_n). \quad (4)$$

3.3. Joint Training and Inverse Optimization

The surrogate and decoder are trained jointly: (1) pre-training with Equations (3) and (4), and (2) joint fine-tuning of $\mathcal{L} = \mathcal{L}_{\text{fwd}} + \mu \mathcal{L}_{\text{flow}}$, $\mu = 0.1$. Given C^{target} , inverse design solves

$$z^* = \arg \min_{z \in \mathcal{Z}} \left\{ \left\| \hat{\mathcal{H}}_\theta(\mathcal{F}_\phi(z)) - C^{\text{target}} \right\|_F^2 + \lambda_\rho (\hat{\rho} - \rho^*)^2 + \mathcal{R}(\mathcal{F}_\phi(z)) \right\}. \quad (5)$$

Theorem 2 (Inverse-design consistency). Let $\hat{\mathcal{H}}_\theta$ achieve error ε_{fwd} on $\mathcal{D}_{\text{train}}$ and $\mathcal{F}_\phi$ a log-likelihood gap $\varepsilon_{\text{flow}}$. Then,

$$\left\| \hat{\mathcal{H}}_\theta(\chi^*) - \mathcal{H}(\chi^*) \right\|_F \leq \varepsilon_{\text{fwd}} + C_{\text{lip}} \varepsilon_{\text{flow}}^{1/2}, \quad (6)$$

where C_{lip} is the Lipschitz constant of $\hat{\mathcal{H}}_\theta$ (estimated in Section 4.2).

A proof is given in Appendix A; the corresponding symmetry classes and notation are summarized in Appendix B, and the full set of training hyperparameters in Appendix C.

Figure 4 demonstrates that this guarantee holds in practice. The decoder and its optimization loop are sketched in panels (a) and (b), while panel (c) contrasts the two inverse

procedures directly: for the reproduced fixed-decoder baseline, the surrogate objective and the FEH re-evaluation diverge, leaving a persistent gap above 30%, whereas the jointly trained flow keeps the two curves aligned and drives the gap below 6%. The parity plots in panels (d–f) confirm that the designs achieve their targets in E_x^{eff} , ν^{eff} , and ρ with $R^2 > 0.92$ when re-evaluated by FEH—the operational meaning of the bound in Equation (6). Having established that the flow keeps designs on the training manifold, we next examine the assumptions and numerical safeguards that make the forward map itself reliable.




Figure 4. Normalizing-flow decoder and inverse design. (a) Flow architecture ($K = 8$ coupling layers). (b) Inverse-design loop (solid arrows indicate deterministic forward evaluation; dashed arrows indicate the latent-update feedback path). (c) Surrogate objective vs. FEH re-evaluation (normalized): reproduced baseline (grey/red) shows a persistent gap $>30\%$; the present method (teal/green) closes it. (d–f) Inverse-design parity for E_x^{eff} , ν^{eff} , ρ (FEH-achieved vs. target), $R^2 > 0.92$, $N = 180$.

4. Assumptions and Numerical Enforcement

Two ingredients of the forward surrogate were stated above without analysis: the positive-definite projection that guarantees physically valid tensors, and the Lipschitz constant that controls the consistency bound of Equation (6). We examine each here, since both determine how far the guarantees can be trusted.

4.1. Positive-Definite Projection

The symmetrized output $\hat{C}$ is mapped to the PD cone by

$$\hat{C} \leftarrow \hat{C} + \max(0, -\lambda_{\min}(\hat{C}) + \varepsilon)I, \quad \varepsilon = 10^{-6}, \quad (7)$$

with $\lambda_{\min}$ the smallest eigenvalue.

Well-posedness: Since $\hat{C}$ is symmetric, its spectrum is real; the shift adds $\delta = \max(0, -\lambda_1 + \varepsilon)$ to every eigenvalue, giving $\lambda'_i \geq \varepsilon > 0$, so $\hat{C}' \succ 0$, and the map is idempotent on the cone.

Closeness: The output changes by at most $\|\hat{C}' - \hat{C}\|_F = \delta\sqrt{3}$. When well-trained, $\lambda_1 \geq 0$ for most inputs and $\delta = 0$. Empirically, the projection was active on 3.1% of the

held-out test inputs, with a mean shift $\delta = 4.2 \times 10^{-3}$ and a maximum of $\delta = 2.1 \times 10^{-2}$ (in units of E_0 -normalized stiffness), confirming that the projection has a small practical footprint and acts only as a safety correction in the rare near-singular cases.

Differentiability: The eigenvalue shift is differentiable almost everywhere; non-differentiability occurs only where $\lambda_1 = 0$ or eigenvalues coincide. We use the eigendecomposition-based gradients with a small ridge, equivalent to the Cholesky correction in Algorithm 1.

Algorithm 1 Forward surrogate training

Require: Dataset $\mathcal{D}_N = \{(\chi_n, C_n^{\text{eff}}, \rho_n)\}$, epochs T

- 1: Initialize CNN weights θ (Xavier)
- 2: **for** epoch = $1, \dots, T$ **do**
- 3: Sample mini-batch $\mathcal{B} \subset \mathcal{D}_N$
- 4: $\hat{C}_n = f_{\text{FC}}(f_{\text{pool}}(f_{\text{conv}}(\chi_n; \theta)))$
- 5: Symmetry: $\hat{C}_n \leftarrow (\hat{C}_n + \hat{C}_n^\top)/2$
- 6: PD project: $\hat{C}_n \leftarrow \hat{C}_n + \max(0, -\lambda_{\min} + \varepsilon)I$
- 7: $\mathcal{L} = \frac{1}{|\mathcal{B}|} \sum_n (\|\hat{C}_n - C_n^{\text{eff}}\|_F^2 + \lambda \|\hat{\rho}_n - \rho_n\|^2)$
- 8: $\theta \leftarrow \theta - \eta \nabla_\theta \mathcal{L}$ (Adam)
- 9: **end for**

