

ORIGINAL ARTICLE

Robust inverse material design with physical guarantees using the Voigt–Reuss Net

Sanath Keshav | Felix Fritzen

¹Institute of Applied Mechanics, Data Analytics in Engineering, University of Stuttgart, Universitätsstr.32, 70569 Stuttgart, Baden-Württemberg, Germany

Correspondence

Corresponding author: Felix Fritzen
Email: fritzen@mib.uni-stuttgart.de

Present address

Universitätsstr. 32, 70569 Stuttgart, Germany

Abstract

We propose a spectrally normalized surrogate model for forward and inverse mechanical homogenization that guarantees physical admissibility of effective stiffness tensors throughout training, inference, and optimization. Building on classical Voigt–Reuss bounds, we construct a universal reparameterization of symmetric positive definite constitutive tensors: the difference between the Voigt and Reuss bounds is factorized into a Cholesky-like operator, and each effective tensor is represented by a dimensionless, symmetric positive semi-definite matrix with spectrum confined to the unit interval. Learning is performed in this bounded spectral space, and the inverse mapping ensures that all predictions are symmetric positive definite and lie between the Voigt and Reuss bounds in the Löwner sense. We first apply this framework to 3D linear elasticity using an open dataset of stochastic biphasic microstructures. A fully connected Voigt–Reuss net is trained on over 750K FFT-based homogenization labels using 236 isotropic morphological descriptors and three non-dimensional elastic contrast parameters. Because the inputs are rotationally invariant, the surrogate recovers the isotropic projection of the effective stiffness tensor with near-perfect fidelity: isotropy-related components achieve $R^2 > 0.99$, while anisotropy-revealing entries are correctly identified as unpredictable. At the tensor level, the relative Frobenius error has a median of $\approx 1.7\%$ and mean $\approx 3.4\%$ across train/validation/test, operating close to the irreducible error floor set by isotropic projection and far below classical Hill-type estimates. We then couple spectral normalization with a differentiable microstructure renderer and a convolutional network for 2D plane-strain elasticity of highly anisotropic, high-contrast biphasic composites generated by thresholded trigonometric fields. The resulting end-to-end differentiable convolutional Voigt–Reuss net maps amplitude matrices and thresholds to effective stiffness tensors, achieving $R^2 > 0.99$ on all components and subpercent Voigt–Reuss normalized losses, while remaining strictly within the bounds. The surrogate captures percolation-induced eigenvalue jumps along threshold sweeps and generalizes robustly to out-of-distribution, diffusion-generated microstructures. Treating the parametric microstructure as design variables, we solve inverse problems via batched, first-order optimization using a single surrogate, matching target tensors with relative errors of only a few percent and producing diverse near-optimal designs. Overall, the Voigt–Reuss net provides a physically grounded surrogate for multiscale elasticity that unifies accurate forward prediction with large-batch, constraint-consistent inverse materials design. The approach is generic to elliptic problems and readily extensible to other transport and coupled-physics settings where rigorous bounds are available.

KEY WORDS

inverse materials design, data-driven models with physical guarantees, physics-augmented neural networks, homogenization

1 | INTRODUCTION

1.1 | Motivation

The mechanical response of heterogeneous materials is governed by a subtle interplay between phase properties, volume fractions, and microstructural geometry across multiple scales. From fiber-reinforced polymers to architected lattices and metamaterials, designers are increasingly seeking to engineer microstructures to tailor effective stiffness, anisotropy, and load transfer pathways for specific applications. Classical homogenization theory provides a rigorous framework for relating microstructure and local constitutive laws to effective properties, e.g., via asymptotic expansions, variational principles, and empirical estimates and bounds^{1,2,3,4}. Yet, the corresponding *G-closure*^{5,6}, i.e., the set of all effective tensors realizable by mixing given phases in all admissible microgeometries, is typically high-dimensional, non-convex, and analytically intractable.

Computational homogenization and direct numerical simulation have become the workhorses to explore the material design space in practice. Finite element (FE) and FFT-based schemes^{7,8,9,10} allow one to solve corrector problems on intricate unit cells and extract homogenized tensors. These methods are robust and systematically improvable, but they are also computationally expensive: A single 3D simulation on a 192^3 voxel grid with strong phase contrast involves ≈ 21 million degrees of freedom. Therefore, it can already be costly, and design tasks that require exploring tens of thousands of candidate microstructures or solving nested optimization problems are often prohibitive if each evaluation relies on high-fidelity simulation. More strikingly, differentiating these results with respect to design parameters is challenging, if not impossible, in many situations. These bottlenecks have motivated a surge of interest in surrogate models that learn microstructure property relationships from data and can be queried at negligible cost once trained.

1.2 | Machine learned effective property forecasting

Machine-learning-based surrogates for homogenization are now widely explored. Approaches range from regression on hand-crafted descriptors (e.g., morphological statistics, correlation measures, or spectral features)^{11,12,13} to deep convolutional networks ingesting microstructure images or volumes directly^{14,15,16}. They have been used to predict scalar effective properties¹², full elasticity tensors¹⁷, and even local fields^{18,19}. In parallel, data-driven paradigms in computational mechanics^{20,21,22} have demonstrated that learned constitutive maps can replace or augment traditional phenomenological models, enabling faster simulations and design. Furthermore, these works have demonstrated the need to incorporate physical constraints into the surrogates to leverage data efficiency, robustness, and out-of-data prediction accuracy. On the design side, several works now combine surrogates with generative models²³, optimization^{24,25,26}, learned inverse mappings²⁷, or autoencoders²⁸ to tackle inverse problems, where one seeks microstructures or process parameters that achieve targeted effective responses.

Despite these advances, most existing surrogates for effective tensors suffer from a critical limitation: they treat the components of the effective tensor as generic regression targets, with physical constraints imposed (if at all) through post-processing or soft penalization. In linear elasticity, the effective stiffness tensor must be symmetric, positive definite, and satisfy basic variational bounds, most notably the first-order Voigt and Reuss bounds, which constrain the admissible tensors in the Löwner sense⁴. Naive regression does not respect such structural constraints: predicted tensors can be indefinite, violate ellipticity, or lie outside the Voigt-Reuss envelope, especially when queried far from the training distribution. Such violations are more than just aesthetic flaws: they can destabilize downstream FE simulations, break thermodynamic consistency, and render inverse-design loops unreliable. Common ad hoc remedies include projecting predictions onto the cone of symmetric positive definite (SPD) matrices, learning a Cholesky decomposition²⁹, clamping eigenvalues, or adding penalty terms to the loss in the spirit of physics-infused neural networks (PINNs³⁰). None of these approaches guarantees adherence to Voigt-Reuss bounds, and often this leads to ill-conditioned optimization landscapes.

At the same time, homogenization theory provides a much richer structure than simple SPD constraints. For a given set of phase tensors and volume fractions, the effective tensor must lie in the intersection of the SPD cone with a set of variational bounds such as the classical Voigt and Reuss bounds^{31,32,33}. Under some assumptions, the Hashin-Shtrikman bounds³⁴, and more refined and problem-specific estimates^{35,36} can apply, as well. In Löwner order, the basic Voigt-Reuss inequality states that the effective tensor is bounded above and below by those of the weighted arithmetic and harmonic averages, respectively. Geometrically, this defines a compact spectral simplex in the space of SPD matrices where the difference between the Voigt

and Reuss bounds is positive semi-definite and encodes the admissible variation of effective responses. Any surrogate used for forward prediction or inverse design *must* remain in the thereby defined feasible set. In particular, any configurations explored along optimization trajectories must be inside this envelope.

Another source of difficulty is the intrinsic complexity of the G-closure, especially for high-contrast and anisotropic composites. For strongly heterogeneous materials, small changes in geometry can cause large, non-smooth changes in effective response, particularly near percolation thresholds or when load-bearing paths reconfigure^{1,37}. Effective stiffness eigenvalues can exhibit jumps as microstructure parameters (e.g., volume fraction, connectivity, or orientation) cross critical values. Mean-field approaches, such as e.g., the Mori-Tanaka scheme⁴ and self-consistent schemes, which implicitly assume dilute or weakly interacting inclusions, cannot capture these transitions. Standard neural surrogates can approximate them, but when unconstrained, they may also generate unphysical responses, e.g., in poorly sampled regions of the design space. Robust inverse design thus demands surrogates that are both expressive enough to capture non-smooth, geometry-driven behavior and structured enough to enforce physical admissibility by construction.

In a recent work³⁸, we proposed a spectral normalization framework for learning effective conductivities resulting from the solutions of elliptic problems. The core idea is to exploit the Voigt-Reuss gap as a natural scaling for the admissible response and to reparameterize the effective tensor via a normalized SPD matrix with spectrum confined to the unit interval. Learning is then carried out in the bounded unit-cube space of eigenvalues and orthogonal factors of the normalized tensor. The reconstruction of the constitutive tensor based on the normalized representation automatically enforces symmetry, positive definiteness, and Voigt-Reuss admissibility. That prior study focused on scalar-valued elliptic problems such as thermal conduction.

1.3 | Objectives and scope.

In this paper, we generalize the Voigt-Reuss net framework to linear elasticity and to demonstrate how a spectrally normalized surrogate can be used not only for fast, physically admissible forward prediction of effective stiffness tensors, but also as a robust engine for gradient-based inverse material design. Building on our earlier study for scalar elliptic problems³⁸, we first formalize a spectral normalization for generic symmetric positive definite constitutive tensors (of arbitrary order), including 3D small-strain elasticity as a special case. Any effective tensor is mapped through a Cholesky-like factor of the Voigt-Reuss difference into a dimensionless matrix with spectrum in $[0, 1]$, yielding a task-agnostic target space with scale-free, well-conditioned losses and the guarantee that, after inverse mapping, all surrogate predictions are symmetric positive definite and confined to the Voigt-Reuss envelope.

Within this general framework, we consider two representative elasticity settings. First, we construct a Voigt-Reuss net for 3D biphasic composites based on the open microstructure dataset^{39,12}, where each sample is described by isotropic averaged morphology descriptors and non-dimensional phase contrast parameters. This case probes the performance of spectral normalization in a purely descriptor-based, rotationally invariant forward problem. Second, we couple spectral normalization with a differentiable microstructure renderer and a convolutional neural network for 2D plane-strain composites generated from thresholded trigonometric fields^{40,37}. In this setting, the surrogate operates directly on parametric microstructure representations, enabling end-to-end differentiability of the map from design parameters to effective stiffness and, consequently, first-order inverse design with respect to full target tensors or tensor-valued functionals. Throughout, we restrict attention to linear elastostatics, but emphasize that the proposed spectral normalization is generic and applicable to other elliptic operators (e.g., thermal conductivity, diffusion, poroelasticity) and further.

The remainder of the paper is organized as follows. In Sections 3.1 and 3.2 we recall the Voigt-Reuss bounds for elliptic problems and derive the associated spectral normalization, including its matrix formulation and numerical implementation. Section 3.3 discusses how this normalization can be combined with generic machine-learning models. Section 4 presents the 3D mechanical homogenization study based on descriptor inputs, while Section 5 introduces the differentiable microstructure renderer, the convolutional Voigt-Reuss net, and a series of forward and inverse design experiments for highly anisotropic, high-contrast 2D composites. The conclusions summarize the main findings and outline opportunities for extending Voigt-Reuss net to richer G-closure problems and multi-physics design.

2 | NOTATION

The spatial average of a quantity over a domain $\mathcal{A}$ with measure $A = |\mathcal{A}|$ is defined as

$$\langle \cdot \rangle_{\mathcal{A}} = \frac{1}{A} \int_{\mathcal{A}} (\cdot) \, dA . \quad (1)$$

In the sequel, boldface lowercase letters denote vectors, boldface uppercase letters denote 2-tensors, and blackboard bold uppercase letters (e.g., $\mathbb{C}$) denote 4-tensors. We use single $(\cdot)$ and double contractions $(:)$. With respect to an orthonormal basis $\mathbf{e}_i$ ($i \in \{1, 2, 3\}$), they are defined as ($\mathbf{A}$: order n tensor; $\mathbf{B}$: order m tensor)

$$\begin{aligned} \mathbf{A} &= A_{i_1 i_2 \dots i_{n-1} i_n} \mathbf{e}_{i_1} \otimes \dots \otimes \mathbf{e}_{i_n}, & \mathbf{B} &= B_{j_1 j_2 \dots j_{m-1} j_m} \mathbf{e}_{j_1} \otimes \dots \otimes \mathbf{e}_{j_m}, \\ \mathbf{A} \cdot \mathbf{B} &= A_{i_1 i_2 \dots i_{n-1} k} B_{k j_2 \dots j_m} \mathbf{e}_{i_1} \otimes \dots \otimes \mathbf{e}_{i_{n-1}} \otimes \mathbf{e}_{j_2} \otimes \dots \otimes \mathbf{e}_{j_m}, \\ \mathbf{A} : \mathbf{B} &= A_{i_1 i_2 \dots i_{n-2} kl} B_{kl j_3 \dots j_m} \mathbf{e}_{i_1} \otimes \dots \otimes \mathbf{e}_{i_{n-2}} \otimes \mathbf{e}_{j_3} \otimes \dots \otimes \mathbf{e}_{j_m}. \end{aligned}$$

using Einstein summation over repeated indices. The tensor product is denoted by $\otimes$. Using the above notation in conjunction with the operator

$$\nabla = \frac{\partial}{\partial x_i} \mathbf{e}_i ,$$

The (right) divergence and the (right) gradient of a general tensor read

$$\mathbf{A} \cdot \nabla = A_{i_1 \dots i_{n-1} j} \mathbf{e}_{i_1} \otimes \dots \otimes \mathbf{e}_{i_{n-1}}, \quad \mathbf{A} \otimes \nabla = A_{i_1 \dots i_n j} \mathbf{e}_{i_1} \otimes \dots \otimes \mathbf{e}_{i_n} \otimes \mathbf{e}_j .$$

Relative Frobenius error is given by,

$$\mathcal{E}_F(\underline{\underline{A}}, \underline{\underline{B}}) = \frac{\|\underline{\underline{A}} - \underline{\underline{B}}\|_F}{\|\underline{\underline{A}}\|_F} = \frac{\sqrt{\sum_{i,j}^{m,n} (A_{ij} - B_{ij})^2}}{\sqrt{\sum_{i,j}^{m,n} A_{ij}^2}}, \quad \text{where } \underline{\underline{A}}, \underline{\underline{B}} \in \mathbb{R}^{m \times n}. \quad (2)$$

Generally, $\hat{(\cdot)}$ denotes predicted quantities, $\tilde{(\cdot)}$ denotes normalized quantities, and $\overline{(\cdot)}$ denotes homogenized quantities.

