A Low-Parametric Rhombic Microstructure Family for Irregular Lattices

DAVI COLLI TOZONI, New York University

JÉRÉMIE DUMAS, Adobe Research, nTopology

ZHONGSHI JIANG, New York University

JULIAN PANETTA, EPFL

DANIELE PANOZZO, New York University

DENIS ZORIN, New York University

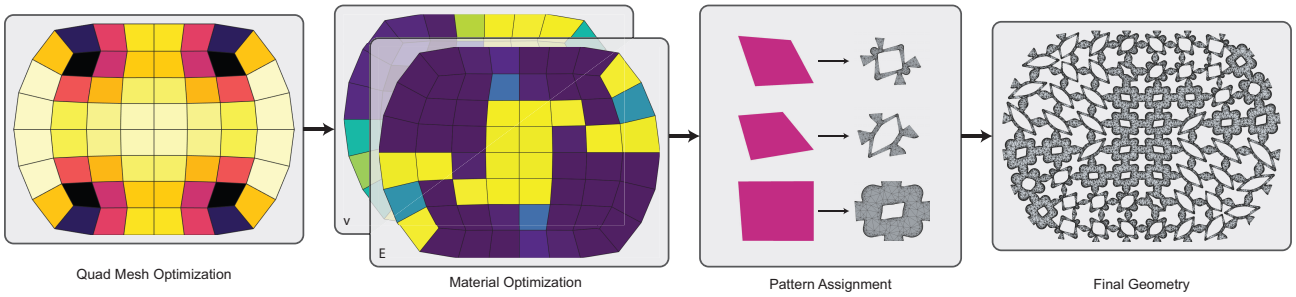

Fig. 1. The cells of a quadrilateral mesh are optimized to become quasi-rhombic; then material properties (variable Young's moduli and Poisson's ratio) are assigned to the cells. The assigned material properties are used to evaluate the geometric parameters of a tileable microstructure, encoded with a smooth spline map.

New fabrication technologies have significantly decreased the cost of fabrication of shapes with highly complex geometric structure. One important application of complex fine-scale geometric structures is to create variable effective elastic material properties in shapes manufactured from a single material. Modification of material properties has a variety of uses, from aerospace applications to soft robotics and prosthetic devices. Due to its scalability and effectiveness, an increasingly common approach to creating spatially varying materials is to partition a shape into cells and use a parametric family of small-scale geometric structures with known effective properties to fill the cells.

We propose a new approach to solving this problem for extruded, planar microstructures. Differently from existing methods for two-scale optimization based on regular grids with square periodic cells, which cannot conform to an arbitrary boundary, we introduce cell decompositions consisting of (nearly) rhombic cells. These meshes have far greater flexibility than those with square cells in terms of approximating arbitrary shapes, and, at the same time, have a number of properties simplifying small-scale structure construction. Our main contributions include a new family of 2D cell geometry structures, explicitly parameterized by their effective Young's moduli E, Poisson's ratios ν, and rhombic angle α with the geometry parameters expressed directly as smooth spline functions of E, ν, and α. This family leads to smooth transitions between the tiles and can handle a broad range of rhombic cell shapes. We introduce a complete material design pipeline based on this microstructure family, composed of an algorithm to generate rhombic tessellation from quadrilateral meshes and an algorithm to synthesize the microstructure geometry. We fabricated a number of models and experimentally demonstrated how our method, in combination with material optimization, can be used to achieve the desired deformation behavior.

CCS Concepts: • **Computing methodologies → Mesh models**.

Additional Key Words and Phrases: Additive Fabrication, Microstructures, Deformable Objects, Homogenization, Shape Optimization

ACM Reference Format:

Davi Colli Tozoni, Jérémie Dumas, Zhongshi Jiang, Julian Panetta, Daniele Panozzo, and Denis Zorin. 2020. A Low-Parametric Rhombic Microstructure Family for Irregular Lattices. *ACM Trans. Graph.* 39, 4, Article 1 (July 2020), 20 pages. https://doi.org/10.1145/3386569.3392451

Authors' addresses: Davi Colli Tozoni, New York University, davi.tozoni@nyu.edu; Jérémie Dumas, Adobe Research, nTopology, jeremie.dumas@ens-lyon.org; Zhongshi Jiang, New York University, jiangzs@nyu.edu; Julian Panetta, EPFL, julian.panetta@epfl.ch; Daniele Panozzo, New York University, panozzo@nyu.edu; Denis Zorin, New York University, dzorin@cs.nyu.edu.

0730-0301/2020/7-ART1 $15.00
https://doi.org/10.1145/3386569.3392451

1 INTRODUCTION

Advances in fabrication of highly complex geometry using additive manufacturing and other technologies resulted in new opportunities for shape design. In particular, *small-scale*, topologically and geometrically complex structures make it possible to achieve variable effective material properties using a single material for fabrication, including material properties not easily obtained by other means (such as negative Poisson's ratio) or material properties needed for precise control of deformation behavior of a shape. Realizing the potential of microstructures requires automatic generation and high-level control of the geometry, absent from commonly used geometric modeling tools.

A variety of approaches were developed for generating small-scale structures. Global methods, like topology optimization, are flexible but require very expensive computations to obtain high-quality results for fine-scale structures (cf. [Aage et al. 2017]). Other approaches partition the problem into different scales, using a variety of methods to generate a small-scale structure locally from a coarse-scale assignment of spatially-varying target properties. In this paper we present a method for design of small-scale structure families supporting an approach of the second type; specifically,

- Partition an input shape into quadrilateral cells, possibly with irregular connectivity; each cell is assigned target elastic properties;
- Assign to each cell a geometric microstructure, chosen from a family of such structures directly parameterized by their effective elastic properties and cell shapes.

The main goals for the choice of the microstructure family include: (1) cover a broad range of material properties, (2) be simple enough for fabrication at a small scale (i.e. avoiding thin features and small holes), (3) handle a range of cell shapes, (4) be easily tileable without significant modifications, (5) depend smoothly on the target effective elasticity tensor to avoid discontinuities in transitioning between varying material properties, and (6) be efficient to compute to enable tiling of large lattices.

The question of designing such microstructures is extensively studied in the literature, although advances in design of practical families covering large ranges of material properties are more recent. Existing techniques fall somewhat short of meeting the requirements above. First, these methods are exclusively based on periodic cell tilings with squares, regular triangles, or hexagons: only a single cell shape is used. This restricts the shapes that can be tiled to those constructed of this single cell type. Other shapes require either cutting cells, or deforming cells; both changes to the cell shape lead to a significant deviation from the intended deformation behavior (Section 8).

Most closely related previous work constructs large libraries of cell geometries [Panetta et al. 2015; Schumacher et al. 2015], with a separate shape or topology optimization performed for a dense set of values of material properties (Young's modulus and Poisson's ratio combinations for isotropic materials), and a specific base material. While this is generally adequate, if variable cell shapes are used, the size of the needed library will become prohibitively large (and still be limited to a discrete set of shapes). Furthermore, a different set of microstructures needs to be computed if some aspects of the properties of the base material change. These approaches do not provide the guarantee of smooth dependence of the resulting geometry on the material properties, and, as a consequence, a large number of samples are needed and interpolation between samples cannot be used.

In this paper, we address these problems for extruded, planar microstructures, aiming to meet the requirements listed above. The main aspects of our approach include the following.

