Differentiable FFT Homogenization for CTE-Driven Design of Thermoelastic Metamaterials with Double-Negative Response

Chenchen CHU
chuchenchen@live.com

Abstract

We present a differentiable topology-optimization framework for the inverse design of three-phase (void plus two solids) periodic thermoelastic metamaterials, with the primary objective of achieving a negative mean effective coefficient of thermal expansion (CTE). The framework combines a matrix-free FFT-based conjugate-gradient (FFT-CG) homogenization solver with JAX-native automatic differentiation: the equilibrium solve is wrapped in an implicit linear-solve boundary, so reverse-mode gradients follow from adjoint solves without hand-coding the full sensitivity, and a staged continuation of the SIMP exponent and projection sharpness guides multiphase designs from a gray field to a near-binary layout. The solver is verified against homogeneous limits, symmetry and positive-definiteness checks, finite-difference gradient comparisons, and an independent finite-element reference. For a 100×100 plane-stress cell, CTE-driven optimization produces a design with a negative mean in-plane homogenized CTE at feasible modulus floors. Postprocessing of the same homogenized stiffness shows that the CTE-optimized design also has a negative effective Poisson ratio; this auxetic response is an observed outcome rather than an optimization target. Five of the six screened initializations yield designs with negative mean in-plane CTE, and a $40 \times 40 \times 40$ three-dimensional run reaches a negative mean CTE as budget-limited feasibility evidence.

Keywords: topology optimization; FFT-based homogenization; automatic differentiation; negative thermal expansion; negative Poisson ratio; thermoelastic metamaterials

1 Introduction

Metamaterials — microstructured materials whose effective response is engineered by internal geometry rather than bulk constituents — offer access to extreme effective properties. Among the most striking targets are materials with a negative effective coefficient of thermal expansion ($\alpha_H^{\text{mean}} < 0$) and materials with negative effective Poisson ratio ($\nu_H < 0$, auxetic solids). Realizing *both* in a single periodic cell — a "double-negative" thermoelastic metamaterial — is an appealing yet challenging inverse-design target because the two properties arise from distinct and often competing mechanisms. Negative Poisson ratio is typically produced by reentrant or chiral architectures that convert axial stretch into lateral expansion [19, 14]. Negative thermal expansion, by contrast, is a multiphase effect: a large CTE contrast between two solids, arranged so that differential expansion is converted into bending or rotation of connecting members, can make the periodic cell contract upon heating [20, 4, 6]. Realizing both responses in one microstructure therefore requires at least two solids plus void, together with a geometry that can host both the reentrant kinematics and the bimaterial bending mechanism [16, 17, 22].

Inverse design of the representative volume element (RVE) whose homogenized thermoelastic tensors are obtained through computational homogenization [5, 1] proceeds along three main routes, each with characteristic strengths and limitations:

- **Topology optimization** treats the phase layout as a free field and seeks extreme homogenized properties by gradient-based search [5, 6, 7, 21, 22, 18]. It is the only route that can explore beyond a prescribed cell family, but it relies on two ingredients remaining reliable throughout the search: a stable homogenization forward model and a continuation strategy that can drive a multiphase design from a gray field to a near-binary layout without stalling.

- **Parametric design** restricts the search to a preselected family of unit cells, typically reentrant or chiral lattices defined by a handful of geometric parameters [16, 17]; the kinematics built into the parameterization limit the topologies that can be discovered.

- **Data-driven inverse design** learns a map from geometry to effective properties, or from target properties back to geometry, from a precomputed library of cells [23, 15]; it inherits the same restriction unless the training library already spans the relevant topologies, and it shifts the burden onto the quality and coverage of that library.

FEM-based multiphase topology optimization has already produced double-negative CTE–Poisson designs, including three-phase constructions [4, 5], later energy-based homogenization studies [6, 7], additively manufactured double-negative cells [21], a multiobjective analysis of the conflict between the two properties [22], and VAE-driven initial-dataset generation with an Alternative Active Phase and Objective (AAPO) method [15]. Automatic differentiation has recently been used to differentiate FFT-based homogenization solves for solid mechanics, with topology optimization demonstrated for mechanical objectives only [11, 8, 9, 10]. Existing negative-CTE topology optimization is predominantly FEM-based, whereas differentiable FFT homogenization has mainly been demonstrated for mechanical objectives. A reliable differentiable FFT framework for three-phase negative-CTE topology optimization has not yet been established.

An alternative is to evaluate the homogenized tensors on a regular pixel (2D) or voxel (3D) grid by discrete Fourier transforms. The method originates from Fourier-based homogenization [1, 2] and has matured into a broad family of solvers [3, 2]; it avoids a global finite-element assembly and scales as $O(N \log N)$ per iteration. For the purpose of this study, the decisive feature is that the entire topology optimization chain — material interpolation, periodic filter, projection, and the matrix-free equilibrium solve — is expressed in a differentiable form, so that the homogenized response can be differentiated with respect to the design fields. We combine a matrix-free FFT-based conjugate-gradient (FFT-CG) solver with a spectral reference preconditioner, a gauge penalty that removes the rigid-body singularity without altering the zero-mean physics, and a hierarchical two-field design parameterization whose raw fields pass through a periodic filter and a tanh projection. The equilibrium solve is wrapped in an implicit linear-solve boundary, so reverse-mode gradients follow from adjoint solves without hand-coding the full sensitivity. On the design side, the present implementation targets the mean in-plane homogenized CTE as the primary objective, with modulus floors (Young's and shear) as constraints, and the Poisson ratio is evaluated from the same effective stiffness obtained from the mechanical load cases.

1.1 Contributions

The contributions of this work are:

- A differentiable, three-phase thermoelastic FFT-CG homogenization framework with implicit-adjoint reverse-mode gradients (Secs. 2, 3.1).

- CTE-driven topology optimization with a staged continuation strategy and its numerical verification, demonstrated by a two-dimensional design reaching a negative mean homogenized CTE (Sec. 3).

- A CTE-driven design that also exhibits a negative effective Poisson ratio in post-processing, together with initialization screening (five of six guesses reach negative mean in-plane CTE) and three-dimensional negative-CTE feasibility evidence (Sec. 3).

2 Materials and methods

2.1 Framework overview

The computational pipeline of the framework is a single differentiable chain:

design fields → filter/projection → phase properties → FFT-CG → homogenized response → adjoint gradients → optimizer,

as illustrated in Fig. 1. Two raw design fields are mapped through a periodic filter and a tanh projection to three phase fractions representing void and two distinct solids; a SIMP interpolation converts them into periodic coefficient fields; an FFT-based conjugate-gradient solve computes the equilibrium fluctuation for each load case; volume averaging yields the homogenized thermoelastic tensors that define the objective and constraints; and reverse-mode differentiation through the implicit solve boundary provides the gradients that drive the optimizer.

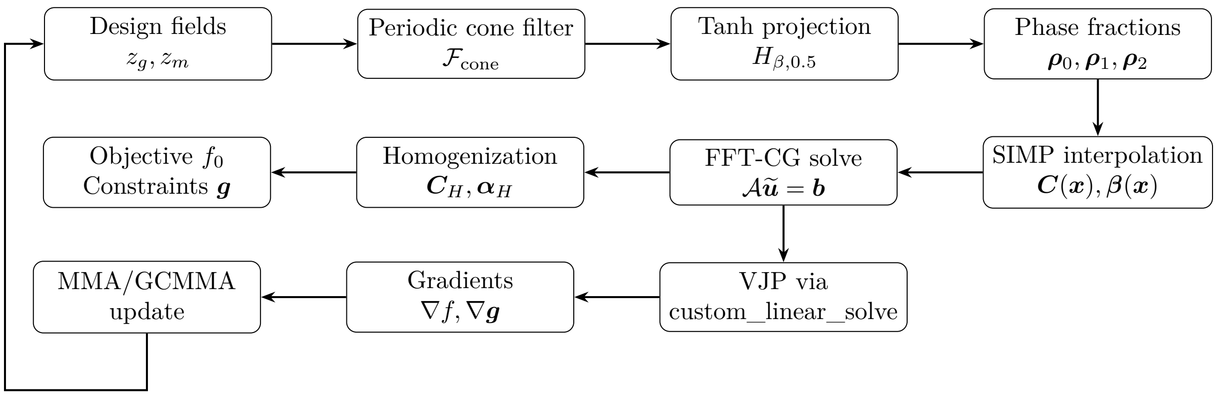

Figure 1: Computational pipeline of the differentiable FFT-homogenization topology-optimization framework. The forward pass (top row) maps design fields to homogenized properties; the reverse pass (bottom row) computes gradients via the implicit differentiation boundary.

2.2 Design and optimization problem

The design is described by two independent periodic fields $z_g, z_m \in [0,1]^{n_1\times\cdots\times n_d}$, flattened as $\boldsymbol{x} = [\text{vec}(z_g), \text{vec}(z_m)] \in [0,1]^{2N}$. Each raw field is passed through a normalized periodic cone filter and a normalized tanh projection:

$$\tilde{z_g} = \mathcal{F}_{\text{cone}} z_g, \qquad \tilde{z_m} = \mathcal{F}_{\text{cone}} z_m, \qquad s = H_{\beta_g,0.5}(\tilde{z_g}), \qquad m = H_{\beta_m,0.5}(\tilde{z_m}). \tag{1}$$

The cone filter kernel is

$$K(\boldsymbol{q}) = \frac{\max(0, r - d(\boldsymbol{q}))}{\sum_{\boldsymbol{q}} \max(0, r - d(\boldsymbol{q}))}, \tag{2}$$

where $d(\boldsymbol{q})$ is the minimum-image distance and r is the filter radius in physical length units (reported with the case setup in Table 9). The projection uses threshold θ and sharpness β with the normalized form

$$H_{\beta,\theta}(z) = \frac{\tanh(\beta\theta) + \tanh(\beta(z-\theta))}{\tanh(\beta\theta) + \tanh(\beta(1-\theta))}. \tag{3}$$

The filtered and projected fields define three phase fractions on the simplex:

$$\boldsymbol{\rho}_0 = 1 - s, \qquad \boldsymbol{\rho}_1 = s\,(1-m), \qquad \boldsymbol{\rho}_2 = s\,m, \tag{4}$$

with phase order $(0,1,2) = (\text{void}, \text{solid 1}, \text{solid 2})$. Throughout, the isotropic source law uses the engineering-shear Voigt order of the ambient dimension; the two-dimensional plane-stress specialization of this work uses Voigt order $(11,22,12)$, with the phase pair obtained by Schur condensation of the out-of-plane block of the full three-dimensional tensor. Material coefficients (the stiffness tensor $\boldsymbol{C}$ and the thermal-stress coefficient $\boldsymbol{\beta}$) are interpolated with a common SIMP power $p \geq 1$ applied to the solid occupancy (the projected geometry field s):

$$\boldsymbol{C}(\boldsymbol{x}) = \boldsymbol{C}_0 + s^p\Big[\sum_{k=1}^{2} q_k\,\boldsymbol{C}_k - \boldsymbol{C}_0\Big], \qquad \boldsymbol{\beta}(\boldsymbol{x}) = \boldsymbol{\beta}_0 + s^p\Big[\sum_{k=1}^{2} q_k\,\boldsymbol{\beta}_k - \boldsymbol{\beta}_0\Big], \tag{5}$$

where $q = (1-m, m)$ are the solid-simplex weights. For $p = 1$ the interpolation (5) is exactly the affine mixture of the reduced phase tables; higher powers penalize intermediate solid occupancy and are advanced by the continuation strategy of Sec. 2.5.

The reference scales are defined using the reference solid: E_{ref} is the solid Young modulus, ν_{ref} the solid Poisson ratio, $G_{\mathrm{ref}} = E_{\mathrm{ref}}/[2(1+\nu_{\mathrm{ref}})]$ the corresponding shear modulus, and α_{ref} the larger solid-phase CTE. The physical objective is the dimensionless normalized mean diagonal CTE:

$$f_0(\boldsymbol{x}) = \frac{1}{n_d}\sum_{i=1}^{n_d} \frac{\boldsymbol{\alpha}_{H\,ii}}{\alpha_{\mathrm{ref}}}, \tag{6}$$

where $n_d = d$ is the number of diagonal Voigt entries. A negative effective CTE corresponds to $f_0 < 0$; in the two-dimensional plane-stress setting the diagonal Voigt entries are $(11,22)$, so $n_d = 2$, f_0 is the normalized mean in-plane CTE, and the physical quantity is $\alpha_H^{\mathrm{mean}} = f_0\,\alpha_{\mathrm{ref}}$. The total objective adds explicit grayness penalties:

$$f(\boldsymbol{x}) = f_0(\boldsymbol{x}) + \lambda_g\,G_g(\boldsymbol{x}) + \lambda_m\,G_m(\boldsymbol{x}), \tag{7}$$

where $\lambda_g, \lambda_m \geq 0$ are the geometry and material penalty weights (stage-dependent; Sec. 2.5) and the grayness metrics of the projected fields are

$$G_g = \langle 4s(1-s)\rangle_\Omega\,, \qquad G_m = \frac{\langle 4s\,m(1-m)\rangle_\Omega}{\langle s\rangle_\Omega + \epsilon_m}, \tag{8}$$

with a nonnegative regularization ϵ_m that keeps the metric finite for an empty solid phase. Both metrics vanish on a binary two-solid design and equal one on the maximally gray field.