For convenience, a six-dimensional orthonormal basis $\{\mathbf{B}_1, \dots, \mathbf{B}_6\}$ of the space of symmetric second-order tensors $Sym(\mathbb{R}^3)$ in a Mandel-like notation is proposed:

$$\begin{aligned} \mathbf{B}_1 &= \mathbf{e}_1 \otimes \mathbf{e}_1, & \mathbf{B}_2 &= \mathbf{e}_2 \otimes \mathbf{e}_2, & \mathbf{B}_3 &= \mathbf{e}_3 \otimes \mathbf{e}_3, \\ \mathbf{B}_4 &= \frac{\sqrt{2}}{2} (\mathbf{e}_1 \otimes \mathbf{e}_2 + \mathbf{e}_2 \otimes \mathbf{e}_1), & \mathbf{B}_5 &= \frac{\sqrt{2}}{2} (\mathbf{e}_1 \otimes \mathbf{e}_3 + \mathbf{e}_3 \otimes \mathbf{e}_1), & \mathbf{B}_6 &= \frac{\sqrt{2}}{2} (\mathbf{e}_2 \otimes \mathbf{e}_3 + \mathbf{e}_3 \otimes \mathbf{e}_2) . \end{aligned} \quad (3)$$

Then any symmetric second-order tensor $\mathbf{U} \in Sym(\mathbb{R}^3)$ and any fourth-order tensor $\mathbb{Y}$ with left and right sub-symmetry can be expressed by a vector $\underline{u} \in \mathbb{R}^6$ and a matrix $\underline{\underline{Y}} \in \mathbb{R}^{6 \times 6}$ via[‡]

$$u_i = \mathbf{U} : \mathbf{B}_i, \quad Y_{ij} = \mathbf{B}_i : \mathbb{Y} : \mathbf{B}_j .$$

The following identities hold for symmetric tensors of order 2 ($\mathbf{U}, \mathbf{V}$) and 4 ($\mathbb{C}, \mathbb{S}$) and their respective vector representations $\underline{u}, \underline{v}, \underline{\underline{C}}, \underline{\underline{S}}$:

$$\mathbf{U} \cdot \mathbf{V} \leftrightarrow \underline{u}^T \underline{v} = \underline{u} \cdot \underline{v} = \underline{v}^T \underline{u}, \quad \mathbb{C} : \mathbf{U} \leftrightarrow \underline{\underline{C}} \underline{u}, \quad \mathbb{C}^{-1} : \mathbf{V} \leftrightarrow \underline{\underline{C}}^{-1} \underline{v} .$$

[‡] In two dimensions vectors are of size 3 and matrices are 3×3 , with basis tensor $\tilde{\mathbf{B}}_1 = \mathbf{B}_1, \tilde{\mathbf{B}}_2 = \mathbf{B}_2, \tilde{\mathbf{B}}_3 = \mathbf{B}_4$.

3 | FOUNDATIONS

3.1 | Löwner bounds for elliptic problems

We build on our recent work³⁸ and extend the framework to 3D mechanical problems. We consider spectral normalization of a symmetric positive definite (SPD) constitutive tensor $\mathbb{Y}$ (of arbitrary order) that maps an input $\mathbf{X}$ and a general tensor $\mathbf{Z}$, i.e.

$$\mathbf{Z} \cdot \mathbb{Y} \cdot \mathbf{X} = \mathbf{X} \cdot \mathbb{Y} \cdot \mathbf{Z};, \quad \mathbf{X} \cdot \mathbb{Y} \cdot \mathbf{X} > 0; . \quad (4)$$

Let $\mathbf{X}$ be a kinematic quantity (e.g., strain or temperature gradient). Many stationary physical problems can be written as the elliptic equation

$$(\mathbb{Y} \cdot \mathbf{X}) \cdot \nabla = \mathbf{0}; . \quad (5)$$

This covers, for instance, elasticity, steady-state heat conduction, and diffusion. In a homogenization setting, we set the boundary condition $\langle \mathbf{X} \rangle = \mathbf{e}_i$, where $\mathbf{e}_i$ denotes the i -th unit tensor and i runs over all admissible loadings (e.g., $i \in 1, 2, 3$ for thermal problems or $i \in 1, \dots, 6$ for small-strain elasticity). The corresponding solutions are denoted by $\mathbf{X}_i$. With the effective fluxes $\bar{\mathbf{Q}}_i = \langle \mathbb{Y} \cdot \mathbf{X}_i \rangle$, the effective constitutive tensor $\bar{\mathbb{Y}}$, with respect to an orthonormal basis $\mathbf{e}_1, \dots, \mathbf{e}_m$ [§], admits the matrix representation

$$\bar{\mathbb{Y}} \leftrightarrow \underline{\bar{\mathbb{Y}}}, \quad (\underline{\bar{\mathbb{Y}}})_{ij} = \bar{\mathbf{Q}}_i \cdot \mathbf{e}_j \quad (i, j = 1, \dots, m). \quad (6)$$

Using standard variational principles for the primal and dual problems (see, e.g.,^{4,3}), the Voigt and Reuss bounds for $\bar{\mathbb{Y}}$ are given by the weighted arithmetic and harmonic means of $\mathbb{Y}$, respectively,

$$\bar{\mathbb{Y}}_V = \langle \mathbb{Y} \rangle, \quad \bar{\mathbb{Y}}_R = (\langle \mathbb{Y}^{-1} \rangle)^{-1}. \quad (7)$$

The ordering is understood in the Löwner sense:

$$\bar{\mathbb{Y}}_R \preceq \bar{\mathbb{Y}} \preceq \bar{\mathbb{Y}}_V \quad \Leftrightarrow \quad \bar{\mathbf{X}} \cdot \bar{\mathbb{Y}}_R \cdot \bar{\mathbf{X}} \leq \bar{\mathbf{X}} \cdot \bar{\mathbb{Y}} \cdot \bar{\mathbf{X}} \leq \bar{\mathbf{X}} \cdot \bar{\mathbb{Y}}_V \cdot \bar{\mathbf{X}}. \quad (8)$$

In the $m \times m$ matrix representation (e.g., $m = 3$ for 3D thermal problems and $m = 6$ for 3D small-strain elasticity), this reads

$$\underline{\bar{\mathbb{Y}}}_R \preceq \underline{\bar{\mathbb{Y}}} \preceq \underline{\bar{\mathbb{Y}}}_V \quad \Leftrightarrow \quad \bar{\mathbf{x}} \cdot \underline{\bar{\mathbb{Y}}}_R \cdot \bar{\mathbf{x}} \leq \bar{\mathbf{x}} \cdot \underline{\bar{\mathbb{Y}}} \cdot \bar{\mathbf{x}} \leq \bar{\mathbf{x}} \cdot \underline{\bar{\mathbb{Y}}}_V \cdot \bar{\mathbf{x}} \quad \forall \bar{\mathbf{x}} \in \mathbb{R}^m. \quad (9)$$

In the sequel, we use the matrix form for convenience.

3.2 | Spectral normalization

The difference between the upper and lower bounds is symmetric and positive semi-definite and thus admits a diagonalization

$$\underline{\underline{\mathbf{Q}}}_0 \underline{\underline{\Lambda}}_0 \underline{\underline{\mathbf{Q}}}_0^T = \underline{\underline{\mathbf{Y}}}_V - \underline{\underline{\mathbf{Y}}}_R. \quad (10)$$

In most situations, $\underline{\underline{\Lambda}}_0$ is positive definite. If not, truncating to eigenvalues greater than a numerical tolerance $\epsilon > 0$ [¶] yields

$$\tilde{\underline{\underline{\mathbf{Q}}}}_0 \tilde{\underline{\underline{\Lambda}}}_0 \tilde{\underline{\underline{\mathbf{Q}}}}_0^T = \underline{\underline{\mathbf{Y}}}_V - \underline{\underline{\mathbf{Y}}}_R + \mathcal{O}(\epsilon) = \underline{\underline{\mathbf{L}}} \underline{\underline{\mathbf{L}}}^T. \quad (11)$$

From that, we can define a matrix $\underline{\underline{\mathbf{L}}}$ closely related to Cholesky factorization and its pseudo-inverse $\underline{\underline{\mathbf{L}}}^+$ via

$$\underline{\underline{\mathbf{L}}} = \tilde{\underline{\underline{\mathbf{Q}}}}_0 \sqrt{\tilde{\underline{\underline{\Lambda}}}_0}, \quad \underline{\underline{\mathbf{L}}}^+ = \sqrt{\tilde{\underline{\underline{\Lambda}}}_0^{-1}} \tilde{\underline{\underline{\mathbf{Q}}}}_0^T. \quad (12)$$

[§] For instance, the basis $\mathbf{B}_1, \dots, \mathbf{B}_6$ from eq. (3) will be used for $\mathbf{e}_1, \dots, \mathbf{e}_6$ for elastic homogenization.

[¶] Note that ϵ is a purely numerical parameter; in exact algebra, one may set $\epsilon = 0$.

Here $\tilde{\bullet}$ is used to denote the column truncation of $\underline{\underline{Q}}$ and the column- and row-truncation of $\underline{\underline{\Lambda}}_0$ to strictly positive eigenpairs padded by zeros to maintain the dimension $m \times m$ for practical reasons. Note that $\underline{\underline{L}}$ and $\underline{\underline{L}}^+$ depend only on the trivial Voigt and Reuss bounds, respectively; i.e., they are independent of further microstructural features. Symmetric application of the pseudo-inverse[#] to eq. (9) yields

$$0 \leq \underline{\underline{\tilde{Y}}} = \underline{\underline{L}}^+ \left(\underline{\underline{Y}}_V - \underline{\underline{Y}} \right) \underline{\underline{L}}^T \leq \underline{\underline{L}}^+ \left(\underline{\underline{Y}}_V - \underline{\underline{Y}}_R \right) \underline{\underline{L}}^T = \underline{\underline{I}}. \quad (13)$$

The dedimensionalized linear operator $\underline{\underline{\tilde{Y}}}$ can be expressed through an orthogonal matrix $\underline{\underline{\tilde{Q}}} \in Orth(\mathbb{R}^m)$ and a diagonal matrix $\underline{\underline{\tilde{\Lambda}}} = \text{diag}(\xi_{\lambda_1}, \dots, \xi_{\lambda_m})$ with $\xi_{\lambda} \in [0, 1]^m$ via

$$\underline{\underline{\tilde{Y}}} = \underline{\underline{\tilde{Q}}} \underline{\underline{\tilde{\Lambda}}} \underline{\underline{\tilde{Q}}}^T.$$

Thereby, the effective constitutive output $\underline{\underline{\bar{Y}}}$ is expressed through

$$\underline{\underline{\bar{Y}}} = \underline{\underline{Y}}_V - \underline{\underline{L}} \left(\underline{\underline{\tilde{Q}}} \underline{\underline{\tilde{\Lambda}}} \underline{\underline{\tilde{Q}}}^T \right) \underline{\underline{L}}^T = \underline{\underline{Y}}_V - \underline{\underline{L}} \underline{\underline{\tilde{Y}}} \underline{\underline{L}}^T. \quad (14)$$

In conclusion, $\underline{\underline{\tilde{Y}}}$ and the trivial bounds $\underline{\underline{Y}}_V$ and $\underline{\underline{Y}}_R$ suffice to reconstruct $\underline{\underline{\bar{Y}}}$. As a result of the spectral normalization, it suffices to predict $\underline{\underline{\tilde{Y}}}$, being a symmetric positive semi-definite matrix devoid of physical dimension and with a spectrum confined to the unit interval.