Ensure: Surrogate $\hat{\mathcal{H}}_\theta, R^2 > 0.99$ on test set

4.2. Lipschitz Constant and Validity of the Bound

The bound in Equation (6) depends on C_{lip} . For a feed-forward CNN with 1-Lipschitz activations, a standard upper bound is the product of spectral norms, $C_{\text{lip}} \leq \prod_\ell \|W_\ell\|_2$, which is typically loose. A tighter empirical estimate samples pairs near the training manifold and takes the maximal ratio $\|\hat{\mathcal{H}}_\theta(\chi_i) - \hat{\mathcal{H}}_\theta(\chi_j)\|_F / \|\chi_i - \chi_j\|_{L^2}$. For the trained network, the spectral-norm product bound evaluates to $\prod_\ell \|W_\ell\|_2 \approx 8.4 \times 10^2$, which (as expected) is loose; the empirical estimate from 10^4 finite-difference pairs sampled near the training manifold is $C_{\text{lip}}^{\text{emp}} \approx 2.3$ (E_0 -normalized stiffness per unit L^2 image distance), about two orders of magnitude smaller, so the practically relevant constant entering Equation (6) is $O(1)$.

Validity: Equation (6) is most informative when $\varepsilon_{\text{flow}}$ is small (the flow accurately models the data). In the well-sampled regime, this holds, whereas in sparse regions, the $\varepsilon_{\text{flow}}^{1/2}$ term dominates and the bound becomes conservative—consistent with the larger inverse error observed for low-density and strongly anisotropic targets (Section 8).

5. Dataset Design Rationale, Diversity, and Active Learning

The reliability of both the forward surrogate and the inverse decoder ultimately rests on the data on which they are trained. We therefore turn to how the dataset is constructed—which microstructure families it spans and how active learning allocates the labeling budget.

5.1. Why Five Microstructure Families

The five families span the qualitatively distinct regions of the 2D plane-stress property space most relevant to inverse design, rather than enumerating all possible topologies: **elliptic void** (positive ν , near-solid baseline); **re-entrant auxetic** ($\nu^{\text{eff}} \in [-0.5, 0]$); **chiral lattice** (nonzero C_{16}, C_{26}); **hierarchical** (multi-scale stiffness); and **banded** (arbitrary anisotropy). Together they cover the four corners of design interest—sign of ν^{eff} , chiral coupling, multi-scale response, and anisotropy magnitude—which no single parametric family spans.

Extensibility.

Adding families (triply periodic minimal surfaces, lattice trusses, stochastic spinodal topologies) is expected to *improve* coverage of regions not spanned by the present five at the cost of larger labeling budgets. Because the decoder is generative and the loop is acquisition-driven, the pipeline accommodates new families without architectural change; the main limiting factor is the FEH labeling cost, which active learning mitigates. We observe no saturation of predictive quality with the present families, suggesting that additional families would help. A systematic study is left to future work.

5.2. Active Learning with Expected Improvement

Definition 4 (Expected Improvement). For a GP surrogate $\hat{f} \sim \mathcal{GP}(\mu, \sigma^2)$ and current best f^* , $\text{EI}(\chi) = (f^* - \mu)\Phi(\frac{f^* - \mu}{\sigma}) + \sigma\phi(\frac{f^* - \mu}{\sigma})$, with Φ, ϕ the standard normal CDF and PDF.

Each iteration queries the EI-maximizing microstructure, evaluates it by FEH, and adds it to the dataset, concentrating labeling where uncertainty is highest—about $3\times$ better coverage than random sampling for equal budget.

The combined effect of the family design and the active-learning loop is quantified in Figure 5. The principal-component projection in panel (a) and the density histogram in panel (b) show that the present dataset fills the property space and covers the full density range $\rho \in [0.10, 0.95]$, breaking the near-solid concentration of the reproduced original dataset. Panel (c) shows that EI active learning reduces prediction uncertainty roughly three times faster than random sampling, and panels (d–f) show representative members of three families with their FEH-computed properties, confirming that the sampled designs are geometrically distinct rather than minor perturbations of one another. A diverse, well-covered dataset is necessary but insufficient for deployable designs; the designs must also be manufacturable, which is the subject of the next section.




Figure 5. Dataset diversity and active learning. (a) PCA of the property space (legend above panel). (b) Density histogram. (c) Active-learning convergence ($\sim 3\times$ faster than random). (d–f) Representative microstructures with FEH-computed properties (normalized by E_0).

6. Manufacturing Constraints

FDM-Aware Regularization

A morphological density filter with radius $r = \ell_{\min}/\Delta x$ removes sub-resolution features ($\ell_{\min} \approx 0.5$ mm for a 0.4 mm nozzle) [48]. Connectivity is enforced using a differen-

tiable penalty $\mathcal{R}_{\text{conn}} = \lambda_{\text{conn}} \max(0, N_{\text{components}} - 1)$, and a near-binary state is encouraged using the Heaviside projection [49]

$$\tilde{\chi} = \frac{\tanh(\beta_H(0.5 - \hat{\chi})) + \tanh(\beta_H/2)}{2 \tanh(\beta_H/2)}. \quad (8)$$

The full regularizer is given by $\mathcal{R} = \mathcal{R}_{\text{conn}} + \lambda_{\text{bin}} \|\chi(1 - \chi)\|^2$. Physical validation of these constraints against printed specimens is reported in Section 8.5, after the designs of Section 8 have been introduced, so that the figure numbering follows the order in which figures are first cited.

7. Baseline Reproduction Protocol

Table 2 states the provenance of every comparison value: **(L)** literature; **(R)** baselines reproduced under identical data and training conditions; **(O)** the authors' own implementation. "Conventional surrogate approaches" in Table 3 refers to category **(R)**: a scalar-proxy CNN re-implemented on the same 10,000-microstructure dataset, at the same 32×32 resolution, using the same optimizer (Adam, $\eta = 10^{-4}$), and under the same epoch budget, with the output head reduced to a single scalar.

Table 2. Provenance of comparison values. L = literature; R = reproduced under identical dataset/training; O = authors' own implementation.

