



HAL
open science

Computational Design of Flexible Planar Microstructures

Zhan Zhang, Christopher Brandt, Jean Jouve, Yue Wang, Tian Chen, Mark
Pauly, Julian Panetta

► **To cite this version:**

Zhan Zhang, Christopher Brandt, Jean Jouve, Yue Wang, Tian Chen, et al.. Computational Design of Flexible Planar Microstructures. ACM Transactions on Graphics, 2023, 42 (6), pp.1-16. 10.1145/3618396 . hal-04425030

HAL Id: hal-04425030

<https://inria.hal.science/hal-04425030>

Submitted on 27 Aug 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

Computational Design of Flexible Planar Microstructures

ZHAN ZHANG, University of California, Davis, USA

CHRISTOPHER BRANDT, 1000shapes GmbH, USA

JEAN JOUVE, University Grenoble Alpes Inria, CNRS, Grenoble INP, LJK, France

YUE WANG, University of Houston, USA

TIAN CHEN, University of Houston, USA

MARK PAULY, EPFL, Switzerland

JULIAN PANETTA, University of California, Davis, USA




Fig. 1. We introduce an efficient, accurate method for periodic homogenization of nonlinear elasticity and develop a computational framework for designing microstructures that best reproduce an input hyperelastic constitutive law over a *finite region* of applied macroscopic strains. For validation, we design a catalog of metamaterials best fitting a broad range of isotropic Hooke’s laws (E, ν) at finite strain. For each highlighted sample, we show the design produced by a traditional linear-elasticity-based method (left) and our optimization result (right), visualizing the accuracy with which the target stress-strain map is reproduced using surfaces in stress space $(\sigma_{xx}, \sigma_{yy}, \sigma_{xy})$ as described in Figure 5. The colormap visualizes the relative error in fitting to the target stress (indicated by the translucent ellipse). Right: an assortment of structures that we fabricated and tested to validate our design optimization.

Mechanical metamaterials enable customizing the elastic properties of physical objects by altering their fine-scale structure. A broad gamut of effective material properties can be produced even from a single fabrication material by optimizing the geometry of a periodic microstructure tiling. Past work has extensively studied the capabilities of microstructures in the small-displacement regime, where periodic homogenization of linear elasticity yields computationally efficient optimal design algorithms. However, many applications involve flexible structures undergoing large deformations for which the accuracy of linear elasticity rapidly deteriorates due to geometric nonlinearities. Design of microstructures at finite strains involves a massive increase in computation and is much less explored; no computational tool yet exists to design metamaterials emulating target hyperelastic laws *over finite regions of strain space*.

Authors’ addresses: Z. Zhang, UC Davis, zzzhang@ucdavis.edu; C. Brandt, 1000shapes GmbH, christopher.brandt.n@gmail.com; J. Jouve, U. Grenoble Alpes Inria, jean.jouve@inria.fr; Y. Wang, UH, ywang298@cougarnet.uh.edu; T. Chen, UH, tianchen@uh.edu; M. Pauly, EPFL, mark.pauly@epfl.ch; J. Panetta, UC Davis, jpanetta@ucdavis.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2023 Copyright held by the owner/author(s).

0730-0301/2023/12-ART185

<https://doi.org/10.1145/3618396>

We make an initial step in this direction, developing algorithms to accelerate homogenization and metamaterial design for nonlinear elasticity and building a complete framework for the optimal design of planar metamaterials. Our nonlinear homogenization method works by efficiently constructing an accurate interpolant of a microstructure’s deformation over a finite space of macroscopic strains likely to be endured by the metamaterial. From this interpolant, the homogenized energy density, stress, and tangent elasticity tensor describing the microstructure’s effective properties can be inexpensively computed at any strain. Our design tool then fits the effective material properties to a target constitutive law over a region of strain space using a parametric shape optimization approach, producing a directly manufacturable geometry. We systematically test our framework by designing a catalog of materials fitting isotropic Hooke’s laws as closely as possible. We demonstrate significantly improved accuracy over traditional linear metamaterial design techniques by fabricating and testing physical prototypes.

CCS Concepts: • **Computing methodologies** → **Shape Modeling; Simulation**.

Additional Key Words and Phrases: metamaterials, homogenization, physics-based simulation, computational design, fabrication

ACM Reference Format:

Zhan Zhang, Christopher Brandt, Jean Jouve, Yue Wang, Tian Chen, Mark Pauly, and Julian Panetta. 2023. Computational Design of Flexible Planar Microstructures. *ACM Trans. Graph.* 42, 6, Article 185 (December 2023), 16 pages. <https://doi.org/10.1145/3618396>

1 INTRODUCTION

Modern high-resolution digital fabrication technologies have opened up the possibility of creating custom-tailored *elastic metamaterials*, fine-scale layouts of deformable material whose geometric structure can give rise to fine-tuned, spatially graded elastic properties. With the right geometry, metamaterials can exhibit strikingly different behaviors from any natural material, like extreme positive or negative Poisson's ratios (*pentamode* [Kadic et al. 2012] or *auxetic* materials [Ren et al. 2018], which have been the subject of significant recent research activity). A large body of past work has focused on the inverse homogenization problem: designing a microstructure geometry achieving a desired elastic material behavior. However, the vast majority of this work relies on the assumption of *infinitesimal displacements* to replace the nonlinear governing physics with linear elasticity. The linearized theory brings dramatic simplifications, permitting a metamaterial geometry's effective small-strain material properties to be concisely summarized by a fourth-order *homogenized elasticity tensor*, calculated efficiently by solving a small number of sparse linear systems.

Although linear elasticity is reasonable for stiff structures, flexible metamaterials used in, *e.g.*, compliant mechanisms, soft robots, and prosthetics, quickly enter the large deformation regime where geometric nonlinearities cannot be neglected and designs can deviate strongly from their linearly elastic behaviors. This issue is illustrated in Figure 2, where an isotropic auxetic metamaterial optimized to produce a Poisson's ratio of $\nu = -0.75$ under linear elasticity is stretched horizontally. Its tangent Poisson's ratio $-\frac{\partial \varepsilon_{yy}}{\partial \varepsilon_{xx}}$ immediately deviates from the designed value, even changing sign before reaching a horizontal strain of $\varepsilon_{xx} = 10\%$. We note that nonlinear properties like these *could* be harnessed for certain applications, but linear-elasticity-based design fundamentally lacks awareness or control of them.

The few past efforts to extend microstructure design to finite strains have employed severe simplifications: controlling behavior in only a single loading scenario [Nakshatrala et al. 2013], at a few sampled loads [Behrou et al. 2021], or along one deformation path (*e.g.*, uniaxial stretching) [Clausen et al. 2015; Schumacher et al. 2018a; Wang et al. 2014]. These limitations are acceptable in some situations, like for bistable auxetics [Chen et al. 2021a] known to admit a *single* low-energy deformation path. However, many applications call for tailoring properties over *finite regions*, \mathcal{E} , of *strain space* (see Figure 4). The distinction is visualized in the inset sketch: control is needed over a full volume (translucent ball) instead of the sparse and low-dimensional samplings of past work, *e.g.* uniaxial strains with varying magnitude or direction (brown and blue dots, respectively). Characterizing and optimizing material properties over this finite region poses several computational challenges that our work seeks to address.

Applications of finite-strain elastic metamaterial design abound. At the most basic level, microstructures can be designed to reproduce simple analytical models, like linear Hooke's laws or neo-Hookean materials. Such target materials are intuitive for engineers to work with and are defined by a concise set of parameters that can be systematically swept (*e.g.*, Young's modulus and Poisson's ratio in




Fig. 2. Finite-strain properties of a microstructure designed using linear elasticity: the tangent Poisson's ratio quickly deviates from the requested value of $\nu = -0.75$.

the isotropic case). This enables generating catalogs of elastic metamaterials covering a wide range of material properties to support efficient, *fully decoupled multiscale design*, as has proven effective in the setting of linear elasticity [Panetta et al. 2015; Tozoni et al. 2020; Zhu et al. 2017].

Single-scale optimization of large microstructured parts like soft robots [Jin et al. 2020; Mark et al. 2016; Rafsanjani et al. 2019], prosthetics and medical implants [Kim et al. 2021; Pirozzi et al. 2022], or compliant mechanisms [Tozoni et al. 2020; Zhu et al. 2020] directly at the resolution of modern additive manufacturing technology (exceeding the teravoxel level) is intractable. But nonlinear homogenization enables a two-scale optimization where the full part is represented by a coarsened, solid mesh whose micro-scale detail is encoded by a collection of scalar fields. In past approaches, these fields encode the microstructure geometry, which ultimately requires solution and sensitivity analysis of nonlinear homogenization cell problems to be performed for each quadrature point in each coarse element concurrently with macro-scale design [Nakshatrala et al. 2013; Raju et al. 2021]. With the general-purpose material catalogs we generate, the scalar design fields on the coarsened mesh can instead represent the material parameters of achieved effective constitutive laws (like Young's moduli and Poisson's ratios), fully decoupling the coarsened design optimization from the microstructure optimization and dramatically improving efficiency.

Flexible microstructures can also be designed to exploit inherently nonlinear elastic material behaviors like multistability and jamming. Multistability is the mechanism underlying a number of fascinating emerging technologies like mechanical logic gates [Song et al. 2019] and neural networks [Lee et al. 2022], shape-reconfigurable materials [Haghpanah et al. 2016], reprogrammable materials [Chen et al. 2021b], and deployable surface structures [Chen et al. 2021a]. It is characterized by having multiple states at which elastic energy is locally minimized and frequently arises from buckling instabilities in slender beam structures. Jamming refers to a mechanical phase transition in which a metamaterial abruptly shifts from deforming flexibly into a near-rigid configuration upon application of certain strains [Wang et al. 2021]. This has applications in impact absorption [Wang et al. 2019] and robotics [Brown et al. 2010]. While we do not address the design of such materials in this work, we believe that our algorithms will provide a useful foundation for future efforts.

To summarize, the need for full nonlinear elasticity models in the analysis and design of flexible elastic metamaterials is twofold: it ensures *correctness* of designs and unlocks fundamentally nonlinear behaviors that are impossible within the framework of linear elasticity. Yet despite this clear need, little progress has been made, and developing optimal design algorithms for nonlinear metamaterials “remains fiendishly difficult” [Bertoldi et al. 2017] due to the immense computations involved. We develop algorithms to ease this computational burden, making the following main contributions:

- We develop an adaptive, data-accelerated homogenization algorithm that efficiently achieves a requested level of accuracy over a finite region of strain space.
- We introduce a material design algorithm optimizing for geometries best fitting a target hyperelastic model.
- We propose a robust scheme for promoting collision-free microstructures, avoiding the need for differentiable contact simulation in the inner loop of the design optimization.
- We generate a catalog of intuitive, as-linear-as-possible isotropic planar metamaterials to support applications in finite-strain multiscale design.

2 RELATED WORK

We contribute to a large and active area of research in the computer graphics community on fabrication-aware design of mechanical objects, which we briefly survey here, referring to [Bermano et al. 2017] for a recent comprehensive overview. Several methods have been proposed for enhancing the structural stability of additively manufactured objects by automatically hollowing, thickening, or inserting supports [Lu et al. 2014; Schumacher et al. 2018b; Stava et al. 2012; Ulu et al. 2017]. Finding the structurally weakest parts of an object with worst-case or stochastic analysis has been addressed by [Langlois et al. 2016; Zhou et al. 2013], as has *designing* under these loads [Cui et al. 2020; Langlois et al. 2016; Panetta et al. 2017; Schumacher et al. 2018b]. Many techniques have been developed to optimize the large-deformation behaviors of elastic objects [Bickel et al. 2010; Chen et al. 2014; Skouras et al. 2013; Xu et al. 2015]. Other works have optimized objects’ balance [Musialski et al. 2015; Prévost et al. 2013], spinnability [Bächer et al. 2017], natural frequencies [Bharaj et al. 2015; Musialski et al. 2016], rate of dissolution in a solvent [Panetta et al. 2022], and appearance while retaining structural soundness [Martínez et al. 2015; Zehnder et al. 2016].