Remark 1. For thermal problems, the conductivities in isotropic multi-phasic materials are always different, yielding strictly positive eigenvalues in $\underline{\underline{\Lambda}}_0$. In mechanical problems, the situation can become technically more intricate: Considering a biphasic isotropic elastic medium, the material parameters of the phases can be expressed through the bulk modulus K_i and the shear modulus G_i of phase $i \in \{0, 1\}$. In case either of the two matches, the difference eq. (11) will contain eigenvalues. For instance, a matching bulk modulus $K_0 = K_1$ induces

$$\overline{\mathbb{C}}_V - \overline{\mathbb{C}}_R = 2\overline{\Delta G} \mathbb{P}_2^{\text{iso}}, \quad \overline{\Delta G} = G_V - G_R, \quad \mathbb{P}_2^{\text{iso}} = \mathbb{I}^{\text{sym}} - \frac{\mathbf{I} \otimes \mathbf{I}}{3}. \quad (15)$$

The matrices $\underline{\underline{L}}$ and $\underline{\underline{L}}^+$ then get

$$\underline{\underline{L}} = \sqrt{2\overline{\Delta G} \mathbb{P}_2^{\text{iso}}}, \quad \underline{\underline{L}}^+ = \frac{1}{\sqrt{2\overline{\Delta G} \mathbb{P}_2^{\text{iso}}}} \mathbb{P}_2^{\text{iso}}, \quad \mathbb{P}_2^{\text{iso}} = \begin{bmatrix} 2/3 & -1/3 & -1/3 & 0 & 0 & 0 \\ -1/3 & 2/3 & -1/3 & 0 & 0 & 0 \\ -1/3 & -1/3 & 2/3 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}. \quad (16)$$

While theoretically straightforward, the robust computation of $\underline{\underline{L}}$ and $\underline{\underline{L}}^+$ must be carefully implemented and tested for such limit cases. Note that the restriction to isotropy in the above is for illustrative purposes only, i.e., the code should be able to cope with general anisotropy. For technical reasons, a numerical threshold $\epsilon > 0$ should be used to filter "near 0" eigenvalues in $\underline{\underline{\Lambda}}_0$.

3.3 | Implementation of the spectral normalization

The reparameterization through $\underline{\underline{\tilde{Q}}}$, $\underline{\underline{\tilde{\Lambda}}}$ can be combined with machine learning to predict $\underline{\underline{\tilde{Y}}} \approx \underline{\underline{\tilde{Y}}}$. Therefore, two distinct ingredients are needed (see also Figure 1):

- a vector-valued surrogate $\hat{\xi}_{\lambda} \in [0, 1]^m$ to model the eigenvalues in $\underline{\underline{\tilde{\Lambda}}}$ and
- a parameterized orthogonal matrix $\hat{\underline{\underline{\tilde{Q}}}}(\hat{\xi}_q) \in Orth(\mathbb{R}^m)$.

[#] PseudoInverse refers to the Moore-Penrose pseudo inverse⁴¹.

The first part is trivial to enforce, e.g., through a simple sigmoid function (or similar alternatives). Parameterization of the orthogonal matrix $\hat{\underline{\underline{Q}}}$ is more challenging and less commonly available. In our implementation using `pytorch`⁴², we deploy `torch.nn.parameterizations.orthogonal`^{||}. In practice, the black box surrogate model must provide $\hat{\underline{\underline{\xi}}}_q$, matching the parameterization used.

A relative error with respect to the range of admissible responses $\underline{Y}_V - \underline{Y}_R$ of an approximation $\hat{\underline{Y}}$ of $\underline{Y}$ can directly be obtained from the absolute error in the spectrally normalized $\hat{\underline{Y}}$:

$$\phi(\hat{\underline{Y}}, \underline{Y}) = \frac{1}{\sqrt{m}} \|\hat{\underline{Y}} - \underline{Y}\|_F, \quad (17)$$

where $\|\cdot\|_F$ denotes the Frobenius norm. This loss is insensitive to changes of unit (e.g., SI vs. imperial), and it can be combined with any other dedimensionalized loss. Therefore, we strongly recommend the use of a loss that is exclusively based on the normalized quantity $\hat{\underline{Y}}$ instead of the actual constitutive target $\underline{Y}$. Notably, the error is interpretable for each sample $\phi(\hat{\underline{Y}}, \underline{Y}) \in [0, 1]$, which gives a relative error of the stiffness normalized to the range of the bounds. This procedure is also in line with the problem-specific weighting of losses in⁴³, which was also shown to be crucial for success. Therefore, the loss will be based on eq. (17).

Remark 2. In our experience, it is prohibitive to compute the loss either based on the parameters $\hat{\underline{\xi}}_q$, or on $\hat{\underline{\xi}}_\lambda$ for the following reasons:

- Redundancy is possible, i.e. different $\hat{\underline{\xi}}_q$ can lead to the same $\hat{\underline{\underline{Q}}}$.
- Depending on the chosen parameterization, errors in the parameters are hard to interpret since the usual norms of $\hat{\underline{\xi}}_q$ do not correspond straightforwardly to errors in the orthogonal matrix.
- As for $\hat{\underline{\xi}}_\lambda$, the ordering of eigenvalues in the surrogate model could be different. The situation is even worse if eigenvalues show algebraic multiplicity, which renders the columns of $\hat{\underline{\underline{Q}}}$ nonunique.

These problems are eliminated when using the suggested loss eq. (17) or any other loss relating $\hat{\underline{Y}}$ to $\underline{Y}$.

As mentioned before, spectral normalization is independent of the type of machine-learning model to be used, i.e., it can be combined with neural networks (in their different flavors), kernel methods, random forests, and many more. The only two essential ingredients are that the outputs must be the vector $\hat{\underline{\xi}}_\lambda \in [0, 1]^m$ and the orthogonal matrix $\hat{\underline{\underline{Q}}} \in \text{Orth}(\mathbb{R}^m)$.

4 | APPLICATION TO 3D LINEAR MECHANICAL HOMOGENIZATION

Owing to the origins of the spectral normalization, we refer to our surrogate as Voigt-Reuss net and apply it in a 3D mechanical setting. The goal is to predict the effective fourth-order stiffness tensor $\hat{\underline{\underline{C}}} \in \text{Sym}^+(\mathbb{R}^6)$ of periodic, biphasic composites from geometric descriptors of the unit cell and phase moduli. All microstructures are drawn from the open 3D dataset of^{39,12}, which comprises nine stochastic model classes designed to span a broad morphological envelope (volume fraction, connectivity, length scales); see Figure 2. Each microstructure is represented by a 192^3 voxel image, and it is labeled by 236 image-derived scalar descriptors.**

We consider the isotropic bulk and shear moduli (K_1, G_1) for phase 1 and (K_2, G_2) for phase 2. To reduce dimensionality, we dedimensionalize the remaining parameters by K_1 , leading to three distinct parameters:

$$\pi_1 = \frac{K_2}{K_1}, \quad \pi_2 = \frac{G_1}{K_1}, \quad \pi_3 = \frac{G_2}{K_1}. \quad (18)$$

^{||} <https://pytorch.org/docs/stable/generated/torch.nn.utils.parameterizations.orthogonal.html>

** As in¹², these include porosity, specific surface area, constrictivity, chord-length and spherical contact distributions, geodesic tortuosity, two-point correlations, and size/shape statistics. They are predominantly isotropy-averaged descriptors; we return to this point below.




FIGURE 1 Two approaches for constructing the effective elasticity tensor $\bar{\underline{\underline{C}}}$ from microstructure: (a) direct numerical simulation (DNS) workflow, and (b) a machine-learning surrogate model employing spectral normalization to enforce two-sided L wner bounds.

Due to the linearity of the elasticity problem, any global scaling by K_1 can be applied a posteriori. To cover a wide range of contrasts, we sample

$$\pi_i = 10^{\pi_i^*}, \quad \pi_i^* \sim \mathcal{U}(-2, 2), \quad i = 1, 2, 3, \quad (19)$$

i.e., $\pi_i \in [10^{-2}, 10^2]$ is uniformly sampled in the logarithmic space.

For each microstructure, we draw triples (π_1, π_2, π_3) . We then solve periodic linear elasticity problems using our FFT-based homogenization code (FANS^{44,8}) to obtain the effective elasticity tensor $\bar{\underline{\underline{C}}}$ for $K_1 = 1$ (physical units are recovered by subsequent scaling by K_1). Overall, 1.18 million microstructure-material combinations were considered. The data were split into train/validation/test, see also table 1.

In the fig. 3 (Left), an Ashby-style scatter of individual engineering materials (elastomers, polymers, woods/foams, concretes, composites, metals, glasses/ceramics, ultra-stiff refractories) in the (K, G) plane. The cloud is narrowly confined between lines $G \approx \alpha K$ with $\alpha = 3(1 - 2\nu)/(2(1 + \nu))$ induced by the isotropic relation $G/K = 3(1 - 2\nu)/(2(1 + \nu))$, i.e., material points fall on a low-dimensional manifold parameterized mainly by Poisson's ratio ν . Typical ranges ($\nu \approx 0.15$ – 0.35 for metals/ceramics, higher for many polymers, very high for near-incompressible elastomers) generate α spanning $\sim 10^{-2}$ (rubbery, nearly incompressible) to ~ 1 (low ν ceramics), which explains the visually linear alignment. In the fig. 3 (Right), the plotted binary material pairs (each shown twice by swapping phase roles) inherit this proportional structure: both K - and G -contrasts largely move along similar $G = \alpha K$ rays because each constituent's Poisson ratio sits in the physically admissible band. Expressed in our non-dimensional parameters, this panel maps directly to the contrast plane $(1/\pi_1, \pi_2/\pi_3) = (K_1/K_2, G_1/G_2)$. The real pairs cluster within a corridor already densely covered by our log-uniform sampling $\pi_i \in [10^{-2}, 10^2]$; hence the training set spans (i) the physically prevalent




FIGURE 2 Examples of 3D microstructures taken from³⁹ for eight of the nine classes (spatial stochastic graphs omitted here). The inclusion phase is shown in red, and the matrix is shown as a void for visibility. Each class contributes 10,000 samples.

region occupied by established engineering materials, (ii) high-contrast extrapolations (extreme K or G mismatch), and (iii) regimes approaching near-incompressibility (very small G/K).

TABLE 1 Summary of the datasets considered for the 3D mechanical homogenization problem.

Dataset	Material parameters	Microstructures	Samples
Train	$\pi_i = 10^{\pi_i^*}, \pi_i^* \sim \mathcal{U}(-2, 2)$	63,000	751,089
Validation		13,500	263,188
Test		13,500	170,753
Total		90,000	1,185,030

TABLE 2 Typical bulk and shear moduli for some common composite microstructured materials at room temperature

No.	Composite			Phase 1 (GPa)		Phase 2 (GPa)	
				Bulk K_1	Shear G_1	Bulk K_2	Shear G_2
1	Steel	–	Aluminum	156 - 165	79.3	68 - 70	25.5
2	Steel	–	Epoxy	156 - 165	79.3	5.2	1.5
3	Epoxy	–	E-Glass	5.2	1.5	38	31
4	Aluminum	–	Silicon carbide	68 - 70	25.5	220	150
5	Cement	–	Aggregate	10 - 30	7	45 - 85	20
6	Polyethylene	–	E-Glass	1	0.1	38	31
7	Copper	–	Epoxy	123	44.7	5.2	1.5
8	Tungsten	–	Epoxy	307 - 314	160	5.2	1.5

The Voigt-Reuss net surrogate for the 3D mechanical problem operates on a 242-dimensional feature vector where the first 236 scalars are morphology descriptors¹², and the remaining six encode the (dimensionless) elastic phase contrasts π_1, π_2, π_3 together with their reciprocals $1/\pi_1, 1/\pi_2, 1/\pi_3$ to symmetrize the contrast representation. All input features were range-normalized prior to training. As in Section 3.2, we predict normalized degrees of freedom $\hat{\underline{\xi}}$ comprising of six eigenvalues $\lambda\{\hat{\underline{C}}\}_i^6 \in (0, 1)$ and an




FIGURE 3 Phase contrasts of the bulk and shear moduli for the material pairs from table 2. Each material pair is represented twice, considering the possibility of the two materials being in either phase 1 or phase 2.

orthogonal factor parameterized by 15 scalars $(\tilde{q}_1, \dots, \tilde{q}_{15})$ also in $(0, 1)$, leading to a total of 21 outputs. The spectral inverse normalization $\mathcal{S}^{-1}$ maps $\hat{\underline{\underline{C}}}$ to $\hat{\underline{\underline{C}}}$.

The regressor is realized as a fully connected feed-forward neural network with four hidden layers of widths [1024, 512, 256, 128], taking the 242-dimensional input $\underline{\chi}$ and outputting the 21 normalized DOFs $\hat{\underline{\xi}}$. Each hidden layer is followed by batch normalization and a mixed activation (combination chosen for stable convergence; identity components act as linear bypass channels). A final sigmoid activation confines all 21 outputs to $(0, 1)$.