Comparison Quantity	Prov.	Notes
Scalar-proxy forward R^2	R	Re-implemented scalar CNN with identical data and training setup
Full-tensor forward R^2	O	This work
Surrogate-PDE gap (>30%)	R	Reproduced fixed-decoder inverse design
Inverse-design R^2	O	This work
Conventional/auxetic foam ν^{eff}	L	Gibson–Ashby [32], Lakes [2]
GAN/chiral comparison points	R	Reproduced under present dataset
Honeycomb anisotropy baseline	L/R	Literature value, re-checked by FEH
Experimental measurements (10)	O	This work; Section 8.5

The fixed-decoder inverse baseline reproduces the *original study* on which this work builds convolutional-neural-network homogenization surrogates and preliminary inverse design for polymer metamaterials, which reported high forward accuracy but only partially successful inverse design. All reproduced baselines use the released code and dataset described in the Availability statements.

Table 3. Quantitative comparison of the present framework vs. the reproduced scalar-proxy baseline (category R of Table 2) on property prediction and inverse-design targets. "Conventional surrogate approaches" denotes the reproduced baseline, not literature values.

Metric	Reproduced Baseline (R)	OPERA (O)	Improvement
<i>Forward surrogate (R^2, held-out test set)</i>			
E_x^{eff}	0.996	0.998	+0.002
E_y^{eff}	0.995	0.997	+0.002
ρ	0.998	0.999	+0.001
ν^{eff}	0.710	0.941	+0.231
C_{12}	0.000	0.912	+0.912
C_{16}	0.000	0.887	+0.887

Table 3. Cont.

Metric	Reproduced Baseline (R)	OPERA (O)	Improvement
<i>Inverse design</i>			
Surrogate–PDE gap	>30%	<6%	$\sim 5\times$
Inverse R^2	≈ 0	0.942	+0.942
<i>Design problems (target error)</i>			
Auxetic	Failed	7.9%	solved
Anisotropy	Failed	4.0%	solved
Isotropic	Failed	10%	solved
Chiral	Failed	achieved	solved
Hierarchical	Failed	9.2%	solved
Average	>100%	6.8%	5/5
<i>Efficiency</i>			
Inference time	0.8 ms	1.2 ms	comparable
Labels for $R^2 > 0.95$	5000	1700 (AL)	$\sim 3\times$ fewer
Density coverage	$\rho \in [0.8, 1.0]$	$\rho \in [0.1, 1.0]$	wider

8. Applications: Five Design Problems

With the forward surrogate, the flow decoder, the diverse dataset, and the manufacturing constraints all in place, we now apply the framework to five inverse-design problems chosen to exercise different regions of the property space: negative Poisson ratio, extreme anisotropy, isotropic low density, chiral coupling, and multi-scale stiffness. Each problem specifies a property target, and we report both the surrogate-predicted result and its FEH re-evaluation.

8.1. Problem 1: Auxetic ($\nu^{\text{eff}} = -0.38$)

Target: $\nu^{\text{eff}} = -0.38$, $\rho \in [0.35, 0.50]$, $E_x^{\text{eff}} \geq 0.4$. The framework converged to $\nu^{\text{eff}} = -0.41$ (7.9% deviation from the target magnitude), $\rho = 0.42$, $E_x^{\text{eff}} = 0.46$. The reproduced fixed-decoder baseline did not produce an auxetic microstructure (the predicted ν^{eff} remained positive). The re-entrant geometry (inner angle $\approx 110^\circ$) expanded laterally under axial tension (DIC-confirmed).

Figure 6 documents this result across data types. Panel (b) shows the surrogate-predicted Poisson ratio decreasing monotonically toward the target during optimization while the baseline remains positive; panel (c) places the FEH-computed ν^{eff} of the auxetic family well below zero, in contrast to the elliptic family; and the FEH stress field and lateral-expansion plots in panels (d,e) confirm the auxetic kinematics mechanically. Panel (f) compares the achieved ν^{eff} against the literature and the reproduced baselines in Table 2, with the present design yielding the most negative value.

8.2. Problem 2: Extreme Anisotropy ($E_x/E_y = 4.8$)

Target: $E_x/E_y = 5.0$, $\rho \in [0.40, 0.60]$; achieved: $E_x/E_y = 4.8$ (4.0% deviation) via a banded microstructure at $\theta = 12^\circ$. As Figure 7 shows, the anisotropy ratio depends strongly on band orientation (panel b), and the polar stiffness diagram in panel (c) develops an elongated butterfly shape characteristic of extreme orthotropy. The directional stress–strain curves (panel d) make the $4.8\times$ stiffness contrast explicit, and panel (e) indicates that the present designs reach into the high-ratio region of the density–anisotropy map that the baselines (panel f) do not.




Figure 6. Problem 1: auxetic design. (a) Designed unit cell. (b) ν^{eff} evolution during inverse optimization (surrogate-predicted). (c) ν^{eff} vs. void fraction (FEH-computed). (d) Von Mises stress field under y -tension (FEH), normalized by the applied stress σ_{applied} . (e) Lateral expansion (FEH-simulated). (f) Achieved ν^{eff} across methods; provenance per Table 2, with literature comparison points from Gibson and Ashby [32] and Lakes [2].




Figure 7. Problem 2: extreme anisotropy. (a) Designed banded microstructure. (b) Anisotropy ratio vs. orientation (FEH; legend above panel). (c) Polar stiffness $E(\theta)$ (FEH). (d) Directional stress–strain (FEH). (e) Achievable anisotropy vs. density (FEH). (f) Cross-method comparison; provenance per Table 2.

8.3. Problems 3 and 4: Isotropic Low-Density and Chiral

Problem 3 (isotropic)—Target: $K^{\text{eff}} = 0.30$, $E_x/E_y \in [0.95, 1.05]$, $\rho < 0.30$; achieved: $K^{\text{eff}} = 0.27$ (10% deviation), $\rho = 0.28$, $E_x/E_y = 1.02$ (hierarchical family).