- We choose to use *rhombic* cells, among different classes of tileable cells, since they have the flexibility to approximate arbitrary boundaries and can be connected in arbitrary orientations while being sufficiently simple to define microstructures for all members of the family in a compact way.
- We develop a parametric family of structures (Figure 2) completely described by: (1) Eight geometric parameters, defined as smooth spline functions $\mathcal{P}(E, \nu; \alpha)$ of material properties (Young's modulus E and Poisson's ratio ν in the isotropic case) and the cell shape parameter (rhombus angle α). (2) The domain of this parameterization in (E, ν) space is defined by three simple linear constraints $\ell(\alpha)$ dependent on the angle. This enables efficient optimization of material properties.
- Our family is tileable, i.e., the structures for adjacent cells connect with little or no modifications. Our tiles have matching topology and positions in the boundary. In addition, due to the smoothness of our mapping functions, if the transition of materials is smooth, the change in the connection region will also be smooth and almost no difference exists in the shape between neighboring cells.
- We demonstrate that the set of geometries is *universal*, i.e., can be used for any base material (although the ranges of realizable material properties do change).
- We present an algorithm for optimizing a quadrangulation of a planar domain to minimize the deviation of quads from the rhombic shape and demonstrate that this approach yields nearly-rhombic cells so that our microstructure family can be used.

We will make the data for computing microstructure geometric parameters, and constructing microstructure geometry from these parameters, publicly available.

2 RELATED WORK

We build on the foundation of previous work on construction of microstructure families and mesh quality optimization.

Periodic homogenization. Homogenization is a central tool in our construction. We use an extension of the FEM-based formulation used in [Panetta et al. 2017, 2015], which, in turn, goes back to [Allaire 2002], and is widely used in the literature. [Schumacher et al. 2018] presents an efficient homogenization approach tailored to rod structures. In these papers, regular lattices are used. We instead

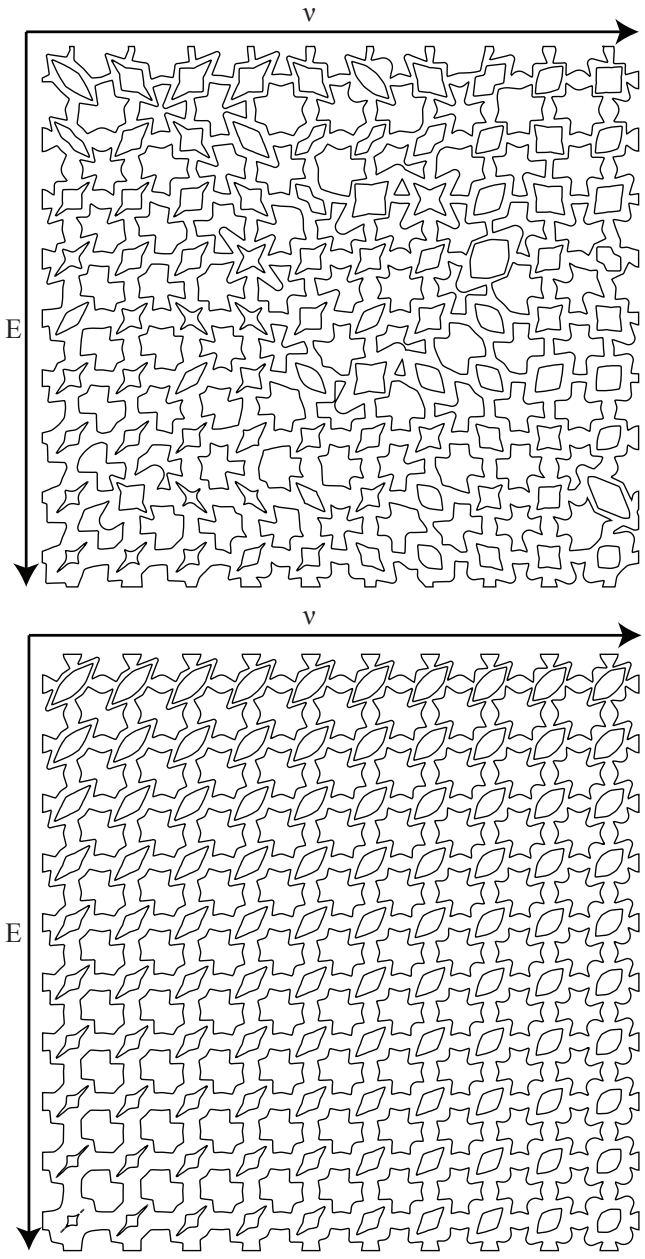

Fig. 2. An array of microstructures with Young's modulus E varying from 0.005 to 0.1 and ν varying from -0.15 to 0.5. Top: obtained by pointwise inverse homogenization of [Panetta et al. 2017]; Bottom: our family.

perform homogenization on a range of rhombic cell shapes, by transforming the problem to a domain where the cell is square, and replacing the target tensor with a transformed one. Transformations of the homogenized material properties when the base material is changed are considered in [Cherkaev et al. 1992].

Microstructure design and optimization. Many papers considered different aspects of microstructure design, see e.g., books [Allaire 2002; Cherkaev 2000; Cioranescu and Donato 1999; Milton 2002; Torquato 2002], and references in [Schumacher et al. 2015] and [Panetta et al. 2015], which represent topology and shape optimization approaches to microstructure optimization respectively. We refer to these methods as *pointwise inverse homogenization*, as they construct microstructures separately for each combination of effective material parameters.

Other examples of topology-optimization based work can be found in [Bendsøe 1989; Bendsøe and Sigmund 2003; Chen et al. 2018; Nakasone and Silva 2010]. Initially, mostly the problem of identifying extremal microstructures are considered (i.e., microstructures with properties at the boundary of the ranges that can be achieved). Recent works, e.g., [Ostanin et al. 2018; Panetta et al. 2015; Schumacher et al. 2015; Zhu et al. 2017] consider the question of constructing families spanning a broad range of elastic properties. In particular, [Ostanin et al. 2018] shows that near-optimal ranges of isotropic material behaviors can be achieved in 2D with hexagonal and triangular cells and a small number of microstructure parameters. We discuss the differences from this work in Section 4. [Milton et al. 2017] provides a characterization of achievable elastic tensors in terms of energies in 2D and 3D.

Most of the works designing microstructure families use an ad hoc approach to connecting structures corresponding to adjacent cells. [Garner et al. 2019] takes a more systematic approach for structures obtained using topology optimization by adding additional terms to the functional. A recent concurrent work [Martínez et al. 2019] introduces a metric that can be used to interpolate between a variety of microstructures on 2D regular grids, allowing to create smooth variation of material properties similar to our construction (but limited to regular lattices). [Konaković-Luković et al. 2018] uses a special type of 2D triangular auxetic structure to effect conformal surface deformations. This method requires domain meshing with triangles close to regular. Similarly, a recent paper [Malomo et al. 2018] uses 2D spiral microstructures for controlling deformation of sheets into a target shape. Varying geometric properties of spirals allows for a restricted control over material properties.

As an alternative to periodic microstructures, [Martínez et al. 2016; Martínez et al. 2017] construct randomized printable structures with control over Young's moduli both for isotropic and anisotropic target properties, but cannot independently control the Poisson's ratio. [Ion et al. 2016, 2019] describe a simple set of small-scale two-dimensional structures that can be assembled into mechanisms, and a computational design tool creating this type of mechanisms. While a range of properties can be obtained by changing the basic structure parameters their properties are difficult to control precisely.

A number of works apply the two-scale approach in a different way in 2D, using a simple rectangular cell structure and obtaining a directional field and scalar fields for geometric parameters defined on a regular grid. A field-aligned coarse mesh is extracted from the directional field and then filled with microstructures with parameters determined by the scalar fields [Gil-Ureta et al. 2019; Groen and Sigmund 2017; Groen et al. 2019]. Compared to these works,

we do not require the mesh cell orientations to be physically meaningful: in our case the effective material distribution and meshing are independent.

Global topology optimization. In [Aage et al. 2017; Liu et al. 2018; Wu et al. 2016], topology optimization was scaled up to high-resolution uniform and adaptive 3D grids. [Wu et al. 2018] performs high-resolution topology optimization with additional constraints to create an evenly distributed porous small-scale structure minimizing compliance for specific loading scenarios. The two-scale methods based on microstructures can always be combined with topology optimization to improve efficiency or resolution achievable in a given time.