Training uses the AdamW optimizer with a ReduceLROnPlateau scheduler (patience-based learning rate halving) until validation stagnation. The loss is the mean of the relative error measure eq. (17) over all samples. It is dimensionless and confined to $[0, 1]$. The model is trained on $\sim 7.5 \times 10^5$ samples (train split). Every prediction is, by design, physically admissible and bounded between the Voigt and Reuss bounds, independent of extrapolation in contrast space.

Figure 4 assembles a component-wise view of predictive accuracy for the effective tensor $\bar{\underline{\underline{C}}}$. The lower triangular panels display the ground truth versus prediction (hexbin densities) for all individual components, along with their corresponding R^2 scores. Two distinct regimes emerge sharply. First, the isotropic block i.e., the three normal moduli $(\bar{C}_{11}, \bar{C}_{22}, \bar{C}_{33})$, the three shear moduli $(\bar{C}_{44}, \bar{C}_{55}, \bar{C}_{66})$, and the symmetric normal couplings $(\bar{C}_{21}, \bar{C}_{31}, \bar{C}_{32})$ exhibit near-perfect agreement with $R^2 \geq 0.998$. Second, all remaining (anisotropy-revealing) entries that encode departures from isotropy (shear-normal couplings and inter-shear couplings) show $R^2 \approx 0$. This pattern is expected given the 236 microstructure descriptors provided in ^{12,39} are exclusively isotropic: Because the input feature map is invariant under $SO(3)$, no information is available to break rotational degeneracy and predict anisotropic tensor parts that depend on orientation or directionality. Consequently, the surrogate can—at best—recover the isotropic projection of the effective stiffness $\bar{\underline{\underline{C}}}_{\text{iso}}$, where,

$$\bar{\underline{\underline{C}}}_{\text{iso}} = \arg \min_{\bar{\underline{\underline{C}}} \in \text{Sym}^+(6)} \|\bar{\underline{\underline{C}}} - \bar{\underline{\underline{C}}}_{\text{iso}}\|_F, \quad \bar{\underline{\underline{C}}}_{\text{iso}} = 3 \bar{K} \frac{\mathbf{I} \otimes \mathbf{I}}{3} + 2 \bar{G} \mathbb{I}_2^{\text{iso}}, \quad (20)$$

$$\bar{K} = \frac{1}{9} \sum_{i,j=1}^3 \bar{C}_{ij}, \quad \bar{G} = \frac{1}{10} \left(\sum_{i=1}^6 \bar{C}_{ii} - 3 \bar{K} \right), \quad (21)$$

which fixes exactly those nine lower-triangular entries with high R^2 (including $C_{21}, C_{31}, C_{32}, C_{12}, C_{13}, C_{23}$). All anisotropy revealing components lie in the orthogonal complement, i.e., $\bar{\underline{\underline{C}}} - \bar{\underline{\underline{C}}}_{\text{iso}}$ and are a priori unpredictable from isotropic descriptors, hence $R^2 \approx 0$.

To quantify accuracy at the tensor level, fig. 5 (left) shows the distribution of $\mathcal{E}_F$ for train/validation/test; the histograms nearly overlap with essentially identical mean/median, indicating strong generalization across the datasets. The mean error is $\approx 3.4\%$,




FIGURE 4 Lower triangle: Ground truth vs. prediction for all the stiffness components of $\bar{\mathbb{C}}$ (hexbin density). Each tile lists per-component R^2 . All entries corresponding to isotropy achieve $R^2 \geq 0.998$: the axial terms $\{\bar{\mathbb{C}}_{11}, \bar{\mathbb{C}}_{22}, \bar{\mathbb{C}}_{33}\}$, shear terms $\{\bar{\mathbb{C}}_{44}, \bar{\mathbb{C}}_{55}, \bar{\mathbb{C}}_{66}\}$, and the normal couplings $\{\bar{\mathbb{C}}_{21}, \bar{\mathbb{C}}_{31}, \bar{\mathbb{C}}_{32}\}$. The remaining (anisotropy-revealing) components have $R^2 \approx 0$.

and the median is $\approx 1.7\%$, evidencing a consistent geometry and material parameter to the effective tensor mapping learned by Voigt-Reuss net.

Figure 5 (right) plots empirical CDFs of the relative Frobenius error for the Voigt-Reuss net together with two analytical references: the isotropic projection and the classical Hill average,

$$\mathcal{E}_{\text{iso}} := \mathcal{E}_{\text{F}} \left(\bar{\mathbb{C}}, \bar{\mathbb{C}}_{\text{iso}} \right), \quad \mathcal{E}_{\text{Hill}} := \mathcal{E}_{\text{F}} \left(\bar{\mathbb{C}}, \frac{\bar{\mathbb{C}}_{\text{V}} + \bar{\mathbb{C}}_{\text{R}}}{2} \right). \quad (22)$$

The isotropic projection error $\mathcal{E}_{\text{iso}}$ (median 1.03%) sets a natural lower floor for any surrogate that only receives isotropy-invariant descriptors as input, as discussed above. In contrast, the Hill average incurs substantially larger discrepancies (median 29.12%), reflecting that simple mixing rules are not competitive once high material contrast and rich morphology are present. The Voigt-Reuss net predictions (median 1.71%) lie very close to the isotropic floor and far below the Hill curve, demonstrating that the network has learned to approximate the effective stiffness tensor with high accuracy within the information constraints imposed by the descriptors.

Importantly, the residual gap of the Voigt-Reuss net to the isotropic limit reflects both the limited information content of the current descriptor set and the finite capacity and specific hyperparameter choices of the network; in principle, more aggressive architectural and training optimization could move the Voigt-Reuss net predictions even closer to this isotropic baseline, even

without changing the inputs. If anisotropy-sensitive inputs (e.g., band-limited features, directional two-point correlations, Minkowski tensors, or convolutional image features) were supplied, the achievable error floor of the current isotropic baseline would disappear, enabling $R^2 > 0$ for anisotropy-revealing components and further reductions in $\mathcal{E}_F$ at the tensor level.

To further examine performance over a wide range of elastic contrasts, fig. 6 shows the dependence of the relative error on phase contrast for the training data. The plot shows errors up to 20% covering approximately 98.18% of the data. The $\approx 1.8\%$ extreme outliers, which are not represented herein, typically appear when the phase contrast is near the limits of the logarithmic sampling (i.e., when moduli ratios are either very large or very small) or when the microstructure lies at the boundary of the considered morphological diversity. In addition, composite pairs representing real-world materials from table 2 fall within the region of relatively low error, which supports the model’s capacity to achieve accurate predictions for practical composites such as, e.g., steel-aluminum, aluminum-SiC, and epoxy-glass systems.

These results confirm that the Voigt-Reuss net achieves strong interpolation and generalization across an expansive space of 3D microstructures and elastic phase contrasts. It also maintains physically admissible predictions, adhering to the fundamental Voigt-Reuss bounds in the stiffness space, and consistently delivers accurate performance across the training, validation, and testing sets, with a median error of a few percent. Crucially, because the network was trained on a large and diverse dataset that includes real-world composite parameter ranges, it is well positioned to handle practical engineering scenarios where bulk and shear moduli vary within similar bounds.




FIGURE 5 **Left:** Relative Frobenius error histograms $\mathcal{E}_F$ for train/val/test splits (nearly coincident; mean $\approx 3.4\%$, median $\approx 1.7\%$). **Right:** ECDFs of $\mathcal{E}_F$ and two references: isotropic projection, Voigt-Reuss net and the Hill average. Voigt-Reuss net surrogate using isotropy-only descriptors operates very close to the lower (isotropy) limit.

5 | APPLICATION TO HIGHLY ANISOTROPIC AND HIGH CONTRAST 2D COMPOSITES

In this section, we describe the application of the Voigt-Reuss net to a 2D mechanical problem, specifically the prediction of the effective elasticity tensor $\bar{\mathbb{C}} \in \text{Sym}_+(\mathbb{R}^3)$ of periodic, biphasic composites under plane strain. The data was generated and obtained from a prior study³⁷, which constructed a large database of unit cells $\Omega = [-0.5, 0.5] \times [-0.5, 0.5] \subset \mathbb{R}^2$ and their homogenized elastic tensors, particularly targeting structured materials with extreme elastic anisotropy. The data are generated with a previously documented procedure to sample large families of periodic, single-scale unit cells composed of two phases that can realize a wide range of anisotropic effective elastic tensors⁴⁰. The construction is based on thresholded trigonometric fields with few spatial modes, which (i) guarantees periodicity by design, (ii) offers direct control over geometric symmetry through simple algebraic constraints on amplitudes, and (iii) yields a monotonic handle on the volume fraction over a wide range (1–99% volume fraction of phase 1; see below for details) through the threshold parameter $\tau \in [0, 1]$.




FIGURE 6 Relative error in the predicted effective stiffness tensor for different phase contrasts. The left plot shows the lower errors (0 - 20%), while the right plot.




FIGURE 7 Schematic of the procedure to generate periodic biphasic unit cells from thresholded trigonometric fields³⁷

Let $\{\mathbf{e}_1, \mathbf{e}_2\} \subset \mathbb{R}^2$ be the orthonormal basis vectors of the 2D unit cell ($\mathbf{e}_i \cdot \mathbf{e}_j = \delta_{ij}$). For two odd integers $M_1, M_2 \in \{1, 3, 5, \dots\}$ (the numbers of modes in the $\mathbf{e}_1$ and $\mathbf{e}_2$ directions, respectively), define the finite index sets

$$\mathcal{I}_1 = \left\{ -\frac{M_1-1}{2}, \dots, 0, \dots, \frac{M_1-1}{2} \right\}, \quad \mathcal{I}_2 = \left\{ -\frac{M_2-1}{2}, \dots, 0, \dots, \frac{M_2-1}{2} \right\}. \quad (23)$$

Given an amplitude matrix $\underline{\underline{A}} = (A_{mn})_{m \in \mathcal{I}_1, n \in \mathcal{I}_2} \in \mathbb{R}^{M_1 \times M_2}$ with entries sampled independently and identically from a uniform distribution $A_{mn} \sim \mathcal{U}[-1, 1]$, we form the periodic field $\psi(\mathbf{x}; \underline{\underline{A}})$,

$$\psi(\mathbf{x}; \underline{\underline{A}}) = \sum_{m \in \mathcal{I}_1} \sum_{n \in \mathcal{I}_2} A_{mn} \cos(2\pi (m x_1 + n x_2)), \quad \mathbf{x} \in \Omega := \left\{ x_1 \mathbf{e}_1 + x_2 \mathbf{e}_2 \mid x_1, x_2 \in \left[-\frac{1}{2}, \frac{1}{2}\right] \right\}. \quad (24)$$

We then apply range normalization on $\psi(\mathbf{x}; \underline{\underline{A}})$ such that the rescaled field $\tilde{\psi}(\mathbf{x}; \underline{\underline{A}}) \in [0, 1]$,

$$\psi_{\min} = \min_{\mathbf{x} \in \Omega} \psi(\mathbf{x}; \underline{\underline{A}}), \quad \psi_{\max} = \max_{\mathbf{x} \in \Omega} \psi(\mathbf{x}; \underline{\underline{A}}), \quad \tilde{\psi}(\mathbf{x}; \underline{\underline{A}}) = \frac{\psi(\mathbf{x}; \underline{\underline{A}}) - \psi_{\min}}{\psi_{\max} - \psi_{\min}} \in [0, 1]. \quad (25)$$




FIGURE 8 Examples of unit cells with different symmetry classes obtained by imposing simple algebraic constraints on the amplitude matrix $\underline{A}$. (a) No symmetry (random $\underline{A}$); (b) diagonal symmetry ($A_{m,n} = A_{n,m}$); (c) orthotropic symmetry ($A_{m,n} = A_{-m,n} = A_{m,-n}$); (d) square symmetry ($A_{m,n} = A_{n,m} = A_{-m,n} = A_{m,-n}$).