Problem 4 (chiral)—Target: $|C_{16}| > 0.05$; achieved: $C_{16} = 0.063$. Explicit surrogate-based targeting of C_{16} has not, to the best of our knowledge, been reported under conditions comparable to those of our reproduced baseline. We state this relative to that baseline rather than as an absolute first.

These two problems are summarized in Figure 8. For the isotropic target, panel (b) shows the designed bulk modulus tracking the target line and remaining within the Voigt–Reuss bounds, while the hierarchical geometry of panel (a) achieves near-isotropy at low density. For the chiral target, panel (d) shows the coupling coefficient C_{16} varying smoothly and sign-reversibly with the chirality parameter, vanishing in the achiral case, and panel (c) shows the corresponding spiral geometry. The manufacturability matrix in panel (e) indicates that, among the methods considered, only the present designs satisfy all four FDM criteria simultaneously, and panel (f) collects the target deviations for the first four problems.




Figure 8. Problems 3–4. (a) Hierarchical microstructure (isotropic target). (b) Bulk modulus vs. density with Voigt–Reuss bounds (FEH). (c) Chiral microstructure ($C_{16} \neq 0$). (d) Chiral coupling C_{16} vs. chirality parameter (FEH). (e) FDM manufacturability check; checkmarks (✓) indicate that a design satisfies the corresponding manufacturability criterion and crosses (×) indicate that it does not. (f) Accuracy summary, problems 1–4.

8.4. Manufacturing Validation

The five designs were finally fabricated and tested to verify that the predicted properties survive the printing process. Figure 9 compares the ideal design (panel a) with its naive FDM discretization (panel b), whose layer artifacts degrade the response, and with the constraint-aware design (panel c), which prints faithfully; panel (d) shows that the constraint recovers about 97% of the ideal stiffness. Panel (e) reports the print-speed/accuracy trade-off, and panel (f) presents the central experimental result—predicted versus measured properties for the ten specimens, with error bars and pooled $R^2 = 0.987$ (quantified in Section 8.5).

8.5. Experimental Validation

Ten specimens (two per design problem) were fabricated on a Creality Ender-5 (Shenzhen Creality 3D Technology Co., Shenzhen, China; PLA, 0.4 mm nozzle, 0.2 mm layer height, 100% infill, 40 mm) at $20 \times 20 \times 5 \text{ mm}^3$. Uniaxial tensile tests (ZwickRoell Z100 universal testing machine, ZwickRoell GmbH & Co. KG, Ulm, Germany; 10 kN cell, crosshead speed 1 mm/min^{-1}) yielded E_x^{eff} , E_y^{eff} , and ν^{eff} via digital image correlation (GOM Aramis 5M system, GOM GmbH (Carl Zeiss GOM Metrology), Braunschweig, Germany; facet size $15 \times 15 \text{ px}$, step 5 px).




Figure 9. FDM manufacturing validation. (a) Ideal design. (b) After FDM discretization without constraint (artifacts degrade the response). (c) With FDM constraint (manufacturable). (d) Property retention with/without constraint (legend above panel). (e) Print-speed vs. accuracy trade-off. (f) Prediction vs. experiment on $n = 10$ specimens with per-point error bars; pooled $R^2 = 0.987$ over E^{eff} and ν^{eff} . Error bars are 95% confidence intervals from three repeated tensile tests per specimen (per-property $R^2 = 0.985$ for E^{eff} and 0.972 for ν^{eff}).

Statistical validation and uncertainty.

The reported $R^2 = 0.987$ is computed across the pooled measured E_x^{eff} and ν^{eff} for the ten specimens (Figure 9f); it is a property-pooled coefficient of determination rather than a per-property value. The two specimens per design problem (ten in total) characterize print-to-print variability, and each specimen was tested three times within the elastic range to assess measurement repeatability. The resulting per-specimen relative standard deviations are 4.1% for E_x^{eff} , 4.6% for E_y^{eff} , and 0.03 (absolute) for ν^{eff} ; the print-to-print relative standard deviation across the two specimens of each problem is 5.8% on average; and the DIC strain-measurement uncertainty is $\pm 0.3\%$ strain (facet size 15×15 px). Propagating these uncertainties into the prediction comparison yields the 95% confidence intervals shown as error bars in Figure 9f; the pooled agreement ($R^2 = 0.987$) is stable across repeated tests, with a per-property breakdown of $R^2 = 0.985$ for E^{eff} and $R^2 = 0.972$ for ν^{eff} . With ten independent prints and thirty tensile measurements, the sample size is sufficient to establish the reported agreement while remaining modest. We therefore present it as a robust *indicative* validation and note that a larger campaign would further tighten the intervals.

9. Benchmarks, Algorithms, and Scalability

9.1. Algorithm Summary

For completeness and reproducibility, we present the two core procedures—forward surrogate training and inverse design with active learning—in Algorithms 1 and 2, after which we report the fifth design problem and the aggregate benchmarks.

Algorithm 2 Inverse design with active learning**Require:** Target $C^{\text{target}}, \rho^*, \nu^*$; trained $\hat{\mathcal{H}}_\theta, \mathcal{F}_\phi$ 1: Initialize $z^{(0)} \sim \mathcal{N}(0, I)$ 2: **for** $t = 1, \dots, T_{\text{inv}}$ **do**3: $\hat{\chi}^{(t)} = \mathcal{F}_\phi(z^{(t-1)})$; FDM filter $\rightarrow \tilde{\chi}^{(t)}$ 4: $\hat{C}^{(t)} = \hat{\mathcal{H}}_\theta(\tilde{\chi}^{(t)})$ 5: $\mathcal{L}(z) = \|\hat{C}^{(t)} - C^{\text{target}}\|_F^2 + \mathcal{R}(\tilde{\chi}^{(t)})$ 6: $z^{(t)} = z^{(t-1)} - \eta_z \nabla_z \mathcal{L}$ 7: **end for**8: $\chi^* = (\mathcal{F}_\phi(z^{(T_{\text{inv}})}) > 0.5)$ 9: Evaluate $\mathcal{H}(\chi^*)$ by FEH; if gap $> \epsilon_{\text{tol}}$: query $\arg \max_\chi \text{EI}(\chi)$, add to $\mathcal{D}$, re-train**Ensure:** χ^* with $\|\hat{C}^* - C^{\text{target}}\| / \|C^{\text{target}}\| < 0.10$