The design problem is

$$\min_{\boldsymbol{x}\in[0,1]^{2N}} f(\boldsymbol{x}) \qquad \text{s.t.} \quad \boldsymbol{g}(\boldsymbol{x}) \leq \boldsymbol{0}, \tag{9}$$

with the constraint vector $\boldsymbol{g} \in \mathbb{R}^{n_g}$ (feasibility requires all entries to be nonpositive) collecting a two-sided band on each of the two independent projected phase fractions,

$$g_{2p+1} = V_p - V_p^* - \epsilon_v, \qquad g_{2p+2} = V_p^* - V_p - \epsilon_v, \qquad p \in \{0,1\}, \tag{10}$$

where $V_p = \langle \boldsymbol{\rho}_p\rangle_h$ are the projected phase-volume fractions and $V_0^* + V_1^* + V_2^* = 1$, together with modulus floors on the normalized engineering constants: one Young-modulus floor per diagonal direction and one shear-modulus floor per shear entry,

$$g_{4+i} = e_{\mathrm{floor}} - \hat{E}_i/E_{\mathrm{ref}}, \qquad i = 1,\ldots,d, \tag{11}$$

$$g_{4+d+j} = g_{\mathrm{floor}} - \hat{G}_j/G_{\mathrm{ref}}, \qquad j = 1,\ldots,d(d-1)/2, \tag{12}$$

where $\hat{E}_i = 1/S_{ii}$ are the diagonal Young moduli and $\hat{G}_j$ the shear moduli $1/S_{kk}$ of the reduced Voigt space; the floors require the homogenized moduli to remain above the prescribed fractions of the reference scales. In the two-dimensional plane-stress setting, $d = 2$ with Voigt order $(11, 22, 12)$, the engineering constants are $\hat{E}_1 = E_1$, $\hat{E}_2 = E_2$, $\hat{G}_1 = G_{12}$, and the constraint vector is the seven-row system

$$g_1 = V_0 - V_0^* - \epsilon_v, \qquad g_2 = V_0^* - V_0 - \epsilon_v, \tag{13}$$
$$g_3 = V_1 - V_1^* - \epsilon_v, \qquad g_4 = V_1^* - V_1 - \epsilon_v, \tag{14}$$
$$g_5 = e_{\text{floor}} - E_1/E_{\text{ref}}, \qquad g_6 = e_{\text{floor}} - E_2/E_{\text{ref}}, \qquad g_7 = g_{\text{floor}} - G_{12}/G_{\text{ref}}, \tag{15}$$

with $n_g = 4 + 2 + 1 = 7$. The demonstration values of the volume targets, the band half-width ϵ_v, and the floors are reported with the case setup (Table 9).

2.3 Periodic homogenization and FFT-CG

We consider a periodic three-phase representative volume element in d spatial dimensions, occupying $\Omega = [0, L_1) \times \cdots \times [0, L_d)$ on a regular pixel grid with spacings $h_i = L_i/n_i$ $(i = 1, \ldots, d)$. The displacement field is $u(\boldsymbol{x}) = \bar{\boldsymbol{E}}\boldsymbol{x} + \widetilde{\boldsymbol{u}}(\boldsymbol{x})$, where $\bar{\boldsymbol{E}}$ is the macroscopic strain and $\widetilde{\boldsymbol{u}}$ is periodic with $\langle \widetilde{\boldsymbol{u}} \rangle_h = 0$. Under small strains, the periodic equilibrium problem for a uniform temperature increment ΔT is

$$\operatorname{div}_h \boldsymbol{\sigma} = \boldsymbol{0} \quad \text{in } \Omega, \tag{16a}$$
$$\boldsymbol{\sigma}(\boldsymbol{x}) = \boldsymbol{C}(\boldsymbol{x})[\bar{\boldsymbol{E}} + \boldsymbol{\varepsilon}(\widetilde{\boldsymbol{u}})] - \boldsymbol{\beta}(\boldsymbol{x})\,\Delta T, \tag{16b}$$

with sign convention $\boldsymbol{\sigma} = \boldsymbol{C}\boldsymbol{\varepsilon} - \boldsymbol{\beta}\,\Delta T$ and $\boldsymbol{\beta} = \boldsymbol{C}\boldsymbol{\alpha}$.

Effective properties are obtained from one periodic equilibrium problem (16) per load case. With the engineering-shear Voigt order of dimension d, the $n_{\text{mech}} = d(d+1)/2$ mechanical load cases impose the macroscopic strain bases $\bar{\boldsymbol{E}} = a\,\boldsymbol{e}^{(i)}$ $(i = 1, \ldots, n_{\text{mech}})$ with $\Delta T = 0$, and one constrained thermal load case imposes $\bar{\boldsymbol{E}} = \boldsymbol{0}$, $\Delta T = t$; in the two-dimensional plane-stress setting of this work this gives three mechanical load cases plus one constrained thermal load case (four in total). All load cases share the same coefficient fields $\boldsymbol{C}(\boldsymbol{x}), \boldsymbol{\beta}(\boldsymbol{x})$, the same zero-mean periodic fluctuation $\langle \widetilde{\boldsymbol{u}} \rangle_h = 0$, and the same discrete operator; they differ only through the imposed pair $(\bar{\boldsymbol{E}}, \Delta T)$, which enters the right-hand side $\boldsymbol{b} = \operatorname{div}_h[\boldsymbol{C}\bar{\boldsymbol{E}} - \boldsymbol{\beta}\Delta T]$ of the operator form (20). Each load case is solved by the FFT-CG solver described below, which returns the periodic fluctuation $\widetilde{\boldsymbol{u}}_J(\boldsymbol{x})$; the raw effective stiffness and the thermal-stress coefficient follow by volume averaging:

$$\boldsymbol{C}^{\text{raw}}_{:,J} = \langle \boldsymbol{\sigma}_J \rangle_h / a, \qquad \boldsymbol{C}_H = \tfrac{1}{2}(\boldsymbol{C}^{\text{raw}} + \boldsymbol{C}^{\text{raw}\,\mathsf{T}}), \tag{17a}$$
$$\boldsymbol{\beta}_H = -\left\langle \boldsymbol{\sigma}^{\text{th}} \right\rangle_h / t, \tag{17b}$$
$$\boldsymbol{\alpha}_H = \boldsymbol{S}_H\,\boldsymbol{\beta}_H, \qquad \boldsymbol{S}_H = \boldsymbol{C}_H^{-1}. \tag{17c}$$

Here a is the macroscopic engineering-strain amplitude and t is the temperature increment of the constrained thermal load; both are finite and nonzero, and the solver verification of Sec. 3.1 shows the extracted properties to be insensitive to their magnitude (load-amplitude invariance). The engineering constants follow from the compliance matrix ($\hat{E}_i = 1/S_{ii}$, $\hat{G}_j = 1/S_{kk}$, $\nu_{12} = -S_{21}/S_{11}$); they enter the modulus floors of the design problem (9).

On the periodic grid a forward difference for the gradient and the adjoint (backward) difference for the divergence make the discrete integration-by-parts identity exact. In 2D engineering-shear Voigt order $(11, 22, 12)$:

$$\boldsymbol{B}_h\,\widetilde{\boldsymbol{u}} := [D_x^+\widetilde{\boldsymbol{u}}_1,\ D_y^+\widetilde{\boldsymbol{u}}_2,\ D_y^+\widetilde{\boldsymbol{u}}_1 + D_x^+\widetilde{\boldsymbol{u}}_2]^{\mathsf{T}}, \tag{18}$$

$$\mathrm{div}_h\, \boldsymbol{\sigma} = \left[D_x^- \sigma_{11} + D_y^- \sigma_{12},\ D_x^- \sigma_{12} + D_y^- \sigma_{22}\right]^{\mathsf{T}}, \tag{19}$$

so the operators are exact spectral adjoints; the Fourier multiplier forms and the treatment of the single unresolved zero mode are given in Appendix B.

The equilibrium operator and right-hand side are

$$\mathcal{A}\,\tilde{\boldsymbol{u}} = -\,\mathrm{div}_h\big[\boldsymbol{C}\,\boldsymbol{B}_h\,\tilde{\boldsymbol{u}}\big] + \zeta\,\mathcal{P}_0\,\tilde{\boldsymbol{u}}, \tag{20a}$$

$$\boldsymbol{b} = \mathrm{div}_h\big[\boldsymbol{C}\,\bar{\boldsymbol{E}} - \boldsymbol{\beta}\,\Delta T\big], \tag{20b}$$

$$\mathcal{A}\,\tilde{\boldsymbol{u}} = \boldsymbol{b}, \tag{20c}$$

where $\mathcal{P}_0$ projects onto the zero-frequency mode and $\zeta > 0$ is a dimensional gauge penalty:

$$\zeta = \langle C_{1111}\rangle_h\, \max_i \Big(\min_{|\delta_i|>0} |\delta_i|\Big)^2. \tag{21}$$

The gauged operator is invertible and self-adjoint on the full space, and its restriction to the zero-mean subspace is the physical equilibrium operator. The Krylov solve is preconditioned with a spectral inverse of the homogeneous-reference operator. In each Fourier mode:

$$\mathcal{A}_0(k) = \boldsymbol{B}_h(k)^{\mathsf{H}}\, \boldsymbol{C}_{\mathrm{ref}}\, \boldsymbol{B}_h(k), \qquad \text{zero mode: } \mathcal{A}_0(0) = \zeta\,\boldsymbol{I}, \tag{22}$$

where $\boldsymbol{C}_{\mathrm{ref}}$ is the voxel arithmetic mean stiffness (the homogeneous reference medium). The PCG iteration starts from the zero field and stops when

$$\|\boldsymbol{r}_k\|_2 \le \mathrm{rtol}\,\|\boldsymbol{b}\|_2, \qquad \boldsymbol{r}_k = \boldsymbol{b} - \mathcal{A}\,\tilde{\boldsymbol{u}}_k, \tag{23}$$

where rtol is the relative residual tolerance and maxiter is the iteration budget; the demonstration values are reported with the case setup (Table 9).

2.4 Implicit differentiation

Three properties of the differentiation boundary matter for this framework. First, the equilibrium equations constrain the design problem (9): the objective depends on the design both directly and through the converged fluctuations of the load cases. Second, the primal solve is wrapped in JAX's `lax.custom_linear_solve`, which differentiates the converged equation implicitly: reverse-mode gradients follow the converged gauged equation $\mathcal{A}\,\tilde{\boldsymbol{u}} = \boldsymbol{b}$ via the adjoint equation $\mathcal{A}\,\boldsymbol{\lambda} = \boldsymbol{c}$ with the same operator, preconditioner, and relative stop rule. Third, the reverse pass therefore solves one adjoint linear system per load case instead of unrolling the CG history, so the computational graph is fixed regardless of the design-dependent iteration count. Summing over the n_{lc} load cases,

$$\frac{df}{d\boldsymbol{x}} = \frac{\partial f}{\partial \boldsymbol{x}}\Big|_{\mathrm{direct}} - \sum_{J=1}^{n_{\mathrm{lc}}} \boldsymbol{\lambda}_J^{\mathsf{T}} \frac{\partial \boldsymbol{R}_J}{\partial \boldsymbol{x}}, \qquad \mathcal{A}\,\boldsymbol{\lambda}_J = \left(\frac{\partial \tilde{f}}{\partial \tilde{\boldsymbol{u}}_J}\right)^{\mathsf{T}}, \tag{24}$$

where $\boldsymbol{R}_J = \mathcal{A}\,\tilde{\boldsymbol{u}}_J - \boldsymbol{b}_J$ is the converged residual of load case J, the solve step contributes the implicit derivative $\mathcal{A}^{-1}$ of the converged gauged equation (20), and the direct term collects the explicit grayness penalties of (8) on the projected fields. The per-factor derivatives of the forward chain — filter convolution, projection derivative, simplex map, SIMP power law, and the residual derivative $\partial \boldsymbol{R}_J/\partial \boldsymbol{x}$ — are derived in Appendix A.

2.5 Optimization and continuation

The constrained problem (9) is solved with the method of moving asymptotes (MMA) and its globally convergent variant GCMMA [12, 13], using gradients obtained by reverse-mode differentiation through the filter, projection, interpolation, and the implicit FFT-CG solve. The

Algorithm 1 Balanced stage-profile continuation

Input: initial design $\boldsymbol{x}_0$, stage profile (Table 7), caps, DR factor
Output: optimized design $\boldsymbol{x}^*$
for each stage s in profile **do**
  Set $p, \beta_g, \beta_m, \lambda_g, \lambda_m$ from stage s ▷ SIMP (5); projection (3); objective (7)
  $k \leftarrow 0$
  while $k < \text{cap}_s$ **do**
    Evaluate $f(\boldsymbol{x}_k)$, $\boldsymbol{g}(\boldsymbol{x}_k)$ ▷ problem (9) via FFT-CG (16)
    Evaluate $\nabla f(\boldsymbol{x}_k)$ via implicit adjoint ▷ Sec. 2.4
    Update $\boldsymbol{x}_{k+1}$ via GCMMA ▷ Sec. 2.5
    $k \leftarrow k + 1$
    if ready condition met **or** DR condition met **then**
      break
    end if
  end while
  Probe next stage: evaluate frozen design at higher p or β
  if probe passes hard gates **then**
    Advance to next stage
  else
    Hold or reject (no change to $\boldsymbol{x}$, p, or β)
  end if
end for

SIMP exponent p and the projection sharpnesses β_g, β_m must be advanced gradually rather than fixed at their final values: at small p the interpolation is soft and the objective is nearly insensitive to the layout in the gray regime, while sharpening the projection increases the CTE-contrast right-hand side of (20) and degrades the conditioning of the equilibrium problem; raising both too quickly can stall the design in a gray state or destabilize the solver. The continuation therefore advances p, β_g, and β_m through discrete stages, and each stage transition is controlled by observables of the run rather than by fixed iteration counts: a stage exits when the trailing committed records form a stable feasible tail (readiness), when the per-operator-application objective gain collapses (diminishing return), or when the stage's committed-step cap is reached, and a transition additionally requires the frozen design to pass hard probe gates at the next level (load-case validity, phase-volume drift, operator work ratio, and grayness non-worsening). Algorithm 1 summarizes the generic stage loop.