The design of kinetic, shape-morphing, and deployable structures is especially relevant, as these often leverage metamaterials to actuate or control their transformations. Several works in this direction design patterns of rigid material that can be attached to pre-stretched elastic sheets, guiding their contraction motion to actuate towards a desired curved surface [Guseinov et al. 2017; Jourdan et al. 2022; Pérez et al. 2017]. Similar transformations to target curved surfaces have been achieved by inflatable sheets [Ou et al. 2016; Panetta et al. 2021] and networks of elastic beams [Baek et al. 2018; Panetta et al. 2019; Pillwein and Musialski 2021; Ren et al. 2022]. The spiral structures of [Malomo et al. 2018], while not deployable per se, are optimized to produce the right bending and stretching stiffness so that patches deform into a desired curvature upon constraining their boundaries. Ion et al. [2016] develop an interactive

design tool that can be used to create compliant mechanisms from a few metamaterial building blocks. Celli et al. [2018] and Chen et al. [2021a] design incision patterns that, when cut into a flat sheet of material, induce metric frustration under stretching and cause out-of-plane deformation. The latter is perhaps the most closely related, as it employs a similar nonlinear periodic homogenization framework. However, it could restrict its analysis to uniaxial trajectories due to the constraints imposed by the material system and did not need sensitivity-analysis-based design optimization due to the small numbers of structure parameters considered.

Past literature on the design of mechanical metamaterials using linear elasticity is vast. We mention only a few representative works from the graphics community, starting with [Panetta et al. 2015; Schumacher et al. 2015] that optimized periodic microstructures to achieve target elasticity tensors and satisfy tileability constraints to support creation of 3D printable parts with spatially graded properties. Microstructures have been used to support two-scale design for linear elasticity [Panetta et al. 2017; Zhu et al. 2017], including conforming to the boundary of the coarsened design mesh (and accounting for the ensuing unit cell distortion during homogenization) [Tozoni et al. 2020]. Extremal microstructures with properties close to the theoretical bounds have also been designed using topology optimization techniques [Chen et al. 2018; Ostanin et al. 2018]. We adopt and improve upon the parametric metamaterial shape optimization framework of [Panetta et al. 2017], which was originally developed to produce microstructures with low worst-case stresses.

The homogenization and design of elastic microstructures for finite elasticity is a much less mature field [Bertoldi et al. 2017]. As discussed in more depth in Section 1, the few works that have tackled large-deformation metamaterial design have restricted their focus to small subsets or trajectories through strain space [Behrou et al. 2021; Clausen et al. 2015; Schumacher et al. 2018a; Wang et al. 2014; Zehnder et al. 2017], rather than fitting over a finite region. Homogenization has been used to accelerate simulation in computer graphics, initially for linear elasticity [Kharevych et al. 2009], but later for nonlinear solids [Chen et al. 2015] and shells [Sperl et al. 2020]. Emerging techniques like proper generalized decompositions [Chinesta et al. 2013; Lamari et al. 2010] and neural networks [Le et al. 2015] have also been applied to homogenization of nonlinear mechanical metamaterials, but so far only in simplified settings as proofs of concept. We note that our interpolation-accelerated homogenization approach bears close resemblance to the recent method of [Chan-Lock et al. 2022] in that it constructs high-order interpolants of unstructured samples in strain space (in their case, using RBF interpolation). However, that work interpolates the homogenized energy density rather than the microstructure deformation as we do. That leads to more compact storage requirements but sacrifices the native measure of interpolant quality that we leverage to drive adaptive refinement (Section 4.3) and the ability to use partially-built interpolants to accelerate sampling (Section 4.2).

Our collision detection and avoidance approach is inspired by past simulation and modeling techniques that formulate response forces as the gradient of an interpenetration volume [Allard et al. 2010; Faure et al. 2008; Harmon et al. 2011]. Our method robustly accounts for multiple overlaps and is incorporated in the *outer loop* of the design optimization using adjoint sensitivity analysis.




Fig. 3. We use nonlinear periodic homogenization to define the effective properties of a periodically tiled unit cell Y containing microstructure Ω .

3 BACKGROUND AND METHODS OVERVIEW

In this section, we provide some background material on how the effective material properties of a microstructure geometry can be defined in the finite-strain setting and give an overview of our computational framework.

3.1 Periodic Homogenization for Hyperelastic Materials

We specifically design *periodic* microstructures that are produced by tiling a periodic *unit cell* patch $Y \subset \mathbb{R}^d$ (often referred to as a representative volume element, or RVE, in the mechanics literature) into a regular grid. This enables the use of periodic homogenization to rigorously define effective material properties in the limit of an infinite tiling. We focus on single-material structures, which are defined by the periodic geometry $\Omega \subset Y$ of the subdomain to be filled with fabrication material (see Figure 3), although our formulas hold also for heterogeneous fabrication materials.

Tiling Deformations. The fundamental assumption of periodic homogenization is that deformations of the tiling preserve its translational symmetry. In other words, all periodic copies of cell Y are assumed to deform into an identical shape that still periodically tiles space. It follows that the deformed position of each material point $\mathbf{X} \in \Omega$ takes the form:

$$\Phi(\mathbf{X}) = \omega(\mathbf{X}) + \bar{F}\mathbf{X},$$

where $\bar{F} \in \mathbb{R}^{d \times d}$ is an applied macroscopic deformation gradient determining how the tiling grid is distorted by the deformation (the linear map relating the lattice of cell corners in the rest and deformed configurations), and $\omega : \Omega \rightarrow \mathbb{R}^d$ is a periodic *fluctuation displacement field* indicating how the microstructure relaxes into equilibrium under this distortion. Note that this assumption can be violated by symmetry-breaking buckling modes whose periods span multiple unit cells, a caveat that we discuss in Section 8.

Homogenized Energy Density. In this work, we further assume that the fabrication material is hyperelastic, being described by a function $\psi : \mathbb{R}^{d \times d} \rightarrow \mathbb{R}$ that associates an elastic energy density with each micro-scale deformation gradient. This hyperelastic law enables us to compute the average energy density over each unit of the deformed periodic tiling, which is minimized at equilibrium:

$$\bar{\psi}(\bar{F}) := \min_{\omega \text{ periodic}} \frac{1}{|Y|} \int_{\Omega} \psi(\nabla \omega + \bar{F}) \, d\mathbf{X}. \quad (1)$$




Fig. 4. Homogenization process: given an applied macroscopic Biot strain $\bar{\epsilon}$ in *strain space* (left), we solve (3) to compute the deformation of the periodic microstructure tiling. Averaging the resulting PK1 stress over the unit cell (2) obtains the homogenized Biot stress $\bar{\sigma} = \bar{\psi}'$, a point in *stress space* (right).

This nonlinear map from the applied average deformation gradient \bar{F} to the resulting average energy density $\bar{\psi}(\bar{F})$ is called the *homogenized energy density function*, and serves as a complete description of the effective material properties of geometry Ω .

We note that $\bar{\psi}$ inherits rotational invariance from ψ : applying a rigid rotation \bar{R} to the equilibrium deformation yields a new, rotated equilibrium configuration $\bar{R}\Phi(\mathbf{X}) = \bar{R}\omega(\mathbf{X}) + \bar{R}\bar{F}\mathbf{X}$ for which $\bar{\psi}(\bar{R}\bar{F}) = \frac{1}{|Y|} \int_{\Omega} \psi(\bar{R}\nabla\omega + \bar{R}\bar{F}) \, d\mathbf{X} = \bar{\psi}(\bar{F})$. Therefore, it suffices to consider only *symmetric* \bar{F} in (1), and these can be thought of as the symmetric part of the polar decomposition ($\bar{F} = \bar{R}\bar{S} = \bar{S}$). For these symmetric deformation gradients \bar{F} , the *macroscopic Biot strain* is given simply by $\bar{\epsilon} = \bar{F} - I$.

Biot Stress and Stiffness. The homogenized Biot stress $\bar{\sigma}$ resulting from an applied macroscopic Biot strain $\bar{F} - I$, can be obtained by averaging the micro-scale stress at equilibrium:

$$\bar{\sigma}(\bar{F}) := \bar{\psi}'(\bar{F}) = \frac{1}{|Y|} \int_{\Omega} \psi'(\nabla \omega^*(\mathbf{X}; \bar{F}) + \bar{F}) \, d\mathbf{X}, \quad (2)$$

where $\psi' := \frac{\partial \psi}{\partial \bar{F}}$ denotes the first Piola-Kirchhoff stress at the micro-scale, and ω^* denotes the equilibrium fluctuation displacement:

$$\omega^*(\bar{F}) := \operatorname{argmin}_{\omega \text{ periodic}} \int_{\Omega} \psi(\nabla \omega + \bar{F}) \, d\mathbf{X}. \quad (3)$$

The stress formula follows directly from the *envelope theorem*: integrand terms of the form $\psi' : \nabla(\delta\omega)$ must integrate to zero for admissible perturbations $\delta\omega$ due to optimality of ω^* . Tangent elasticity tensors \bar{C} can be obtained by differentiating $\bar{\sigma}(\bar{F})$ once more with respect to \bar{F} , and in this case derivatives of ω^* cannot be neglected:

$$\bar{C}(\bar{F}) : e^{ij} = \frac{1}{|Y|} \int_{\Omega} \psi''(\nabla \omega^*(\mathbf{X}; \bar{F}) + \bar{F}) : (\nabla \omega^{ij} + e^{ij}) \, d\mathbf{X}, \quad (4)$$

where $\omega^{ij} := \frac{\partial \omega^*}{\partial \bar{F}} : e^{ij}$ are equilibrium derivatives that we provide formulas for in Section 4.1, and $e^{ij} := \frac{1}{2}(\mathbf{e}_i \otimes \mathbf{e}_j + \mathbf{e}_j \otimes \mathbf{e}_i)$ is the canonical basis for symmetric matrices.

Homogenization Process. To probe the material's behavior at a single applied Biot strain tensor $\bar{\epsilon}$, one first must solve the *nonlinear, nonconvex* optimization (3) with $\bar{F} = I + \bar{\epsilon}$, after which the homogenized energy and stress can be calculated via (1) and (2). This process is visualized in Figure 4. However, we recall that our goal is to characterize the behavior over a finite strain region \mathcal{E} , which for

planar microstructures is a three dimensional space with coordinates $(\varepsilon_{xx}, \varepsilon_{yy}, \varepsilon_{xy})$. Therefore, separate instances of (3) must in theory be solved for each of *infinitely many* $\bar{F} \in \mathcal{F}$, where $\mathcal{F} = \{I + \bar{\varepsilon} | \bar{\varepsilon} \in \mathcal{E}\}$. In practice, an approximation to $\bar{\psi}$ or ω^* must be constructed with suitable accuracy *over all of* \mathcal{E} , and this process must be computationally efficient to run in the inner loop of a design optimization.

3.2 Overview

The first core technical contribution of our work is an accelerated scheme for constructing an interpolant of the equilibrium fluctuation displacement ω^* with controlled accuracy over a suitably chosen strain region \mathcal{E} , effectively solving the homogenization problem. This approach is described in Section 4. The second contribution, described in Section 5, is an optimal design algorithm for generating metamaterials best-fitting target hyperelastic laws $\bar{\psi}_{\text{tgt}}$ over \mathcal{E} . This section also describes our approach to eliminate self-contact at the level of design rather than augmenting (3) with collision constraints or a smoothed contact model like Incremental Potential Contact (IPC) [Li et al. 2020]; this approach circumvents the large increase in nonlinearity inherent in accurate contact simulation and is effective when $\bar{\psi}_{\text{tgt}}$ can be achieved without harnessing self contact. Next, we provide the details of our implementations and explain our choice of validations in Section 6 before reporting results in Section 7. Finally we conclude in Section 8 by discussing a few of the many avenues for further research to build on our methods.

4 INTERPOLANT-ACCELERATED HOMOGENIZATION

Given a microstructure geometry Ω and a strain region \mathcal{E} , our ultimate goal is to characterize the effective properties $\bar{\psi}(\bar{F})$ to controlled accuracy at all macroscopic deformation gradients $\bar{F} \in \mathcal{F}$. However, our insight is that rather than approximating $\bar{\psi}(\bar{F})$ directly, it is advantageous to construct an approximant of the equilibrium fluctuation displacement $\omega^*(\bar{F})$ (from which $\bar{\psi}(\bar{F})$ and $\bar{\sigma}(\bar{F})$ can be inexpensively computed): (i) the optimality conditions of cell problem (3) provide a natural, physically meaningful measure of approximation quality with no need for ground-truth solutions; (ii) a partially built approximant yields high-quality initializations for solving at new sample points; and (iii) the data defining our specific approximant can be reused in the shape derivative formulas of Section 5.