For a threshold $\tau \in [0, 1]$, the binary indicator χ of the microstructure is

$$\chi(\mathbf{x}; \underline{A}, \tau) = 1_{\{\tilde{\psi}(\mathbf{x}; \underline{A}) \geq \tau\}}, \quad c_1(\tau; \underline{A}) = \frac{1}{|\Omega|} \int_{\Omega} \chi(\mathbf{x}; \underline{A}, \tau) \, d\mathbf{x}, \quad c_0(\tau; \underline{A}) = 1 - c_1(\tau; \underline{A}). \quad (26)$$

By construction c_0 and c_1 vary monotonically with τ , that is, $c_0(\tau) \leq c_0(\tau + \delta)$ for $\delta > 0$, and the inverse inequality for c_1 holds by construction. Symmetry classes of the microstructure can be enforced by simple constraints on $\underline{A}$ (see fig. 8). Increasing (M_1, M_2) introduces finer features and generally reduces extreme anisotropy by distributing the power spectral density over more spatial frequencies³⁷. Each unit cell $\chi(\mathbf{x}; \underline{A}, \tau)$ represents a biphasic composite. Phase 0 (stiff - DM 8530 Grey60) and Phase 1 (soft - Tango Black) are modeled as isotropic linear elastic solids with $(E_0, \nu_0) = (1 \text{ GPa}, 0.3)$, $(E_1, \nu_1) = (1 \text{ MPa}, 0.49)$, as chosen in³⁷ are representative of materials from a commercial multi-material Connex 3D printer. Note the impressive Young's modulus contrast of 10^3 and a large contrast in bulk/shear response due to the near incompressibility of Tango black due to $\nu_1 \approx 0.5$. The data from³⁷ (2D plane strain homogenized elasticity tensors given in Voigt notation) is used to rewrite the linear stress-strain relation in Mandel notation (see section 2):

$$\underline{\bar{\sigma}} = \begin{bmatrix} \bar{\sigma}_{11} \\ \bar{\sigma}_{22} \\ \sqrt{2} \bar{\sigma}_{12} \end{bmatrix} = \begin{bmatrix} \bar{C}_{1111} & \bar{C}_{1122} & \sqrt{2} \bar{C}_{1112} \\ \bar{C}_{2211} & \bar{C}_{2222} & \sqrt{2} \bar{C}_{2212} \\ \sqrt{2} \bar{C}_{1211} & \sqrt{2} \bar{C}_{1222} & 2 \bar{C}_{1212} \end{bmatrix} : \begin{bmatrix} \bar{\varepsilon}_{11} \\ \bar{\varepsilon}_{22} \\ \sqrt{2} \bar{\varepsilon}_{12} \end{bmatrix} = \underline{\bar{C}} : \underline{\bar{\varepsilon}}. \quad (27)$$

The homogenized elasticity tensor $\underline{\bar{C}}$ is obtained by a standard two-scale first-order computational homogenization, i.e., by solving the corrector problems for three orthonormal macroscopic loadings. Thermodynamic stability requires $\bar{C}_{1111}, \bar{C}_{2222}, \bar{C}_{1212} > 0$, while off-diagonal couplings $\bar{C}_{1122}, \bar{C}_{1112}, \bar{C}_{2212}$ may be negative. Symmetry and positive definiteness of the plane strain stiffness $\underline{\bar{C}}$ (27) enforces the full set of positivity constraints (see section 3.2; also³⁸) automatically. In the study³⁷, a large database was constructed by sampling many amplitude matrices $\underline{A}$ for different values of M_1, M_2 and by sweeping the threshold τ ; see³⁷ for details. In summary, the dataset construction involves:

- **Mode grids.** Five sets of spatial modes $(M_1 \times M_2) \in \mathcal{P}$, with $\mathcal{P} = \{3 \times 3, 5 \times 5, 7 \times 7, 9 \times 9, 11 \times 11\}$.
- **Amplitude sampling.** For each mode set $p \in \mathcal{P}$, $N_A = 2,000$ independent amplitude matrices $\underline{A}^{(p,i)}$ with entries $A_{mn} \sim \mathcal{U}[-1, 1]$ are drawn. Fully random $\underline{A}$ are drawn to encourage low symmetry.
- **Threshold sweep.** For each $\underline{A}^{(p,i)}$, we evaluate $N_\tau = 25$ equi-spaced thresholds $\{\tau_k\}_{k=1}^{N_\tau} \subset (0, 1)$.
- **Filtering.** In the resulting $2,000 \times 25 = 250,000$ structures, to avoid trivial geometries (all phase 0; all phase 1), we retain only realizations whose phase volume fraction satisfies $c_0(\tau_k; \underline{A}^{(p,i)}) \in [0.01, 0.99]$.
- **Homogenized elasticity tensors.** Unit cells are rasterized at 100×100 pixels and homogenized in plane strain with periodic boundary conditions, using the phase pairs $(E_1, \nu_1), (E_2, \nu_2)$ to obtain $\underline{\bar{C}}$.

The data from³⁷ is used in the subsequent study, with additional simulations being performed using our FFT-based homogenization FANS⁴⁴. For the resulting structures, we define a random split (80% / 20%) into training and validation sets for each mode set $(M_1, M_2) \in \mathcal{P}$ (see table 3). For each retained unit cell we store $(\underline{A}^{(p,i)}, \tau_k)$ ($i \in \{1, \dots, 2,000\}; k \in \{1, \dots, 25\}$) and




FIGURE 9 Examples of unit cells generated by thresholded trigonometric fields with a fixed amplitude matrix $\underline{A}$ and increasing threshold τ from left to right, showing the monotonic increase in the volume fraction of phase 0 (stiff phase). The microstructures exhibit a wide range of microstructure topologies, demonstrating the challenging nature of the dataset.

the homogenized elasticity tensor $\underline{\underline{C}}_{(M_1, M_2)}^{(i, k)}$. After removal of inadmissible samples ($c_0 < 0.01$ or $c_0 > 0.99$), 219,378 samples remained; see table 3. Note that each sample is characterized entirely by the parameter vector $\underline{\mu} = [M_1, M_2, \underline{A}_{(M_1, M_2)}^{(i, k)}, \tau]$.^{††}

TABLE 3 Summary of the 2D plane-strain elasticity dataset constructed by³⁷ from cosine-thresholded unit cells (post-filtering by $c_0 \in [0.01, 0.99]$)

	3×3	5×5	7×7	9×9	11×11	Total
Train	36,444	35,917	35,179	34,246	33,718	175,504
Validation	9,111	8,979	8,795	8,562	8,427	43,874
Total	45,555	44,896	43,974	42,808	42,145	219,378

The prediction of effective elastic properties for periodic composites is fundamentally constrained by the classical G-closure problem. The G-closure describes the set of all possible effective elasticity tensors that can be realized by mixing two given phases in all possible microgeometries, subject to prescribed volume fractions and symmetry constraints. For high-contrast composites, where the elastic moduli of the two phases differ by several orders of magnitude, the G-closure set becomes extremely large and nontrivial. The variability of the microstructures in this dataset is substantial: the use of random trigonometric fields with varying numbers of modes and amplitude matrices generates a vast diversity of geometries, ranging from nearly isotropic to highly anisotropic and from simple to complex topologies. This diversity is further amplified by sweeping the threshold parameter τ , which tunes the volume fraction and can induce topological transitions in the microstructure (see fig. 9). These factors render the learning problem particularly challenging. First, the mapping from microstructure to effective tensor is highly nonlinear and nonlocal: small changes in geometry can lead to large changes in the effective response, especially near percolation thresholds or in the high-contrast regime. Second, the admissible set of effective tensors (the G-closure) is a highly complicated set of symmetric positive definite matrices, strictly bounded by the upper Voigt and lower Reuss bounds. Third, the presence of extreme anisotropy and high contrast means that the model must generalize across a wide range of physical behaviors, including cases where subtle geometric features dominate the effective response. Consequently, accurately predicting the effective elasticity tensor in this space requires a surrogate model that is highly expressive and physically constrained, motivating the use of spectral normalization and the proposed generalization of the Voigt-Reuss net³⁸ for mechanical problems.

5.1 | Neural surrogate architecture

Next, we describe the end-to-end differentiable surrogate used to map the parametric unit cell specification $(\underline{A}, \tau)$ to the effective plane-strain elasticity tensor $\underline{\underline{C}}$. The model has three stages:

- (i) A fixed-shape embedding $\mathcal{T}$ that pads amplitude grids of various mode sets in $\mathcal{P}$ into a common $K \times K$ hyper mode set; Let the mode grid of a sample be $\underline{A} \in \mathbb{R}^{m \times n}$, where $m, n \in \{3, 5, 7, 9, 11\}$. We define a centered embedding into $\underline{A}^\top \in \mathbb{R}^{K \times K}$

^{††} Matrices for $M_1 = M_2 < 11$ can be represented by 11×11 matrices padded with zeros, rendering the dimension of $\underline{\mu}$ constant for convenience (see eq. (28)).

with K chosen such that $K \geq \max\{m, n\}$ for all possible m, n (here, $K = 11$):

$$\mathcal{T} : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{K \times K}, \quad \underline{\underline{A}}_{ij}^\top = \begin{cases} \underline{\underline{A}}_{i-i_0, j-j_0}, & i_0 \leq i < i_0 + m, \quad j_0 \leq j < j_0 + n, \\ 0, & \text{otherwise,} \end{cases} \quad \text{where } \begin{cases} i_0 = \frac{(K-m)}{2}, \\ j_0 = \frac{(K-n)}{2}. \end{cases} \quad (28)$$

Thus, every sample is represented by a fixed $K \times K$ matrix $\underline{\underline{A}}^\top$, enabling batched GPU training while preserving the spectral layout (modes remain centered).

- (ii) A differentiable microstructure renderer $\mathcal{R}$ that synthesizes a soft microstructure from $\underline{\mu} = (\underline{\underline{A}}^\top, \tau)$; Given the grid resolution $(N_{x_1}, N_{x_2}) = (100, 100)$, the scalar field is given by,

$$\underline{\underline{\psi}}(x_1, x_2; \underline{\underline{A}}^\top) = \sum_{p=-(K-1)/2}^{(K-1)/2} \sum_{q=-(K-1)/2}^{(K-1)/2} A_{pq}^\top \cos(2\pi(px_1 + qx_2)), \quad (x_1, x_2) \in [-\frac{1}{2}, \frac{1}{2}]^2 \cap \Pi_{N_{x_1}, N_{x_2}}, \quad (29)$$

where $\Pi_{N_{x_1}, N_{x_2}}$ denotes the pixel lattice. We then apply range normalization eq. (25) such that $\underline{\underline{\tilde{\psi}}} \in [0, 1]$. The scalar field $\underline{\underline{\tilde{\psi}}}$ is thresholded with a *soft* threshold τ and temperature $T > 0$,

$$\underline{\underline{\chi}}(x_1, x_2; \underline{\underline{A}}^\top, \tau) = \sigma \left(\frac{\underline{\underline{\tilde{\psi}}}(x_1, x_2; \underline{\underline{A}}^\top) - \tau}{T} \right), \quad \sigma(t) = \frac{1}{1 + e^{-t}}. \quad (30)$$

By construction, the map $\mathcal{R}(\underline{\underline{A}}^\top, \tau) \mapsto \underline{\underline{\chi}}$ is smooth; therefore, the gradients can flow from the loss to both $\underline{\underline{A}}^\top$ and τ (allowing inverse design later). We use $T = 0.0001$ to approximate a hard threshold while maintaining numerical stability and some degree of differentiability.

- (iii) A regressor predicting the normalized degrees of freedom $\underline{\underline{\hat{\xi}}}$ (see section 3.2). Let the rendered microstructure be a monochrome image $\underline{\underline{\chi}} \in [0, 1]^{N_{x_1} \times N_{x_2}} \subset \mathbb{R}^2$ and let $\underline{\underline{s}} = [c_0, \tau]^\top \in [0, 1]^2$ be an auxiliary feature vector, where c_0 is the phase volume fraction of phase 0 in the rendered image $\underline{\underline{\chi}}$ and τ is the threshold used. A convolutional trunk $\mathcal{C}_\theta$ extracts an image descriptor $\underline{\underline{h}} = \mathcal{C}_\theta(\underline{\underline{\chi}}) \in \mathbb{R}^{d_h}$; see fig. 10. A multilayer perceptron (MLP) with a sigmoid head $\mathcal{M}_\theta$ then regresses the normalized target from the concatenation $[\underline{\underline{h}}; \underline{\underline{s}}]$, such that $\underline{\underline{\hat{\xi}}} = \mathcal{M}_\theta([\underline{\underline{h}}; \underline{\underline{s}}]) = [\underline{\underline{\hat{\xi}}}_\lambda, \underline{\underline{\hat{\xi}}}_q]^\top$. Here, $\underline{\underline{\xi}}_\lambda \in (0, 1)^3$ are the normalized eigenvalues and $\underline{\underline{\xi}}_q \in (0, 1)^3$ parameterize an orthogonal matrix; both are produced with sigmoid activations to confine them to the unit cube. This is followed by the unique and deterministic spectral inverse transform $\mathcal{S}^{-1}$ to $\underline{\underline{\hat{\underline{\underline{C}}}}}$.

The final surrogate can be written as the composition,

$$\mathcal{M}_\theta \circ \mathcal{C}_\theta \circ \mathcal{R} \circ \mathcal{T} : (\underline{\underline{A}}, \tau) \mapsto \underline{\underline{\hat{\xi}}} \in (0, 1)^6, \quad (31)$$

where subscript $(\bullet)_\theta$ indicates a dependence on learnable parameters. This construction ensures automatic feature detection from microstructure images, as well as end-to-end differentiability in $(\underline{\underline{A}}, \tau)$ for inverse design. Note that by replacing $\mathcal{R} \circ \mathcal{T}(\underline{\underline{A}}, \tau)$ with any other image source, it is equally possible to generalize this approach to arbitrary microstructure sources. Furthermore, this construction can be straightforwardly extended to 3D by embedding 3D amplitude grids, rendering volumetric microstructures via trigonometric fields, and applying a 3D convolutional regressor, enabling end-to-end differentiable surrogates for 3D systems. However, the computational requirement for the 3D setting will increase considerably compared to the 2D setting.