The full sixteen-stage schedule, thresholds, windows, caps, and the operator-work ledger are reported in Appendix C (Tables 7 and 8).

3 Results

3.1 Numerical verification of the solver and gradients

Before presenting the optimization cases, four checks establish the numerical basis used below: an analytical homogeneous limit, a directional finite-difference audit of the reverse-mode gradients, an independent finite-element comparison, and a frozen-design resolution and void-ratio study. The operator and gradient evidence is reported in Appendix B (Tables 4 and 5); the discretization and void-sensitivity evidence is reported in Appendix C (Table 10 and Figs. 10–14).

The homogeneous test reaches a maximum relative error of 3.28×10^{-14}. Under the production setting, the largest equilibrium relative residual is 9.87×10^{-9}, the homogenized stiffness is symmetric and positive definite, and the directional reverse-mode gradients agree with centered finite differences to worst-case relative errors from 9.09×10^{-6} down to 9.07×10^{-8} as the step size is reduced. Against an independent JAX-FEM solution of a synthetic three-phase voxel RVE, the differences decrease overall with refinement and, at 20^3, are 1.4% for effective stiffness, 0.2% for the diagonal effective CTE components, and 0.08% for effective Poisson ratio.

For the frozen case-1 design, the 80^2 and 100^2 evaluations of the mean CTE agree within 1.8%, whereas the 40^2 grid under-resolves the reentrant-bimetal mechanism. The production void stiffness ratio 10^{-4} introduces an approximately 7% shift in mean CTE relative to the softer-void limit, but softer voids become increasingly difficult to solve on the finer grids. These results set the production resolution and void ratio used in the optimization cases.

Taken together, the checks support the consistency of the discrete operators and implicit gradients and show agreement with the independent FEM reference within the reported tests and tolerances. They do not replace a design-level FEM cross-validation, which remains open. The 2D optimization results below are therefore interpreted with the observed 1.8% resolution sensitivity and an approximately 7% void-regularization shift.

3.2 Two-dimensional CTE-driven design

The demonstration case is a 100×100 plane-stress cell with three phases: $E_1 = E_2 = 1.0$, $\nu = 0.3$, $\alpha = (10^{-5}, 2 \times 10^{-4})$ K^{-1}, and void stiffness ratio 10^{-4} (phase data in Table 6). The phase-volume targets are $(0.40, 0.30, 0.30)$, the filter radius is 0.05, and the modulus floors are $e_{\text{floor}} = 0.05$ and $g_{\text{floor}} = 0.01$. The design starts from a deterministic D4-symmetric reentrant-bimetal seed, and the equilibrium solves use PCG with relative tolerance 10^{-8} and an iteration budget of 2000. The full parameter table is given in Appendix C (Table 9).

The final accepted design and its optimization history are shown in Fig. 2, and its homogenized properties are reported in Table 1. The run stopped at the Dm_3 committed-step cap without meeting the terminal convergence criterion, so the reported result is the final accepted (cap-terminated) design; over the final 20-record window the physical objective varies by only $\pm 0.2\%$. This terminal-window variation is reported as a stability diagnostic, not as evidence that the unmet terminal convergence criterion can be disregarded.

Table 1: Results for the final accepted design of the two-dimensional demonstration case.

Quantity	Value
Physical objective f_0	-0.526020
Mean in-plane CTE	$\alpha_H^{\text{mean}} = -1.05204 \times 10^{-4}$ K^{-1}
In-plane CTE components	$(\alpha_{11}^H, \alpha_{22}^H) = (-1.05204, -1.05204) \times 10^{-4}$ K^{-1}
Effective Poisson ratio	$\nu_{12} = \nu_{21} = -0.4107$
$E_1,\ E_2$	$0.0500,\ 0.0500$
G_{12}	0.0300
Phase fractions (V_0, V_1, V_2)	$(0.380, 0.320, 0.300)$
Grayness $G_g,\ G_m$	$0.053,\ 0.072$
Max constraint violation $g_{\max}$	-1.21×10^{-6}
Max solver residual (four load cases)	9.87×10^{-9}
Termination	Dm_3 committed-step cap; terminal convergence criterion not met (status terminal_stage_capped)

The physical objective changes from positive at the seed to negative at the final design (Fig. 2), and the final effective Poisson ratio is negative. The modulus floors are active but feasible ($g_{\max} < 0$), the design is near-binary (Figs. 6 and 8), the phase fractions evolve to their targets (Fig. 7), and the run stops at the committed-step cap without satisfying the terminal convergence criterion (Fig. 9). The reentrant-bimetal topology of the seed persists in the final design and supports both bending-induced negative thermal expansion and auxetic deformation:

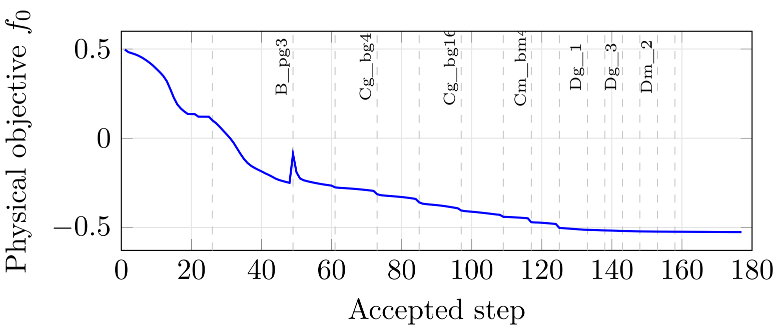

(a) physical objective f_0 versus accepted step

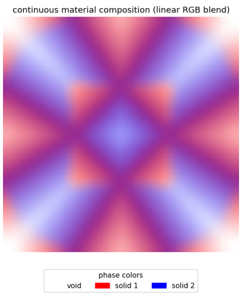

(b) initial design

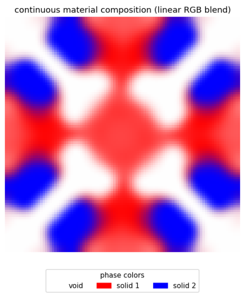

(c) stage A, step 20

(d) stage B, step 10

(e) stage B, step 40

(f) Cg4 final

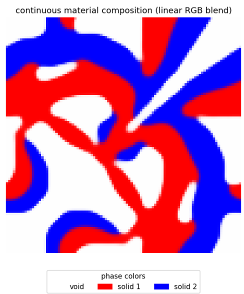

(g) final accepted design

Figure 2: Objective history and chronological structural evolution of the two-dimensional demonstration. (a) The physical objective decreases from +0.497 to −0.526 over 177 accepted records. Panels (b)–(g) follow the design from the filtered initial seed through stages A, B, and Cg to the final accepted design. Continuous composition: void = white, solid 1 = red, solid 2 = blue.

negative mean in-plane CTE is the optimized response, whereas the negative Poisson ratio emerges in post-processing; together they form an observed double-negative combination in this two-dimensional design.

3.3 Initialization sensitivity

To test the dependence of the framework on the starting field, six structurally distinct initial guesses — center circle, chiral hint, layered hint, random circles (two seeds), and reentrant hint — are run with identical case-1 settings through the full continuation schedule. The final designs are shown in Fig. 3, with the final mean in-plane CTE reported in each panel, and the quantitative outcome is summarized in Table 2. The full normal CTE components are reported in Appendix D (Table 11).

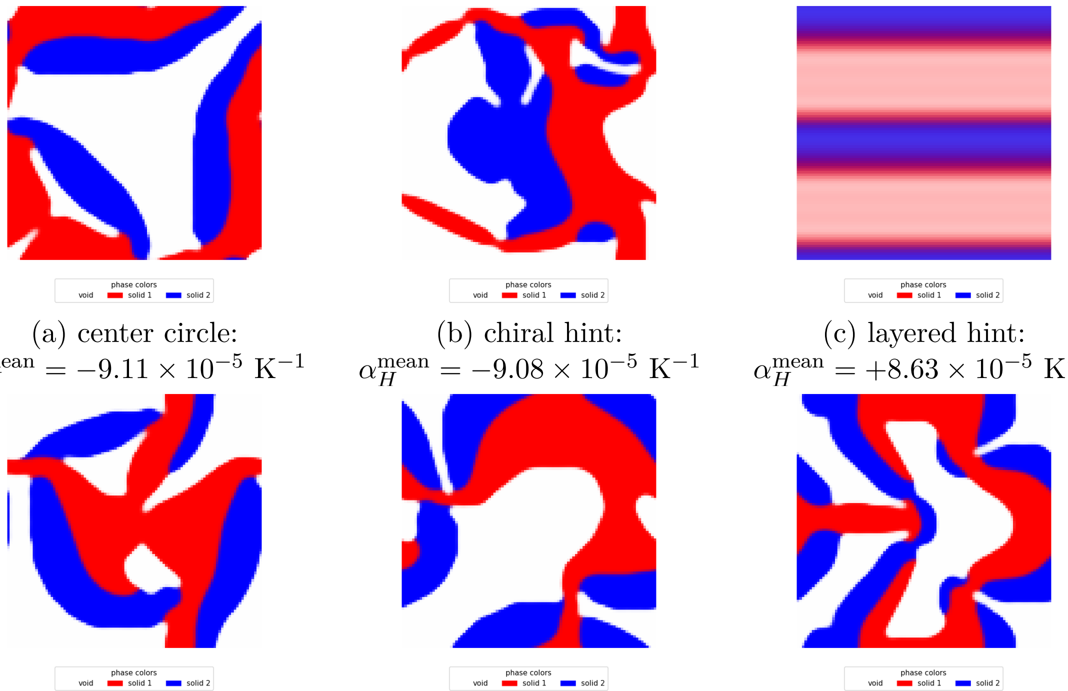

(a) center circle: $\alpha_H^{\text{mean}} = -9.11 \times 10^{-5}\ \text{K}^{-1}$

(b) chiral hint: $\alpha_H^{\text{mean}} = -9.08 \times 10^{-5}\ \text{K}^{-1}$

(c) layered hint: $\alpha_H^{\text{mean}} = +8.63 \times 10^{-5}\ \text{K}^{-1}$

(d) random circles s1: $\alpha_H^{\text{mean}} = -10.01 \times 10^{-5}\ \text{K}^{-1}$

(e) random circles s2: $\alpha_H^{\text{mean}} = -8.14 \times 10^{-5}\ \text{K}^{-1}$

(f) reentrant hint: $\alpha_H^{\text{mean}} = -9.46 \times 10^{-5}\ \text{K}^{-1}$

Figure 3: Final designs of the six screening initializations after the full continuation schedule. Each panel lists the final mean in-plane homogenized CTE; the layered hint (c) remains trapped in a layered topology with positive CTE. Continuous composition: void = white, solid 1 = red, solid 2 = blue.

Five of the six tested initializations reached designs with negative mean in-plane CTE under the same continuation schedule. The single failure, the layered hint, terminates early at stage B_pg3 with high grayness and positive CTE, trapped in a layered topology incompatible with the phase-volume targets. The objective trajectories of all six runs are shown in Fig. 15, and the per-guess audit table, including entry objectives, steps, and operator work, is given in Appendix D (Table 12).

3.4 Three-dimensional feasibility

The three-dimensional formulation is identical to the two-dimensional one, with the same isotropic source law, spectral operators, gauge penalty, and reference-medium preconditioner; the derivation is therefore not repeated. Two structured seeds at the $40 \times 40 \times 40$ production

Table 2: Initial-guess screening summary: final mean homogenized CTE ($\alpha_{\mathrm{ref}} = 2 \times 10^{-4}$ K^{-1}; units 10^{-5} K^{-1}), exit grayness, and termination. Final Poisson ratios were not recorded for the screening runs.

Initial guess	α_H^{mean}	G_g exit	G_m exit	Termination
Center circle	-9.11	0.033	0.052	Dm_3
Chiral hint	-9.08	0.051	0.053	Dm_3
Layered hint	$+8.63$	0.648	0.534	B_pg3
Random circles (s1)	-10.01	0.037	0.051	Dm_3
Random circles (s2)	-8.14	0.029	0.049	Dm_3
Reentrant hint	-9.46	0.044	0.070	Dm_3

resolution — a chiral helix hint and a center-sphere hint — are run through the full schedule. The final designs are shown in Fig. 4 and summarized in Table 3.