4.1 Hermite Interpolation of the Micro-Scale Equilibrium

A simple approach would be to sample ω^* at a number of macroscopic deformation gradients $\bar{F} \in \mathcal{F}$ and interpolate the vector fields. However, each solve of (3) is expensive, making it desirable to extract as much data for each \bar{F} as can be obtained and stored efficiently. Fortunately, the *derivatives* ω^{ij} of ω^* with respect to the three independent coordinates defining symmetric \bar{F} are efficient to compute, enabling the use of higher-order C^1 Hermite interpolation with substantially improved approximation power. Implicit in our approach is an assumption that ω^* depends smoothly on \bar{F} . While ω^* can change non-smoothly in the presence of self-contact, or even discontinuously at the onset of buckling, we do expect ω^*

to be smooth for optimal designs reproducing the simple hyperelastic laws we target in this paper. We leave more sophisticated, discontinuity-aware interpolation for future work.

To obtain this derivative data, we consider the optimality conditions characterizing the solution of (3):


$$\int_{\Omega} \nabla \phi : \psi'(\nabla \omega^* + \bar{F}) dX = 0,$$

for all admissible perturbations ϕ of ω^* (those preserving periodicity). Differentiating both sides of this equation with respect to \bar{F} (computing a directional derivative along each symmetric basis matrix e^{ij}) obtains an equation for ω^{ij} :

$$\int_{\Omega} \nabla \phi : \psi'' : (\nabla \omega^{ij} + e^{ij}) dX = 0 \quad \forall \phi. \quad (5)$$

After discretization, this is a sparse linear system featuring the same Hessian matrix used to solve (3) with Newton's method. We also note that at the undeformed configuration $\bar{F} = I$, this equation reduces to the cell problems solved for periodic homogenization of linear elasticity.

To further boost sampling efficiency, we construct a trivariate C^1 Powell-Sabin interpolant [Worsey and Piper 1988] on a tetrahedral mesh \mathcal{T} adapted to tightly contain \mathcal{E} rather than sampling on a regular grid. Beyond the efficiencies of better conformance to \mathcal{E} , we note that adaptively refining a tetrahedron of this grid by a standard edge-splitting scheme only requires inserting 6 new points (see inset), more surgical than the 19 new sample points involved in subdividing a voxel into 8 subvoxels.



The Powell-Sabin scheme achieves a piecewise quadratic C^1 interpolant by splitting each tetrahedron of the interpolation mesh into 24 subelements and defining a quadratic polynomial on each consistent with the data ω^* and ω^{ij} at the original tet corners. Strict geometric requirements must hold to ensure all of these quadratic interpolants meet with C^1 continuity; a sufficient condition suggested by Worsey and Piper [1988] is for the scheme's split points (one per edge, face, and interior of the original tetrahedron) to be placed at the circumcenters of each associated simplex—which works provided the circumcenter falls within the simplex. However, ensuring each tetrahedron's circumcenter lies within, even after the subdivisions applied by adaptive refinement, turns out to be difficult. Surprisingly, the regular tetrahedra that meshing algorithms typically aspire to produce are unacceptable: four of the eight refined tetrahedra generated by edge subdivision have circumcenters lying on an edge rather than the interior. We therefore developed a custom mesh optimization algorithm detailed in the supplement that minimizes the distance between $\partial\mathcal{T}$ and $\partial\mathcal{E}$ along with a distortion measure based on a St. Venant Kirchhoff energy, all while ensuring tet circumcenters fall within the tet at all levels of subdivision with a stiff penalty term.

4.2 Accelerated Interpolant Construction

Constructing our interpolant requires solving the nonlinear elasticity simulation problem (3) once for each vertex of the interpolation mesh, probing the microstructure tiling with the vertex's associated




Fig. 5. Visualization of our homogenization result: given a microstructure (bottom), we construct an accurate interpolant of ω^* over the entire ellipsoid \mathcal{E} , which enables us to efficiently sample the strain-stress map $\bar{\psi}'$ at any $\bar{\epsilon} \in \mathcal{E}$. This map is visualized by showing the images of concentric ellipses. We observe significant nonlinearity in this structure: the stress surfaces deviate from ellipses even at this mild $\bar{\epsilon}_{\max} = 0.05$.

\bar{F} . This simulation will converge much faster if started with a high-quality initial guess, which we note can be obtained from existing interpolation data.

We construct the interpolant with a breadth-first search (BFS) approach: we begin at the vertex closest to the origin of \mathcal{E} for which we expect the solver to converge most rapidly. All interpolation meshes used in our results contain a vertex directly at the origin $\bar{\epsilon} = 0$, in which case the exact solution $\omega^* = 0$ is immediately returned with no iteration. We next solve the tangent cell problems (5) to obtain the full interpolation data for the initial vertex. Our BFS algorithm then iteratively expands outwards from the initial vertex along the mesh edges. Every time a new vertex is encountered, an initial guess is obtained from the first-order Taylor expansion defined by the data at a processed neighbor. If multiple processed neighbors exist, the guess yielding the lowest elastic energy is picked. As we demonstrate in Figure 15, this strategy achieves a measurable speedup over a naïve $\omega = 0$ initialization.

4.3 Adaptive Refinement

Our aim is to characterize the homogenized properties to controlled accuracy, which the user specifies in terms of a tolerance ϵ for the residual norm of the optimality conditions for (3) when evaluated on our ω^* interpolant. Physically speaking, we require that residual forces are negligible in the interpolated deformation of the microstructure tiling.

We enforce this condition approximately throughout \mathcal{E} by checking the interpolant quality at the barycenter of each tetrahedron of \mathcal{T} : we obtain the interpolated fluctuation displacement field and evaluate its residual norm. If the residual norm exceeds ϵ , we subdivide the tetrahedron into 8 subtetrahedra, inserting new points at each edge midpoint. For each of these six new vertices created, we solve (3) using the existing interpolant as a high-quality initial guess, and then obtain the required derivative data from (5). We repeat this process on all subdivided tetrahedra until the requested subdivision tolerance is met by all query points (or a maximum number of subdivision levels is hit).




Fig. 6. Our automatic construction of \mathcal{E} and its discretization \mathcal{T} for the structure depicted in the upper left. The tangent elasticity tensor at the undeformed configuration, \bar{C} , is applied to stresses within the unit ball in Mandel space (middle), producing an ellipsoid (left) of likely low-energy strains. Right: a tetrahedral discretization of this ball that was optimized for compatibility with our adaptive Powell-Sabin technique, and the subdivisions made during adaptive refinement of ω^* .

The result is a highly accurate interpolant defined on a tetrahedral mesh hierarchy $\tilde{\mathcal{T}}$ that can be used to evaluate $\bar{\psi}$ and $\bar{\sigma}$ throughout \mathcal{E} , as visualized in Figure 5. Given a query point $\bar{\epsilon}_q \in \mathcal{E}$, we locate the *finest-level* tetrahedron containing $\bar{\epsilon}_q$ and evaluate the Hermite interpolant on that tetrahedron. This point location is done recursively, first locating the coarse tetrahedron containing $\bar{\epsilon}_q$ and then, if the containing tetrahedron was refined, repeating the search process over its 8 subtetrahedra.

4.4 Picking the Interpolant Domain

The question remains what region to pick for \mathcal{E} ; while this could be selected by users with knowledge of the deformations they intend to apply, an automatic choice is desired for generating material catalogs for multiscale design. Picking a strain radius $\bar{\epsilon}_{\max}$ and taking the ball $\|\bar{\epsilon}\|_F \leq \bar{\epsilon}_{\max}$ might seem like the most natural choice, but it is often problematic. The metamaterial is likely to resist some strains much more stiffly than others, and forcing the structure into those strains not only fails to reflect practical use cases but also leads to contorted, high-energy deformations for which (3) can be slow to converge. An example would be an isotropic auxetic material with Poisson's ratio $\nu = -1$: such materials freely expand and contract uniformly but strongly resist shearing. Picking a reasonable domain \mathcal{E} is even more important for material design to prevent the optimizer from attempting to fit properties at extreme strains that never will be exercised in practice at the expense of more typical strains.

The automated selection strategy we propose picks an ellipsoid in macroscopic strain space that approximately corresponds to a ball $\|\bar{\sigma}\|_F \leq \bar{\sigma}_{\max}$ in macroscopic stress space; the resulting ellipsoid will naturally stretch along the structure's flexible directions and contract in its stiff directions as desired. Specifically, we use the tangent elasticity tensor at the undeformed configuration to map each point within the stress ball into strain space, obtaining \mathcal{E} as the image. When running our homogenization as a standalone process, this is the homogenized elasticity tensor $\bar{C}(I)$ given by (4), while for material fitting we use the tangent elasticity tensor of the target constitutive law, $\bar{\psi}'_{\text{tgt}}$. The user still specifies a strain bound $\bar{\epsilon}_{\max}$, from which we determine $\bar{\sigma}_{\max} = \bar{\epsilon}_{\max}/\lambda_{\max}$, where λ_{\max} is the

largest eigenvalue of the tangent elasticity tensor. This choice scales the ellipsoid \mathcal{E} so that the strain bound holds tightly.

The resulting strain domain is visualized in Figure 6, along with our tetrahedral discretization. We wish for its vertices to be roughly uniformly distributed in stress space, and thus construct this discretization in the coordinate system where the ellipsoid is a unit sphere. This can be interpreted as meshing a sphere in the “Mandel stress space” $(\bar{\sigma}_{xx}, \bar{\sigma}_{yy}, \sqrt{2}\bar{\sigma}_{xy})$ in which the Euclidean norm corresponds to the Frobenius norm $\|\bar{\sigma}\|_F$.

An alternative to constructing ellipsoid \mathcal{E} as the (approximate) image of a ball in stress space would be to simply use that stress-space ball as the native sampling domain. This means reformulating the homogenization problem (1) in terms of an applied macroscopic stress $\bar{\sigma}$, which can be implemented simply by adding a linear load term $-\bar{\sigma} : \bar{F}$ to the potential energy in (3), and optimizing over \bar{F} in addition to ω . The material design problem formulated in Section 5 likewise can be recast as fitting strains at prescribed stresses. We implemented this version of our framework but found it less robust than the strain sampling approach, particularly in the presence of buckling.

4.5 Exploiting Reflectional Symmetry

Especially when designing isotropic materials, it is often desirable to restrict microstructure geometries to have reflectional symmetry across the period cell faces. In this case, by assuming this additional symmetry imposed on the tiling deformation is not broken by buckling, we can both reduce the number of displacement variables and speed up the simulation by roughly a factor of two. This is done by posing the homogenization problem over the smaller symmetry base cell—the positive quadrant in 2D or octant in 3D. The way shear strains transform under reflection ends up requiring multiple copies of the ω field to be defined over the symmetry base cell (two in 2D, four in 3D), hence the only $2\times$ reduction in problem size. The details of our orthotropic homogenization method, which we used for all examples in the paper, are provided in the supplement

5 MATERIAL DESIGN

Given a target hyperelastic material model $\bar{\psi}_{\text{tgt}}$, our goal is to construct a microstructure geometry Ω^* whose homogenized properties computed using the methods of Section 4 best fit to it. Ideally we would solve for a geometry satisfying the constraints $\bar{\psi}(\bar{F}) = \bar{\psi}_{\text{tgt}}(\bar{F})$ for every $\bar{F} \in \mathcal{F}$. However, we propose the following more amenable nonlinear least squares optimization formulation:

$$\min_{\Omega} \frac{1}{|\mathcal{F}|} \int_{\mathcal{F}} w_e (\bar{\psi}/\bar{\psi}_{\text{tgt}} - 1)^2 + w_\sigma \|\bar{\psi}' - \bar{\psi}'_{\text{tgt}}\|^2 / \|\bar{\psi}'_{\text{tgt}}\|^2 d\bar{F}, \quad (6)$$

where the objective penalizes relative deviations from the target energy values $\bar{\psi}_{\text{tgt}}$ and Biot stresses $\bar{\psi}'_{\text{tgt}}$. Although both terms express a similar goal, including both improves robustness when the integral in (6) is approximated by quadrature and enables the user to trade off between different forms of error by adjusting the weights w_e and w_σ when a perfect fit is impossible.