The fifth design problem targets a hierarchical lattice with prescribed multi-scale stiffness; the resulting geometry and its scale-dependent response are shown in Figure 10a,b, with a target error of 9.2%. The remaining panels of Figure 10 aggregate performance across all problems: panel (c) summarizes the five target errors against the unsolved baseline, panel (d) shows that accuracy improves with dataset size and that active learning reaches a given R^2 with fewer labels, panel (e) quantifies the inference cost (1.2 ms per design, versus 3200 ms for FEM homogenization), and panel (f) confirms $R^2 > 0.88$ on all six Voigt components, where the scalar baseline attains zero on the coupling terms. Figure 11 complements this with a per-component view: the pseudo-code boxes in the top row restate the algorithms graphically, while panels (d–f) show that active learning lowers the error on new targets about three times faster than random sampling, accelerates property-space coverage, and yields component-wise R^2 between 0.94 and 0.99 for the full pipeline.




Figure 10. Problem 5 and comprehensive benchmarks. (a) Hierarchical microstructure. (b) Scale-dependent stiffness (FEH). (c) Five-problem accuracy summary (legend above panel). (d) R^2 vs. training samples. (e) Computational cost per design. (f) R^2 for all six Voigt components.




Figure 11. Algorithms and active-learning performance. (a–c) Pseudo-code for the three algorithms (top row). (d) Active-learning efficiency, present method (solid line) vs. random sampling (dashed line) ($\sim 3 \times$ faster than random). (e) Design-space coverage vs. labeled samples. (f) Per-component R^2 .

9.2. Quantitative Benchmark

Table 3 consolidates these comparisons in one place: the present framework matches the reproduced baseline on the directional moduli and density, while gaining the coupling terms, the inverse-design accuracy, and the wider property-space coverage that the baseline lacks entirely. We now interpret why these gains arise.

10. Discussion

10.1. Why the Prior Inverse Design Failed

High forward accuracy is necessary but insufficient for inverse design. In the reproduced baseline, the fixed decoder creates an inverse manifold disjoint from the training manifold, so the surrogate’s accuracy provides no guarantee about inverse-design quality—a train–test distribution shift arising from the inverse procedure. The flow decoder resolves this by construction: any latent z decodes to a microstructure with high probability under the training distribution, so the surrogate’s accuracy transfers (Theorem 2; Appendix A).

10.2. Importance of the Full Elastic Tensor

A scalar proxy is appropriate when targets are only directional stiffnesses and density. For auxetic, chiral, or anisotropic design, however, the Poisson ratio and coupling stiffnesses are the primary targets. The physics-embedded CNN predicts all six independent components ($R^2 > 0.88$), which is possible because the symmetry layer prevents wasted capacity on asymmetric outputs and the PD projection ensures physically meaningful results.

10.3. Active Learning as a Dataset-Design Principle

The near-solid bias of the original dataset arises naturally from the enforced solid boundary frame (physically motivated for connectivity), making the surrogate accurate in the high-density regime and weaker in the low-density, anisotropic regimes where design is most valuable. EI active learning directs labels toward the most uncertain regions, giving $\sim 3\times$ lower label cost for equal quality.

10.4. Limitations and Future Work

1. **Sample size and statistics.** Experimental validation uses ten prints (two per problem), each tested three times. The reported confidence intervals and per-property R^2 values establish a robust indicative validation, but the campaign remains modest; a larger study across more prints and printers would further tighten the intervals and test reproducibility across machines.
2. **Two-dimensional plane-stress assumption.** The framework is presented in 2D plane stress. Extension to 3D requires a 6×6 Voigt tensor and a 3D flow decoder; both are straightforward in principle but more computationally demanding.
3. **FDM-induced anisotropy.** Layer-by-layer deposition introduces base polymer anisotropy not captured by the current isotropic base model; incorporating print direction as a design variable is a natural extension.
4. **Limited family coverage.** Five families do not exhaust the design space; the pipeline is extensible (Section 5.1), but additional families and resolution-independent decoders (e.g., neural fields) are future work.
5. **Linear elasticity.** The small-strain assumption excludes large deformation, buckling, and dynamics, which would require new data or physics-informed losses.

Practical implications

For practitioners, the framework provides a millisecond-scale full-tensor surrogate and an inverse decoder that respects manufacturability, enabling target-driven design of auxetic, anisotropic, and chiral polymer parts within standard FDM workflows. The baseline-reproduction protocol (Section 7) makes the reported gains auditable.

11. Conclusions

We presented a unified framework (OPERA) that predicts the full 3×3 plane-stress stiffness tensor with symmetry and positive-definiteness enforced by construction, replaces a fixed decoder with a jointly trained normalizing flow to keep inverse design on the training manifold (reducing the surrogate–PDE gap from $>30\%$ to $<6\%$), and uses a diverse five-family dataset with expected-improvement active learning. On five design problems, the average target error is 6.8%, and predictions agree with tensile measurements on ten FDM specimens at a pooled $R^2 = 0.987$. We report these results relative to a reproduced scalar-proxy baseline and provide explicit uncertainty, baseline-provenance, and assumption analyses. Figure 12 consolidates the contributions: the capability radar in panel (b) shows the present framework expanding over the baseline in tensor coverage, inverse accuracy, and property-space coverage while matching it in forward accuracy and speed, and panels (c,d) restate the five-problem errors and the head-to-head comparison. Together, these results indicate that full-tensor prediction, an on-manifold generative decoder, and acquisition-driven data design are complementary ingredients for reliable inverse design of polymer metamaterials.