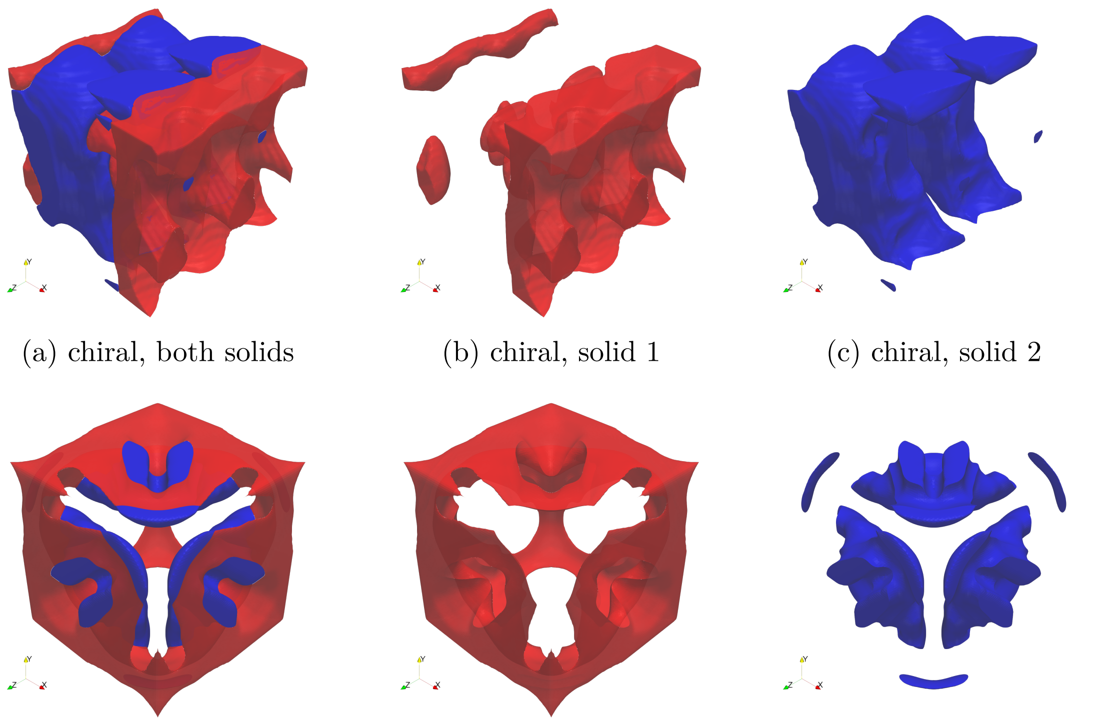

(a) chiral, both solids (b) chiral, solid 1 (c) chiral, solid 2

(d) center sphere, both solids (e) center sphere, solid 1 (f) center sphere, solid 2

Figure 4: Final designs of the two 40^3 runs: (a)–(c) chiral hint, (d)–(f) center sphere; (a),(d) both solids, (b),(e) solid 1 (red) only, (c),(f) solid 2 (blue) only, with void transparent.

Table 3: Summary of the two 40^3 runs.

	Chiral hint	Center sphere
Exit objective f_0	$+0.1178$	-0.09196
Final CTE tensor	anisotropic, principal values $(-6.87, 2.16, 8.24) \times 10^{-5}$ K^{-1}	nearly isotropic $\approx -8.99 \times 10^{-6}$ K^{-1}
Termination	budget-limited at Dm_3 cap	budget-limited at Dm_3 cap

The center-sphere run reaches a negative mean CTE, while the chiral run retains a positive mean CTE with only one negative principal direction; both terminate at the Dm_3 committed-step cap, budget-limited, and neither meets the terminal convergence criterion. Both are budget-limited feasibility evidence, not converged optima. The per-stage trajectory and run detail are given in Appendix E (Tables 13 and 14).

Directional-objective demonstration. To examine how objective directionality changes the three-dimensional morphology, a third 40^3 run replaces the mean-CTE objective with the single-axis objective $f_0 = \alpha_{11}^H/\alpha_{\text{ref}}$ ($\alpha_{\text{ref}} = 10^{-4}$ K^{-1}), keeping the center-sphere seed and all other settings unchanged; the design therefore optimizes the 11-direction CTE alone. The resulting design is shown in Fig. 5: solid 1 (red) forms a connected open-cell framework of thick walls and large openings, with solid 2 (blue) embedded as rounded pockets inside the openings and void as the through channels. The homogenized CTE response is strongly directional: the 11-direction reaches $\alpha_{11}^H = -1.08 \times 10^{-4}$ K^{-1} ($f_0 = -1.08$), while the two transverse components remain positive at $\alpha_{22}^H = \alpha_{33}^H = +1.77 \times 10^{-4}$ K^{-1}. The phase volumes $(0.499, 0.251, 0.250)$ lie within the target band of $(0.5, 0.25, 0.25)$ and all ten constraints are feasible ($g_{\max} < 0$). In contrast to the mean-CTE center-sphere case, the single-axis objective concentrates the negative response in the selected direction and produces a visibly different open-cell architecture. This case is included as a directional design demonstration rather than as evidence of negative mean CTE.

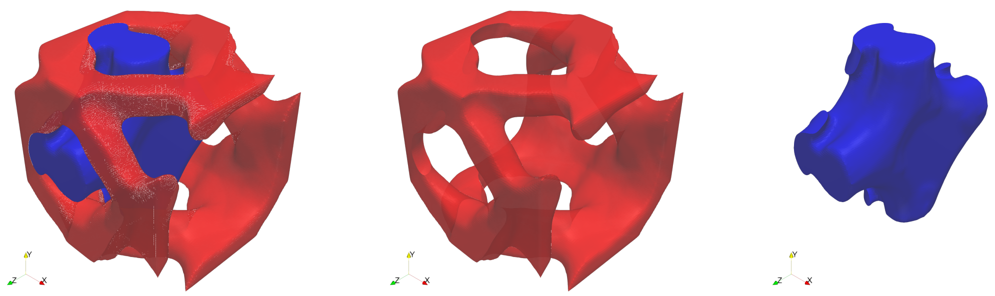

(a) both solids (b) solid 1 (c) solid 2

Figure 5: Design of the 40^3 center-sphere run under the single-axis α_{11} objective. (a) both solids, (b) solid 1 (red) only, (c) solid 2 (blue) only, with void transparent. The homogenized CTE components of the shown design are $\alpha_{11} = -1.08 \times 10^{-4}$, $\alpha_{22} = +1.77 \times 10^{-4}$, and $\alpha_{33} = +1.77 \times 10^{-4}$ K^{-1} (objective $f_0 = \alpha_{11}/\alpha_{\text{ref}} = -1.08$).

4 Discussion

Significance of differentiable FFT-CG for periodic multiphase design. The combination of a matrix-free FFT-CG homogenization solver with an implicit adjoint boundary makes the homogenized response of a periodic three-phase cell differentiable with respect to the design fields at $O(N \log N)$ cost per FFT-based operator application, with exact periodicity built into the Fourier basis. Because the reverse pass solves one adjoint linear system per load case instead of unrolling the CG history, the computational graph stays fixed as the design-dependent conditioning varies, enabling FFT-based homogenization to be embedded in a gradient-based topology-optimization loop.

Why CTE-driven design retains the auxetic mechanism. The reentrant-bimetal mechanism that produces bending-induced negative thermal expansion also induces lateral expansion under axial tension; the two target responses are carried by overlapping geometric features rather than by disjoint substructures. The final two-dimensional design of Sec. 3 illustrates this: the Poisson ratio was not part of the optimization objective, yet the CTE-optimized design exhibits a negative effective Poisson ratio, an observed outcome rather than an imposed target.

Practical tradeoffs between FFT and FEM. The choice between FFT-based and FEM-based homogenization for topology optimization involves a fundamental tradeoff. The advantages of FFT-based homogenization are decisive for large-scale periodic problems: no matrix assembly, natural periodicity, and $O(N \log N)$ scaling. Its disadvantages become acute inside topology optimization: the continuous design evolution through filtering, projection, and grayness penalties produces material distributions with smooth but high-contrast interfaces, degrading the spectral properties of the equilibrium operator and slowing CG convergence.

Limitations. (i) The independent finite-element reference of Sec. 3.1 is a three-dimensional synthetic voxel RVE, not the two-dimensional case-1 design; a design-level two-dimensional FEM cross-validation remains open. (ii) The void regularization at stiffness ratio 10^{-4} carries a bias of about 7% in the reported mean CTE relative to the saturated softer-void limit (Sec. 3.1). (iii) The mean-CTE three-dimensional runs of Sec. 3 are budget-limited feasibility evidence rather than converged optima. (iv) Manufacturing constraints — length-scale control, overhang restrictions, and multi-material printing constraints — are not yet included in the filter/projection chain. Regarding the observed robustness of the initialization screening, we state a hypothesis rather than a conclusion: the observed initialization behavior may result from the combined effects of spectral discretization, filtering, projection, and continuation; controlled FFT–FEM comparisons are required to isolate these contributions. The elastic moduli in the demonstrations are normalized rather than tied to a specific engineering material pair; the reported property values therefore demonstrate numerical design behavior, not a material-specific manufacturing prediction.

5 Conclusion

We presented a differentiable FFT-homogenization topology-optimization framework for periodic three-phase thermoelastic metamaterials. The framework combines a matrix-free FFT-CG primal solver with reverse-mode gradients through an implicit linear solve (wrapped adjoint), a hierarchical two-field design with periodic filter and tanh projection, and a safety-gated staged projection-continuation schedule with an explicit operator-work ledger. The solver is verified against homogeneous limits, symmetry and positive-definiteness checks, finite-difference gradient comparisons, an independent finite-element reference, and a frozen-design resolution sensitivity study.

The two-dimensional demonstration reached a negative mean in-plane homogenized CTE, and postprocessing of the same homogenized stiffness reveals a negative effective Poisson ratio.

The negative CTE is the optimized response, whereas the negative Poisson ratio is an additional observed outcome; together they constitute a double-negative combination in this two-dimensional design without implying a double-objective optimization. Five of the six screened initializations yield designs with negative mean in-plane CTE under the same continuation schedule, and a $40 \times 40 \times 40$ three-dimensional run reaches a negative mean CTE as budget-limited feasibility evidence.

References

[1] J. C. Michel, H. Moulinec, and P. Suquet, "A computational scheme for linear and non-linear composites with arbitrary phase contrast," *International Journal for Numerical Methods in Engineering*, vol. 52, pp. 139–160, 2001. DOI: 10.1002/nme.275.

[2] M. Schneider, "A review of nonlinear FFT-based computational homogenization methods," *Acta Mechanica*, vol. 232, pp. 2051–2100, 2021. DOI: 10.1007/s00707-021-02962-1.

[3] J. Vondřejc, J. Zeman, and I. Marek, “An FFT-based Galerkin method for homogenization of periodic media,” *Computers & Mathematics with Applications*, vol. 68, pp. 156–173, 2014. DOI: 10.1016/j.camwa.2014.05.014.

[4] O. Sigmund and S. Torquato, “Composites with extremal thermal expansion coefficients,” *Applied Physics Letters*, vol. 69, pp. 3203–3205, 1996. DOI: 10.1063/1.117961.

[5] O. Sigmund and S. Torquato, “Design of materials with extreme thermal expansion using a three-phase topology optimization method,” *Journal of the Mechanics and Physics of Solids*, vol. 45, pp. 1037–1067, 1997. DOI: 10.1016/S0022-5096(96)00114-7.

[6] L. Wu, B. Li, and J. Zhou, “Isotropic negative thermal expansion metamaterials,” *ACS Applied Materials & Interfaces*, vol. 8, pp. 17721–17727, 2016. DOI: 10.1021/acsami.6b05717.

[7] Y. Guo, H. Wang, W. Wang, C. Chen, and Y. Wang, “Topology optimization for metamaterials with negative thermal expansion coefficients using energy-based homogenization,” *Advances in Engineering Software*, vol. 189, 103794, 2024. DOI: 10.1016/j.advengsoft.2024.103794.

[8] I. Jödicke, R. J. Leute, T. Junge, and L. Pastewka, “Efficient topology optimization using compatibility projection in micromechanical homogenization,” *arXiv preprint arXiv:2107.04123*, 2021.

[9] Z. Chen, B. Wu, Y. M. Xie, X. Wu, and S. Zhou, “FFT-based inverse homogenization for cellular material design,” *International Journal of Mechanical Sciences*, vol. 231, 107572, 2022. DOI: 10.1016/j.ijmecsci.2022.107572.

[10] M. Matsui, H. Hoshiba, K. Nishiguchi, H. Ogura, and J. Kato, “Multiscale topology optimization applying FFT-based homogenization,” *International Journal for Numerical Methods in Engineering*, vol. 126, e70009, 2025. DOI: 10.1002/nme.70009.

[11] M. Pundir and D. S. Kammer, “Simplifying FFT-based methods for solid mechanics with automatic differentiation,” *Computer Methods in Applied Mechanics and Engineering*, vol. 435, 117572, 2025. DOI: 10.1016/j.cma.2024.117572.

[12] K. Svanberg, “The method of moving asymptotes — a new method for structural optimization,” *International Journal for Numerical Methods in Engineering*, vol. 24, pp. 359–373, 1987. DOI: 10.1002/nme.1620240207.

[13] K. Svanberg, “A class of globally convergent optimization methods based on conservative convex separable approximations,” *SIAM Journal on Optimization*, vol. 12, pp. 555–573, 2002. DOI: 10.1137/S1052623499362822.

[14] T.-C. Lim, “Negative-Poisson’s-ratio materials: auxetic solids,” *Annual Review of Materials Research*, vol. 47, pp. 63–93, 2017. DOI: 10.1146/annurev-matsci-070616-124118.