Fabrication. [Hollister 2005; Kang 2010; Lin et al. 2004a,b] have demonstrated fabrication of optimized microstructures in the context of bone scaffold and fusion cage design. [Andreassen et al. 2014; Bückmann et al. 2012; Greaves et al. 2011; Schwerdtfeger et al. 2011] have shown the possibility of manufacturing auxetic materials. The idea of manufacturing objects with spatially varying properties using tileable structures also appears in [Hiller and Lipson 2009]. [Bickel et al. 2010] designs and fabricates objects satisfying an input deformation by optimizing for the best combination of stacked layers of their multi-material 3D printer's base materials. [Skouras et al. 2013] applies discrete material optimization to achieve desired deformations of complex characters with actuation, fabricating the results with multi-material printing.

Fabrication constraints. [Gaynor and Guest 2016; Langelaar 2016; Qian 2016] have made recent progress in incorporating undercut/overhang angle constraints in the topology optimization framework. However, these method enforce the constraints only approximately, requiring parameter tuning, and add nonlinearities to the problem that hinder convergence [Gaynor and Guest 2016]. [Panetta et al. 2015] enforces approximate printability via constraints on the skeleton defining the structure. In our planar case, there is no need to impose fabrication constraints other than requiring that the geometry should stay connected, and the minimal feature size is large enough for the chosen fabrication process.

Mesh optimization. The optimization of the element shape of discrete meshes has been studied in many disciplines. In the context of finite element simulations, the shape is optimized to reduce the distortion introduced by the geometric map [Brewer et al. 2003; Livesu et al. 2015]. Similarly, for texture mapping and quadrilateral meshing applications the distortion of a map from a surface to a plane is minimized by either evolving a Tutte's parameterization [Aigerman et al. 2014; Degener et al. 2003; Fu et al. 2015; Hormann and Greiner 2000; Kovalsky et al. 2016; Poranne and Lipman 2014; Sander et al. 2001; Schüller et al. 2013; Shtengel et al. 2017; Smith and Schaefer 2015; Sorkine et al. 2002; Zhu et al. 2017] or recovering from a possibly inverted initial guess [Aigerman and Lipman 2013; Fu and Liu 2016; Kovalsky et al. 2015; Lipman 2012]. These maps are also commonly used for mesh deformation applications [Bouaziz et al. 2012; Sorkine and Alexa 2007], and special constraints are often used in architectural geometry to generate meshes with planar faces [Bouaziz et al. 2014; Deuss et al. 2015; Jiang et al. 2015; Poranne et al. 2013; Tang et al. 2014].

We introduce an algorithm to optimize a quadrilateral mesh to have rhombic elements, a requirement that, to the best of our knowledge, has never been studied before. Our algorithm follows the paradigm introduced in ShapeUp [Bouaziz et al. 2012], alternating the projection to the space of rhombic shapes and a continuous optimization, and it relies on the solver introduced in [Rabinovich et al. 2017] to speed up convergence.

3 PROBLEM FORMULATION AND SOLUTION OVERVIEW

Problem. Our goal is to approximate a distribution of (potentially variable) material properties given by a spatially variable symmetric elasticity 4-tensor $\bar{C}(\mathbf{x})$ on a two-dimensional polygonal domain Ω, by partitioning it into constant-material cells, and assigning a microstructure to each cell. We assume that a single *base material* with known elastic properties given by the tensor C^{base} is used for all small-scale geometry.

The geometry of each cell Q is chosen so that its *homogenized* or *effective* elasticity tensor matches that of the assigned material $\bar{C}(Q)$.

Informally, the effective elasticity tensor can be understood as follows. If the cell Q were repeated periodically, a block of material consisting of sufficiently large number of such cells would behave as if it were made of a homogeneous material with elasticity tensor $\bar{C}(Q)$. In this work, we focus on isotropic target material distributions, for which the elasticity tensor is defined by a pair of parameters $M = (E, \nu)$, Young's modulus and Poisson's ratio.

We focus on isotropic materials since the effective orthotropic elasticity tensors that rhombic cells can produce have principal axes aligned with the diagonals of the rhombi; while for some purposes this broader space can still be used, the space of possible materials becomes mesh-dependent, complicating and restricting material optimization. In contrast, considering isotropic materials makes the result much more mesh-independent.

As mentioned in Section 1, our goal is to design the geometry in individual cells to have the following properties: (1) cover a broad range of material properties, (2) be simple enough for fabrication at small scale (i.e. avoiding thin features and small holes), (3) handle a range of cell shapes, (4) be easily tilable without significant modifications, (5) depend smoothly on the target effective elasticity tensor to avoid discontinuities in transitioning between varying material properties, and (6) be efficient to compute to enable tiling of large lattices.

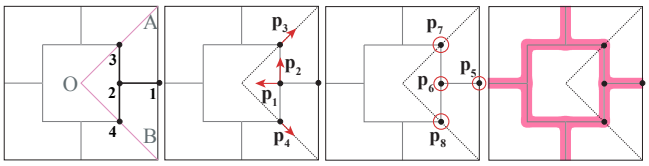

Fig. 3. Parameters defining our elementary cell geometry: 4 displacement parameters (two nodes 3,4 have only one free displacement and the node 1 is fixed, due to symmetry), and 4 radii, one per node. The displacements are specified in barycentric coordinates with respect to the triangle OAB.

Notation. We use *rhombic cells* R and a specific parametric cell geometry structure, which respects the two reflection symmetries of the rhombus. The shape of the rhombic cell is described by a single parameter, $\alpha \leqslant \pi/2$, the smaller angle of the rhombus.

The geometry of each cell is defined by a vector of parameters $\mathbf{p} \in \mathbb{R}^n$, $n = 8$, (Figure 3); the geometry is generated from parameter values using a 2D version of the method of [Panetta et al. 2017], described in Appendix A. The reasons for this choice are summarized in Section 4. We assume that for each parameter a range $[p_i^{min}, p_i^{max}]$ is provided. The n-dimensional box of admissible parameter values is denoted B_p. In particular, for *any* choice of parameters, the resulting homogenized elasticity tensor is *orthotropic*, and completely determined by four components, e.g., Young's moduli E_1, E_2 along two directions, shear modulus G, and one of two Poisson's ratios ν_{12}.

The function $H(\mathbf{p}; \alpha) : B_p \rightarrow \mathbb{R}^4$, the *homogenization function*, computes the effective elasticity tensor components of a cell with angle α from its geometry parameters. We describe a way to compute this function on arbitrary parallelogram cells, including rhombic, in Section 5. We use $H_E(\mathbf{p})$ and $H_\nu(\mathbf{p})$ to denote the Young's modulus and Poisson's ratio corresponding to $H(\mathbf{p})$ when it is isotropic.

Solution overview. Our approach is composed of two main components:

- We construct a parametric family of cell structures solving the *inverse homogenization problem* for a range $D(\alpha) \subset H(B_p; \alpha)$ of material properties $(E, \nu) \in D(\alpha)$, where E is the Young's modulus and ν is the Poisson's ratio, for a range of cell angles α, $[\alpha_{min}, \pi/2]$. This family is described by the *material-to-geometry* map $\mathcal{P}(E, \nu; \alpha) : D(\alpha) \rightarrow B^p$.
- We optimize a quadrangulation to obtain quasi-rhombic cells, so that the map $\mathcal{P}(E, \nu, \alpha)$ can be used to fill in the small-scale geometry.

The difficulty of the problem of inverse homogenization is due to several factors. The homogenization map $H(\mathbf{p})$ is straightforward, but expensive, to compute, because it requires several finite-element solves as a part of its computation. More fundamentally, for a fixed α, we need to invert the map $H\colon B_p \subset \mathbb{R}^n \rightarrow \mathbb{R}^4$ on a two-dimensional subspace Iso $\subset \mathbb{R}^4$ of isotropic materials. This is challenging, since the inverse is far from unique. Last but not least, the map depends on the base material properties C^{base}, in addition to the homogenized parameters and cell geometry.