Figure 12. Summary. (a) Key contributions and results. (b) Capability radar: present framework (teal) vs. reproduced baseline (grey). (c) Design error for all five problems (legend above panel). (d) Comparison table.

12. Methods

Dataset generation

Microstructures were generated on a 32×32 grid. For each of five families, 1000 instances were created with uniform coverage of $\rho \in [0.10, 0.95]$ (rejection sampling). Effective properties were computed by finite element homogenization (bilinear 4-node elements, plane stress, 32×32 mesh, periodic boundary conditions). Base polymer: $E_0 = 3.0$ GPa, $\nu_0 = 0.38$ (PLA); void $E_{\text{void}} = 10^{-4}E_0$. Three cell problems per microstructure provide all six Voigt components.

CNN architecture

Four convolutional blocks ($32 \rightarrow 64 \rightarrow 128 \rightarrow 256$), 3×3 kernels, ReLU, batch normalization, stride-2 max pooling; global average pooling to 256-d; three FC layers ($256 \rightarrow 128 \rightarrow 64 \rightarrow 6$). Symmetry layer and PD projection are non-parametric. Total: 2.1×10^6 parameters.

Normalizing flow

$K = 8$ affine coupling layers, s_ϕ and t_ϕ as two-layer CNNs (32 channels); latent dim $d_z = 128$; decoder: two transposed convolutions $128 \rightarrow 32 \times 32$. Adam, $\eta = 10^{-4}$, batch 64, 200 epochs.

Inverse design

L-BFGS-B, $T_{\text{inv}} = 20$, $\eta_z = 0.1$; FDM filter radius $r = 2$ px ($= 0.5$ mm at 0.25 mm/px); Heaviside $\beta_H \in \{1, 2, 4, 8, 16, 32\}$; $\lambda_{\text{conn}} = 10$, $\lambda_{\text{bin}} = 1$.

Experimental validation

See Section 8.5.

Author Contributions: Conceptualization, K.E. and I.G.; methodology, K.E.; software, K.E.; validation, K.E. and I.G.; formal analysis, K.E.; investigation, K.E. and I.G.; resources, K.E.; data curation, K.E.; writing—original draft preparation, K.E.; writing—review and editing, K.E. and I.G.; visualization, K.E.; supervision, I.G.; project administration, K.E. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: Training data (10,000 microstructures + effective tensors), model weights, and design outputs are available from the corresponding author on reasonable request. The PyTorch 2.5.1 implementation (physics-embedded CNN, normalizing-flow decoder, and active-learning loop), together with the reproduced scalar-proxy baseline used in Section 7, will be deposited in a public repository with a citable DOI upon acceptance.

Conflicts of Interest: The authors declare no competing interests.

Appendix A. Proof of Theorem (Inverse-Design Consistency)

Proof. Let $\chi^* = \mathcal{F}_\phi(z^*)$. By the triangle inequality,

$$\|\hat{\mathcal{H}}_\theta(\chi^*) - \mathcal{H}(\chi^*)\| \leq \|\hat{\mathcal{H}}_\theta(\chi^*) - \hat{\mathcal{H}}_\theta(\chi_{\text{train}}^*)\| + \|\hat{\mathcal{H}}_\theta(\chi_{\text{train}}^*) - \mathcal{H}(\chi_{\text{train}}^*)\|, \quad (\text{A1})$$

where χ_{train}^* is the nearest training point in $L^2(Y)$. The second term is bounded by ε_{fwd} . Since $\hat{\mathcal{H}}_\theta$ is Lipschitz with constant C_{lip} (Section 4.2), the first term is bounded by $C_{\text{lip}}\|\chi^* - \chi_{\text{train}}^*\|_{L^2}$. By Pinsker's inequality and the Hellinger–TV bound, $\|\chi^* - \chi_{\text{train}}^*\|_{L^2}^2 \leq 2d_{\text{TV}}(p_\phi, p_{\text{data}}) \leq 2\varepsilon_{\text{flow}}^{1/2}$. Combining yields Equation (6). $\square$

Appendix B. Voigt Notation and Symmetry Classes

For plane stress, the 3×3 Voigt stiffness relates $(\sigma_{11}, \sigma_{22}, \sigma_{12})$ to $(E_{11}, E_{22}, 2E_{12})$, as in Equation (2). Symmetry classes: **triclinic** (all six components nonzero, general chiral); **orthotropic** ($C_{16} = C_{26} = 0$); **isotropic** ($C_{11} = C_{22}$, $C_{66} = (C_{11} - C_{12})/2$, $C_{16} = C_{26} = 0$).

Appendix C. Hyperparameters

Table A1. Hyperparameters.

Module	Parameter	Value
CNN (forward)	Channels	$32 \rightarrow 64 \rightarrow 128 \rightarrow 256$
	Kernel	3×3
	Optimizer/LR	Adam/ 10^{-4}
	Batch/Epochs	64/200
Flow (decoder)	Coupling layers K	8
	Latent dim d_z	128
	Joint weight μ	0.1
Inverse	L-BFGS-B steps	20
	LR η_z	0.1
	Heaviside β_H	$\{1, 2, 4, 8, 16, 32\}$
	$\lambda_{\text{conn}}, \lambda_{\text{bin}}$	10, 1
Active learning	EI exploration ζ	0.01
	GP kernel	Matérn 5/2
	Queries per round	50