[15] C. Chu, A. Leichner, F. Wenz, and H. Andrä, “Exploring VAE-driven implicit parametric unit cells for multiscale topology optimization,” *Materials & Design*, vol. 244, 113087, 2024. DOI: 10.1016/j.matdes.2024.113087.

[16] X. Li, L. Gao, W. Zhou, Y. Wang, and Y. Lu, “Novel 2D metamaterials with negative Poisson’s ratio and negative thermal expansion,” *Extreme Mechanics Letters*, vol. 30, 100498, 2019. DOI: 10.1016/j.eml.2019.100498.

[17] L. Ai and X.-L. Gao, “Metamaterials with negative Poisson’s ratio and non-positive thermal expansion,” *Composite Structures*, vol. 162, pp. 70–84, 2017. DOI: 10.1016/j.compstruct.2016.11.056.

[18] D. Akamatsu, Y. Noguchi, K. Matsushima, Y. Sato, J. Yanagimoto, and T. Yamada, "Two-phase topology optimization for metamaterials with negative Poisson's ratio," *Composite Structures*, vol. 311, 116800, 2023. DOI: 10.1016/j.compstruct.2023.116800.

[19] R. Lakes, "Foam structures with a negative Poisson's ratio," *Science*, vol. 235, pp. 1038–1040, 1987. DOI: 10.1126/science.235.4792.1038.

[20] R. Lakes, "Cellular solid structures with unbounded thermal expansion," *Journal of Materials Science Letters*, vol. 15, pp. 475–477, 1996. DOI: 10.1007/BF00275406.

[21] Z. Han and K. Wei, "Multi-material topology optimization and additive manufacturing for metamaterials incorporating double negative indexes of Poisson's ratio and thermal expansion," *Additive Manufacturing*, vol. 54, 102742, 2022. DOI: 10.1016/j.addma.2022.102742.

[22] C. Almeida, P. Coelho, J. M. Guedes, A. L. Custódio, F. Conde, and A. Velhinho, "Multiobjective topology optimization of metamaterials to analyse the conflicting nature of negative Poisson's ratio and negative thermal expansion," *Structural and Multidisciplinary Optimization*, vol. 68, 202, 2025. DOI: 10.1007/s00158-025-04131-7.

[23] M. W. Cho, K. Ko, M. Mohammadhosseinzadeh, J. H. Kim, D. Y. Park, D. S. Shin, and S. M. Park, "Inverse design of Bézier curve-based mechanical metamaterials with programmable negative thermal expansion and negative Poisson's ratio via a data augmented deep autoencoder," *Materials Horizons*, vol. 11, pp. 2615–2627, 2024. DOI: 10.1039/d4mh00302k.

[24] T. Xue, S. Liao, Z. Gan, C. Park, X. Xie, W. K. Liu, and J. Cao, "JAX-FEM: A differentiable GPU-accelerated 3D finite element solver for automatic inverse design and mechanistic data science," *Computer Physics Communications*, vol. 291, 108802, 2023. DOI: 10.1016/j.cpc.2023.108802.

A Adjoint sensitivity analysis

This appendix presents the abstract adjoint sensitivity formulation for the differentiable FFT-CG framework. The derivation follows the implicit function theorem applied to the converged equilibrium equation.

A.1 Abstract formulation

Let $\boldsymbol{x} \in \mathbb{R}^{2N}$ be the design vector and let $\widetilde{\boldsymbol{u}}(\boldsymbol{x}) \in \mathbb{R}^{2N}$ be the converged displacement fluctuation satisfying the discrete equilibrium equation

$$\boldsymbol{R}(\widetilde{\boldsymbol{u}}, \boldsymbol{x}) = \mathcal{A}(\boldsymbol{x})\,\widetilde{\boldsymbol{u}} - \boldsymbol{b}(\boldsymbol{x}) = \boldsymbol{0}, \tag{25}$$

where $\mathcal{A}(\boldsymbol{x})$ is the equilibrium operator and $\boldsymbol{b}(\boldsymbol{x})$ is the right-hand side, both depending on the design through the material interpolation.

The objective function $f(\boldsymbol{x})$ depends on the design directly and indirectly through the homogenized properties, which are functions of the converged displacement:

$$f(\boldsymbol{x}) = \tilde{f}(\boldsymbol{x}, \widetilde{\boldsymbol{u}}(\boldsymbol{x})). \tag{26}$$

The total derivative is

$$\frac{df}{d\boldsymbol{x}} = \frac{\partial \tilde{f}}{\partial \boldsymbol{x}} + \frac{\partial \tilde{f}}{\partial \widetilde{\boldsymbol{u}}} \frac{d\widetilde{\boldsymbol{u}}}{d\boldsymbol{x}}. \tag{27}$$

A.2 Implicit differentiation

Differentiating the equilibrium equation with respect to $\boldsymbol{x}$ gives

$$\frac{\partial \boldsymbol{R}}{\partial \boldsymbol{x}} + \frac{\partial \boldsymbol{R}}{\partial \tilde{\boldsymbol{u}}}\frac{d\tilde{\boldsymbol{u}}}{d\boldsymbol{x}} = \boldsymbol{0}, \tag{28}$$

hence

$$\frac{d\tilde{\boldsymbol{u}}}{d\boldsymbol{x}} = -\left(\frac{\partial \boldsymbol{R}}{\partial \tilde{\boldsymbol{u}}}\right)^{-1}\frac{\partial \boldsymbol{R}}{\partial \boldsymbol{x}} = -\mathcal{A}^{-1}\frac{\partial \boldsymbol{R}}{\partial \boldsymbol{x}}, \tag{29}$$

where we used $\partial \boldsymbol{R}/\partial \tilde{\boldsymbol{u}} = \mathcal{A}$.

Substitution into the total derivative gives

$$\frac{df}{d\boldsymbol{x}} = \frac{\partial \tilde{f}}{\partial \boldsymbol{x}} - \frac{\partial \tilde{f}}{\partial \tilde{\boldsymbol{u}}}\mathcal{A}^{-1}\frac{\partial \boldsymbol{R}}{\partial \boldsymbol{x}}. \tag{30}$$

Define the adjoint variable $\boldsymbol{\lambda}$ as the solution of

$$\mathcal{A}^{\mathsf{T}}\boldsymbol{\lambda} = \left(\frac{\partial \tilde{f}}{\partial \tilde{\boldsymbol{u}}}\right)^{\mathsf{T}}. \tag{31}$$

Since $\mathcal{A}$ is self-adjoint (symmetric), $\mathcal{A}^{\mathsf{T}} = \mathcal{A}$, and

$$\frac{df}{d\boldsymbol{x}} = \frac{\partial \tilde{f}}{\partial \boldsymbol{x}} - \boldsymbol{\lambda}^{\mathsf{T}}\frac{\partial \boldsymbol{R}}{\partial \boldsymbol{x}}. \tag{32}$$

A.3 Wrapped adjoint in JAX

In the JAX implementation, the primal solve is wrapped in `lax.custom_linear_solve`, which automatically:

1. Solves the primal system $\mathcal{A}\tilde{\boldsymbol{u}} = \boldsymbol{b}$ using PCG with the spectral preconditioner;
2. During reverse-mode differentiation, solves the adjoint system $\mathcal{A}\boldsymbol{\lambda} = \boldsymbol{c}$ with the same operator, preconditioner, and stopping criterion;
3. Returns the gradient $\nabla_{\boldsymbol{x}} f = \partial \tilde{f}/\partial \boldsymbol{x} - \boldsymbol{\lambda}^{\mathsf{T}}(\partial \boldsymbol{R}/\partial \boldsymbol{x})$ without unrolling the CG iteration.

The key advantage is that the computational graph is fixed: the reverse pass solves a single linear system regardless of how many PCG iterations the primal solve required. This is essential for FFT-based topology optimization, where the iteration count varies with the design-dependent conditioning.

A.4 Explicit gradient components

The gradient of the objective f_0 with respect to the design fields decomposes as

$$\frac{df_0}{d\boldsymbol{x}} = \frac{\partial f_0}{\partial \boldsymbol{x}} - \sum_{J=1}^{4}\boldsymbol{\lambda}_J^{\mathsf{T}}\frac{\partial \boldsymbol{R}_J}{\partial \boldsymbol{x}}, \tag{33}$$

where J indexes the four load cases (three elastic, one thermal), and $\boldsymbol{\lambda}_J$ is the adjoint solution for load case J.

The partial derivative $\partial \boldsymbol{R}_J/\partial \boldsymbol{x}$ involves the derivative of the material interpolation:

$$\frac{\partial \boldsymbol{R}_J}{\partial \boldsymbol{x}} = \frac{\partial \mathcal{A}}{\partial \boldsymbol{x}}\tilde{\boldsymbol{u}}_J - \frac{\partial \boldsymbol{b}_J}{\partial \boldsymbol{x}}, \tag{34}$$

where

$$\frac{\partial \mathcal{A}}{\partial \boldsymbol{x}} = -\,\mathrm{div}_h \left[\frac{\partial \boldsymbol{C}}{\partial \boldsymbol{x}} \boldsymbol{B}_h \tilde{\boldsymbol{u}}_J \right], \qquad \frac{\partial \boldsymbol{b}_J}{\partial \boldsymbol{x}} = \mathrm{div}_h \left[\frac{\partial \boldsymbol{C}}{\partial \boldsymbol{x}} \bar{\boldsymbol{E}}_J - \frac{\partial \boldsymbol{\beta}}{\partial \boldsymbol{x}} \Delta T_J \right]. \tag{35}$$

The material derivatives follow from the chain rule through the filter, projection, and phase-fraction mapping:

$$\frac{\partial \boldsymbol{C}}{\partial \boldsymbol{x}} = \sum_p \frac{\partial \boldsymbol{C}}{\partial \boldsymbol{\rho}_p} \frac{\partial \boldsymbol{\rho}_p}{\partial \boldsymbol{x}}, \qquad \frac{\partial \boldsymbol{\beta}}{\partial \boldsymbol{x}} = \sum_p \frac{\partial \boldsymbol{\beta}}{\partial \boldsymbol{\rho}_p} \frac{\partial \boldsymbol{\rho}_p}{\partial \boldsymbol{x}}. \tag{36}$$

A.5 Composition of the forward chain

The complete forward map is the composition of the filter, projection, phase-fraction, SIMP, equilibrium, and homogenization maps of Section 2,

$$\begin{aligned} &\boldsymbol{x} \xrightarrow{\mathcal{F}_{\mathrm{cone}}} \tilde{\boldsymbol{z}} \xrightarrow{H_{\beta,\theta}} (s, m) \xrightarrow{\text{simplex}} \boldsymbol{\rho} \xrightarrow{\text{SIMP}} (\boldsymbol{C}, \boldsymbol{\beta}) \xrightarrow{\mathcal{A}^{-1}} \tilde{\boldsymbol{u}}, \\ &\xrightarrow{\text{homogenization}} (\boldsymbol{C}_H, \boldsymbol{\alpha}_H) \xrightarrow{\text{objective/constraints}} (f, \boldsymbol{g}), \end{aligned} \tag{37}$$

and reverse-mode differentiation propagates the objective cotangent through the same chain in reverse order. Each factor of the chain (37) has an elementary derivative: the constant convolution $\mathcal{F}_{\mathrm{cone}}$, the elementwise projection derivative $\mathrm{diag}(H'_{\beta,\theta})$, the linear simplex map, the SIMP power law with

$$\frac{\partial \boldsymbol{C}}{\partial s} = p\, s^{p-1} \Big[\sum_{k=1}^{2} q_k\, \boldsymbol{C}_k - \boldsymbol{C}_0 \Big], \tag{38}$$

and the volume averages in Eq. (17); the residual derivative $\partial \boldsymbol{R}_J / \partial \boldsymbol{x}$ is given explicitly in the previous subsection.

B Discrete operators and numerical verification

This appendix collects the full discrete-operator definitions and the complete numerical data behind the verification summary of Section 3.1.

B.1 Discrete operators

On the periodic grid we use a forward difference for the gradient and the adjoint (backward) difference for the divergence, so that the discrete integration-by-parts identity

$$(\boldsymbol{B}_h\, \tilde{\boldsymbol{u}}, \boldsymbol{\sigma})_h = -\,(\tilde{\boldsymbol{u}}, \mathrm{div}_h\, \boldsymbol{\sigma})_h \tag{39}$$

holds exactly. In 2D engineering-shear Voigt order $(11, 22, 12)$:

$$\boldsymbol{B}_h\, \tilde{\boldsymbol{u}} := \left[D_x^+ \tilde{\boldsymbol{u}}_1,\ D_y^+ \tilde{\boldsymbol{u}}_2,\ D_y^+ \tilde{\boldsymbol{u}}_1 + D_x^+ \tilde{\boldsymbol{u}}_2 \right]^{\mathsf{T}}, \tag{40}$$

$$\mathrm{div}_h\, \boldsymbol{\sigma} = \left[D_x^- \sigma_{11} + D_y^- \sigma_{12},\ D_x^- \sigma_{12} + D_y^- \sigma_{22} \right]^{\mathsf{T}}. \tag{41}$$

The Fourier multipliers satisfy $\delta_i^+(k) = (e^{2\pi \mathrm{i} k_i / n_i} - 1)/h_i$ and $\delta_i^-(k) = -\overline{\delta_i^+(k)}$, so the operators are exact spectral adjoints.