We introduce a novel, smooth, closed form material-to-geometry map $\mathcal{P}(E, \nu; \alpha, C^{\text{base}})$ covering a broad range of material properties and rhombus angles α, that can be used for an *arbitrary* isotropic base material with tensor C^{base}. This map is uniquely determined by selecting a suitable four-dimensional subspace of the geometric parameter space $\mathbb{R}^n$, and can be represented in a very compact form by a set of 3D spline.

In the next sections, we discuss our choice of cell shape (Section 4), review homogenization (Section 5) and show how it can be extended to rhombic cells. Then we explain how we solve the inverse homogenization problem (Section 6), and how to compute a tiling with quasi-rhombic cells (Section 7).

4 CHOICE OF CELL SHAPE AND STRUCTURE

The choice of cell structure (i.e., topology of small-scale geometry, and parameters defining the geometry) is not unique: many choices may have similar behavior. We outline the heuristics we have used to select our structure for quadrilateral cells and briefly compare to alternatives.

Our choice of structure follows the general procedure outlined in [Panetta et al. 2015], applied to planar square cells. On the one hand, we want to minimize the number of parameters and topological complexity of the structure, as complex cells are difficult or impossible to manufacture. On the other hand, we want to maximize the *coverage* that can be achieved by the structure, that is, the range of material properties (E, ν). To obtain a rough prediction of coverage, we run a coarse sampling of geometric parameter space, evaluating the material parameters for each sample. In these sweeps, we enforce square symmetry (Figure 5), so that the resulting elastic tensor has only three free parameters (E, ν, G), and we look at the coverage we get in the projection to the (E, ν) plane as an indicator of what one can expect for inverse homogenization.

Among simplest topology patterns, the specific one we have chosen has the largest area covered. Expanding this area requires increasingly complex topology which negatively affects manufacturability.

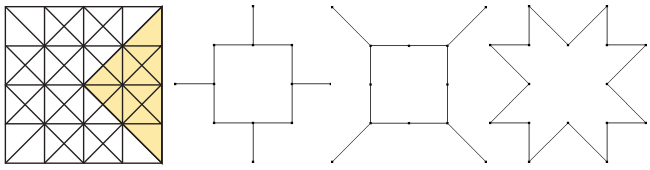

Fig. 4. We consider structures consisting of edges of the graph depicted on the left, with degrees of freedom consisting of radii at nodes and node displacements. Right: several examples of structures with few edges/simple topology.

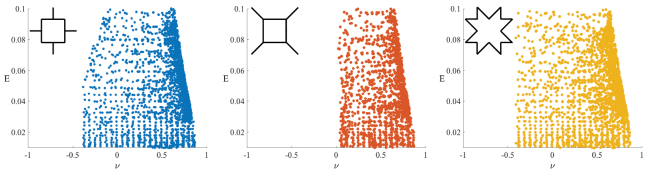

Fig. 5. Coverage in (E, ν) space for three cell structures, with smaller and larger coverage.

Choice of cell shape. As we would like to partition arbitrary shapes into cells, using square cells is not possible: in general, we cannot conform to an arbitrary boundary without introducing cells of other shapes. At the same time, the cells need to be close to periodically tileable, i.e., if we use quadrilateral cells, close to parallelograms. Parallelograms have only central symmetry, which does

not restrict the possible homogenized (averaged) elasticity tensors much; on the other hand, *rhombic* cells have two reflectional symmetries with respect to diagonals, and, as a result (Section 6), its elasticity tensor is orthotropic, as long as the small-scale geometry satisfies the same constraint. This considerably simplifies the construction of the material-to-geometry properties map, by decreasing the dimension of the space of possible materials.

As an alternative, regular triangular and hexagonal cells as shown in [Ostanin et al. 2018] can span a large range of material properties, and have an important advantage of being isotropic *by construction*, i.e., *any* geometry can be used on a cell as long as it has the symmetries of the cell (regular triangle or regular hexagon). While this significantly simplifies the problem of constructing geometries with target properties for regular cells, using distorted triangular or hexagonal cells as we need to tile an arbitrary shape, nullifies this advantage.

For triangular cells, we can enforce an additional symmetry by requiring the cells to be isosceles; in a periodic tiling, this is equivalent to using rhombic cells. At the same time, triangular cells are more restrictive in terms of choosing structure topology, if it needs to have the symmetry of the regular triangle. Finally, and perhaps most importantly, due to non-existence of regular tetrahedral tilings in 3D, methods based on quadrilateral tilings generalize better to 3D. For these, reasons, we choose rhombic cells.

To define the actual microstructure geometry, we use an implicit function parametrized by an edge skeleton vertices and radii, using the method of [Panetta et al. 2017], This method, on the one hand, allows for explicit control of topology through the skeleton connectivity, on the other hand, allows for merges of parts of the structure when these run into each other due to parameter choices, which is critical for robustness of the optimization process.

5 HOMOGENIZATION AND SHAPE OPTIMIZATION

We briefly review the standard homogenized tensor computation on square cells and show how it extends to rhombic cells. We also discuss how the inverse homogenization problem can be solved using optimization.

5.1 Homogenization and shape optimization on square cells

Our formulation follows [Panetta et al. 2015]. Suppose that a planar domain Ω is tiled by identical square cells Q, each with the same microstructure ω. At a large scale, we can consider averaged deformations: in the limit of infinitesimal cells, these deformations can be viewed as linear on each cell, and the actual deformation as a sum of the averaged one and a local fluctuation on the cell. The averaged deformation is the elastic deformation when the solid Ω is viewed as a completely filled volume with variable material properties defined by cell small-scale geometry, in the limit of zero cell size. The averaged deformation $\bar{\mathbf{u}}$ satisfies the *macroscopic* elasticity equation

$$-\nabla \cdot [\bar{C} : \varepsilon(\bar{\mathbf{u}})] = \bar{\mathbf{f}} \quad \text{in } \Omega, \tag{1}$$

where the elasticity tensor $\bar{C}$ is the effective elasticity 4-tensor with entries $\bar{C}_{ijkl}$ and the $\varepsilon(\mathbf{u}) := \frac{1}{2}(\nabla\mathbf{u}+(\nabla\mathbf{u})^T)$ is the linearized Cauchy strain tensor of the averaged deformation. The colon notation is used for contraction over the last two indices: $A : B = A_{ijkl}B_{kl}$, in the Einstein summation notation.

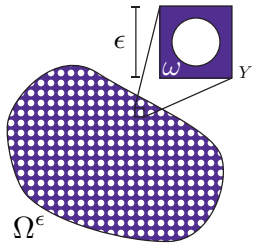

Fig. 6. A tiling of a domain Ω with base cell Y having geometry ω [Panetta et al. 2015].