We evaluate the integral in (6) using linear Newton-Cotes quadrature over the tetrahedral discretization \mathcal{T} of \mathcal{E} : the quadrature points of this rule are simply the vertices of \mathcal{T} , so the sampled ω^*

and ω^{ij} data can be used directly without interpolation. If the origin $\bar{\varepsilon} = 0$ is a vertex of \mathcal{T} (as in all our examples), that quadrature point is dropped: its relative error is undefined, but its absolute error is zero. The resulting discrete objective is in the form of a standard nonlinear least squares problem with a finite number of residuals (four per vertex of \mathcal{T}). We minimize it using a Gauss-Newton solver [Wright 2006] after defining a finite set of variables parametrizing the design geometry Ω (Section 5.3). We note that each step of this design optimization requires solving the homogenization problem for the design candidate (Section 4) and then differentiating these homogenized properties with respect to the microstructure’s geometry to determine how best to improve the design.

5.1 Shape Derivatives of the Residuals

The Gauss-Newton algorithm needs only the gradients of each residual function (which are just the derivatives of $\bar{\psi}$ and $\bar{\psi}'$ evaluated at vertices of \mathcal{T}). We employ shape optimization in this work, preferring to explore different topologies by enumeration as discussed in Section 5.3. Therefore, we need formulas for the *shape derivatives* of these properties: how they change when the rest shape Ω is advected. As in [Panetta et al. 2017], we prefer the “volume integral form” of the shape derivative expressions (also known as a “distributed shape derivative” [Laurain and Sturm 2016]) rather than the more common boundary integral form. This is because the former yields the exact discrete gradient when discretized directly.

We show in the supplement that the initial rate of change of each residual when the rest shape is advected by a velocity field \mathbf{v} (undergoing the map $\mathbf{X} \mapsto \mathbf{X} + \mathbf{v}(\mathbf{X})$) can be written in the form:

$$\left\langle \frac{\partial \bar{\psi}}{\partial \Omega}, \mathbf{v} \right\rangle = \int_{\Omega} G_{\bar{\psi}} : \nabla \mathbf{v} d\mathbf{X}, \quad \left\langle \frac{\partial \bar{\sigma}_{ij}}{\partial \Omega}, \mathbf{v} \right\rangle = \int_{\Omega} G_{\bar{\sigma}_{ij}} : \nabla \mathbf{v} d\mathbf{X}, \quad (7)$$

with the tensor fields:

$$G_{\bar{\psi}} = \psi I - (\nabla \omega^*)^\top \psi', \quad G_{\bar{\sigma}_{ij}} = ((\nabla \omega^{ij} + e^{ij}) : \psi') I - (\nabla \omega^{ij})^\top \psi' - (\nabla \omega^*)^\top \psi'' : (\nabla \omega^{ij} + e^{ij}).$$

We discretize Ω with quadratic *subparametric* elements. Despite having mid-edge nodes that allow the elements to curve in the deformed configuration, the rest shapes of these elements are straight-edged, defined only by their vertex positions. We can therefore use the formulas in (7) to obtain the exact discrete gradient with respect to each rest vertex position \mathbf{X}_k efficiently as $\int_{\Omega} G \nabla \lambda_k d\mathbf{X}$, where λ_k is the piecewise linear shape function associated with vertex k , and G stands for either $G_{\bar{\psi}}$ or $G_{\bar{\sigma}_{ij}}$. The resulting discrete shape derivative can be visualized as a vector field as in Figure 9.

5.2 Collision Detection and Avoidance

At finite strain, and particularly under compression, self-collisions of the microstructure are likely. A straightforward though expensive approach for capturing collision behavior in our framework would be to incorporate a robust, differentiable contact model like IPC [Li et al. 2020] into our simulation. In particular, the IPC barrier terms can be added to the potential energy in (3) and accounted for in sensitivity analysis. However, this not only increases the cost of each simulation step (requiring continuous collision detection), but the resulting more complex physical behaviors take more iterations




Fig. 7. Collision area visualization: the collision area \mathcal{A} computed for each strain $\bar{\varepsilon} \in \mathcal{E}$ is visualized as a scalar field on the left. For the marked point, the microscopic deformation is visualized, with the self-colliding regions highlighted in red. This example features multiple levels of overlap as indicated by the winding numbers annotated for each polygon in the positive quadrant. We note that in this example, the collision area was computed within reflectional symmetry base cell.

to simulate, are more difficult to approximate with our interpolants of Section 4.1, and greatly increase the nonconvexity of the design optimization.

We note that self-collisions induce discontinuous changes in the material properties that are counterproductive for fitting a smooth target model $\bar{\psi}_{\text{tgt}}$. Consequently, in this work we focus on the goal of *avoiding* collisions rather than attempting to harness them (as might be needed to produce, e.g., jamming behaviors). Accordingly, we propose to preserve the simplicity and robustness of the simulation (3) by permitting interpenetration at the simulation level but then adding a penalty term to (6) that guides the design algorithm to eliminate it.

Our collision penalty term is a global measure of how much interpenetration exists in the microstructure tiling across the full strain domain \mathcal{E} :

$$\frac{w_c}{|\mathcal{F}|} \int_{\mathcal{F}} \mathcal{A}(\omega^*(\bar{F}) + \bar{F}\mathbf{X}; \Omega)^2 d\bar{F}. \quad (8)$$

This penalty is added to the nonlinear least squares problem (6) by including a weighted version of \mathcal{A} as an additional residual function. Function \mathcal{A} evaluates the “collision area,” which we define as:

$$\mathcal{A}(\Phi; \Omega) := \int_{\Phi(\Omega)} (\text{wind}(\mathbf{x}, \Phi(\partial\Omega)) - 1)_+ dx,$$

where $\text{wind}(\mathbf{x}, \Gamma)$ evaluates the integer *winding number* of a closed curve Γ around deformed point \mathbf{x} , and $(\cdot)_+$ clamps its argument to positive values. This collision area is visualized as a function of $\bar{\varepsilon}$ in Figure 7; note how the winding number expression accounts for the multiple layers of overlaps occurring in this example.

We implement this term in 2D using CGAL’s 2D Arrangements package [Wein et al. 2023], inserting all boundary curves of the deformed microstructure tiling that intrude into the symmetry base cell, along with the borders of this base cell. Specifically, we insert *piecewise linear* approximations of these curves, ignoring the edge midpoint nodes. The resulting arrangement decomposes the base cell into polygons whose constant winding numbers we can determine using a BFS approach starting from the unbounded face, as suggested in [Zhou et al. 2016]. The desired integral \mathcal{A} is obtained




Fig. 8. Top: the six topologies selected as described in Section 6.1 along with the number of parameters in each associated parametric design space (which is constrained to preserve reflectional symmetry). Bottom: a few geometries sampled from the design space of the upper-left topology.

simply by summing each *collision polygon* area times its corresponding winding number expression.

We compute the shape derivative of \mathcal{A} in multiple steps, employing the chain rule. First, we compute the derivative with respect to each vertex of the collision polygon, which is straightforward: the winding number of each polygon is constant, and only polygon areas must be differentiated. We note that collision polygon vertices are either mesh boundary vertices directly or are defined by an intersection—either between two mesh boundary edges or one edge and the base cell boundary. Using derivatives of the intersection point positions, we obtain a gradient with respect to deformed mesh node positions by chain rule. Finally, we use the adjoint method to obtain a gradient with respect to the rest vertex positions.

5.3 Parametric Shape Optimization

For improved robustness, ease of enforcing manufacturability constraints, and to benefit from higher-order optimization algorithms, we choose to solve (6) using *parametric* shape optimization: rather than perturbing individual vertices according to the shape derivative formulas from Section 5.1, we work within a parametric design space. We specifically use the framework of [Panetta et al. 2017], which defines organic truss-like structures composed of smoothly blended beams. These microstructures are specified by first picking a lattice structure topology, after which the geometry is defined by a smooth implicit function controlled by a vector of shape parameters $\mathbf{p} \in \mathbb{R}^{n_p}$; example geometries are visualized in Figure 8. Each vertex of the lattice is converted to a ball specified by position and radius parameters. Each edge becomes a beam defined by the convex hull of the endpoint balls. Finally, beams are blended together smoothly according to the smoothness parameter associated with each vertex.




Fig. 9. Each component of the design objective gradient is computed as the inner product of two vector fields: the *shape derivative* computed using (7) (represented on the left as a steepest-ascent vector field) and the *shape velocity* for parameter p_i (visualized at the right for a horizontal position parameter) computed from the parametric geometry model using automatic differentiation.

In this parametric design subspace, the discretized version of (6) becomes a finite-dimensional nonlinear least squares problem:

$$\min_{\mathbf{p}} \frac{1}{|\mathcal{T}|} \sum_{\bar{\varepsilon} \in \mathcal{T}} \|\mathbf{r}(\bar{\varepsilon}, \mathbf{p})\|^2 \text{Vol}(\bar{\varepsilon}),$$

where in 2D, $\mathbf{r} : \mathbb{R}^3 \times \mathbb{R}^{n_p} \rightarrow \mathbb{R}^4$ evaluates the weighted residuals (the weighted relative error in energy and each stress component), and $\text{Vol}(\bar{\varepsilon})$ evaluates the volume of the barycentric region of vertex $\bar{\varepsilon}$ in tetrahedral mesh \mathcal{T} . Constraints are additionally imposed on this problem to ensure manufacturability and to keep the lattice nodes within the unit cell. In 2D, this entails minimum bounds on the radius parameters; for 3D structures, a self-supporting constraint can be formulated using linear inequality constraints as detailed in [Panetta et al. 2017].

The necessary residual Jacobians $\frac{\partial \mathbf{r}}{\partial \mathbf{p}}$ are obtained by computing inner products of the shape derivatives from Section 5.1 with the *shape velocities*, the piecewise linear perturbation velocity on Ω induced by changing pattern parameters (Figure 9). These velocities for boundary vertices are obtained via automatic differentiation of the implicit function and then interpolated throughout the interior by solving a Laplace equation (harmonic interpolation). For efficiency, the adjoint of the harmonic interpolation operator is applied to the shape derivative field to avoid solving for an interpolation for each parameter n_p , and again we refer to [Panetta et al. 2017] for the details.

5.4 Improved Beam Blending

We highlight one improvement we made to the parametric geometry definition of [Panetta et al. 2017] to fix an issue with how beams are blended. The original approach, though striving for smoothness, produces kinks in certain common situations like the one pictured in Figure 10. These kinks are problematic since they lead to stress concentrations and degrade the accuracy of our shape derivatives. Following the original blending algorithm, we first construct a smooth implicit geometry φ_j separately for each joint j by blending together the beam primitive signed distance functions φ^e for each incident edge e using a smoothed union controlled by joint j 's




Fig. 10. We improve the beam blending algorithm of [Panetta et al. 2017], fixing a discontinuity issue that led to sharp corners in the blended geometry (left). These kinks are smoothed away with our new method for blending together the per-joint geometries φ_j (right).

smoothing parameter. Ideally the material added by the smoothed unions performed at different joints do not interfere (overlap), in which case all φ_j can be combined with an exact union (min operation) without introducing creases. However, in the case of overlaps, they must be joined smoothly. The difficulty is that smooth union operations induce *bulging*, adding material wherever the boundaries of the input shapes nearly coincide. Simply smooth-unioning all φ_j will cause every beam to bulge (since each edge geometry φ^e is present in two blended joint geometries). The solution proposed in [Panetta et al. 2017] was to evaluate the final implicit function by (i) finding the closest two joints (j_1, j_2); and (ii) blending φ_{j_1} and φ_{j_2} together with a spatially modulated smoothness determined by an estimate of the degree to which their blends conflict to form a crease (Figure 10 of that paper). Unfortunately, we discovered that the spatial discontinuities induced by selection (i) causes creases.