B.2 The unresolved zero mode

The forward difference has exactly one unresolved mode: the spatial mean translation at $\boldsymbol{k} = (0, 0)$. On the subspace of zero-mean fluctuations the equilibrium operator is invertible, and the gauge penalty $\zeta\, \mathcal{P}_0\, \tilde{\boldsymbol{u}}$ of Eq. (20) (Sec. 2.3) renders the operator invertible and self-adjoint on the full space while leaving the zero-mean physics unchanged: its restriction to the zero-mean subspace is the physical equilibrium operator, and the gauge term acts only on the rigid-body mode it is designed to pin.

B.3 Numerical verification

Homogeneous limit. The homogeneous-limit test evaluates a single homogeneous material on a $6\times6\times6$ cell and compares the extracted stiffness, thermal-expansion tensor, Young moduli, and Poisson ratios against the exact analytical values. The elementwise maximum relative errors over nonzero-expected entries are 2.635×10^{-15} (stiffness), 1.626×10^{-15} (thermal-expansion tensor), 3.283×10^{-14} (Young moduli), and 2.126×10^{-15} (Poisson ratios), giving a maximum over all quantities of 3.28×10^{-14}. The same protocol shows the extracted properties to be insensitive to the magnitude of the load amplitudes a and t (load-amplitude invariance).

Gradient finite-difference audit. Reverse-mode gradients are audited against centered finite differences along random directions for both the objective and the constraint gradients. Table 4 reports the worst-case relative error over the four audited (problem, modulus-floor) combinations, as a function of the step size; the error decreases with the step size as expected for a centered difference on a smooth objective.

Table 4: Gradient finite-difference audit: worst-case relative error over the audited directions (objective and constraint gradients) versus centered finite-difference step size.

Step size h	Max relative error
3×10^{-3}	9.09×10^{-6}
1×10^{-3}	1.01×10^{-6}
3×10^{-4}	9.07×10^{-8}

Independent finite-element reference. The homogenized tensors of a three-phase sphere-plus-void voxel RVE are computed with the JAX-FEM finite-element solver [24] (direct sparse solve) and compared with the FFT-CG solver of this work on voxel grids 8^3, 12^3, 16^3, and 20^3. Table 5 reports the relative differences of the effective stiffness $\boldsymbol{C}_H$, the effective compliance $\boldsymbol{S}_H$, the diagonal entries of the effective thermal-expansion tensor $\boldsymbol{\alpha}^H$, and the effective Poisson ratio; the differences generally decrease with resolution and reach 1.4%, 0.5%, 0.2%, and 0.08%, respectively, at 20^3.

Table 5: Relative differences between FFT-CG and the JAX-FEM reference for the three-phase sphere-plus-void voxel RVE, by grid size.

Grid	8^3	12^3	16^3	20^3
$\boldsymbol{C}_H$ relative difference	3.4%	2.8%	2.1%	1.4%
$\boldsymbol{S}_H$ relative difference	1.1%	0.9%	0.8%	0.5%
$\boldsymbol{\alpha}^H$ diagonal relative difference	0.6%	0.3%	0.2%	0.2%
ν relative difference	0.1%	0.2%	0.1%	0.08%

Evaluator validity flags. The frozen case-1 evaluation at the production setting (100×100, void stiffness ratio 10^{-4}) passes all validity gates: the raw and filtered fields are finite and in bounds; the phase fractions are finite, in bounds, and on the simplex; the binarity metrics are finite; all load cases are finite and converged; the raw stiffness is symmetric (`raw_symmetry_valid`); the Kelvin-form stiffness is positive definite (`kelvin_spd`); the compliance and thermal solves are valid; the property bounds are valid; and all derived properties are finite (`properties_finite`). All 15 flags are true and `all_valid` holds. The per-load-case final equilibrium relative residuals are 9.872×10^{-9}, 9.872×10^{-9}, 9.722×10^{-9} (mechanical) and 9.539×10^{-9} (thermal), with PCG iteration counts 802, 802, 793, and 782 respectively.

C Two-dimensional CTE-driven design: continuation schedule and run evidence

This appendix collects run-specific evidence for the two-dimensional demonstration case (Sec. 3). This supporting but secondary material comprises the balanced stage schedule, its per-family profile constants, the full parameter table, the grayness, phase-volume, near-binary, and operator-work trajectory diagnostics, and the frozen-design resolution and void-ratio sensitivity study.

The phase materials of the demonstration case are given in Table 6.

Table 6: Phase materials of the two-dimensional demonstration case. E is in one consistent stress unit, α in K^{-1}. The void stiffness ratio is $E_0/E_{\min}$; the verification reference uses 10^{-6}, the demonstration run uses 10^{-4}.

Phase	E	ν	α (1/K)
0 = void (regularized)	10^{-4}	0.0	0.0
1 = solid 1	1.0	0.3	10^{-5}
2 = solid 2	1.0	0.3	2×10^{-4}

The continuation advances the SIMP exponent p and the projection sharpnesses β_g, β_m through discrete levels, one field at a time, on a staged schedule. The sixteen-stage demonstration ladder is $A(p{=}1) \to B\_p2(p{=}2) \to B\_pg3(p{=}3) \to Cg\_bg2/4/8/16 \to Cm\_bm2/4/8 \to Dg\_1..3 \to Dm\_1..3$, with p raised $1 \to 3$, β_g raised $1 \to 16$, and β_m raised $1 \to 8$ through the Cg/Cm/D stages (Table 7). The balanced profile assigns each stage family a committed-step cap, a stage stability tolerance, and a diminishing-return (DR) factor (Table 8); the generic stage loop is Alg. 1.

Table 7: Balanced stage profile of the demonstration run.

Stage	p	β_g	β_m	λ_g	λ_m
A_physical	1	1	1	0	0
B_pg2	2	1	1	0	0
B_pg3	3	1	1	0	0
Cg_bg2	3	2	1	0	0
Cg_bg4	3	4	1	0	0
Cg_bg8	3	8	1	0	0
Cg_bg16	3	16	1	0	0
Cm_bm2	3	16	2	0	0
Cm_bm4	3	16	4	0	0
Cm_bm8	3	16	8	0	0
Dg_1	3	16	8	0.0826	0
Dg_2	3	16	8	0.2064	0
Dg_3	3	16	8	0.4128	0
Dm_1	3	16	8	0.4128	0.0637
Dm_2	3	16	8	0.4128	0.1592
Dm_3	3	16	8	0.4128	0.3184

Let k index the committed records of the run, with physical objective $f_{0,k}$ and cumulative operator-application count W_k, and let

$$\eta(a, b) = \frac{f_{0,a} - f_{0,b}}{W_b - W_a}, \qquad b > a, \tag{42}$$

be the per-operator-application f_0 gain over committed records $a, \ldots, b$. The safety window is the last five committed records, all load-case-valid with $\max_i g_i \leq 10^{-6}$ and phase volumes

Table 8: Per-family profile constants of the balanced schedule: committed-step caps, stage stability tolerances, and diminishing-return (DR) factors.

Stage family	Committed cap	Stability tolerance	DR factor
A	25	5×10^{-3}	—
B	25	1×10^{-2}	0.25
Cg	12	1×10^{-2}	0.25
Cm	8	5×10^{-3}	0.25
Dg	8	1×10^{-2}	0.25
Dm	20	1×10^{-2}	0.25

inside the target bands. A nonterminal stage exits by the first applicable condition among: (i) *readiness*: the last five committed records satisfy $\mathrm{ptp}(f_0)/|\mathrm{mean}(f_0)| \le \tau_s$ together with the safety window, where $\tau_s = 5 \times 10^{-3}$ for the A and Cm families and 10^{-2} for the B, Cg, Dg, Dm families (Table 8); (ii) *diminishing return*: over the trailing terminal_window = 20 committed records, each of the two most recent complete five-record windows attains $\eta \le \gamma\, \eta_{\mathrm{ref}}$, where η_{ref} is the gain of the oldest five-record window in the trailing window and $\gamma = 0.25$ is the family DR factor (Table 8), together with the safety window; (iii) *cap*: the family committed-step cap of Table 8 is reached and the safety window passes. A transition to the next stage then requires the frozen-design probe at the next level to pass its hard gates: both snapshots load-case-valid, phase-volume drift $\max_p |V_p' - V_p| \le 0.02$, operator-work ratio $\max(W'/W, P'/P) \le 3$, and the grayness of the probed field does not increase when a projection sharpness changes.

The final stage is additionally checked against the terminal convergence criterion over the trailing terminal_window = 20 committed records: all records valid, $\max_i g_i \le 10^{-6}$, objective relative range $\mathrm{ptp}(f_0)/|\mathrm{mean}(f_0)| \le 10^{-4}$, phase changes $\max_p |\Delta V_p| \le 10^{-3}$ and $\mathrm{mean}_p |\Delta V_p| \le 10^{-4}$, and both grayness relative ranges $\le 10^{-3}$ (Table 9). All work is counted on an explicit operator-work ledger; the balanced run sets no operator-work budget. The grayness-penalty weights of the D stages are resolved at family entry as $\lambda = r\, |f_0| / \max(G, 1/N)$ with contribution ratios $r \in \{0.01, 0.025, 0.05\}$, where G is the family grayness at entry and $N = n_1 n_2$ the pixel count; this yields the Dg weights $\lambda_g \in \{0.0826, 0.2064, 0.4128\}$ and Dm weights $\lambda_m \in \{0.0637, 0.1592, 0.3184\}$ of Table 7.

The full parameter table of the demonstration run is given in Table 9.

C.1 Resolution convergence and void-ratio sensitivity

The frozen case-1 design is re-evaluated with the production pipeline (cone filter radius 0.05, projection $\beta_g = 16$, $\beta_m = 8$, SIMP $p = 3$, FFT-CG rtol $= 10^{-8}$, $a = 10^{-4}$, $t = 1.0$) at 40×40, 80×80, and the native 100×100 grids — the two coarser grids are nearest-neighbor decimations of the frozen raw fields — and with void stiffness ratios 10^{-4}, 10^{-6}, and 10^{-8} (Table 10). At the production setting (100×100, 10^{-4}) the evaluation reproduces the reported $f_0 = -0.526020$ to 7×10^{-15} relative, validating the protocol. All converged evaluations pass the validity gates.

At a void stiffness ratio of 10^{-4}, the mean CTE is $-3.55 \times 10^{-5}\ \mathrm{K}^{-1}$ on the 40^2 grid, -10.71×10^{-5} on the 80^2 grid, and -10.52×10^{-5} on the native 100^2 grid (Fig. 10): the coarse 40^2 representation loses the reentrant-bimetal mechanism (the thin beam skeleton is under-resolved by the two-pixel filter kernel), while 80^2 and 100^2 agree within 1.8%. The 100×100 grid of the demonstration is therefore adequate, with a residual resolution uncertainty of order 2%. The effective Poisson ratio converges to $\nu_{12} = \nu_{21} \approx -0.41$ (Fig. 11); the diagonal symmetry of the design makes the two values equal exactly. At 80^2 the decimated design's $E_1 = 0.048$ dips slightly below the 0.05 modulus floor satisfied at 100^2, a resolution artifact of the constraint evaluation.

Softer void systematically shifts the prediction: at 100^2, f_0 changes from -0.5260 at void 10^{-4} to -0.5631 at 10^{-6} (about 7% more negative per 100-fold void softening), with ν_{12} moving from

Table 9: Setup of the 100×100 two-dimensional demonstration case.

Setting	Value
Grid	100×100
Initialization	Reentrant bimetal seed
Filter	Periodic FFT cone, radius 0.05
Void stiffness ratio	10^{-4}
Materials	$E_1 = E_2 = 1.0$, $\nu = 0.3$, $\alpha = (10^{-5}, 2 \times 10^{-4})$ K^{-1}
Load amplitudes	$a = 10^{-4}$, $t = 1.0$
Reference scales	$E_{\mathrm{ref}} = 1$, $G_{\mathrm{ref}} = 1/[2(1 + 0.3)]$, $\alpha_{\mathrm{ref}} = 2 \times 10^{-4}$ K^{-1}
Solver	Sequential PCG, rtol $= 10^{-8}$, maxiter $= 2000$ (reference rtol $= 10^{-10}$, maxiter $= 800$)
Phase-volume targets	$(0.40, 0.30, 0.30)$, $\epsilon_v = 0.02$
Modulus floors	$e_{\mathrm{floor}} = 0.05$, $g_{\mathrm{floor}} = 0.01$
Grayness regularization	$\epsilon_m = 0$
Near-binary threshold	$\delta = 0.05$
Projection continuation	Staged levels up to $(\beta_g, \beta_m) = (16, 8)$; threshold $\theta = 0.5$
GCMMA	10 inner resolves, feasibility tolerance 10^{-6}, attempt budget 12,000
Probe gates	load-case validity (both snapshots); volume drift ≤ 0.02; work ratio ≤ 3; grayness non-worsening under β changes
Termination	window 20; objective range $\leq 10^{-4}$, phase $L_\infty \leq 10^{-3}$, phase mean $\leq 10^{-4}$, grayness range $\leq 10^{-3}$
Exit	Balanced stage profile: caps A=25/B=25/Cg=12/Cm=8/Dg=8/Dm=20, DR factors 0.25
JAX	CPU, x64 (Python 3.10.18, JAX 0.6.2)

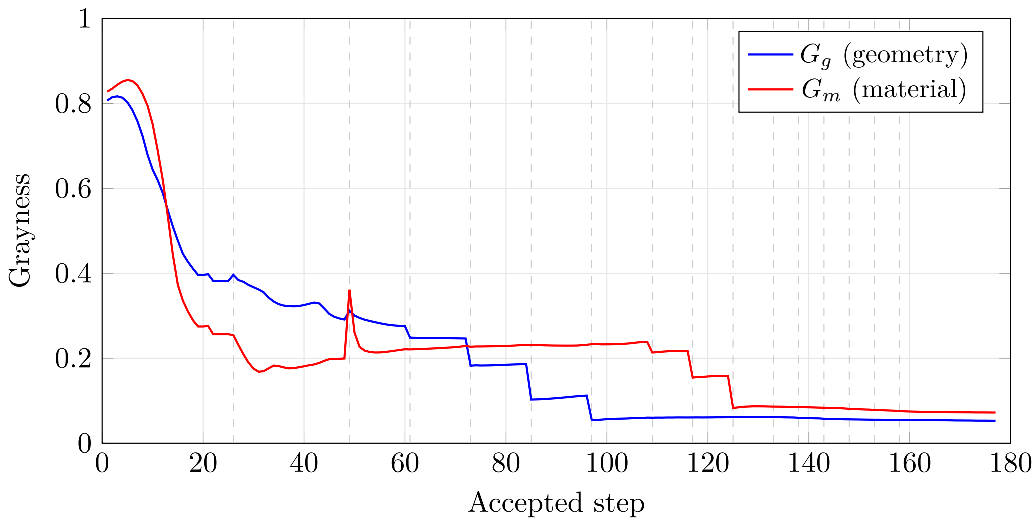

Figure 6: Geometry grayness G_g and material grayness G_m versus accepted step. Both metrics decrease monotonically from $0.81, 0.83$ to $0.05, 0.07$, with geometry grayness leading during early stages (stage A and the Cg continuation, $\beta_g : 2 \to 16$) and material grayness catching up during the Cm continuation ($\beta_m : 2 \to 8$).