On a single cell, in the limit of zero cell size in a periodic tiling, we can view elastic deformations as having a constant averaged strain $\varepsilon(\bar{\mathbf{u}})$ and a periodic "microscopic fluctuation" component $\mathbf{w}$, having zero average strain by periodicity. We consider fluctuations $\mathbf{w}$ for three constant average strains forming a basis for all possible constant strains $e^{11} = \mathbf{e}_1\otimes\mathbf{e}_1$, $e^{22} = \mathbf{e}_2\otimes\mathbf{e}_2$, and $e^{12} = \mathbf{e}_1\otimes\mathbf{e}_2+\mathbf{e}_2\otimes\mathbf{e}_1$. Then the expressions for the components of $\bar{C}$, are given by [Panetta et al. 2017]:

$$\bar{C}_{ijkl} = \frac{1}{|Q|}\int_\omega C^{\text{base}}_{ijpq}[\varepsilon(\mathbf{w}^{kl}) + e^{kl}]_{pq}\ \mathrm{d}\mathbf{x}. \tag{2}$$

To obtain the microscopic fluctuation term, the force balance equation in the cell Q is solved for each of the three basis strains:

$$-\nabla\cdot\left(C^{\text{base}} : [\varepsilon(\mathbf{w}^{kl}) + e^{kl}]\right) = 0 \quad \text{in } \omega, \tag{3a}$$

$$\hat{\mathbf{n}}\cdot\left(C^{\text{base}} : [\varepsilon(\mathbf{w}^{kl}) + e^{kl}]\right) = 0 \quad \text{on } \partial\omega \setminus \partial Q, \tag{3b}$$

$$\mathbf{w}^{kl}(\mathbf{x})\ Q\text{-periodic}, \tag{3c}$$

$$\int_\omega \mathbf{w}^{kl}(\mathbf{x})\,\mathrm{d}\mathbf{x} = \mathbf{0}. \tag{3d}$$

Inverse homogenization and shape derivatives. The conceptually straightforward, but technically difficult, way to solve the inverse homogenization problem for a given target tensor $\bar{C}^*$ is by shape optimization, i.e. by varying the geometry ω using shape parameters $\mathbf{p}$ as variables to minimize a functional. In the case of inverse homogenization, the functional penalizes the difference of the homogenized elasticity tensor and the target $\bar{C}^*$:

$$J(\mathbf{p}) = \|\bar{C}(\mathbf{p}) - \bar{C}^*\|^2 + wE^{reg}(\mathbf{p}), \tag{4}$$

with the norm taken to be e.g., the Frobenius norm. A combination of regularizing terms can be used as E^{reg}, the most common being the volume of the structure defined by $\mathbf{p}$ and the proximity to the initial solution [Panetta et al. 2017, 2015].

Unfortunately, due to non-uniqueness of the solution, resulting geometry parameters typically do not depend smoothly on the input material properties which has a number of negative consequences. We compare direct pointwise optimization to our approach in Section 8.

To minimize the functional (4) efficiently, we need derivatives of $J(\mathbf{p})$ with respect to the geometry parameters $\mathbf{p}$. Note that this requires differentiating $\bar{C}(\mathbf{p})$; the computation of $\bar{C}$ requires solving elasticity PDEs on a domain depending on $\mathbf{p}$. This type of derivatives are called *shape derivatives*; these are computed by solving an *adjoint* PDE, which in our case is identical to elasticity, but with redefined volume forces. The details of the computation can be found in [Panetta et al. 2017].

Discretization. To solve the PDEs (3), as well as similar equations needed for shape derivative computation, we use a standard FEM discretization. The domain is remeshed at each iteration using marching squares to extract the boundary and Triangle [Shewchuk 1996] to mesh the domain.

We discretize the cell problems with quadratic triangle elements, which we found essential for accurate stress and homogenized tensor evaluation. We use straight-edged elements (subparametric FEM) for representing the geometry to simplify meshing and the shape derivative formulas (edge nodes are placed at the edge midpoints).

In our experiments, running 2D homogenization usually took under 1s for meshes with an order of 4000 vertices. The mesh resolution (and hence the number of vertices) is determined by the maximal triangle area constraint (flag `-a` in Triangle) and the marching squares grid size (used to compute the boundary polygon), which we chose to be 5×10^{-4} and 256 respectively for most of our experiments. Figure 7 shows an example of the resolution we used.

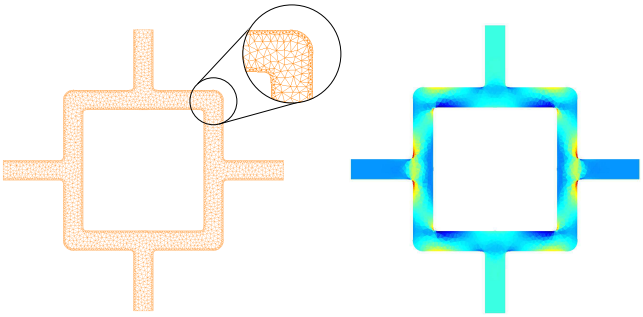

Fig. 7. Left: An example of the cell geometry discretization. Right: fluctuation strain norm obtained from (3) for the constant strain e^{11}.

5.2 Deformed structure homogenization

Computing the homogenization map $H(\mathbf{p})$ is the essential step in our process for obtaining a parametric family of structures. We use a modification of the standard homogenization method on square cells to handle arbitrary rhombi. Homogenization for an arbitrary rhombic cell (more generally parallelogram) can be transformed to homogenization on a square cell by a change of basis for elasticity tensors C^{base} and $\bar{C}$.

Suppose we have a cell R (rhombus with unit sides) and we want to write all our equations on a domain Q (unit square). Let F be the affine map $Q \to R$ and let its inverse be G.

Proposition 1. *Let $(\mathbf{e}_1, \mathbf{e}_2)$ be a nonorthogonal basis aligned with the sides of R; in this coordinate system, the domain is a unit square, i.e., the map F maps $(\mathbf{e}_1, \mathbf{e}_2)$ to the unit coordinate vectors. If displacements $\mathbf{u}$ satisfy the elasticity equation on domain R, $-\nabla \cdot [C : \varepsilon(\mathbf{u})] = \mathbf{f}$, then if $\tilde{\mathbf{u}}$ and $\tilde{\mathbf{f}}$ are the displacements and forces expressed in the coordinate system $(\mathbf{e}_1, \mathbf{e}_2)$, they satisfy the elasticity equation*

$$-\nabla \cdot [C' : \varepsilon(\tilde{\mathbf{u}})] = \tilde{\mathbf{f}}, \tag{5}$$

where the transformed tensor C' components are given by

$$C'_{ijkl} = G_{pi} G_{qj} G_{rk} G_{sl} C_{pqrs}. \tag{6}$$

The proposition is verified directly by the change of variables in (1) and (3).

Note that the components of C can be expressed in terms of the components of C' using the inverse relation:

$$C_{ijkl} = F_{pi} F_{qj} F_{rk} F_{sl} C'_{pqrs}. \tag{7}$$

This observation applies both to the microscopic equation on the domain in R filled with the base material, as well as to the homogenized equation. This leads to the following 3-step procedure for determining the homogenized properties of a rhombic cell:

- Transform the base material tensor C^{base} to the coordinate system $(\mathbf{e}_1, \mathbf{e}_2)$ using (6).
- Compute the homogenized tensor $\bar{C}'$ on Q, as described in Section 5.1.
- Transform $\bar{C}$ back to the original coordinate system, using (7).

This allows us to solve the *forward* problem: given the geometry in a rhombic cell, compute the homogenized elasticity tensor. As the problem is reduced to square-cell problems, shape derivatives are computed exactly in the same way, and the problem is discretized using the same approach.

We describe our solution to the inverse problem in the next section.

6 MATERIAL-TO-GEOMETRY MAP CONSTRUCTION

In this section, we explain the steps for constructing a map $\mathcal{P}$ from isotropic material parameters to geometry parameters. We assume in most of the exposition that C^{base} is fixed, and show how this assumption can be removed in Section 6.3. We use $E = 1$ and $\nu = 0$ for the base material, but the specific choice does not matter.

6.1 Orthotropy

The fact that the rhombic cells always have orthotropic elasticity tensor is of critical importance in our construction, since it determines the dimension of the image of $H(\mathbf{p}, \alpha)$. We consider this fact in more detail and define the measures of deviation from *isotropy* tailored for orthotropic materials.