We propose a simpler approach that eliminates these creases. We blend all φ_j smoothly using the Kreisselmeier-Steinhauser (KS) but subtract away the “double counting” committed by this operation that would lead to bulging. Specifically, we obtain the full implicit function via the formula:

$$\varphi = m - s \log \left(\sum_j \exp((m - \varphi_j)/s) - \sum_e \exp((m - \varphi^e)/s) \right),$$

where $m = \min_j \varphi_j$ is an implicit function for the exact union, and s is a global smoothing parameter that we set to 0.05. The subtracted sum over all edges is motivated by the observation that each edge is contained in *two* different joints whose implicit functions contribute to the first sum. There exist certain degenerate configurations (like evaluating at the center of a triangle arrangement of edges) where the edges are not fully double-counted and the subtraction can end up eroding material. To prevent this, we clamp the argument of log to remain above 1, guaranteeing the operator is additive. This clamp operation is the only potential source of nonsmoothness in our improved definition, and in practice we noticed no creases in our results—a substantial improvement over the frequent kinks in the original geometries, as demonstrated in Figure 10.

6 IMPLEMENTATION DETAILS

We generate all designs using a plane stress neo-Hookean model with Poisson’s ratio $\nu = 0.35$ for the base material ψ . We note that the Young’s modulus parameter simply globally scales the generated metamaterials’ stiffnesses without fundamentally changing the results. We investigate the influence of the Poisson’s ratio in Figure 12. Our modular implementation supports other hyperelastic models like corotated linear elasticity and St. Venant Kirchhoff, and adding additional more realistic models is straightforward. All of these fabrication material models differ only in their large-strain behaviors, which we argue does not significantly impact the behavior of flexible microstructures: deformations of flexible metamaterials feature primarily bending deformations that do not highly strain the base material, keeping strain magnitudes within the regime that a simple neo-Hookean model is reasonably accurate.

Simulations and shape derivative calculations are implemented using our custom quadratic triangle finite element simulation code developed in C++. We use SciPy’s nonlinear least squares solver [Viratanen et al. 2020] for (6). Our implementation is fully parallelized using TBB apart from collision area calculation, yielding multifold speedups for individual material design tasks on a multi-core system. However, for the large-scale material design reported in Section 7.1, we restricted each material design task to a single thread and ran them in parallel on an HPC cluster. In this single-threaded mode, the 6137 design optimizations reported on in the top row of Figure 14 each took 42 minutes on average on an Intel Xeon 8280 CPU. Of this time, an average of 27.5 minutes were spent solving (3), 3 minutes on computing collision areas, and 3.5 minutes calculating shape derivatives. Designing a single material for $\bar{\epsilon}_{\max} = 0.05$ takes roughly 13 minutes when run multithreaded on a Mac Studio with an M1 Ultra CPU.

As initial designs often behave poorly when stretched to large strains, we run our design optimization incrementally, starting at strain magnitude $\bar{\epsilon}_{\max} = 0.05$ and rerunning at increments of 0.05 until reaching strain magnitude $\bar{\epsilon}_{\max} = 0.15$. We use $w_e = 0.1$ and $w_\sigma = 1$ for all design optimizations. When penalizing collisions, we fix $w_c = 10^5$ for the results presented here in the main paper, and show in the supplement results for the higher penalty weight $w_c = 10^6$; increasing this weight further decreases the number and severity of collisions with only minor sacrifices to fitting quality.

6.1 Validation Methodology: Isotropic Hooke’s Laws

We systematically test our pipeline on the task of designing as-linear-as-possible isotropic materials. This is a challenging test case that can be attacked comprehensively and one for which performance is easy to visualize and interpret. The resulting material catalog also has direct applications for multi-scale design problems, e.g., for compliant mechanisms or soft robots.

We first ran the pipeline of [Panetta et al. 2017] to generate linearly elastic designs and determine a feasible set of target Hooke’s laws. That code optimizes designs on a regular grid in the $(\log(E), \nu)$ plane. For each feasible grid point, the linearly elastic design tool produces multiple designs with different topologies and shape parameters, and we use these both as the baseline for the comparisons




Fig. 11. Our experimental setup for both compression and tensile testing of small tilings of our planar microstructures.




Fig. 12. We rehomogenize our structures (which were optimized for a fabrication material with $\nu = 0.35$) at $\bar{\epsilon}_{\max} = 0.05$ using different settings for the base material’s Poisson’s ratio. We note the resulting homogenized properties are essentially unchanged, achieving similar closeness to the target properties (compare to the top left of Figure 13).

of Section 7.1 and as initial guesses for our finite-strain design optimization. To reduce the total number of design optimizations run, we manually picked a few topologies with high coverage—those visualized in Figure 8—and for each grid point, clustered the initial designs of the same topology based on their parameter vectors to remove redundancies (keeping each parameter vector only if it differs by more than 1% from the other kept vectors). Furthermore, because our focus is on designing *flexible* metamaterials, we removed target properties exceeding one tenth of the fabrication material’s stiffness. The resulting set of target Young’s moduli and Poisson’s ratios can be seen in Figure 1 and the scatter plots in Section 7.1

6.2 Fabrication and Testing Setup

We fabricated our structures using a multi-material 3D printer (Stratasys J35). Analogous to inkjet paper printers, this 3D printing technology jets up to three photopolymeric resins at the same time in a layer-by-layer fashion. A UV light cures the photopolymers after each layer of deposition. The novel rotational printing platform of the J35 reduces the degree of anisotropy in the elastic properties of the printed structures. As a result, we assume that the fabricated materials are transversely isotropic. All the physical specimens are fabricated using RGD8530, which is a digital material blend between VeroultraWhite (stiff polymer) and ElasticoBlack (elastomer). This material has a Young’s modulus of 1.1–1.7 GPa [Stratasys 2021] and a Poisson’s ratio of 0.4 [Zhang et al. 2018]. For the tensile testing,

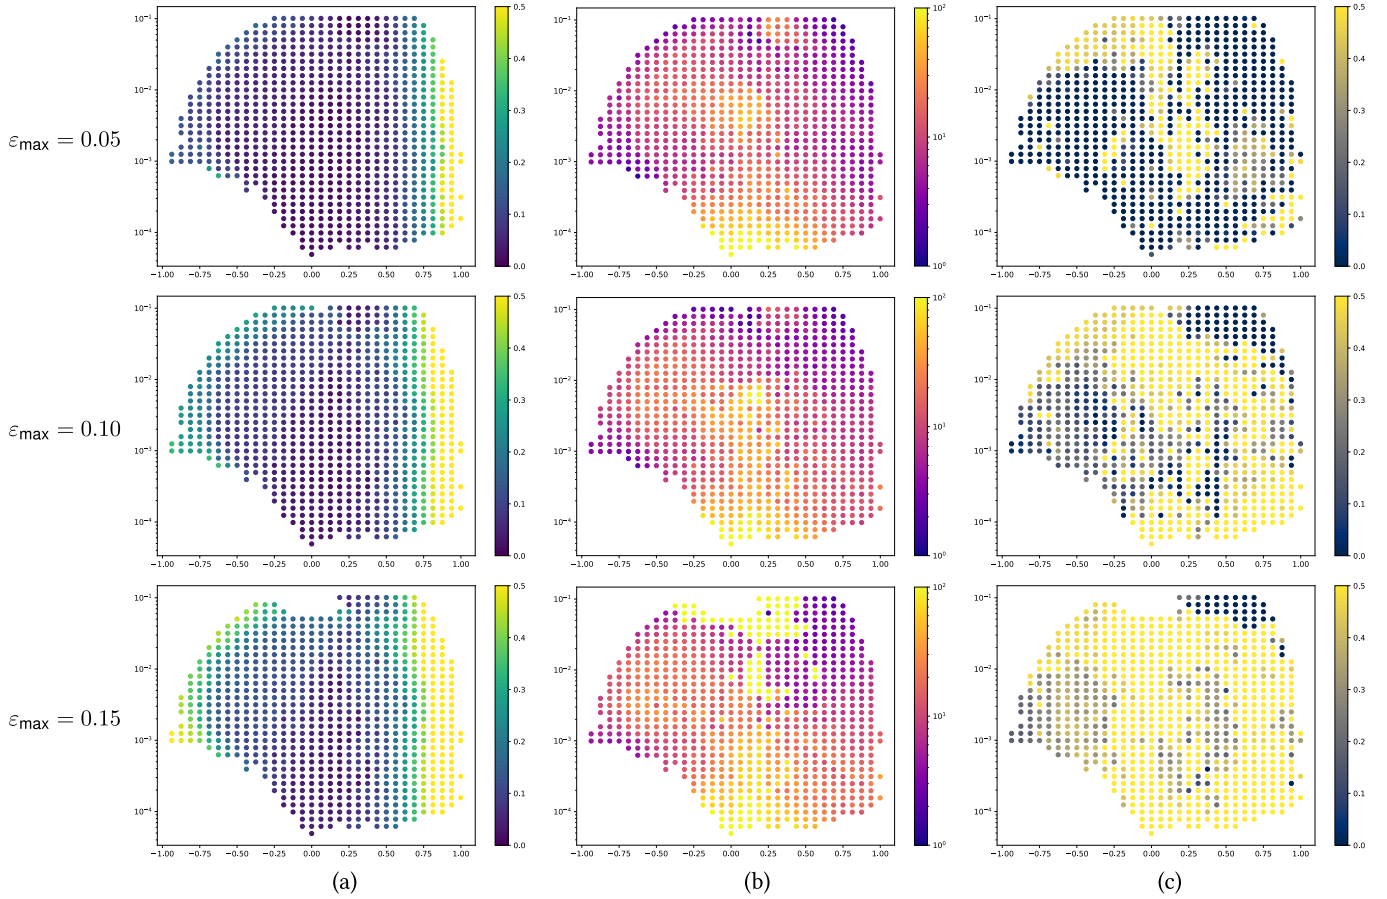


Fig. 13. Statistics for metamaterials designed *without penalizing collisions*. (a): maximum relative error in stress; (b): *improvement* in maximum relative stress error vs. linearly elastic design; (c) fraction of the strain ellipsoid volume exhibiting collisions. The vertical axis reports the target *relative* Young's moduli (with 1.0 corresponding to the full fabrication material's stiffness), and the horizontal axis reports target Poisson's ratio.

additional solid blocks are printed to both ends of the specimens for gripping. The microstructured unit cells are tiled either in a 3×3 or 5×5 arrangement. For these two arrangements, the edge lengths of the unit cells are set to 30 mm and 20 mm respectively. The thickness is set to 3 mm for all specimens. This approximates the plane stress condition while ensuring that the specimens do not buckle easily under compression.

We test the microstructures under uniaxial tension or compression using a universal testing system (Instron® 68SC-2) with a 500 N load transducer. The testing system allows the application of a prescribed displacement and the measurement of the corresponding reaction force (or vice versa). For tensile specimens, clamps are used to grip both ends of the specimens. Prescribed displacements are applied at these boundary edges parallel to the sheets while the lateral displacements are constrained to zero. For compression testing, a sleeve is fabricated using acrylic sheets to sandwich the specimens in between. Two additional acrylic sheets are used to apply prescribed compressive displacements to the specimens (Figure 11). Liquid lubricant is used to minimize the influence of friction. Here, the lateral displacements at the boundaries are unconfined.

All experiments are conducted under ambient conditions, and the loading rate is set to 3 mm/min. Given that room temperature is significantly below the glass transition temperature of the polymer ($\approx 60^\circ\text{C}$), we assume that the material is in its glassy phase.

7 RESULTS AND VALIDATIONS

7.1 Material Design

The results of our large-scale material design experiments described in Section 6.1 are summarized in Figure 13 and Figure 14. The first figure reports the results of running design optimization (6) without including our collision avoidance term. Consequently, a significant number of the designs feature collisions over a large fraction of their respective strain domains \mathcal{E} as seen in column (c). We clipped this collision fraction colormap to 0.5 as we expect collision to occur primarily in the compressive half-space (although surprisingly we found some patterns for which nearly 100% of the strain domain was inaccessible due to collisions).