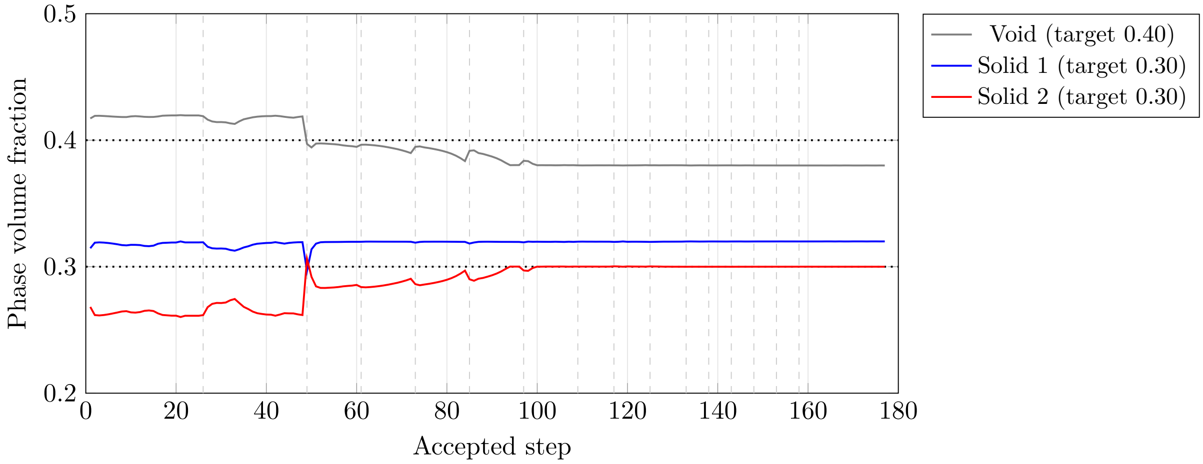

Figure 7: Phase volume fractions versus accepted step. Dotted horizontal lines mark the targets $(0.40, 0.30, 0.30)$; all three phases converge within $\epsilon_v = 0.02$ by stage Cm_bm2 (step 109) and remain stable at $(0.380, 0.320, 0.300)$ thereafter.

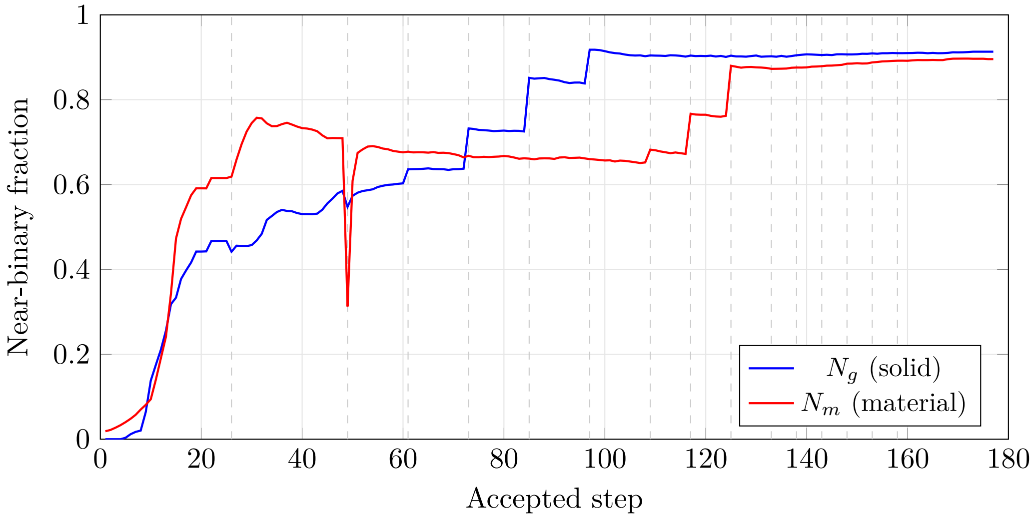

Figure 8: Near-binary fractions N_g (solid) and N_m (material) versus accepted step. The design evolves from fully gray ($N_g = 0.00$, $N_m = 0.02$) to $N_g = 0.913$ and $N_m = 0.896$ at the final step, with the solid fraction increasing during stage A and the material fraction increasing during the Cm continuation.

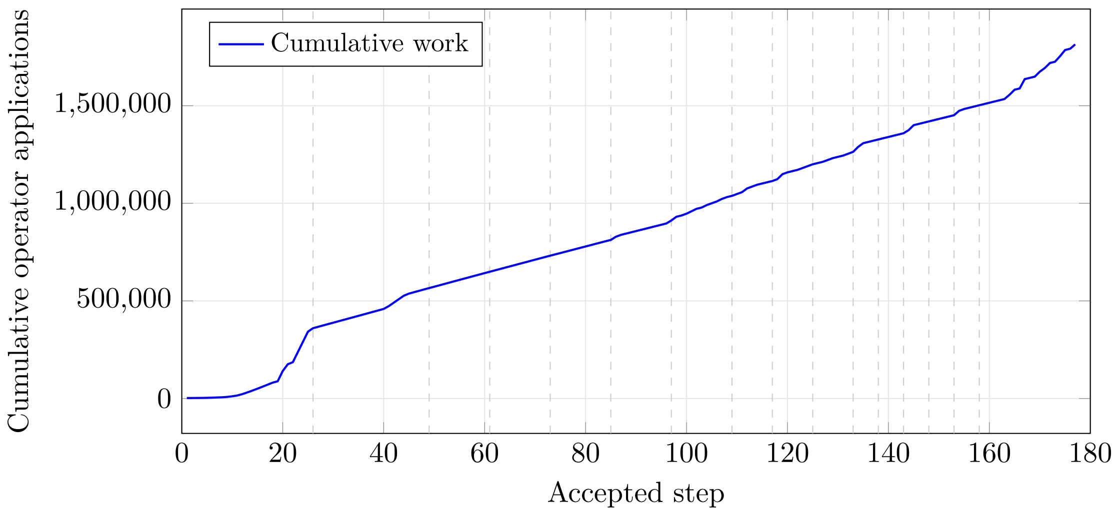

Figure 9: Cumulative operator applications versus accepted step. The run totals 1,814,061 applications. Stage A (340k, 19%) and the final stage Dm_3 (312k, 17%) account for the largest shares because their committed-step caps are 25 and 20, respectively; intermediate stages consume 25k–120k each.

-0.4107 to -0.4225; the trend is monotone, similar at 40^2 and 80^2, and saturates below 10^{-6}: the $10^{-6} \to 10^{-7}$ change in α_H^{mean} is below 0.1% at 100^2 (Fig. 12, Fig. 13). The demonstration's void choice 10^{-4} therefore introduces a void-regularization bias of order 7% in f_0 relative to the softer-void limit. The effective Poisson ratio of the final design is $\nu_{12} = \nu_{21} \approx -0.41$ at the production setting — the value obtained here is negative as expected from the reentrant mechanism.

The void-softening limit is governed by solver convergence: at void 10^{-8}, the FFT-CG does not reach rtol $= 10^{-8}$ within 20,000 iterations on the 80^2 and 100^2 grids (final relative residuals 1.6×10^{-7} and 2.5×10^{-7}; Table 10), while operator work grows about 2.6-fold per decade of void softening at 100^2 (3,183 applications at 10^{-4}, 8,380 at 10^{-5}, 21,699 at 10^{-6}, and 53,725 at 10^{-7}; Fig. 14). The production iteration budget (2,000) supports void ratios down to about 10^{-6} at 40^2 but only 10^{-4} at 80^2 and 100^2; the demonstration's 10^{-4} choice lies at the practical conditioning limit of the production solver.

Table 10: FFT-only resolution and void-ratio study of the frozen case-1 design (production pipeline; $\alpha_H^{\mathrm{mean}} = f_0\,\alpha_{\mathrm{ref}}$ with $\alpha_{\mathrm{ref}} = 2 \times 10^{-4}$ K^{-1}; "conv." = converged to rtol $= 10^{-8}$ within the stated iteration budget). Unconverged evaluations are reported with their final relative residuals; no derived quantities are claimed for them.

Grid	Void ratio	f_0	α_H^{mean} (10^{-5} K^{-1})	ν_{12}	E_1	G_{12}	conv.
40^2	10^{-4}	-0.1776	-3.553	-0.330	0.054	0.028	yes
40^2	10^{-6}	-0.1899	-3.798	-0.345	0.053	0.027	yes
40^2	10^{-8}	-0.1899	-3.799	-0.345	0.053	0.027	yes
80^2	10^{-4}	-0.5353	-10.707	-0.417	0.048	0.029	yes
80^2	10^{-6}	-0.5747	-11.494	-0.429	0.047	0.029	yes
80^2	10^{-8}	—	—	—	—	—	no (1.6×10^{-7})
100^2	10^{-4}	-0.526020	-10.520	-0.411	0.050	0.030	yes
100^2	10^{-6}	-0.5631	-11.262	-0.422	0.049	0.030	yes
100^2	10^{-8}	—	—	—	—	—	no (2.5×10^{-7})

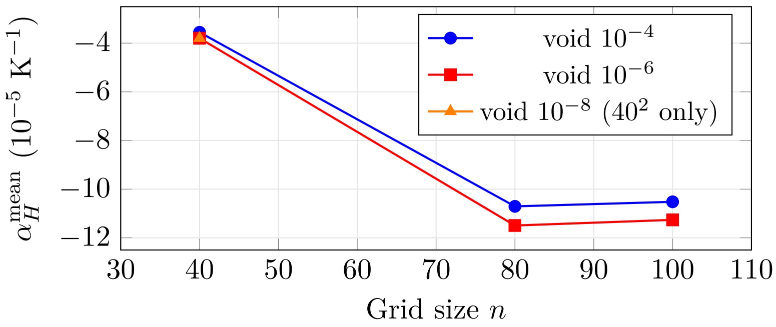

Figure 10: Mean homogenized CTE of the frozen case-1 design versus grid resolution for void stiffness ratios 10^{-4}, 10^{-6}, and 10^{-8} (converged evaluations only). The 40^2 grid under-resolves the reentrant-bimetal mechanism; 80^2 and 100^2 agree within 1.8%. The 10^{-8} evaluation converges only at 40^2.

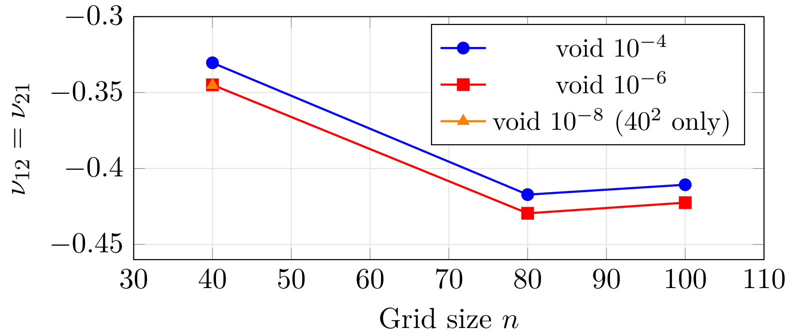

Figure 11: Effective Poisson ratio $\nu_{12} = \nu_{21}$ of the frozen case-1 design versus grid resolution (converged evaluations only). The first measured value of the case-1 design, $\nu_{12} \approx -0.41$ at the production setting, is negative as expected from the reentrant mechanism and shifts by about 3% when the void is softened by two decades.