5.2 | Numerical experiments and results

The convolutional trunk $\mathcal{C}_\theta$ is a deep residual CNN with periodic convolutions and four stages. Each stage consists of a ResidualBlock with two periodic convolutional layers with increasing kernel sizes ($k_1 = 3, k_2 = 5, k_3 = 7, k_4 = 9$), batch normalization, and ReLU activations, followed by 2×2 average pooling. The number of output channels C_{out} increases as $18 \rightarrow 24 \rightarrow 30 \rightarrow 36$ across the four blocks, which implies $C_{\text{in}} = 1 \rightarrow 18 \rightarrow 24 \rightarrow 30$. After the final block, the feature map is flattened and concatenated with the auxiliary scalars $[c_0, \tau]$. This is passed to a multi-layer perceptron (MLP) with five hidden layers of sizes 256, 128, 64, 64, and 64, each with batch normalization and a mixed activation ($[\text{SELU}, \text{Tanh},$




FIGURE 10 End-to-end differentiable surrogate. The pipeline maps the parametric unit cell specification $(\underline{A}^\top, \tau)$ to the effective plane-strain elasticity tensor $\underline{\underline{C}}$. **(Top):** The center-padded amplitude matrix $\underline{A}^\top$ together with threshold τ , renders a soft microstructure image $\underline{\chi}$ via a differentiable renderer $\mathcal{R}$. The image is processed by a convolutional trunk $\mathcal{C}_\theta$ with four convolutional residual blocks^{45,46} $\mathcal{B}_i^k(C_{\text{in}}, C_{\text{out}})$. The resulting feature map is flattened to $\underline{h}$ and concatenated with auxiliary scalars $[c_0, \tau]$, then passed to the MLP $\mathcal{M}_\theta$ with four linear layers to predict normalized spectral parameters $\underline{\xi}$. These are spectrally back-transformed to the effective tensor $\underline{\underline{C}}$ via the inverse spectral normalization $\mathcal{S}^{-1}$. **(Bottom):** Structure of a single residual block, two parallel convolutional layers with kernel size k and a skipped connection along with batch normalization, ReLU activations, and average pooling.

Sigmoid, Identity]), and a final sigmoid-activated output layer of dimension 6. The model uses periodic convolutions throughout to preserve microstructure periodicity. Training uses the AdamW optimizer with an initial learning rate of 10^{-1} , and a ReduceLROnPlateau scheduler (factor 0.5, patience 50 epochs). The loss function is defined as the mean Voigt-Reuss normalized error ϕ (eq. (17)) over each batch,

$$\mathcal{L} = \frac{1}{N_{\text{batch}}} \sum_{i=1}^{N_{\text{batch}}} \phi(\underline{\underline{C}}^{(i)}, \widehat{\underline{\underline{C}}}^{(i)}),$$

where N_{batch} is the batch size, $\underline{\underline{C}}^{(i)}$ is the true normalized stiffness tensor, and $\widehat{\underline{\underline{C}}}^{(i)}$ is the predicted tensor for sample i in the batch.

5.2.1 | Microstructure to property forecasting

Figure 11 (left) compiles ground truth versus predicted entries of homogenized plane-strain stiffness $\underline{\underline{C}}$ for all samples. The six components tightly cluster on the identity line with a coefficient of determination $R^2 > 0.99$, indicating (i) negligible bias across magnitudes and signs (including shear-normal couplings $\bar{C}_{1112}, \bar{C}_{2212}$), and (ii) uniformly low variance of the residuals. Notably, the spectral back-transform ensures that every predicted $\underline{\underline{C}}$ is symmetric positive definite and lies spectrally within Voigt-Reuss bounds; hence, even the few visible outliers remain physically admissible by construction. The right plot in fig. 11

reports the loss of training and validation of the model. Because ϕ for each sample is intrinsically bound to $[0, 1]$, the training and validation losses provide an absolute scale of progress. Training converges smoothly with a small and stable generalization gap to subpercent errors relative to the admissible spectral range: $\mathcal{L}_{\text{train}} \approx 0.586\%$ and $\mathcal{L}_{\text{val}} \approx 0.824\%$ at saturation. In particular, the validation loss shows the same qualitative and nearly the same quantitative behavior as the training loss. The scheduler lowers the learning rate from 10^{-1} to 5×10^{-5} , coincident with the flattening of both curves, which signals that the model has reached its accuracy floor under current capacity and data. Crucially, the bounded, well-conditioned target space (all outputs in $[0, 1]$) produces well-behaved gradients and loss descent without instability, a hallmark advantage of spectral normalization over direct regression in physical units. We also attribute the smooth training to the incorporation of the identity activation in our multi-activation layers, which acts as a sort of bypass inspired by the ResNet architecture.




FIGURE 11 2D plane strain elasticity. **Left:** Ground truth vs. prediction for the six independent components of $\bar{\mathbb{C}}$. **Right:** The Voigt-Reuss normalized loss decays and saturates; the learning rate decreases with training via a ReduceLROnPlateau scheduler.

In fig. 12 we analyze the predictions in more detail. On the left, we plot the three eigenvalues $\{\lambda_i(\bar{\mathbb{C}})\}_{i=1}^3$ of the effective stiffness tensor as a function of the volume fraction for all validation samples. The shaded area indicates the Voigt-Reuss bounds, which collapse at the endpoints $c_0 \rightarrow 0$ and $c_0 \rightarrow 1$ and widen near $c_0 \approx 0.5$, where microstructural variability and anisotropy peak. The ground truth eigenvalues (blue markers) span a wide range from 0 to 2,000 MPa, illustrating the extreme variability and high phase contrast of the dataset. The predicted eigenvalues (green markers) tightly overlap with the ground truth across the entire range, demonstrating the model’s ability to capture the effective stiffness even in challenging regimes accurately. In the middle panel, we show the same eigenvalues after spectral normalization, i.e., $\{\xi_{\lambda_i}\}_{i=1}^3 \in [0, 1]$. The normalization compresses the wide variation in physical units down to a bounded unit interval while preserving the ordering and relative positions of the eigenvalues. The near-perfect overlap between ground truth and predictions in this normalized space illustrates why training in the bounded domain is numerically well-conditioned and stable. Finally, the right panel presents a 2D histogram of the relative Frobenius error in the predicted stiffness tensor as a function of volume fraction. The intensity indicates the number of samples in each 2D bin, with solid and dashed lines showing the mean and median errors, respectively. As expected, errors are smallest near $c_0 \rightarrow 0$ and $c_0 \rightarrow 1$, where the spectral gap collapses, and largest near $c_0 \approx 0.5$. Nevertheless, the mean and median errors remain very modest (around 3% and 2%, respectively), and even the rare worst-case errors (up to around 65%) are still spectrally admissible by construction, i.e., they lie within the bounds.

To stress-test generalization, we consider ten binary microstructures sampled from an independent diffusion-based image generator (completely disjoint from the cosine-threshold family used for training). We can evaluate image-only inference by feeding a binary microstructure image $\underline{\chi} \in \{0, 1\}^{N_{x_1} \times N_{x_2}}$ directly into the convolutional trunk and regressor, followed by the spectral back-transform. The predictor is

$$\hat{\bar{\mathbb{C}}} = \mathcal{S}^{-1} \left(\mathcal{M}_{\theta}(\mathcal{C}_{\theta}(\underline{\chi})) \right), \quad (32)$$




FIGURE 12 2D plane strain elasticity (validation set). **Left:** Distribution of eigenvalues $\{\lambda_i(\bar{\mathbf{C}})\}_{i=1}^3$ and the normalized eigenvalues $\{\xi_{\lambda_i}\}_{i=1}^3 \in [0, 1]$ with respect to the volume fraction. **Right:** 2D histogram of the relative Frobenius error $\mathcal{E}_F(\bar{\mathbf{C}}, \hat{\mathbf{C}})$ against volume fraction.




FIGURE 13 Out-of-distribution image-to-elasticity inference. **Top:** Ten diffusion-generated binary microstructures taken from^{23,47} (family disjoint from training). **Bottom:** Young's modulus polar surfaces $E(\theta)$ from ground-truth homogenization (solid) and direct Voigt-Reuss net predictions from the image (dashed).

where $\mathcal{C}_\theta$ is the periodic convolutional trunk, $\mathcal{M}_\theta$ the MLP head producing the normalized DOFs $\hat{\xi}$, and $\mathcal{S}$ the spectral normalization. This path requires a single forward pass. Figure 13 shows ten diffusion-generated test images (top) and the corresponding plane-strain Young's modulus surfaces $E(\theta)$ (bottom), comparing ground truth (solid) versus direct Voigt-Reuss net predictions (dashed). Despite pronounced out-of-distribution features (e.g., very slender beams, sharp junctions), the predicted $E(\theta)$ curves closely track the ground truth in both shape and level at most angles. The relative Frobenius errors at the tensor level are modestly higher than on the in-distribution data, consistent with the geometric gap between these images and the training manifold. However, all predictions remain strictly admissible by construction. Practically, this enables the instant screening of arbitrary microstructure images with a single model evaluation, while preserving physical guarantees. Additionally, when new input data becomes available, the existing model appears to be a good candidate for transfer learning, as it can be refined at a much lower cost than when starting from a random initialization.




FIGURE 14 Sensitivity of the Voigt-Reuss net to the threshold parameter τ for a fixed amplitude matrix $\underline{A}$. **Top-left:** Relative Frobenius error $\mathcal{E}_F(\underline{\hat{C}}, \underline{\hat{C}})$ of the Voigt-Reuss net (solid) and a Mori-Tanaka estimate with spherical Eshelby tensor (dashed). **Top-middle:** Eigenvalues $\{\lambda_i(\underline{\hat{C}})\}_{i=1}^3$, together with Voigt/Reuss and Hashin-Shtrikman bounds (HS $\pm$). **Top-right:** Normalized gradient of surrogate prediction with respect to the threshold τ . **Bottom:** Ten frames along the τ -path; red boxes mark the two topology changes (island merger/percolation) coinciding with error and gradient spikes, along with the Young's modulus surfaces in logarithmic scale of the ground truth $\underline{\hat{C}}$ (solid) and Voigt-Reuss net prediction $\underline{\hat{C}}$ (dashed).

In order to systematically investigate the sensitivity of the surrogate to the threshold parameter τ , we fix a random amplitude matrix $\underline{A}$ and sweep τ over 100 values from 0 to 1, rendering $\underline{\chi}(\cdot; \underline{A}, \tau)$ and computing both the homogenized tensor $\underline{\hat{C}}$ and the Voigt-Reuss net prediction $\underline{\hat{C}}$. For each τ , we track the relative Frobenius error $\mathcal{E}_F(\underline{\hat{C}}, \underline{\hat{C}})$ and compare it against a classical Mori-Tanaka estimate built from the volume fraction and a spherical Eshelby tensor. Figure 14 (top-left) shows that the Voigt-Reuss net maintains sub-percent errors for almost the entire path; only two prominent spikes appear. In contrast, the Mori-Tanaka estimate begins reasonably at low fractions but degrades sharply as τ increases: once the inclusion topology ceases to resemble well-separated, non-interacting particles (assumptions underlying Mori-Tanaka estimates), its error grows and eventually becomes unacceptable. This behavior is expected in high-contrast media where connectivity and shape effects dominate the effective stiffness.

The two error peaks of the Voigt-Reuss net align exactly with abrupt topology changes: previously disconnected islands coalesce into percolating channels (see the bottom strip of microstructures). These events are level-set percolation transitions of the field $\tilde{\psi}(\mathbf{x}; \underline{A})$ under thresholding; they trigger discontinuous jumps in transport paths and load-bearing skeletons, hence non-smooth changes in $\underline{\hat{C}}$. The top-middle plot of Fig. 14 makes this explicit: the eigenvalues $\{\lambda_i(\underline{\hat{C}})\}_{i=1}^3$ show two clear jumps, and the Voigt-Reuss net tracks both jumps essentially on target. The Hashin-Shtrikman (HS $\pm$) and the Voigt-Reuss envelopes, which depend only on phase moduli and volume fraction, provide a geometry-agnostic benchmark; the ground truth and Voigt-Reuss net eigenvalues remain within these envelopes as well. The fact that the Voigt-Reuss net resolves the spectral jumps is notable: The model only sees the amplitude matrix $\underline{A}$ and the threshold τ , and consequently the rendered microstructure images, and learns the global response, including percolation-driven anisotropy while retaining strict physical admissibility.

To further probe sensitivity, we differentiate the Frobenius norm of the Voigt-Reuss net prediction $\underline{\hat{C}}$ with respect to the threshold τ . The top-right plot of Fig. 14 shows that the normalized gradient exhibits sharp spikes precisely at the two critical thresholds. These spikes signal non-smooth property changes induced by global connectivity flips, which are well-known triggers

for abrupt property changes that are beyond the reach of first-order mean-field closures, such as Mori-Tanaka or self-consistent schemes.