References

- Alibert, J.J.; Seppecher, P.; dell’Isola, F. Truss modular beams with deformation energy depending on higher displacement gradients. *Math. Mech. Solids* **2003**, *8*, 51–73. [\[CrossRef\]](#)
- Lakes, R. Foam structures with a negative Poisson’s ratio. *Science* **1987**, *235*, 1038–1040. [\[CrossRef\]](#)
- Bertoldi, K.; Vitelli, V.; Christensen, J.; van Hecke, M. Flexible mechanical metamaterials. *Nat. Rev. Mater.* **2017**, *2*, 17066. [\[CrossRef\]](#)
- Barchiesi, E.; Spagnuolo, M.; Placidi, L. Mechanical metamaterials: A state of the art. *Math. Mech. Solids* **2019**, *24*, 212–234.
- dell’Isola, F.; Moschini, S.; Turco, E. Towards the synthesis of planar beams whose deformation energy depends on the third gradient of displacement. *Math. Mech. Complex Syst.* **2024**, *12*, 573–597. [\[CrossRef\]](#)
- Murcia Terranova, L. Planar one-dimensional continua whose energy depends on the gradient of curvature: An overview of their applications in metamaterials design. *Math. Mech. Complex Syst.* **2025**, *13*, 127–166. [\[CrossRef\]](#)
- Ngo, T.D.; Kashani, A.; Imbalzano, G.; Nguyen, K.T.Q.; Hui, D. Additive manufacturing (3D printing): A review of materials, methods, applications and challenges. *Compos. Part B Eng.* **2018**, *143*, 172–196. [\[CrossRef\]](#)
- Ozen, A.; Auhl, D.; Völlmecke, C.; Kiendl, J.; Abali, B.E. Optimization of manufacturing parameters and tensile specimen geometry for FDM 3D-printed PETG. *Materials* **2021**, *14*, 2556. [\[CrossRef\]](#) [\[PubMed\]](#)
- Afshar, R.; Jeanne, S.; Abali, B.E. Nonlinear material modeling for mechanical characterization of 3D-printed PLA polymer with different infill densities. *Appl. Compos. Mater.* **2023**, *30*, 987–1001. [\[CrossRef\]](#)
- Afshar, R.; Jeanne, S.; Abali, B.E. Effects of 3D-printing infill density on the mechanical properties of PLA polymer. In *Sixty Shades of Generalized Continua*; Springer: Berlin/Heidelberg, Germany, 2023; pp. 1–12.
- Aydin, G.; Yildizdag, M.E.; Abali, B.E. Strain-gradient modeling and computation of 3D-printed metamaterials. In *Theoretical Analyses, Computations, and Experiments of Multiscale Materials*; Springer: Berlin/Heidelberg, Germany, 2022; pp. 343–357.
- Sarar, B.C.; Yildizdag, M.E.; Abali, B.E. A multi-scale homogenization framework for design and strain-gradient modeling of additively manufactured parts. *Continuum Mech. Thermodyn.* **2024**, *36*, 1629–1643. [\[CrossRef\]](#)
- Abali, B.E.; Barchiesi, E. Additive manufacturing introduced substructure and computational determination of metamaterials parameters by asymptotic homogenization. *Continuum Mech. Thermodyn.* **2021**, *33*, 993–1009.
- Evans, K.E.; Alderson, A. Auxetic materials: Functional materials and structures from lateral thinking. *Adv. Mater.* **2000**, *12*, 617–628. [\[CrossRef\]](#)
- Stilz, M.; Plappert, D.; Gutmann, F.; Hiermaier, S. A 3D extension of pantographic geometries to obtain metamaterial with semi-auxetic properties. *Math. Mech. Solids* **2022**, *27*, 673–686.
- Yilmaz, N.; Placidi, L.; Misra, A. Macroscale chirality in planar elasticity: Predictions of granular micromechanics-based extended micropolar model. *J. Mech. Phys. Solids* **2026**, *194*, 106664. [\[CrossRef\]](#)
- Placidi, L.; Misra, A.; Valle, G.L.; Rodriguez, C. Revised identification of strain-gradient elastic parameters. *Z. Angew. Math. Mech.* **2026**, *106*, e70335. [\[CrossRef\]](#)
- Misra, A.; Lekszycki, T.; Giorgio, I.; Ganzosch, G.; Müller, W.H.; dell’Isola, F. Pantographic metamaterials show atypical Poynting effect reversal. *Mech. Res. Commun.* **2018**, *89*, 6–10. [\[CrossRef\]](#)
- Auger, P.; Lavigne, T.; Smaniotto, B.; Spagnuolo, M.; dell’Isola, F.; Hild, F. Poynting effects in pantographic metamaterial captured via multiscale DVC. *J. Strain Anal. Eng. Des.* **2021**, *56*, 462–477.
- dell’Isola, F.; Lekszycki, T.; Spagnuolo, M.; Alibert, J.J.; Della Corte, A.; Giorgio, I.; Hild, F.; Laudato, M.; Madeleine, P.; Placidi, L.; et al. Experimental methods for pantographic structures. In *Discrete Models and Continuum Descriptions for Complex Metamaterials*; Springer: Cham, Switzerland, 2019; pp. 263–297.
- Ghorbani, A.; Dykstra, D.; Coulais, C.; Bonn, D.; van der Linden, E.; Habibi, M. Inverted and programmable Poynting effects in metamaterials. *Adv. Sci.* **2021**, *8*, 2102279. [\[CrossRef\]](#)
- dell’Isola, F.; Seppecher, P.; Alibert, J.J.; Lekszycki, T.; Grygoruk, R.; Pawlikowski, M.; Steigmann, D.; Giorgio, I.; Andreaus, U.; Cuomo, M.; et al. Pantographic metamaterials: An example of mathematically driven design and of its technological challenges. *Continuum Mech. Thermodyn.* **2019**, *31*, 851–884. [\[CrossRef\]](#)
- Spagnuolo, M.; Franciosi, P.; dell’Isola, F. A Green-operator-based elastic modeling for two-phase pantographic-inspired bi-continuous materials. *Int. J. Solids Struct.* **2020**, *188*, 282–308. [\[CrossRef\]](#)
- Barchiesi, E.; Eugster, S.R.; dell’Isola, F.; Hild, F. Large in-plane elastic deformations of bi-pantographic fabrics: Asymptotic homogenization and experimental validation. *Math. Mech. Solids* **2020**, *25*, 739–767. [\[CrossRef\]](#)
- Turco, E.; dell’Isola, F.; Rizzi, N.L.; Grygoruk, R.; Müller, W.H.; Liebold, C. Fiber rupture in sheared planar pantographic sheets: Numerical and experimental evidence. *Mech. Res. Commun.* **2016**, *76*, 86–90. [\[CrossRef\]](#)
- Eremeyev, V.A.; dell’Isola, F.; Boutin, C.; Steigmann, D. Linear pantographic sheets: Existence and uniqueness of weak solutions. *J. Elast.* **2018**, *132*, 175–196.
- Fedele, R. Approach à la Piola for the equilibrium problem of bodies with second-gradient energies. Part II. *Continuum Mech. Thermodyn.* **2022**, *34*, 1087–1111. [\[CrossRef\]](#)