The rhombus has two reflection symmetries with respect to its diagonals, and we consider geometries that by construction have the same symmetries. This means that the periodic structure obtained by repeating the elementary cell R is invariant with respect to these

two transformations. As a consequence, the elasticity tensor $\bar{C}$ is also invariant with respect to these transformations. This implies *orthotropy*:

Proposition 2. *[Love 1944] If an elasticity tensor is invariant with respect to reflections about two orthogonal axes, then it is orthotropic with respect to these axes and in the coordinate system aligned with these axes has the form (in Voigt notation)*

$$\begin{pmatrix} C_{11} & C_{12} & 0 \\ C_{12} & C_{22} & 0 \\ 0 & 0 & C_{33} \end{pmatrix} \tag{8}$$

This means, in particular, that only two additional constraints need to be satisfied to obtain an isotropic elastic tensor:

$$C_{11} = C_{22}; \quad 2C_{33} = (C_{11} + C_{22})/2 - C_{12} \tag{9}$$

We use two anisotropy measures, corresponding to these two constraints:

$$a_1 = \frac{|C_{11} - C_{22}|}{\max(C_{11}, C_{22})}, \quad a_2 = \frac{|(C_{11} + C_{22})/2 - C_{12} - 2C_{33}|}{\max(C_{11}, C_{22}, C_{33})}. \tag{10}$$

We note that principal directions of the tensor are exactly the diagonals of the rhombus, so cannot be set independently of the mesh directions.

6.2 Map construction for $\alpha = \pi/2$ (squares)

Recall that the material-to-geometry map $\mathcal{P}(E, \nu; \alpha)$ for a given α, is not uniquely defined; we describe a simple way to obtain a unique initial map on a smaller initial domain $D(\alpha)$, for $\alpha = \pi/2$, which we then expand. For clarity, we drop the dependence on α in this section.

We explain the general form of construction, to clarify how it can be used for other types of structures (e.g., 3D), and explain how it applies in the specific case of planar rhombic structures. Recall that $p = 8$ is the number of geometric parameters and we denote $m = 4$ the number of material parameters for 2D orthotropic materials.

We start with the homogenization map $H : B_p \to \mathbb{R}^m$, $p \geqslant m$. In general, a $p - m$-dimensional submanifold of B_p is mapped to each orthotropic elasticity tensor in $H(B_p)$. *Iso* corresponds to the set of possible isotropic materials in the space of material parameters $\mathbb{R}^m$, e.g., it is defined by equations $a_1 = 0$ and $a_2 = 0$ for orthotropic materials.

Our key idea is to restrict H to a carefully chosen m-dimensional affine subspace $V \subset \mathbb{R}^p$, for which H has large coverage, i.e. $H(V \cap B_p) \cap \text{Iso}$, the set of isotropic material properties covered by geometries in V, is not too far from $H(B_p) \cap \text{Iso}$. We refer to V as the *transversal subspace.*

$H' : V \to \mathbb{R}^m$, the restriction of H to V, is locally injective near all points where it is non-degenerate.

The map $\mathbf{p}^{\text{ortho}} : D = H(B_p) \to \mathbb{R}^p$, can be constructed by inverting H' as described below, and its restriction to isotropic materials $\text{Iso} \cap D$ yields the desired material-to-geometry map $\mathbf{p}$. Note that if non-isotropic orthotropic materials are desired, this map also can be used to compute these.

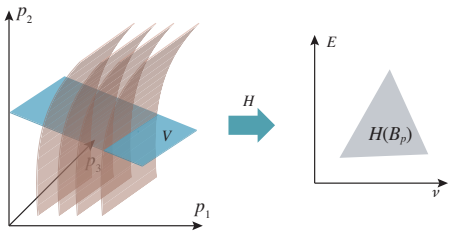

Fig. 8. Transversal subspace V illustrated for the simplified case of a 3-parametric structure and 2-dimensional property space (E, ν) (In our construction, we consider spaces that are 8 and 4-dimensional respectively). The curved surfaces are isosurfaces of $E(\mathbf{p})$. The map H, restricted to V, is bijective.

As the $H(\mathbf{p})$ depends on $\mathbf{p}$ smoothly, almost everywhere we expect $\mathcal{P}(E, \nu)$ to be smooth as well: the only places where a jump is possible is in the areas where the the differential of $H(\mathbf{p})$ is degenerate, or at the boundary of $H(B_p)$, which naturally defines the domain D for $\mathcal{P}(E, \nu)$. In practice, we have not observed the degeneracies of $\mathcal{P}(E, \nu)$ for the shape and material parameter ranges that we have considered, and the smooth approximation we obtain matches the samples closely.

Defining the transversal subspace V. We start with identifying a m-dimensional affine subspace containing a large range of geometric configurations with isotropic elasticity tensor, i.e., has large isotropic coverage. A direct approach to this would be to sample H sufficiently densely on B_p, select points for which the anisotropy measures a_1 and a_2 are small, and fit an m-dimensional hyperplane to this set, e.g., via PCA. Unfortunately, densely sampling the p-dimensional set (even for $p = 8$) is prohibitively expensive, as each evaluation of H requires meshing and solving three elasticity problems at sufficiently high resolution.

By considering a coarse sampling of B_p, we experimentally observed that, for our setting, B_p contains an *axis-aligned* hyperplane subset of dimension p', $p \geqslant p' \geqslant m$, which has good coverage. Intuitively, this corresponds to fixing the geometry parameters that can be fixed without decreasing the range of material properties; we refer to these as *redundant.* This reduction leads to a practical two-stage algorithm: first, remove the redundant parameters that have the least effect on coverage. Once no dimensions can be dropped, use denser sampling and PCA for the final reduction to m dimensions. In our case, $p' = 5$, so only one extra dimension needs to be eliminated.

Next, we discuss these two steps in greater detail:

(1) *Finding redundant parameters.* We compute a coarse regular grid sampling $q^j = H(\mathbf{p}^j)$ of the full homogenization map $H : B_p \to \mathbb{R}^m$. Each geometric parameter p^i is uniformly sampled at values c^i_k, $k = 1 \dots n_i$. In our case, a total of $540k$ data points were collected (see more information in Appendix B). Then

we extract a subset of $\mathbf{p}^j$ with the deviation from anisotropy sufficiently small (in our case 0.005 for both a_1 and a_2). Consider slices π_{ik} consisting of all samples $\mathbf{p}^j$ with a fixed value of $p_i^j = c_k^i$: by direct check of the coverage area of each slice, we identify the geometric parameters i for which there is a k such that coverage of π_{ik} is close to complete coverage, i.e., the coverage area obtained with the initial sweep, where we use all parameters as variables. Figure 9 shows the observed coverage for a sequence of expanding sets of fixed parameters.

For the rhombic structures, we observe that the offset parameters p_1, p_2, and the thickness parameter p_5 all have such values c_k^i (0.5, 0.25, and 0.3 respectively). This is a particularly convenient set of parameters to fix, since this allows us to merge cells *perfectly* as the radius on the boundary is fixed and thus the same for all cells. Fixing p_1, p_2, and p_5 reduces the dimension

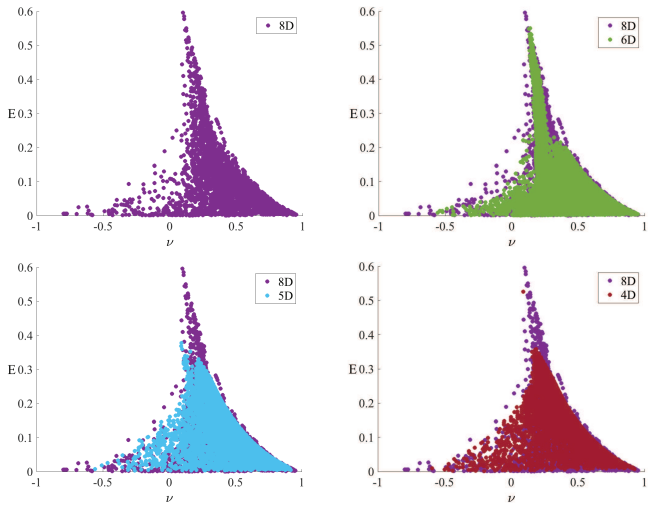

Fig. 9. Material properties coverage for the subset of samples of B_p with constant values of p_1 and p_2 fixed (6D space, samples shown in green), p_1, p_2, p_5 fixed (5D space, samples shown in blue) and V (4D space, shown in red), compared to the full set of samples $\mathbf{p}_j$, shown in purple.

of the geometry space to just 5, leading to a 5D affine subspace V^{5D}. It remains to reduce the dimension by 1.