Overall, Voigt-Reuss net accurately captures the geometry-controlled locally non-smooth dependence of the effective stiffness on the threshold, including (multiple) percolation-induced spectral jumps, while remaining strictly admissible. In particular, the number of such spectral jumps varies among the microstructures studied; that is, it is not explicitly considered in the trained Voigt-Reuss net, but rather implicitly through the data. In contrast, Mori-Tanaka estimates, which hinge on dilute, non-interacting inclusion assumptions and spherical Eshelby kernels, cannot track topology-driven anisotropy and fail as connectivity grows. The differentiable chain $(\underline{A}, \tau) \rightarrow \hat{\underline{C}}$ provides reliable gradients, except for genuine topological events, which appear as informative spikes that could be exploited in inverse design (e.g., to seek or avoid connectivity transitions).

5.2.2 | Inverse materials design




FIGURE 15 Inverse design to prescribed effective stiffnesses in plane strain. For four distinct targets (rows), the first column lists the target tensor $\underline{\underline{C}}^*$; the second overlays the Young's modulus polar surface $E(\theta)$ of the target (black) with the best five candidates (colored); the next five columns show the corresponding candidate microstructures with their relative Frobenius error $\mathcal{E}_F(\underline{\underline{C}}^*, \underline{\underline{C}})$ in %. Across all cases, the recovered $E(\theta)$ nearly overlaps the target, and the tensor errors are only a few percent.

We finally demonstrate the application of our model for the inverse design of microstructures to achieve prescribed effective plane-strain stiffness tensors $\underline{\underline{C}}^*$. Given a physically admissible target plane-strain stiffness $\underline{\underline{C}}^* \in \text{Sym}^+(\mathbb{R}^3)$, we seek unit cell parameters $(\underline{A}, \tau)$ whose predicted effective tensor is as close to the target as possible. Let the differentiable surrogate (section 5.1) be written compactly as

$$\hat{\underline{\underline{C}}}(\underline{A}, \tau) = \mathcal{S}^{-1} \left(\mathcal{M}_\theta \left(\mathcal{C}_\theta \left(\mathcal{R}(\mathcal{T}(\underline{A}), \tau) \right) \right) \right), \quad (33)$$

where $\mathcal{S}$ denotes the spectral normalization. We minimize the (scale-free) distance

$$\mathcal{J}(\hat{\underline{\underline{C}}}) = \min_{\underline{A}, \tau \in (0,1)} \mathcal{E}_F(\underline{\underline{C}}^*, \hat{\underline{\underline{C}}}(\underline{A}, \tau)). \quad (34)$$

Since $\mathcal{T}, \mathcal{R}, \mathcal{C}_\theta, \mathcal{M}_\theta, \mathcal{B}$ are all differentiable end-to-end, the gradients $\partial \mathcal{J} / \partial (\underline{A}, \tau)$ are obtained by automatic differentiation. We solve (34) in batched form: with N_{start} independent initializations, and optimize all $(\underline{A}^{(s)}, \tau^{(s)})$ ($s = 1, \dots, N_{\text{start}}$) in parallel using the AdamW optimizer. At convergence, we verify a few realizations related to the lowest objective function values by full homogenization to report the true relative error, which is inaccessible during the inverse design phase. Figure 15 assembles four representative targets (one per row). For each target, the second column shows the directional Young's modulus surface computed from the plane-strain compliance tensor, overlaid for the target and the five best candidates. In all four cases, the recovered directional Young's modulus $E(\theta)$ almost perfectly matches the target in shape and amplitude, while the relative Frobenius error between the optimized candidates and the target tensor is only a few percent. We also observe a many-to-one effect: different unit cell candidates can realize near-identical effective tensors consistent with the nonuniqueness of G-closure. However, the optimizer reliably lands on physically meaningful, percolating, or laminated motifs when required. Practically, the batched formulation can yield dozens to hundreds of high-quality candidates in a single run; no auxiliary "inverse" network is needed, and one surrogate suffices for both forward prediction and backpropagated design.




FIGURE 16 Inverse design for extreme shear-normal coupling. Left: Young's modulus surfaces $E(\theta)$ of the top candidates produced by the batched optimizer for the coupling objective; Right: seven exemplar microstructures achieving the largest shear-normal coupling while remaining spectrally admissible.

Beyond matching a prescribed tensor, we can pose design goals directly in terms of tensor functionals. To maximize plane-strain shear-normal couplings, we consider the objective

$$\mathcal{K}(\hat{\underline{\underline{C}}}) = \max_{\underline{A}, \tau \in (0,1)} \frac{|\hat{\underline{\underline{C}}}_{1112}| + |\hat{\underline{\underline{C}}}_{2212}|}{\|\hat{\underline{\underline{C}}}\|_F}. \quad (35)$$

We again use a multistart, fully batched AdamW optimization, harvesting the most extreme candidates at convergence. Figure 16 (left) plots $E(\theta)$ for the seven most promising candidates, evidencing strong directional asymmetry consistent with large shear-normal couplings; the right panels show the seven related microstructures uncovered by the optimizer. In particular, among the discovered solutions, the optimizer recovers exactly the theoretical optimum: a laminate oriented at 45° with volume fraction $c_0 \approx 0.5 \approx c_1$, which is known to maximize the shear-normal coupling in plane-strain elasticity.

6 | CONCLUSIONS

6.1 | Summary

Across all settings, and in both 2D and 3D, we demonstrated that spectral normalization realized as the major building block of the Voigt-Reuss net provides a principled scaffold for learning effective constitutive tensors that are guaranteed to be symmetric positive definite and to lie between Voigt and Reuss bounds in the Löwner sense. By mapping every sample into a bounded, unit-cube target space, our surrogates possess robust training dynamics, well-conditioned, scale-free losses, and return physically admissible predictions by construction.

We formalized a universal reparameterization introduced in³⁸ and extended it to mechanical homogenization problems. The learning objective leads to a task-agnostic loss $\phi(\tilde{\underline{Y}}, \hat{\underline{Y}}) \in [0, 1]$ that directly measures errors relative to the admissible Voigt-Reuss gap, decoupling learning from units and improving numerical stability.

In the case of 3D mechanical homogenization as a forward problem, we trained a Voigt-Reuss net surrogate using 236 isotropic microstructure descriptors from^{39,12} along with three non-dimensional phase parameters (π_1, π_2, π_3) , on > 1.18 Mio high-fidelity DNS labels of effective elasticity tensors. Because the inputs are rotationally invariant, the surrogate recovers the isotropic projection of $\underline{\underline{C}}$ nearly perfectly, while anisotropy-revealing entries remain unpredictable: the components related to isotropy attain $R^2 \geq 0.998$, and the remaining components show $R^2 \approx 0$. At the tensor level, histograms of the relative Frobenius error of the training, validation, and testing datasets overlap tightly (mean $\approx 3.4\%$, median $\approx 1.7\%$), indicating strong generalization across microstructures and contrasts. ECDFs reveal that the isotropic projection forms a hard lower floor (median 1.03%), and Voigt-Reuss net sits close to the floor (median 1.71%), outperforming the Hill average (median 29.12%) by far. Real composite pairs (steel-Al, Al-SiC, epoxy-glass, etc.) lie squarely in the low-error regime spanned by our extensive sampling.

In the context of the 2D mechanical homogenization (forward and inverse), we introduced an end-to-end differentiable renderer for parameterized thresholded trigonometric fields, enabling a hybrid convolutional and multi-layer perceptron surrogate that ingests $(\underline{A}, \tau)$ and outputs a spectrally admissible $\underline{\underline{C}}$. Forward prediction reaches $R^2 > 0.99$ across all components of the elasticity tensor and subpercent Voigt-Reuss normalized losses, and the model remains accurate with modest errors across volume fractions and respects the Voigt-Reuss bounds inherently. Crucially, the differentiable chain $(\underline{A}, \tau) \mapsto \underline{\underline{C}}$ exposes extremely informative gradients and resolves non-smooth percolation events (eigenvalue jumps) along threshold sweeps seamlessly, while remaining strictly admissible. Leveraging these properties, we performed inverse design to target full tensors and functionals (e.g., maximizing shear-normal coupling), recovering classical optima (e.g., 45° laminates at $c_0 \approx 0.5$) and diverse near-optimal candidates in a single batched optimization run. This capability is particularly remarkable: by specifying a single target effective stiffness tensor $\underline{\underline{C}}^*$, the optimizer can generate a multitude of structurally distinct microstructures that achieve nearly identical mechanical properties, with relative Frobenius errors as low as a few percent. This many-to-one mapping underscores the richness of the G-closure set, where multiple geometric realizations can produce equivalent effective responses, offering designers flexibility in selecting structures based on additional criteria such as manufacturability, weight, or connectivity. Moreover, the batched optimization approach enables simultaneous exploration of hundreds of candidates, ensuring diversity while guaranteeing that every proposed design lies within the physically admissible Voigt-Reuss bounds, thus avoiding infeasible or unstable configurations. It should also be noted that the same machinery supports inverse tasks without auxiliary networks, enabling large-batch, first-order exploration of G-closure subsets while remaining inside the Voigt-Reuss bounds at every iterate.

Overall, Voigt-Reuss net delivers accurate, admissible, and data-efficient surrogates for multiscale materials, unifying forward prediction and gradient-based design under a single, physically grounded normalization.

6.2 | Discussion

The results in 3D and 2D together highlight two complementary roles of the Voigt-Reuss net framework. First, when coupled with hand-crafted descriptors in 3D, it acts as a physically constrained surrogate for effective stiffness that is robust to large contrast variation but fundamentally limited by the information content of the input features (here: isotropic descriptors). Second, in the 2D setting with a differentiable renderer, it becomes a fully image-based, end-to-end differentiable map from parametric microstructure generators $(\underline{A}, \tau)$ to effective tensors, opening the door to gradient-based inverse design. Both use cases share the same spectral normalization backbone and loss, suggesting that the methodology is not tied to any particular architecture but rather to the Löwner-bounded reparameterization of constitutive tensors.

The 3D study illustrates a key structural limitation that is often implicit in data-driven homogenization: when the feature map is invariant under $SO(3)$, no surrogate can recover orientation-dependent anisotropy. In our case, the descriptors¹² encode only isotropy-averaged information, so the model can learn at best the isotropic projection $\bar{\underline{C}}_{\text{iso}}$ of the effective stiffness. The sharp split in the R^2 -matrix, near-perfect prediction for isotropic components, and $R^2 \approx 0$ for anisotropy-revealing entries is not a failure of the network but a faithful reflection of the available input information. The ECDF analysis makes this explicit: the isotropic projection error forms a hard lower floor for any surrogate using these descriptors, and Voigt-Reuss net operates close to that floor. This has two practical consequences. First, it provides a diagnostic for feature sufficiency: if anisotropic components remain unpredictable while isotropic ones are accurate, one should enrich the descriptor set (e.g., directional correlations, band features, Minkowski tensors, or full CNN-based encoders). Second, even in such limited regimes, spectral normalization still guarantees admissible outputs and stable training, making it a safe drop-in replacement for unconstrained regressors.

The 2D experiments, by contrast, operate with a much richer generative parameterization of microstructures via thresholded trigonometric fields. Here, the bottleneck is no longer missing directional information but the size and diversity of the G-closure subset explored. Embedding the amplitude matrix into a fixed hyper-grid, rendering via a differentiable trigonometric field, and regressing the normalized spectral degrees of freedom together yield a surrogate that is sensitive to connectivity, topology changes, and extreme anisotropy, yet remains numerically stable. The threshold-sweep study is particularly interesting: Voigt-Reuss net tracks sharp eigenvalue jumps at percolation events and exposes these transitions through spikes in $\partial \|\hat{\underline{C}}\|_F / \partial \tau$, while classical Mori-Tanaka estimates deteriorate as their dilute, non-interacting assumptions are violated. This combination of geometric sensitivity and spectral admissibility is precisely what is needed for reliable inverse design in high-contrast media. Also, the straightforward gradient evaluation facilitates the detection of such jumps.

Against this backdrop, the inverse-design studies highlight the distinctive benefits of a spectrally normalized surrogate: (i) **Robust gradients.** All outputs live in $[0, 1]^m$, and the renderer $\mathcal{R}$ is smooth in $(\underline{A}, \tau)$; hence, gradients remain well behaved even at extreme phase contrasts, enabling stable first-order optimization without manual rescaling or ad-hoc clipping. (ii) **Batched optimization.** Optimizing hundreds of candidates in one GPU pass provides a diverse family of designs (useful for screening, multi-criteria trade-offs, or robustness analysis) essentially at the cost of a "single run". (iii) **Single model, forward and backward.** Unlike frameworks that pair a forward surrogate with a separate inverse network or require differentiable simulators, our pipeline uses one spectrally normalized surrogate both for prediction and for backpropagating arbitrary differentiable objectives (tensor matching, eigenvalue ratios, directional moduli, compliance measures, volume constraints, etc.). (iv) **Guaranteed admissibility.** Every iterate and every final design respects Voigt-Reuss bounds in the Löwner sense; the optimizer never leaves the feasible spectral simplex, avoiding unphysical attractors that commonly plague unconstrained neural inverse design. (v) **Empirical coverage of G-closure.** The cosine-threshold parameterization spans a rich, multimodal subset of the G-closure. Empirically, we recover both classical optima (e.g., 45° laminates at $c_0 \approx 0.5$ for maximal shear-normal coupling) and less obvious patterns that realize nearly identical effective property tensors, directly reflecting the intrinsic non-uniqueness of microstructure-to-property mappings.