A worst relative stress error over \mathcal{E} is computed by rehomogenizing the design using adaptive refinement (with $\epsilon = 10^{-2}$) and then

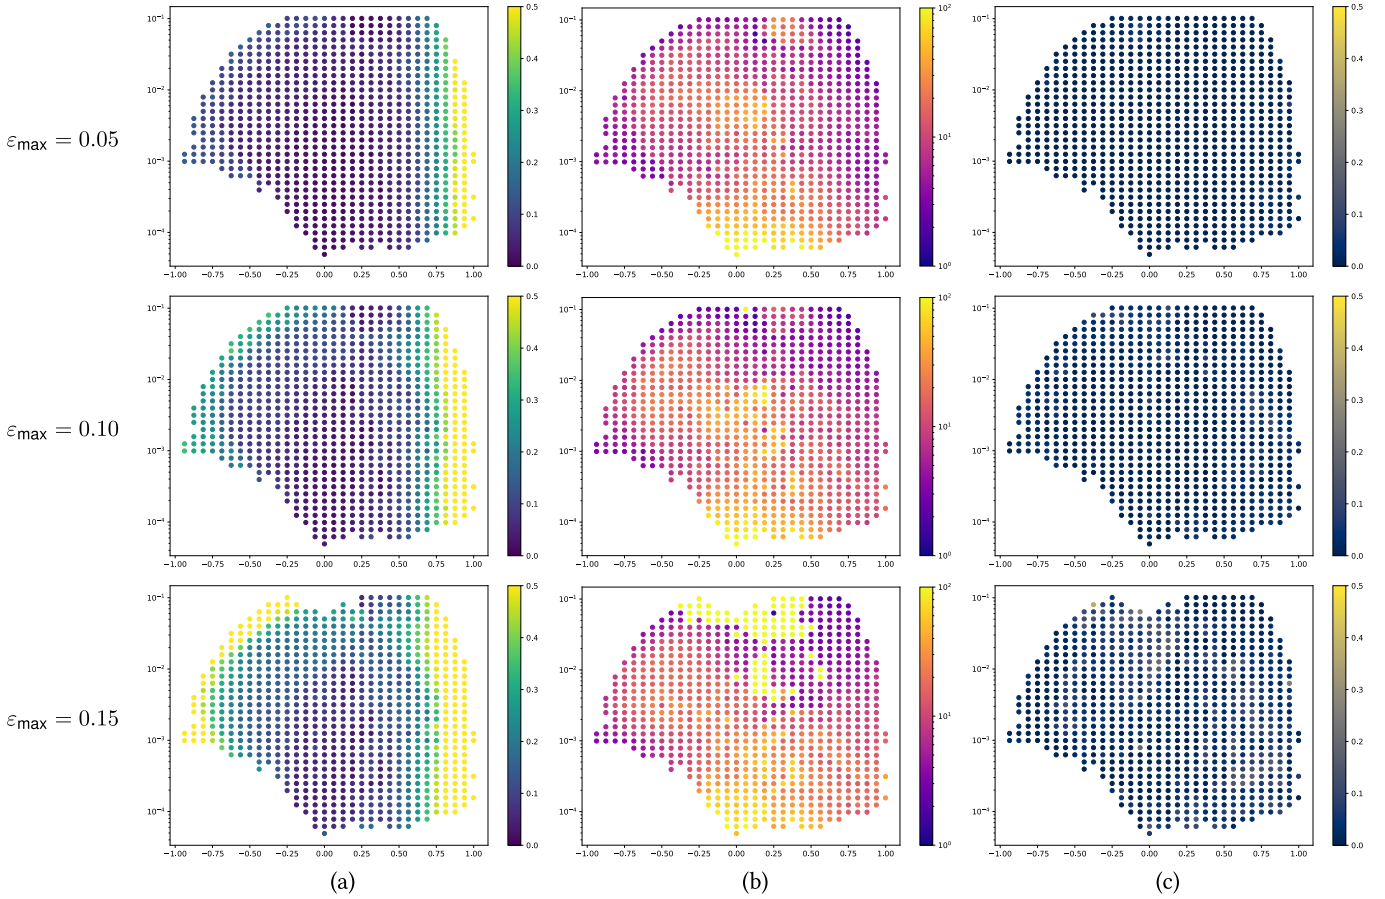


Fig. 14. Metamaterial design results *penalizing collisions*. (a): maximum relative error in stress; (b): *improvement* in maximum relative stress error vs. linearly elastic design; (c) fraction of the strain ellipsoid volume exhibiting collisions.

finding the maximum of $\|\bar{\psi}' - \bar{\psi}'_{\text{tgt}}\|_F / \|\bar{\psi}'_{\text{tgt}}\|_F$ over all vertices of the subdivided strain tetrahedralization $\bar{\mathcal{T}}$; the results are plotted in (a). We note that this measurement differs from the term optimized by (6): it uses the ℓ^∞ norm rather than the ℓ^2 norm. When designing for mild strains ($\bar{\epsilon}_{\text{max}} = 0.05$), we note that most material properties can be fit closely with the notable exception of high Poisson's ratios $\nu > 0.75$. To put this performance in context, we compare with the relative errors of the designs generated using a linear elasticity-based approach ([Panetta et al. 2017]): column (b) reports the ratio of the maximum relative stress error of the baseline designs (averaged over all original designs retained for that target property grid point by clustering) to the error achieved by our optimization. We note dramatic improvements, especially for more flexible materials (lower Young's moduli), where improvements reach 100 \times .

When designing for higher strain ($\bar{\epsilon}_{\text{max}} = 0.15$), we found that attempts began to fail for the stiffest target materials. These failures occur when the Newton solver for (3) is unable to converge within the allotted iteration limit. Homogenization is especially difficult for the poorly behaved baseline designs, and failures occurred even for designs that were able to optimize (due to our strategy of initializing the optimal designs for $\bar{\epsilon}_{\text{max}} = 0.15$ with the $\bar{\epsilon}_{\text{max}} = 0.10$

results); these are reported with the bright dots seen at the top of the improvement ratio plot.

We reran this entire design and analysis process including our collision avoidance term (8) with a fixed weight $w_c = 10^5$, reporting the results in Figure 14. We observe from the collision fractions (c) the high effectiveness of our collision removal strategy. Note that our term is a *soft* penalty, and so does not drive collision areas fully to zero. Therefore, when calculating collision fraction for this figure, we allow a negligible collision area of up to 10^{-5} in a deformation before counting it as colliding: we expect the work done by collision forces in eliminating tiny overlaps like these to be insignificant. Furthermore, since we used a fixed w_c across all optimizations, collision areas can grow commensurately with the fitting term of (6), explaining the mild reemergence of collisions at $\bar{\epsilon}_{\text{max}} = 0.15$.

Importantly, we emphasize that the optimization is generally able to eliminate collisions without substantial sacrifices to fitting quality (with the exception of targets at the upper-left of the material property space). We provide expanded visualizations in the supplement that offer more insight on the severity of remaining collisions as well as the effect of increasing w_c .




Fig. 15. Homogenization cost vs. requested accuracy for our Powell Sabin + BFS approach and alternatives. Especially when the tolerance for error is low, constructing our higher-order interpolant requires orders of magnitude fewer Newton iterations than a piecewise-linear interpolant.

7.2 Homogenization

We demonstrate in Figure 15 the speedup achieved by our use of C^1 interpolation and our BFS strategy when homogenizing to a specified level of accuracy (controlled by subdivision tolerance ϵ). We use the total number of Newton iterations taken in solving instances of (1) as an implementation-agnostic proxy for performance. We implemented a variant of our adaptive interpolant construction with a simpler piecewise linear interpolation (where only values of ω^* are interpolated, not its derivatives) and show that it requires orders-of-magnitude more iterations to converge (and demonstrates asymptotically worse performance with increased accuracy demands).

The improvement due to our BFS initialization strategy is less dramatic, but still reduces Newton iterations by roughly 33% in the lower-accuracy regime used during design optimization, even for the mild strain bound $\bar{\epsilon}_{\max} = 0.1$ used in this experiment. We note that the impact of this initialization speedup shrinks as higher accuracy is requested, when the Newton iterations are predominantly spent on adaptive refinement (for which we use the high-quality C^1 interpolant of the tetrahedron being subdivided in both strategies).

In Figure 12, we confirm stability of the homogenized properties with respect to perturbations of the input fabrication material. Using the designs produced by our optimization run for a fabrication material with $\nu = 0.35$, we rerun homogenization using $\nu = 0.3$ and $\nu = 0.4$ and compute the new deviation from the target properties. Because our flexible metamaterials' properties derive primarily from their geometries, the influence of this base material property is not significant.

7.3 Physical Experiments

We selected a few designs that are amenable to physical experiments, and we selectively apply tensile or compressive loading to a tiled grid of these designs. Since it is infeasible to replicate the periodic boundary condition physically, we choose to approximate uniaxial loading instead. To draw meaningful comparisons, and tease apart errors that might be contributed by the assumptions of periodic homogenization, numerical simulations of the finite tilings




Fig. 16. Results of physical experiments and their simulation counterparts. The simulations (solid curves) indicate significant nonlinearities in the original designs that are virtually eliminated by our optimization. The experimental data (dotted curves) show good agreement with the simulation and clearly demonstrate the improved linearity of our structures.

we fabricated with experimentally correct boundary conditions are performed. Under compression, it is shown that for the two representative designs each featuring an original and an optimized versions, the deformations are captured accurately throughout the entire loading sequence (Figure 16). Quantitatively, we plot the effective stress-strain response of these metamaterials. Since the experiments are displacement-controlled, the strain can be easily derived. The effective stress is simply the reaction forces divided by the enveloping cross-sectional area. The plots show that the initial slope of the curves remain similar. This is expected as the original designs were optimized to match the same target properties using linear elasticity [Panetta et al. 2017]. The curves diverge at larger strains however, with the optimized structures continuing to exhibit a linear response up to a strain of -0.1. The excellent agreement for the optimized designs between the experiment and simulation is made more remarkable as the constituent material (of RGD8530) was selected at random. This demonstrates the desired properties are largely influenced by the geometry of the microstructure rather than the material. For the original designs, a number of non-linear events (such as buckling) cause the experimental response to deviate from simulation.

To demonstrate that the 2D metamaterials are indeed transversely isotropic as designed, we re-apply the same displacements after a 90°




Fig. 17. Our parametric shape optimization framework and solver for (1) all work in 3D, but scaling our ω^* interpolant homogenization approach to 3D will require new algorithmic breakthroughs.

rotation. As seen in Figure 16, the behavior in both these directions are virtually indistinguishable *for the optimized structures*.

We emphasize that these physical tests only probe a single trajectory through \mathcal{E} and are limited due to the boundary conditions that are practical to apply in a lab setting; therefore, they fail to demonstrate the full extent of improvement achieved by our designs, especially since the worst nonlinearities of the original designs often occur under shearing conditions.

8 CONCLUSIONS AND FUTURE WORK

We have developed several new algorithms centered around the design of elastic metamaterials at finite strains, enabling for the first time a large-scale optimization of planar microstructures fitting target hyperelastic laws over entire regions of strains space. While our methods are able to generate performant designs across a wide range of target material properties in our challenging material design tasks, there are many exciting opportunities for improving and extending our work.

The most obvious extension is to tackle the microstructure design problem in 3D. While our framework can already solve for and tune the behavior of individual instances of (1) in 3D as visualized in Figure 17, our finite-region homogenization and design methods cannot be directly applied for volumetric metamaterials. One simple reason is that a Powell-Sabin interpolant has not been developed in dimensions higher than three (as needed for the 6D strain space); alternative interpolation strategies would need to be explored. Even if an off-the-shelf 6D simplicial Hermite interpolant were available, the increased strain dimension paired with the dramatically higher cost of each 3D elasticity simulation means a daunting increase in computational expense. We believe that breakthroughs in sampling and interpolation technologies are needed to make the full 3D design problem tractable.

Certain instances of volumetric metamaterial design may yet be within reach of our methods. For example, the space of *uniaxial strains* in 3D is just three dimensional: these strains can be specified by a unit direction and a signed magnitude. Furthermore, the region of strain space that must be sampled can be reduced by leveraging symmetries. For instance, a structure with cubic symmetry will have identical properties along each coordinate axis, meaning a more




Fig. 18. A visualization of the shape derivative evaluated at an optimal design in our parametric subspace. This vector field indicates the potential for improvement by enriching or departing from our parametric subspace.

compact (yet still six-dimensional!) subregion of strain space can be probed during homogenization. If a parametric geometry space can be constructed that guarantees isotropic behavior *over finite deformations*—as can be done for linear elasticity by enforcing six axes of five-fold symmetry [Christensen 1987]—then the probing domain can be reduced to just three dimensions (parametrized by principal strains).