(2) *Reduce to m dimensions by PCA.* The restriction of the map H to $V' \cap B_p$, where V' is the p'-dimensional subspace obtained by dropping dimensions, can be sampled much more densely (we use 15 samples per dimension for 5-dimensional V'). To obtain the m-dimensional transversal space V, we approximate $H(\mathbf{p})$ with a linear map on B_p: $H(\mathbf{p}) \approx Q\bar{\mathbf{p}} + H_0$, where Q is a $m \times p'$ matrix, and $\bar{\mathbf{p}}$ is obtained from $\mathbf{p}$ by discarding the fixed coordinates. We compute Q and H_0 using PCA. Let W be the matrix spanning the nullspace of Q, i.e., $QW = 0$, and W has dimension $p' \times (p' - m)$ and maximal rank. For all points $\bar{\mathbf{p}} + Wu$, where u is $(p' - m)$-dimensional, the value of the map is constant. Then the condition $W^T(\bar{p} - \bar{\mathbf{p}}_0) = W^T\bar{\mathbf{p}} + c_0 = 0$, where $\bar{\mathbf{p}}_0$ is a point in V', defines an m-dimensional subspace $V \subset V'$, perpendicular to the constant-value affine spaces, on which $Q\bar{\mathbf{p}} + H_0$ is one-to-one. For rhombic structures, we obtain $p' - m = 1$, and W has dimensions $p' \times 1$, i.e., it is a vector in this case. W is approximated well by a vector with nonzero components $v_3 = 1$ and $v_7 = 1$, and with $c_0 = -0.82$.

This gives us a complete description of $V = \{\mathbf{p} \mid p_1 = 0.5, p_2 = 0.25, p_5 = 0.3, p_8 = 0.82 - p_4\}$. Geometrically, V is described as the space of structures where node 2 is fixed at the default position at barycentric coordinates $(0.5, 0.5, 0)$, the radius of node 1 connecting the cell to other cells is fixed at 0.3, and the radius and displacement of node 4 are related linearly.

Inverse of H' on the isotropic subspace. First, we construct a piecewise-linear interpolant of the samples of H'. The map H' is defined as the restriction of H to $V \cap B_p$. We sample H' on a regular grid in V. For the rhombic structure, we use p_3, p_4, p_6, p_8 as coordinates for sampling. We sample at the points of the grid $q_j = (E_j, \nu_j)$ contained inside B_p.

As samples q^j are on a regular m-dimensional grid, they form mD cubes, which we then subdivide into m-simplices. We discard all simplices with one of the two possible orientations (as H' is not guaranteed to be injective everywhere, simplices may have both orientations; we choose the predominant orientation to be the "correct" one). The map is defined by linear interpolation on each simplex.

For each sample q on a regular 2D grid of samples in Iso, we search for m-simplex S_ℓ in V, such that $H(S_\ell)$ contains q. While the simplices found in this way may be not unique, in practice we do not observe this problem. The value of $\mathcal{P}(q)$ is obtained by interpolating the values of $\mathbf{p}$ at the corners of each simplex S_ℓ. Finally, we fit a B-spline to the sampled values of q with Laplacian regularization. We use 20×20 control points for rhombic structures.

Increasing coverage of the material-to-geometry map. Constraining the map H to V makes the map invertible, and we observe the inverse to be smooth. However, this also restricts the coverage, as the inverse image $H^{-1}(Iso \cap H(B_p))$ in B_p is not necessarily contained in V. We can expand the range by the following procedure using optimization-based inverse homogenization, similar to [Zhu et al. 2017].

For each regular-grid sample point q_j in (Iso), for which $\mathcal{P}(q_j)$ is not defined, but which is adjacent to a point q_ℓ where $\mathcal{P}$ is defined, we initialize the shape optimization for the functional (4) using the spline approximation to $\mathcal{P}$ obtained above evaluated at q_j. Then we optimize using all 8 geometry parameters as variables (and not only the 4 independent variables used in the sampling procedure), thus taking the value of $\mathcal{P}(q_j)$ out of V. If the optimization converges to a value sufficiently close, using a tolerance of 0.005 for Poisson's ratio and 0.001 for Young's modulus, we include the additional point in the set, and refit the splines to the new set of points.

We note that, in principle, this process may suffer from the same flaws as the original inverse homogenization of [Panetta et al. 2015] (cf. Figure 10) as the map $\mathcal{P}(E, \nu; \alpha)$ may become non-smooth for lower values of α. However, the procedure defined above prevents

this from happening, and the map $\mathcal{P}(E, \nu; \alpha_{min})$ remains smooth (Figure 18).

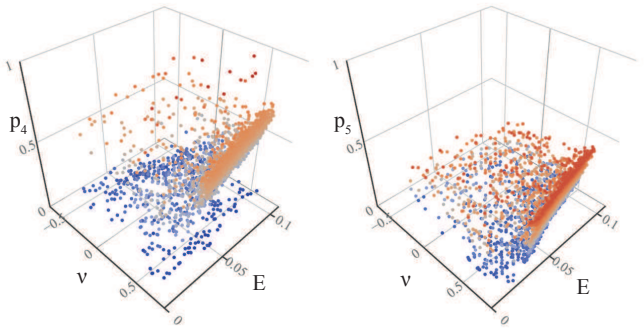

Fig. 10. Plots of parameters p_1 and p_4 for pointwise inverse homogenization. Note the lack of smoothness; other parameters also have a varying degree of noise. The color in the charts represents the geometric parameter value, varying from blue to red.

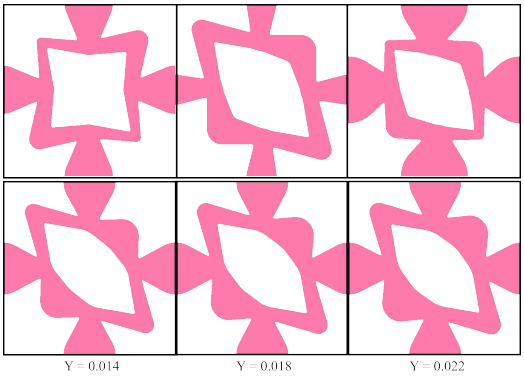

Fig. 11. Changes of geometry as a result of small changes in Young's modulus Top: obtained by pointwise inverse homogenization [Panetta et al. 2017]; Bottom: our family.

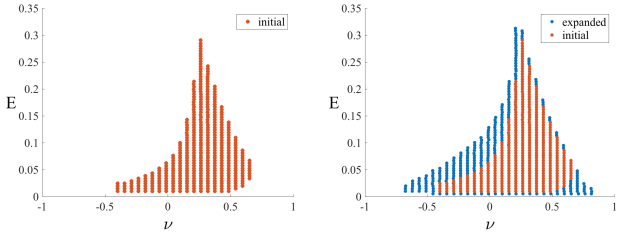

Fig. 12. Expanding the initial material-to-geometry map coverage. Left: original coverage; Right: expanded coverage.

Bounds for the domains $D(\alpha)$. For applications using optimization of material properties on a partitioned domain (Section 8), it is useful to ensure that the values used in the optimization stay within the coverage zone of the material-to-geometry function **p**.

To keep these constraints efficient, we approximate the coverage area by a convex polygon bounded by a set of a maximum of 6 lines: 2 horizontal lines defining minimum and maximum Young's modulus E_{min} and E_{max}, 2 vertical lines for minimum and maximum Poisson's ratio ν_{min} and ν_{max}, and 2 additional slanted lines, s_l and s_r, that allow for larger non-rectangular domains (see Figure 13).