At a higher level, the framework can be viewed as a "spectral layer" that decouples physical admissibility from architectural choice. Any regressor that produces normalized spectral parameters $\hat{\underline{\xi}}$ can be wrapped by $\mathcal{S}^{-1}$ to yield outputs that are symmetric positive definite and automatically bounded by Voigt and Reuss. This separation of concerns should generalize beyond small-strain elasticity to other linear elliptic operators, and to multi-property surrogates where several coupled tensors must remain jointly admissible. Similarly, the 2D image-based architecture is not unique: one could replace the cosine renderer with voxelized tomography, graph-based generators, or diffusion models, while still leveraging the same spectral normalization for the output layer.

There are, however, clear limitations and avenues for improvement. In 3D, the present results are intentionally constrained by isotropic descriptors; to capture full anisotropy, one must either (i) augment the feature set with orientation-sensitive statistics or (ii) transition to 3D CNNs operating directly on voxelized microstructures, at a much higher computational cost. In 2D, the cosine-threshold family, while expressive, samples only a subset of all possible microstructures; enriching the generator (e.g., multi-scale fields, hierarchical patterns, or manufacturing-aware constraints) would expand the reachable portion of the G-closure. Moreover, the current formulation is deterministic: incorporating uncertainty quantification (e.g., Bayesian surrogates or ensemble methods) would allow one to propagate microstructural variability and quantify confidence in inverse designs, which is essential for safety-critical applications.

Finally, the combination of spectral normalization and end-to-end differentiability makes Voigt-Reuss net compatible with existing structural and topology-optimization workflows. Because the surrogate provides fast, differentiable, and physically admissible predictions, it can act as a "constitutive oracle" inside larger multi-scale loops, enabling gradient-based co-design of

architecture and material. The same machinery supports inverse tasks without auxiliary networks, enabling large-batch, first-order exploration of G-closure subsets while staying inside Voigt-Reuss bounds at every iterate. In this sense, Voigt-Reuss net offers a unifying, physics-grounded interface between classical homogenization theory and modern deep-learning-based materials design.

ACKNOWLEDGMENTS

Funded by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2075 - 390740016. Contributions by Felix Fritzen are funded by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) within the Heisenberg program - DFG-FR2702/10 - 517847245. This work is also supported as part of the consortium NFDI-MatWerk, funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under the National Research Data Infrastructure - NFDI 38/1 - 460247524. We acknowledge the support of the Stuttgart Center for Simulation Science (SimTech).

Financial disclosure

None reported.

Conflict of interest

The authors declare no potential conflict of interest.

DATA AVAILABILITY STATEMENT

Upon acceptance of the manuscript, the problem datasets will be published, as well as a simple implementation of the spectral normalization method in Pytorch.

References

1. Torquato S. *Random heterogeneous materials*. Interdisciplinary applied mathematics New York, NY: Springer, 2013.
2. Milton GW. *The theory of composites*. Cambridge Monographs on Applied and Computational Mathematics Cambridge, England: Cambridge University Press (Virtual Publishing), 2009.
3. Zohdi T, Wriggers P. *An Introduction to Computational Micromechanics: Corrected Second Printing*. Springer Berlin Heidelberg, 2005
4. Nemat-Nasser S, Hori M, Achenbach J. *Micromechanics: Overall Properties of Heterogeneous Materials*. ISSN Elsevier Science, 2013.
5. Cherkaev A. *Obtaining G-Closures*:261–277; New York, NY: Springer New York . 2000
6. Allaire G. *The mathematical modeling of composite materials*:91–188; New York, NY: Springer New York . 2002
7. Moulinec H, Suquet P. A numerical method for computing the overall response of nonlinear composites with complex microstructure. *Computer Methods in Applied Mechanics and Engineering*. 1998;69–94.
8. Leuschner M, Fritzen F. Fourier-Accelerated Nodal Solvers (FANS) for homogenization problems. *Computational Mechanics*. 2017;62(3):359–392. doi: 10.1007/s00466-017-1501-5
9. Schneider M. A review of nonlinear FFT-based computational homogenization methods. *Acta Mechanica*. 2021;232(6):2051–2100. doi: 10.1007/s00707-021-02962-1
10. Keshav S, Fritzen F, Kabel M. FFT-based homogenization at finite strains using composite boxels (ComBo). *Computational Mechanics*. 2023(1):191–212. doi: <https://doi.org/10.1007/s00466-022-02232-4>
11. Lißner J, Fritzen F. Data-Driven Microstructure Property Relations. *Mathematical and Computational Applications*. 2019;24(2):57. doi: 10.3390/mca24020057

12. Prifling B, Röding M, Townsend P, Neumann M, Schmidt V. Large-Scale Statistical Learning for Mass Transport Prediction in Porous Materials Using 90, 000 Artificially Generated Microstructures. *Frontiers in Materials*. 2021;8. doi: 10.3389/fmats.2021.786502
13. Lißner J, Fritzen F. Microstructure homogenization: human vs machine. *Advanced Modeling and Simulation in Engineering Sciences*. 2024;11(1). doi: 10.1186/s40323-024-00275-1
14. Aldakheel F, Elsayed ES, Zohdi TI, Wriggers P. Efficient multiscale modeling of heterogeneous materials using deep neural networks. *Computational Mechanics*. 2023;72(1):155–171. doi: 10.1007/s00466-023-02324-9
15. Nguyen PT, Heider Y, Kochmann DM, Aldakheel F. Deep learning-aided inverse design of porous metamaterials. 2025.
16. Heider Y, Aldakheel F, Ehlers W. A multiscale CNN-based intrinsic permeability prediction in deformable porous media. *Appl. Sci. (Basel)*. 2025;15(5):2589.
17. Eidel B. Deep CNNs as universal predictors of elasticity tensors in homogenization. *Computer Methods in Applied Mechanics and Engineering*. 2023;403:115741. doi: <https://doi.org/10.1016/j.cma.2022.115741>
18. Bhattacharya K, Kovachki NB, Rajan A, Stuart AM, Trautner M. Learning Homogenization for Elliptic Operators. *SIAM Journal on Numerical Analysis*. 2024;62(4):1844–1873. doi: 10.1137/23M1585015
19. Nguyen BH, Schneider M. Universal Fourier Neural Operators for periodic homogenization problems in linear elasticity. *Journal of the Mechanics and Physics of Solids*. 2026;206:106418. doi: <https://doi.org/10.1016/j.jmps.2025.106418>
20. As'ad F, Avery P, Farhat C. A mechanics-informed artificial neural network approach in data-driven constitutive modeling. *International Journal for Numerical Methods in Engineering*. 2022;123(12):2738–2759. doi: 10.1002/nme.6957
21. Linden L, Klein DK, Kalina KA, Brummund J, Weeger O, Kästner M. Neural networks meet hyperelasticity: A guide to enforcing physics. *Journal of the Mechanics and Physics of Solids*. 2023;179:105363. doi: 10.1016/j.jmps.2023.105363
22. Klein DK, Hossain M, Kikinov K, Kannapinn M, Rudykh S, Gil AJ. Neural networks meet hyperelasticity: A monotonic approach. *Eur. J. Mech. A Solids*. 2026;116(105900):105900.
23. Bastek JH, Kochmann DM. Inverse design of nonlinear mechanical metamaterials via video denoising diffusion models. *Nat. Mach. Intell.*. 2023;5(12):1466–1475.
24. Raßloff A, Seibert P, Kalina KA, Kästner M. Inverse design of spinodoid structures using Bayesian optimization. *Computational Mechanics*. 2025. doi: 10.1007/s00466-024-02587-w
25. Otto A, Rosenkranz M, Kalina KA, Kästner M. Data-driven inverse design of spinodoid architected materials. *GAMM-Mitteilungen*. 2025;48(4).
26. Rosenkranz M, Kästner M, Sbalzarini IF. Data-efficient inverse design of spinodoid metamaterials. *Integr. Mater. Manuf. Innov.*. 2025.
27. Kumar S, Tan S, Zheng L, Kochmann DM. Inverse-designed spinodoid metamaterials. *Npj Comput. Mater.*. 2020;6(1).
28. Zheng L, Karapiperis K, Kumar S, Kochmann DM. Unifying the design space and optimizing linear and nonlinear truss metamaterials by generative modeling. *Nat. Commun.*. 2023;14(1):7563.
29. Xu K, Huang DZ, Darve E. Learning constitutive relations using symmetric positive definite neural networks. *Journal of Computational Physics*. 2021;428:110072. doi: 10.1016/j.jcp.2020.110072
30. Karniadakis GE, Kevrekidis IG, Lu L, Perdikaris P, Wang S, Yang L. Physics-informed machine learning. *Nature Reviews Physics*. 2021;3(6):422–440. doi: 10.1038/s42254-021-00314-5
31. Voigt W. Ueber Die Beziehung Zwischen Den Beiden Elastizitätsconstanten Isotroper Körper. *Annalen der Physik*. 1889;274(12):573–587. doi: 10.1002/andp.18892741206
32. Reuss A. Berechnung Der Fließgrenze von Mischkristallen Auf Grund Der Plastizitätsbedingung Für Einkristalle .. *ZAMM - Journal of Applied Mathematics and Mechanics*. 1929;9(1):49–58. doi: 10.1002/zamm.19290090104
33. Wiener O. Die Theorie Des Mischkörpers Für Das Feld Der Stationären Strömung 1 Die Mittelwertsätze Für Kraft, Polarisation Und Energie / von Otto Wiener. *Abhandlungen der Königlich Sächsischen Gesellschaft der Wissenschaften ; 61,6*. 1912.
34. Hashin Z, Shtrikman S. A variational approach to the theory of the elastic behaviour of multiphase materials. *J. Mech. Phys. Solids*. 1963;11(2):127–140.
35. Castaneda PP. Bounds and estimates for the properties of nonlinear heterogeneous systems. *Philos. Trans. Phys. Sci. Eng.*. 1992;340(1659):531–567.
36. Talbot DRS, Willis JR. Bounds for the effective constitutive relation of a nonlinear composite. *Proc. Math. Phys. Eng. Sci.*. 2004;460(2049):2705–2723.
37. Boddapati J, Daraio C. Planar structured materials with extreme elastic anisotropy. *Mater. Des.*. 2024;246(113348):113348. doi: 10.1016/j.matdes.2024.113348

38. Keshav S, Herb J, Fritzen F. Spectral Normalization and Voigt-Reuss net: A universal approach to microstructure-property forecasting with physical guarantees. *GAMM-Mitteilungen*. 2025;48(3). doi: 10.1002/gamm.70005
39. Prifling B, Röding M, Townsend P, Neumann M, Schmidt V. Large-scale statistical learning for mass transport prediction in porous materials using 90,000 artificially generated microstructures. Zenodo; 2021
40. Boddapati J, Flaschel M, Kumar S, De Lorenzis L, Daraio C. Single-test evaluation of directional elastic properties of anisotropic structured materials. *J. Mech. Phys. Solids*. 2023;181(105471):105471. doi: 10.1016/j.jmps.2023.105471
41. Penrose R. A generalized inverse for matrices. *Mathematical Proceedings of the Cambridge Philosophical Society*. 1955;51(3):406–413. doi: 10.1017/s0305004100030401
42. Ansel J, Yang E, He H, et al. PyTorch 2: Faster Machine Learning Through Dynamic Python Bytecode Transformation and Graph Compilation. In: ASPLOS '24. ACM. Association for Computing Machinery 2024; New York, NY, USA:929–947
43. Fernández M, Fritzen F, Weeger O. Material modeling for parametric, anisotropic finite strain hyperelasticity based on machine learning with application in optimization of metamaterials. *International Journal for Numerical Methods in Engineering*. 2021;123(2):577–609. doi: 10.1002/nme.6869
44. Keshav S, Rieg F, Desai I, Sigg M, Haag C, Fritzen F. FANS - An open-source, efficient, and parallel implementation of the Fourier-Accelerated Nodal Solvers for microscale multiphysics problems.. Github repository; 2024.
45. He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. In: IEEE. 2016:770–778.
46. He K, Zhang X, Ren S, Sun J. Identity Mappings in Deep Residual Networks. In: Leibe B, Matas J, Sebe N, Welling M., eds. *Computer Vision – ECCV 2016* Springer International Publishing 2016; Cham:630–645.
47. Bastek JH. Inverse-design of nonlinear mechanical metamaterials via video denoising diffusion models: Dataset and model checkpoints. ETH Zurich research collection; 2023