If one wishes to *harness* the self-contact behavior of microstructures, our efficient collision avoidance strategy (Section 5.2) is no longer appropriate. Instead, the true physics of the contacts must be simulated and differentiated—and we anticipate this to be important to achieve highly nonlinear behaviors like jamming.

Despite the apparent simplicity of fitting as-linear-as-possible materials, the design problem we pose (6) is actually highly challenging due to the extreme constraints it imposes on the design. Better fitting quality could be achieved by addressing easier goals (like restricting to uniaxial strains). We expect our contributions, once suitably adapted, will still yield meaningful accelerations in these settings but leave this investigation for future work.

Our parametric shape optimization, though offering many advantages like enabling the use of high-order solvers for (6) and simplifying the imposition of fabrication constraints, is fundamentally limited to the subspace of designs representable by our parametric geometry definition. We note that the shape derivative remains nonzero after converging to an optimum in our parametric space, indicating the possibility of improving our results by enriching or departing from our design space (Figure 18).

We do not capture buckling modes whose periods span multiple unit cells. This is a fundamental challenge in nonlinear homogenization, as the number of spanned cells can be arbitrarily large and cannot be determined *a priori*. In fact, our reduced orthotropic homogenization method potentially worsens the problem since it assumes additional symmetries of the response. We leave systematic study of this phenomenon for future work. However, we have observed empirically that by optimizing structures to be as linear as possible in our single-cell periodic homogenization framework, we reduce the likelihood of buckling in finite tilings.

Finally, we would like to apply the metamaterial catalogs generated using our approach to build a multiscale design framework for nonlinear elasticity.

REFERENCES

- J r mie Allard, Francois Faure, Hadrien Courtecuisse, Florent Falipou, Christian Duriez, and Paul G. Kry. 2010. Volume Contact Constraints at Arbitrary Resolution. *ACM Trans. Graph.* 29, 4, Article 82 (jul 2010), 10 pages. <https://doi.org/10.1145/1778765.1778819>
- Moritz B cher, Bernd Bickel, Emily Whiting, and Olga Sorkine-Hornung. 2017. Spin-it: Optimizing moment of inertia for spinnable objects. *Commun. ACM* 60, 8 (2017), 92–99. <https://doi.org/10.1145/3068766>
- Changyeob Baek, Andrew O Sageman-Furnas, Mohammad K Jawed, and Pedro M Reis. 2018. Form finding in elastic gridshells. *Proceedings of the National Academy of Sciences* 115, 1 (2018), 75–80.
- Reza Behrou, Maroun Abi Ghanem, Brianna C. Macnider, Vimarsh Verma, Ryan Alvey, Jinho Hong, Ashley F. Emery, Hyunsun Alicia Kim, and Nicholas Boechler. 2021. Topology optimization of nonlinear periodically microstructured materials for tailored homogenized constitutive properties. *Composite Structures* 266 (6 2021). <https://doi.org/10.1016/j.compstruct.2021.113729>
- Amit H Bermann, Thomas Funkhouser, and Szymon Rusinkiewicz. 2017. State of the art in methods and representations for fabrication-aware design. In *Computer Graphics Forum*, Vol. 36. Wiley Online Library, 509–535.
- Katia Bertoldi, Vincenzo Vitelli, Johan Christensen, and Martin Van Hecke. 2017. Flexible mechanical metamaterials. <https://doi.org/10.1038/natrevmats.2017.66>
- Gaurav Bharaj, David I.W. Levin, James Tompkin, Yun Fei, Hanspeter Pfister, Wojciech Matusik, and Changxi Zheng. 2015. Computational design of metallophone contact sounds. *ACM Transactions on Graphics* 34, 6 (2015), 1–13. <https://doi.org/10.1145/2816795.2818108>
- Bernd Bickel, Moritz B cher, Miguel A. Otaduy, Hyunho Richard Lee, Hanspeter Pfister, Markus Gross, and Wojciech Matusik. 2010. Design and fabrication of materials with desired deformation behavior. (2010), 1. <https://doi.org/10.1145/1833349.1778800>
- Eric Brown, Nicholas Rodenberg, John Amend, Annan Mozeika, Erik Steltz, Mitchell R. Zakin, Hod Lipson, and Heinrich M. Jaeger. 2010. Universal robotic gripper based on the jamming of granular material. *Proceedings of the National Academy of Sciences* 107, 44 (2010), 18809–18814. <https://doi.org/10.1073/pnas.1003250107> arXiv:<https://www.pnas.org/doi/pdf/10.1073/pnas.1003250107>
- Paolo Celli, Connor McMahan, Brian Ramirez, Anton Bauhofer, Christina Naify, Douglas Hofmann, Basile Audoly, and Chiara Daraio. 2018. Shape-morphing architected sheets with non-periodic cut patterns. *Soft matter* 14, 48 (2018), 9744–9749.
- Antoine Chan-Lock, Jes s P rez, and Miguel A Otaduy. 2022. High-Order Elasticity Interpolants for Microstructure Simulation. In *Computer Graphics Forum*, Vol. 41. Wiley Online Library, 63–74.
- Desai Chen, David IW Levin, Shinjiro Sueda, and Wojciech Matusik. 2015. Data-driven finite elements for geometry and material design. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 1–10.
- Desai Chen, M lina Skouras, Bo Zhu, and Wojciech Matusik. 2018. Computational discovery of extremal microstructure families. *Science advances* 4, 1 (2018), eaao7005.
- Tian Chen, Julian Panetta, Max Schnaubelt, and Mark Pauly. 2021a. Bistable Auxetic Surface Structures. *ACM Trans. Graph.* 40, 4, Article 39 (July 2021), 9 pages. <https://doi.org/10.1145/3450626.3459940>
- Tian Chen, Mark Pauly, and Pedro M Reis. 2021b. A reprogrammable mechanical metamaterial with stable memory. *Nature* 589, 7842 (2021), 386–390.
- Xiang Chen, Changxi Zheng, Weiwei Xu, and Kun Zhou. 2014. An asymptotic numerical method for inverse elastic shape design. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 1–11.
- Francisco Chinesta, Roland Keunings, and Adrien Leygue. 2013. *The proper generalized decomposition for advanced numerical simulations: a primer*. Springer Science & Business Media.
- R. M. Christensen. 1987. Sufficient Symmetry Conditions for Isotropy of the Elastic Moduli Tensor. *Journal of Applied Mechanics* 54, 4 (12 1987), 772–777. <https://doi.org/10.1115/1.3173115> arXiv:https://asmedigitalcollection.asme.org/appliedmechanics/article-pdf/54/4/772/5459546/772_1.pdf
- Anders Clausen, Fengwen Wang, Jakob S. Jensen, Ole Sigmund, and Jennifer A. Lewis. 2015. Topology Optimized Architectures with Programmable Poisson’s Ratio over Large Deformations. *Advanced Materials* 27 (10 2015), 5523–5527. Issue 37. <https://doi.org/10.1002/adma.201502485>
- Qiaodong Cui, Timothy Langlois, Pradeep Sen, and Theodore Kim. 2020. Fast and Robust Stochastic Structural Optimization. In *Computer Graphics Forum*, Vol. 39. Wiley Online Library, 385–397.
- Francois Faure, J r mie Allard, Florent Falipou, and S bastien Barbier. 2008. Image-based collision detection and response between arbitrary volumetric objects. In *ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. Eurographics Association, 155–162.
- Ruslan Guseinov, Eder Miguel, and Bernd Bickel. 2017. CurveUps: Shaping objects from flat plates with tension-actuated curvature. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 1–12.
- Babak Haghpanah, Ladan Salari-Sharif, Peyman Pourrajab, Jonathan Hopkins, and Lorenzo Valdevit. 2016. Multistable Shape-Reconfigurable Architected Materials. *Advanced Materials* 28 (9 2016), 7915–7920. Issue 36. <https://doi.org/10.1002/adma.201601650>
- David Harmon, Daniele Panozzo, Olga Sorkine, and Denis Zorin. 2011. Interference Aware Geometric Modeling. *ACM Transactions on Graphics (proceedings of ACM SIGGRAPH ASIA)* 30, 6 (2011), 137:1–137:10.
- Alexandra Ion, Johannes Frohnhofen, Ludwig Wall, Robert Kovacs, Mirela Alistar, Jack Lindsay, Pedro Lopes, Hsiang-Ting Chen, and Patrick Baudisch. 2016. Metamaterial mechanisms. In *Proceedings of the 29th annual symposium on user interface software and technology*. 529–539.
- Lishuai Jin, Antonio Elia Forte, Bolei Deng, Ahmad Rafsanjani, and Katia Bertoldi. 2020. Kirigami-Inspired Inflatables with Programmable Shapes. *Advanced Materials* 32, 33 (2020), 2001863. <https://doi.org/10.1002/adma.202001863> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1002/adma.202001863>
- David Jourdan, M lina Skouras, Etienne Vouga, and Adrien Bousseau. 2022. Computational Design of Self-Actuated Surfaces by Printing Plastic Ribbons on Stretched Fabric. In *Computer Graphics Forum*, Vol. 41. Wiley Online Library, 493–506.
- Muamer Kadic, Timo B ckmann, Nicolas Stenger, Michael Thiel, and Martin Wegener. 2012. On the practicability of pentamode mechanical metamaterials. *Applied Physics Letters* 100, 19 (2012), 191901.
- Lily Kharevych, Patrick Mullen, Houman Owghadi, and Mathieu Desbrun. 2009. Numerical coarsening of inhomogeneous elastic materials. *ACM Transactions on graphics (TOG)* 28, 3 (2009), 1–8.
- Yujin Kim, Kuk Hui Son, and Jin Woo Lee. 2021. Auxetic structures for tissue engineering scaffolds and biomedical devices. *Materials* 14, 22 (2021), 6821.
- Hajer Lamari, Amine Ammar, Patrice Cartraud, Gr gory Legrain, Francisco Chinesta, and Fr d ric Jacquemin. 2010. Routes for efficient computational homogenization of nonlinear materials using the proper generalized decompositions. *Archives of Computational methods in Engineering* 17, 4 (2010), 373–391.
- Timothy Langlois, Ariel Shamir, Daniel Dror, Wojciech Matusik, and David IW Levin. 2016. Stochastic structural analysis for context-aware design and fabrication. *ACM Transactions on Graphics (TOG)* 35, 6 (2016), 1–13.
- Antoine Laurain and Kevin Sturm. 2016. Distributed shape derivative via averaged adjoint method and applications. *ESAIM: Mathematical Modelling and Numerical Analysis* 50 (7 2016), 1241–1267. Issue 4. <https://doi.org/10.1051/m2an/2015075>
- B. A. Le, J. Yvonnet, and Q. C. He. 2015. Computational homogenization of nonlinear elastic materials using neural networks. *Internat. J. Numer. Methods Engrg.* 104 (12 2015), 1061–1084. Issue 12. <https://doi.org/10.1002/nme.4953>
- Ryan H Lee, Erwin A B Mulder, and Jonathan B Hopkins. 2022. Mechanical neural networks: Architected materials that learn behaviors. <https://www.science.org>
- Minchen Li, Zachary Ferguson, Teseo Schneider, Timothy Langlois, Denis Zorin, Daniele Panozzo, Chenfanfu Jiang, and Danny M. Kaufman. 2020. Incremental Potential Contact: Intersection- and Inversion-free Large Deformation Dynamics. *ACM Trans. Graph. (SIGGRAPH)* 39, 4, Article 49 (2020).
- Lin Lu, Andrei Sharf, Haisen Zhao, Yuan Wei, Qingnan Fan, Xuelin Chen, Yann Savoye, Changhe Tu, Daniel Cohen-Or, and Baoquan Chen. 2014. Build-to-last: Strength to weight 3D printed objects. *ACM Transactions on Graphics* 33, 4 (2014). <https://doi.org/10.1145/2601097.2601168>
- Luigi Malomo, Jes s P rez, Emmanuel Iarussi, Nico Pietroni, Eder Miguel, Paolo Cignoni, and Bernd Bickel. 2018. FlexMaps: Computational design of flat flexible shells for shaping 3D objects. *SIGGRAPH Asia 2018 Technical Papers, SIGGRAPH Asia 2018* 37, 6 (2018), 1–14. <https://doi.org/10.1145/3272127.3275076>
- Andrew G. Mark, Stefano Palagi, Tian Qiu, and Peer Fischer. 2016. Auxetic metamaterial simplifies soft robot design. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*. 4951–4956. <https://doi.org/10.1109/ICRA.2016.7487701>
- Jon s Mart nez, J r mie Dumas, Sylvain Lefebvre, and Li Yi Wei. 2015. Structure and appearance optimization for controllable shape design. *ACM Transactions on Graphics* 34, 6 (2015). <https://doi.org/10.1145/2816795.2818101>
- Przemyslaw Musialski, Thomas Auzinger, Michael Birsak, Michael Wimmer, and Leif Kobbelt. 2015. Reduced-order shape optimization using offset surfaces. *ACM Transactions on Graphics* 34, 4 (2015). <https://doi.org/10.1145/2766955>
- Przemyslaw Musialski, Christian Hafner, Florian Rist, Michael Birsak, Michael Wimmer, and Leif Kobbelt. 2016. Non-linear shape optimization using local sub-space projections. *ACM Transactions on Graphics* 35, 4 (2016), 1–13. <https://doi.org/10.1145/2897824.2925886>
- Praveen Babu Nakshatrala, Daniel A Tortorelli, and KB Nakshatrala. 2013. Nonlinear structural design using multiscale topology optimization. Part I: Static formulation. *Computer Methods in Applied Mechanics and Engineering* 261 (2013), 167–176.
- Igor Ostanin, George Ovchinnikov, Davi Colli Tozoni, and Denis Zorin. 2018. A parametric class of composites with a large achievable range of effective elastic properties. *Journal of the Mechanics and Physics of Solids* 118 (2018), 204–217.
- Jifei Ou, M lina Skouras, Nikolaos Vlavianos, Felix Heibeck, Chin-Yi Cheng, Jannik Peters, and Hiroshi Ishii. 2016. aeroMorph-heat-sealing inflatable shape-change materials for interaction design. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. 121–132.
- Julian Panetta, Florin Ivoranu, Tian Chen, Emmanuel Siefert, Beno t Roman, and Mark Pauly. 2021. Computational inverse design of surface-based inflatables. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 1–14.