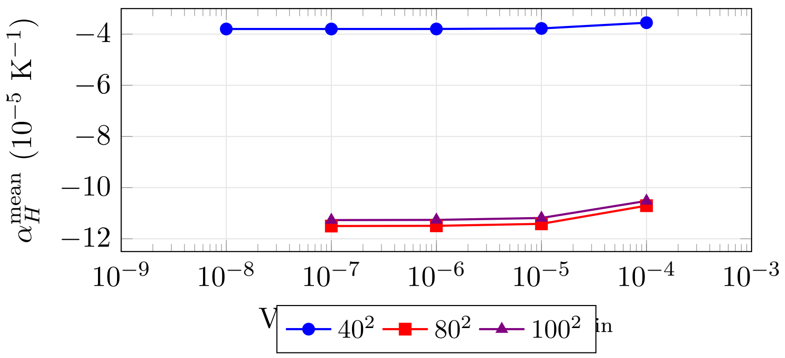

Figure 12: Mean homogenized CTE of the frozen case-1 design versus void stiffness ratio (logarithmic axis) for the three grids. Decreasing the void stiffness shifts the prediction monotonically toward more negative CTE — about 7% from 10^{-4} to the converged limit — and the curves saturate below 10^{-6} (the $10^{-6} \to 10^{-7}$ change is below 0.1% at 100^2); the demonstration's 10^{-4} choice therefore carries a void-regularization bias of about 7% in α_H^{mean}. The 10^{-8} evaluations at 80^2 and 100^2 do not converge and are omitted.

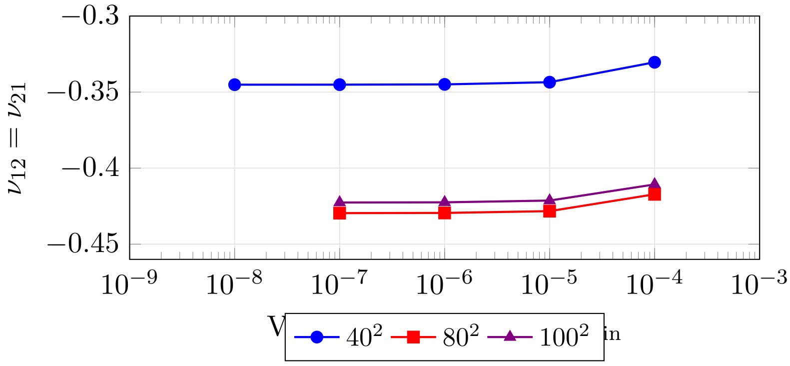

Figure 13: Effective Poisson ratio $\nu_{12} = \nu_{21}$ of the frozen case-1 design versus void stiffness ratio (logarithmic axis) for the three grids. ν_{12} shifts by about 3% from 10^{-4} to the converged limit at 100^2 (-0.411 to -0.423) and saturates below 10^{-6}.

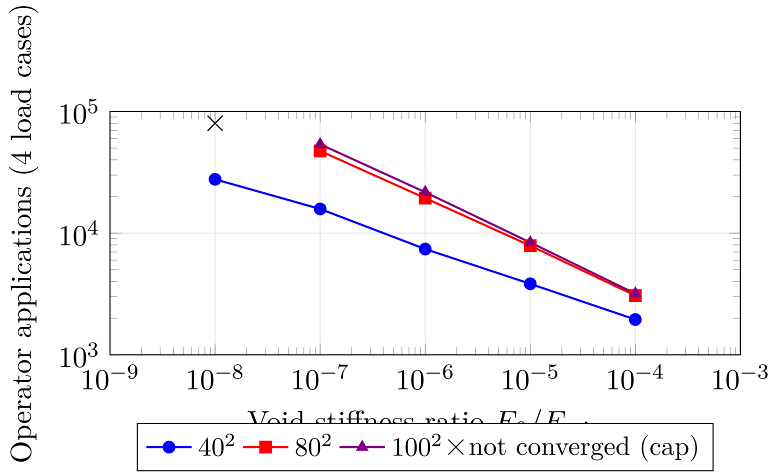

Figure 14: Cumulative operator applications of the four load cases versus void stiffness ratio (logarithmic axes) over the decade sweep 10^{-8}–10^{-4}, per grid. Work grows about 2.6-fold per decade of void softening at 100^2 (3,183 applications at 10^{-4} to 53,725 at 10^{-7}); the two $\times$ markers are the 10^{-8} evaluations on the 80^2 and 100^2 grids, which do not converge within 20,000 iterations.

D Initialization sensitivity: trajectories and audit data

This appendix collects the per-guess audit data from the initialization screening described in Sec. 3.3: the objective trajectories of all six runs and the full per-guess table, which includes entry objectives, steps, and operator work.

Table 11: Final normal CTE components for the six initialization screening runs (units 10^{-5} K^{-1}). The first two components are in-plane; the third is the reconstructed out-of-plane component.

Initial guess	α_{11}^H	α_{22}^H	α_{33}^H
Center circle	−9.11	−9.11	+6.14
Chiral hint	−9.59	−8.58	+5.48
Layered hint	+12.07	+5.19	+6.70
Random circles (s1)	−10.50	−9.52	+6.19
Random circles (s2)	−8.85	−7.43	+6.30
Reentrant hint	−11.63	−7.28	+5.94

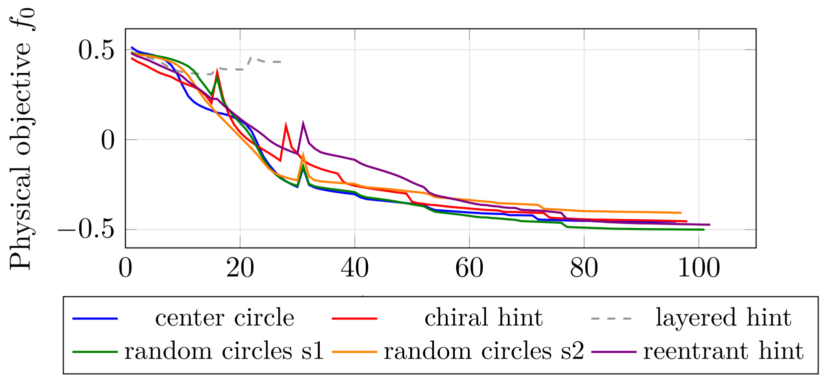

Figure 15: Physical objective f_0 versus accepted step for the six screening initializations of Sec. 3.3. All runs except the layered hint cross zero during stage B_pg2 (steps 21–25) and reach negative f_0; the layered hint terminates at stage B_pg3 with $f_0 = +0.431$ and high grayness. Final geometry and material grayness of the successful runs are $G_g \leq 0.051$ and $G_m \leq 0.070$ (Table 12).

Table 12: Initial-guess screening through the full continuation schedule. f_0 entry = objective at the first committed step; $\alpha_H^{\text{mean}} = f_0^{\text{exit}}\,\alpha_{\text{ref}}$ is the physical homogenized CTE of the final design, with $\alpha_{\text{ref}} = 2 \times 10^{-4}$ K^{-1} (units 10^{-5} K^{-1}); Work = cumulative operator applications. All runs except layered hint reached the final stage Dm_3.

Initial guess	f_0 entry	α_H^{mean}	G_g exit	G_m exit	Steps	Work
Center circle	0.5148	−9.11	0.033	0.052	96	764470
Chiral hint	0.4533	−9.08	0.051	0.053	98	1068017
Layered hint	0.4807	+8.63	0.648	0.534	28	13598
Random circles (s1)	0.4856	−10.01	0.037	0.051	101	821727
Random circles (s2)	0.4840	−8.14	0.029	0.049	97	710853
Reentrant hint	0.4809	−9.46	0.044	0.070	102	975389
Balanced-v2 (reentrant)	0.4972	−10.52	0.053	0.072	177	1814061

E Three-dimensional feasibility: supporting data

This appendix records the run details for the two 40^3 production runs reported in Sec. 3: the full setup, the per-stage objective trajectory, and the termination audit. The three-dimensional homogenization uses seven load cases — six mechanical engineering-strain bases in Voigt order $(11, 22, 33, 12, 23, 31)$ plus one constrained thermal load at $\bar{\boldsymbol{E}} = \mathbf{0}$, $\Delta T = t$ — and the production 3D objective is the normalized mean diagonal CTE $f_0 = (\alpha_{11}^H + \alpha_{22}^H + \alpha_{33}^H)/(3\,\alpha_{\mathrm{ref}})$ with $\alpha_{\mathrm{ref}} = 10^{-4}$ K^{-1}.

Table 13: Setup of the $40 \times 40 \times 40$ three-dimensional production runs.

Setting	Value
Grid	$40 \times 40 \times 40$
Initialization	Structured seed families (center sphere, chiral hint, random spheres, reentrant hint, layered hint)
Filter	Periodic FFT cone, radius 0.125 (five voxels)
Void stiffness ratio	10^{-4}
Materials	$E_1 = E_2 = 1.0$, $\nu = 0.3$, $\alpha = (10^{-5}, 2 \times 10^{-4})$ K^{-1} (Table 6)
Objective	$f_0 = (\alpha_{11}^H + \alpha_{22}^H + \alpha_{33}^H)/(3\,\alpha_{\mathrm{ref}})$, $\alpha_{\mathrm{ref}} = 10^{-4}$ K^{-1}
Reference scales	$E_{\mathrm{ref}} = 1$, $G_{\mathrm{ref}} = 1/[2(1 + 0.3)]$, $\alpha_{\mathrm{ref}} = 10^{-4}$ K^{-1}
Phase-volume targets	$(0.50, 0.25, 0.25)$, band 10^{-3}
Modulus floors	$e_{\mathrm{floor}} = g_{\mathrm{floor}} = 0.10$
Solver	Sequential PCG, rtol $= 10^{-8}$, maxiter $= 2000$
GCMMA	10 inner resolves
Schedule	16-stage ladder: β_g: $1 \to 16$, β_m: $1 \to 8$, p: $1 \to 3$
Budgets	200 attempts; accepted-step budgets 160 (chiral) / 100 (center); wall-time budget 64800 s
JAX	CPU, x64 (Python 3.10.18, JAX 0.6.2)

Table 14: Per-stage objective trajectory of the two 40^3 runs; entry and exit denote $f_0 = \mathrm{mean}(\alpha_{11}^H, \alpha_{22}^H, \alpha_{33}^H)/\alpha_{\mathrm{ref}}$ at the stage's first and last committed records, respectively. Re-evaluation under each stage's $(p, \beta_g, \beta_m, \lambda)$ causes entry/exit discontinuities at stage transitions. All stage exits are cap-triggered except the chiral run's B_pg3 (diminishing return); both runs terminate in Dm_3 at its 10-step committed cap.

Stage	Chiral entry	Chiral exit	Center entry	Center exit
A_physical	1.0366	0.9516	1.1414	1.0682
B_p2	1.0024	0.8826	1.0081	0.9627
B_pg3	0.9040	1.0146	0.8746	0.7904
Cg_bg2	1.0075	0.9828	0.7762	0.7065
Cg_bg4	0.9640	0.9088	0.7263	0.6144
Cg_bg8	0.8352	0.6916	0.6311	0.5243
Cg_bg16	0.6381	0.5305	0.5148	0.4211
Cm_bm2	0.5055	0.4152	0.3914	0.3345
Cm_bm4	0.3863	0.3257	0.3084	0.2575
Cm_bm8	0.2954	0.2571	0.2330	0.1900
Dg_1	0.2500	0.2197	0.1768	0.1429
Dg_2	0.2152	0.1812	0.1315	0.0956
Dg_3	0.1759	0.1532	0.0873	0.0544
Dm_1	0.1517	0.1355	0.0391	0.0121
Dm_2	0.1341	0.1265	−0.0027	−0.0482
Dm_3	0.1258	0.1178	−0.0542	−0.0920

Both runs terminate at their accepted-step budgets: the final stage Dm_3 reaches its 10-step committed cap at 136 and 91 accepted steps against budgets of 160 and 100, and neither satisfies the terminal convergence criterion: over the trailing 10-record window the objective relative ranges are 0.066 (chiral) and 0.50 (center) against the 5×10^{-3} tolerance, and the center-sphere objective is still decreasing at termination. The chiral run consumes 1,113,436 operator applications in 13,406 s wall time (136 attempts, 136 accepted); the center-sphere run consumes 1,038,470 operator applications in 1,860 s (91 attempts, 91 accepted). The 3D runs were resumed after documented intermittent macOS/XLA crashes with zero progress loss. All committed records are feasible ($\max_i g_i < 0$), and the final designs meet the phase-volume targets within the band: $(0.499, 0.251, 0.250)$ (chiral) and $(0.500, 0.251, 0.249)$ (center). Both final designs lie at the modulus floors: the center-sphere design has $E_{11}/E_{\mathrm{ref}} \approx E_{22}/E_{\mathrm{ref}} \approx E_{33}/E_{\mathrm{ref}} \approx 0.1002$ and $G \approx 0.10001$ (all $\approx$ the 0.10 floors); the chiral design has $E_{11}, E_{22} \approx 0.1000$, $E_{33} \approx 0.1181$, $G_{12}, G_{31} \approx 0.10003$, $G_{23} \approx 0.1353$. Both designs percolate in all three axes; solid 1 is a single component in both designs, and solid 2 is one component in the chiral design and seven in the center-sphere design. No 3D Poisson ratio was recorded; double-negative 3D results have not yet been obtained, so the center-sphere result is reported as budget-limited evidence of a negative mean CTE, not a converged optimum.