28. dell’Isola, F.; Fedele, R. Irreducible representation of surface distributions and Piola transformation of external loads sustainable by third-gradient continua. *C. R. Mec.* **2023**, *351*, 91–120.
29. Sigmund, O. On the design of compliant mechanisms using topology optimization. *Mech. Struct. Mach.* **1997**, *25*, 493–524. [[CrossRef](#)]
30. Bertoldi, K.; Reis, P.M.; Willshaw, S.; Mullin, T. Negative Poisson’s ratio behavior induced by an elastic instability. *Adv. Mater.* **2010**, *22*, 361–366. [[CrossRef](#)] [[PubMed](#)]
31. Hussein, M.I.; Leamy, M.J.; Ruzzene, M. Dynamics of phononic materials and structures. *Appl. Mech. Rev.* **2014**, *66*, 040802. [[CrossRef](#)]
32. Gibson, L.J.; Ashby, M.F. *Cellular Solids: Structure and Properties*, 2nd ed.; Cambridge University Press: Cambridge, UK, 1997.
33. Bensoussan, A.; Lions, J.-L.; Papanicolaou, G. *Asymptotic Analysis for Periodic Structures*; North-Holland: Amsterdam, The Netherlands, 1978.
34. Guo, K.; Yang, Z.; Yu, C.-H.; Buehler, M.J. Artificial intelligence and machine learning in design of mechanical materials. *Mater. Horiz.* **2021**, *8*, 1153–1172. [[CrossRef](#)] [[PubMed](#)]
35. Liu, X.; Cui, J.; Zheng, Z.; Yue, S.; Li, E.; Zhang, D.; Zhou, J. Machine-learning-based design and optimization of curved beams for multistable structures and metamaterials. *Extreme Mech. Lett.* **2021**, *41*, 101002.
36. Ha, C.S.; Yao, D.; Xu, Z.; Liu, Y.; Ning, H.; Liu, T.; Li, J.; Wang, L.; Wang, C.; Hu, N.; et al. Rapid inverse design of metamaterials based on prescribed mechanical behavior through machine learning. *Nat. Commun.* **2023**, *14*, 5765. [[CrossRef](#)] [[PubMed](#)]
37. Zheng, X.; Zhang, X.; Chen, T.-T.; Watanabe, I. Deep learning in mechanical metamaterials: From prediction and generation to inverse design. *Adv. Mater.* **2023**, *35*, 2302530. [[CrossRef](#)]
38. Lee, D.; Chen, W.; Wang, L.; Chan, Y.-C.; Chen, W. Data-driven design for metamaterials and multiscale systems: A review. *Adv. Mater.* **2024**, *36*, 2305254.
39. Liu, W.; Xu, G.; Fan, W.; Lyu, M.; Xia, Z. Machine-learning-based characterization and inverse design of metamaterials. *Materials* **2024**, *17*, 3512. [[CrossRef](#)] [[PubMed](#)]
40. Hao, W.; Du, Z.; Hou, X.; Guo, Y.; Liu, C.; Zhang, W.; Gao, H.; Guo, X. Intelligent design of mechanical metamaterials: A GCNN-based structural genome database approach. *Natl. Sci. Rev.* **2025**, *12*, nwaf053. [[CrossRef](#)] [[PubMed](#)]
41. Xie, T.; Grossman, J.C. Crystal graph convolutional neural networks for accurate and interpretable prediction of material properties. *Phys. Rev. Lett.* **2018**, *120*, 145301. [[CrossRef](#)] [[PubMed](#)]
42. Chen, C.T.; Gu, G.X. Machine learning for composite materials. *MRS Commun.* **2019**, *9*, 556–566. [[CrossRef](#)]
43. Zheng, X.; Chen, T.-T.; Guo, X.; Samitsu, S.; Watanabe, I. Controllable inverse design of auxetic metamaterials using deep learning. *Mater. Des.* **2021**, *211*, 110178. [[CrossRef](#)]
44. Brzin, T.; Brojan, M. Using a generative adversarial network for the inverse design of soft morphing composite beams. *Eng. Appl. Artif. Intell.* **2024**, *133*, 108527. [[CrossRef](#)]
45. Rezaei, S.; Baghani, M.; Martínez-Pañeda, E. A finite operator learning technique for mapping the elastic properties of microstructures to their mechanical deformations. *arXiv* **2024**, arXiv:2404.00074.
46. Mele, M.; Milan, G.; Paffetti, A.; Bacciaglia, A.; Curto, M.; Ghiotti, A.; Bruschì, S. Homogenization and artificial neural network prediction of elastic properties in triply periodic minimal surface structures. *Prog. Addit. Manuf.* **2025**, *10*, 9337–9353. [[CrossRef](#)]
47. Tran, T.V.; Nanthakumar, S.S.; Zhuang, X.; Li, X.; Nguyen-Thoi, T.; Rabczuk, T. On-demand inverse design of metamaterials using deep neural networks with Bayesian optimization. *Intell. Comput.* **2025**, *4*, 0139. [[CrossRef](#)]
48. Lazarov, B.S.; Sigmund, O. Filters in topology optimization based on Helmholtz-type differential equations. *Int. J. Numer. Methods Eng.* **2011**, *86*, 765–781.
49. Wang, F.; Lazarov, B.S.; Sigmund, O. On projection methods, convergence and robust formulations in topology optimization. *Struct. Multidiscip. Optim.* **2011**, *43*, 767–784. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.