- Julian Panetta, Mina Konaković-Luković, Florin Isvoranu, Etienne Bouleau, and Mark Pauly. 2019. X-Shells: A New Class of Deployable Beam Structures. *ACM Trans. Graph.* 38, 4, Article 83 (July 2019), 15 pages. <https://doi.org/10.1145/3306346.3323040>
- Julian Panetta, Haleh Mohammadian, Emiliano Luci, and Vahid Babaei. 2022. Shape from Release: Inverse Design and Fabrication of Controlled Release Structures. *ACM Transactions on Graphics (TOG)* 41, 6 (2022), 1–14.
- Julian Panetta, Abtin Rahimian, and Denis Zorin. 2017. Worst-case Stress Relief for Microstructures. *ACM Trans. Graph.* 36, 4, Article 122 (July 2017), 16 pages. <https://doi.org/10.1145/3072959.3073649>
- Julian Panetta, Qingnan Zhou, Luigi Malomo, Nico Pietroni, Paolo Cignoni, and Denis Zorin. 2015. Elastic Textures for Additive Fabrication. *ACM Trans. Graph.* 34, 4, Article 135 (July 2015), 12 pages. <https://doi.org/10.1145/2766937>
- Jesús Pérez, Miguel A Otaduy, and Bernhard Thomaszewski. 2017. Computational design and automated fabrication of kirchhoff-plateau surfaces. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 1–12.
- Stefan Pillwein and Przemyslaw Musialski. 2021. Generalized deployable elastic geodesic grids. *ACM Transactions on Graphics (TOG)* 40, 6 (2021), 1–15.
- Ileana Pirozzi, Ali Kight, Rohan Shad, Amy Kyungwon Han, Seraina A. Dual, Robyn Fong, Allison Jia, William Hiesinger, Paul Yock, and Mark Cutkosky. 2022. RVEX: Right Ventricular External Device for Biomimetic Support and Monitoring of the Right Heart. *Advanced Materials Technologies* 7, 8 (2022), 2101472. <https://doi.org/10.1002/admt.202101472> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1002/admt.202101472>
- Romain Prévost, Emily Whiting, Sylvain Lefebvre, and Olga Sorkine-Hornung. 2013. Make it stand: Balancing shapes for 3D fabrication. *ACM Transactions on Graphics* 32, 4 (2013). <https://doi.org/10.1145/2461912.2461957>
- Ahmad Rafsanjani, Katia Bertoldi, and André R. Studart. 2019. Programming soft robots with flexible mechanical metamaterials. *Science Robotics* 4, 29 (2019), eaav7874. <https://doi.org/10.1126/scirobotics.aav7874> arXiv:<https://www.science.org/doi/pdf/10.1126/scirobotics.aav7874>
- Karthikayen Raju, Tong-Earn Tay, and Vincent Beng Chye Tan. 2021. A review of the FE² method for composites. *Multiscale and Multidisciplinary Modeling, Experiments and Design* 4 (2021), 1–24.
- Xin Ren, Raj Das, Phuong Tran, Tuan Duc Ngo, and Yi Min Xie. 2018. Auxetic metamaterials and structures: a review. *Smart materials and structures* 27, 2 (2018), 023001.
- Yingying Ren, Uday Kusupati, Julian Panetta, Florin Isvoranu, Davide Pellis, Tian Chen, and Mark Pauly. 2022. Umbrella meshes: elastic mechanisms for freeform shape deployment. *ACM Transactions on Graphics* 41, ARTICLE (2022), 1–15.
- Christian Schumacher, Bernd Bickel, Jan Rys, Steve Marschner, Chiara Daraio, and Markus Gross. 2015. Microstructures to Control Elasticity in 3D Printing. *ACM Trans. Graph.* 34, 4, Article 136 (July 2015), 13 pages.
- Christian Schumacher, Steve Marschner, Markus Cross, and Bernhard Thomaszewski. 2018a. Mechanical Characterization of Structured Sheet Materials. *ACM Trans. Graph.* 37, 4, Article 148 (July 2018), 15 pages. <https://doi.org/10.1145/3197517.3201278>
- Christian Schumacher, Jonas Zehnder, and Moritz Bächer. 2018b. Set-in-stone: Worst-case optimization of structures weak in tension. *SIGGRAPH Asia 2018 Technical Papers, SIGGRAPH Asia 2018* 38, 6 (2018). <https://doi.org/10.1145/3272127.3275085>
- Mélina Skouras, Bernhard Thomaszewski, Stelian Coros, Bernd Bickel, and Markus Gross. 2013. Computational design of actuated deformable characters. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 1–10.
- Yuanping Song, Robert M. Panas, Samira Chizari, Lucas A. Shaw, Julie A. Jackson, Jonathan B. Hopkins, and Andrew J. Pascall. 2019. Additively manufacturable micro-mechanical logic gates. *Nature Communications* 10 (12 2019), Issue 1. <https://doi.org/10.1038/s41467-019-08678-0>
- Georg Sperl, Rahul Narain, and Chris Wojtan. 2020. Homogenized Yarn-Level Cloth. *ACM Trans. Graph.* 39, 4, Article 48 (aug 2020), 16 pages. <https://doi.org/10.1145/3386569.3392412>
- Ondrej Stava, Juraj Vanek, Bedrich Benes, Nathan Carr, and Radomír Mvech. 2012. Stress relief: Improving structural strength of 3D printable objects. *ACM Transactions on Graphics* 31, 4 (2012). <https://doi.org/10.1145/2185520.2185544>
- Stratasys. 2021. Digital Materials Data Sheet. https://www.stratasys.com/siteassets/materials/materials-catalog/mds_pj_digitalmaterialsdatasheet_0122a-1.pdf?v=49c186
- Davi Colli Tozoni, Jérémie Dumas, Zhongshi Jiang, Julian Panetta, Daniele Panozzo, and Denis Zorin. 2020. A low-parametric rhombic microstructure family for irregular lattices. *ACM Transactions on Graphics (TOG)* 39, 4 (2020), 101–1.
- Erva Ulu, James McCann, and Levent Burak Kara. 2017. Lightweight structure design under force location uncertainty. *ACM Transactions on Graphics* 36, 4 (2017), 1–13. <https://doi.org/10.1145/3072959.3073626>
- Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, Ilhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, António H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. 2020. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* 17 (2020), 261–272. <https://doi.org/10.1038/s41592-019-0686-2>
- Fengwen Wang, Ole Sigmund, and Jakob Søndergaard Jensen. 2014. Design of materials with prescribed nonlinear properties. *Journal of the Mechanics and Physics of Solids* 69 (2014), 156–174.
- Yifan Wang, Liuchi Li, Douglas Hofmann, José E Andrade, and Chiara Daraio. 2021. Structured fabrics with tunable mechanical properties. *Nature* 596, 7871 (2021), 238–243.
- Yifan Wang, Brian Ramirez, Kalind Carpenter, Christina Naify, Douglas C. Hofmann, and Chiara Daraio. 2019. Architected lattices with adaptable energy absorption. *Extreme Mechanics Letters* 33 (2019), 100557. <https://doi.org/10.1016/j.eml.2019.100557>
- Ron Wein, Eric Berberich, Efi Fogel, Dan Halperin, Michael Hemmer, Oren Salzman, and Baruch Zukerman. 2023. 2D Arrangements. In *CGAL User and Reference Manual* (5.5.2 ed.), CGAL Editorial Board. <https://doc.cgal.org/5.5.2/Manual/packages.html#PkgArrangementOnSurface2>
- A.J. Worsey and B. Piper. 1988. A trivariate Powell-Sabin interpolant. *Computer Aided Geometric Design* 5, 3 (1988), 177–186. [https://doi.org/10.1016/0167-8396\(88\)90001-5](https://doi.org/10.1016/0167-8396(88)90001-5)
- Jorge Nocedal Stephen J Wright. 2006. Numerical optimization.
- Hongyi Xu, Yijing Li, Yong Chen, and Jernej Barbivc. 2015. Interactive Material Design using Model Reduction. *ACM Trans. on Graphics* 34, 2 (2015).
- Jonas Zehnder, Stelian Coros, and Bernhard Thomaszewski. 2016. Designing Structurally-Sound Ornamental Curve Networks. *ACM Trans. Graph.* 35, 4, Article 99 (jul 2016), 10 pages. <https://doi.org/10.1145/2897824.2925888>
- Jonas Zehnder, Espen Knoop, Moritz Bächer, and Bernhard Thomaszewski. 2017. Metasilicone: Design and Fabrication of Composite Silicone with Desired Mechanical Properties. *ACM Trans. Graph.* 36, 6, Article 240 (nov 2017), 13 pages. <https://doi.org/10.1145/3130800.3130881>
- Hang Zhang, Xiaogang Guo, Jun Wu, Daining Fang, and Yihui Zhang. 2018. Soft mechanical metamaterials with unusual swelling behavior and tunable stress-strain curves. *Science advances* 4, 6 (2018), eaar8535.
- Qingnan Zhou, Eitan Grinspun, Denis Zorin, and Alec Jacobson. 2016. Mesh Arrangements for Solid Geometry. *ACM Trans. Graph.* 35, 4, Article 39 (jul 2016), 15 pages. <https://doi.org/10.1145/2897824.2925901>
- Qingnan Zhou, Julian Panetta, and Denis Zorin. 2013. Worst-case Structural Analysis. *ACM Trans. Graph.* 32, 4, Article 137 (July 2013), 12 pages. <https://doi.org/10.1145/2461912.2461967>
- Bo Zhu, Mélina Skouras, Desai Chen, and Wojciech Matusik. 2017. Two-Scale Topology Optimization with Microstructures. *ACM Trans. Graph.* 36, 5, Article 164 (July 2017), 16 pages.
- Benliang Zhu, Xianmin Zhang, Hongchuan Zhang, Junwen Liang, Haoyan Zang, Hai Li, and Rixin Wang. 2020. Design of compliant mechanisms using continuum topology optimization: A review. *Mechanism and Machine Theory* 143 (2020), 103622. <https://doi.org/10.1016/j.mechmachtheory.2019.103622>