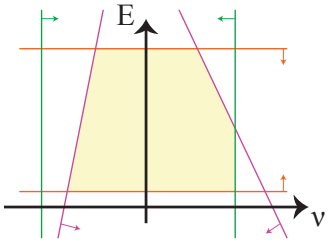

Fig. 13. Lines defining the shape of our convex domain.

The lines are defined to maximize a weighted sum of the area and the height (maximum Young's modulus) of the convex polygon while keeping it completely inscribed in the domain coverage region. Below, we explain this process in more detail.

In our coverage data, we first identify two sets of points on the boundary of our covered area. For each value of Young's modulus, the minimum achieved Poisson's ratio is added to B_l, while the maximum should be in B_r. Then, we iteratively run the following two steps, trying different values for lines ν_{min}, ν_{max}, E_{min} and E_{max} at each time:

(1) *Filtering.* We filter out boundary points in B_l and B_r that are outside the square defined by ν_{min}, ν_{max}, E_{min} and E_{max}, creating new boundary sets B'_l and B'_r.

(2) *Optimization.* We optimize the slanted lines as follows. We create one linear constraint for each boundary point b in B'_l (resp. B'_r), making sure that slanted line s_l (resp. s_r) is to the right (resp. left) of b. After building our constraint matrix, we run the `fmincon()` solver in MATLAB to optimize our objective function (composed of area and height of the convex domain) given our new linear constraints.

At the end, the objective function values obtained for all iterations are compared and we select the 6 lines corresponding to the best value achieved. In our own case, we decided to fix minimum and maximum Young's modulus (E_{min} and E_{max}) to 0.005 and 0.32, due to our observed coverage. However, we tried multiple values for maximum and minimum Poisson's ratio (ν_{min}, ν_{max}). Notice that these values can also be chosen according to the application.

6.3 Arbitrary angles and base materials

We extend the map to all angles in the range $[\alpha_{min}, \pi/2]$. We use $\alpha_{min} = \pi/4$ since the coverage becomes very small below this value. In our experiments, the range $[\pi/4, \pi/2]$ proved to be sufficient, as, after quad optimization, the minimum angle for quads was larger than $\pi/4$. It is however possible that for complex geometries smaller angles are needed.

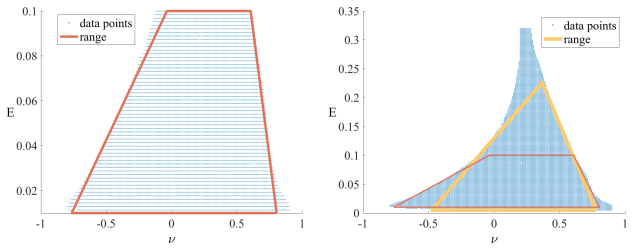

Fig. 14. Domain coverage for $\alpha = \pi/2$, with approximating linear bounds. Left: resulting polygon when restricting minimum and maximum Young's modulus to 0.01 and 0.1; Right: convex polygon (in yellow) when Young's modulus minimum and maximum values are unlimited. We show a comparison with the bounds in the constrained case (in red). Depending on the Young's modulus range the application needs, a different polygon may be obtained in the optimization. If the application requires larger stiffness values, the polygon on the right should be preferred. On the other hand, if a larger range on Poisson's ratio values is needed, the polygon in red is a better option.

We perform the extension incrementally for a sequence of angles $\alpha_i = \pi/2 - i\,d$, where $d = 1.25°$. Using the geometric parameters for α_{i-1} as a starting point, we use shape optimization to obtain a point for the same target value (E, ν) but for α_i, with a regularization term $\|\boldsymbol{p} - \boldsymbol{p_0}\|^2$ penalizing deviation from the initial value of each parameter, where $\boldsymbol{p}$ and $\boldsymbol{p_0}$ represent the current and the initial parameter values during the optimization, respectively.

Once the whole range of angles is covered, we fit a 3D spline in variables (E, ν, α) to the whole set of points. The resulting change of geometry parameters as a function of α is smooth (Figure 15). For spline fitting, we use least-squares on the whole cube. An extra regularization term, the Laplacian of the function at grid points, is added in the cost, which promotes smoothness and eliminates the null space corresponding to the regions with a few or no data points.

Complete map description. The final map $\mathcal{P}$ is defined as:

- a 3D uniform spline approximation on a $12 \times 12 \times 12$ grid of 8 geometric parameters, for a total of 8×12^3 coefficients;
- a 1D splines in α for each of the coefficients defining the lines of the convex polygon approximating the domain coverage region $D(\alpha)$.

The complete description of the map is relatively compact, requiring less than 14k coefficients. The plots for the eight components are shown in Figure 17 for $\alpha = \pi/2$, while Figure 18 shows plots for two of the parameter when $\alpha = \pi/4$. In addition, Figure 11 illustrates (non-)smoothness of the map obtained by pointwise shape optimization, compared to our map, for small changes of E and ν for three choices of (E, ν).

Change of the base material. In principle, we would need to compute a new map $\mathcal{P}$ each time the base material changes, as the homogenization map H depends on the choice of material. However,

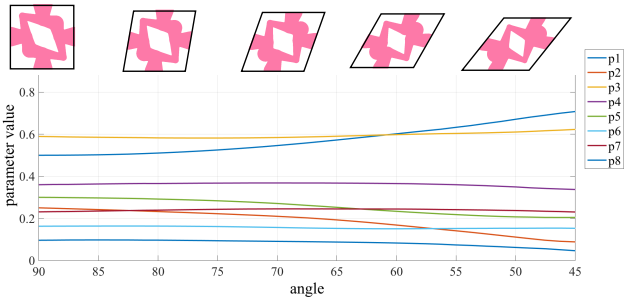

Fig. 15. Dependence of the geometry parameters $\mathbf{p}$ on the angle, for $E = 0.076$ and $\nu = 0.263$ (using base material $E = 1$, $\nu = 0$).

two observations show that $\bar{C}$ for any base material (E', ν') can be obtained from a base material (E, ν) by a simple transformation.

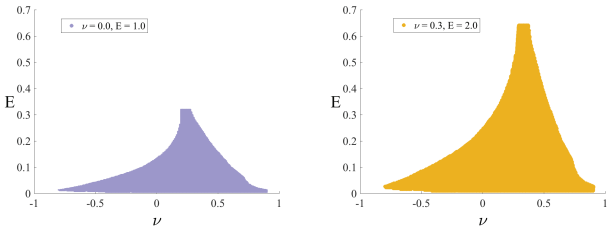

Fig. 16. Left: coverage for $\nu = 0$, $E = 1$. Right: coverage for $\nu = 0.3$, $E = 2$.

The first simple observation is that the expressions for the components of $\bar{C}$ and the solution for (3) depend linearly on the base elastic tensor C^{base}, i.e., if the base Young's modulus E is replaced with sE, the homogenized one E^H becomes sE^H.

The less trivial part are the changes due to changes in the Poisson's ratio. Fortunately, if ν is replaced with ν', the new homogenized elasticity tensor $\bar{C}'$ is characterized by the CLM theorem, a corollary of which for isotropic materials can be formulated as follows.

Theorem 6.1. *[Cherkaev et al. 1992] Suppose a periodic 2D medium with holes is filled with material with Young's modulus E and Poisson's ratio ν, and the effective (homogenized) properties of this medium are E_H and ν_H. Suppose ν is replaced with ν'. Then the effective Young's modulus is preserved, and the Poisson's ratio changes as*

$$\nu'_H = \nu_H - (\nu - \nu')\frac{E_H}{E_b}.$$

This immediately leads to the change of variables in the function $H(E, \nu)$ that corresponds to the change of base material properties $(E, \nu) \to (E', \nu')$:

$$\begin{aligned} E'_H &= \frac{E'}{E} E_H, \\ \nu'_H &= \nu_H - (\nu - \nu')\frac{E' E_H / E}{E'} = \nu_H - (\nu - \nu')\frac{E_H}{E}. \end{aligned} \tag{11}$$

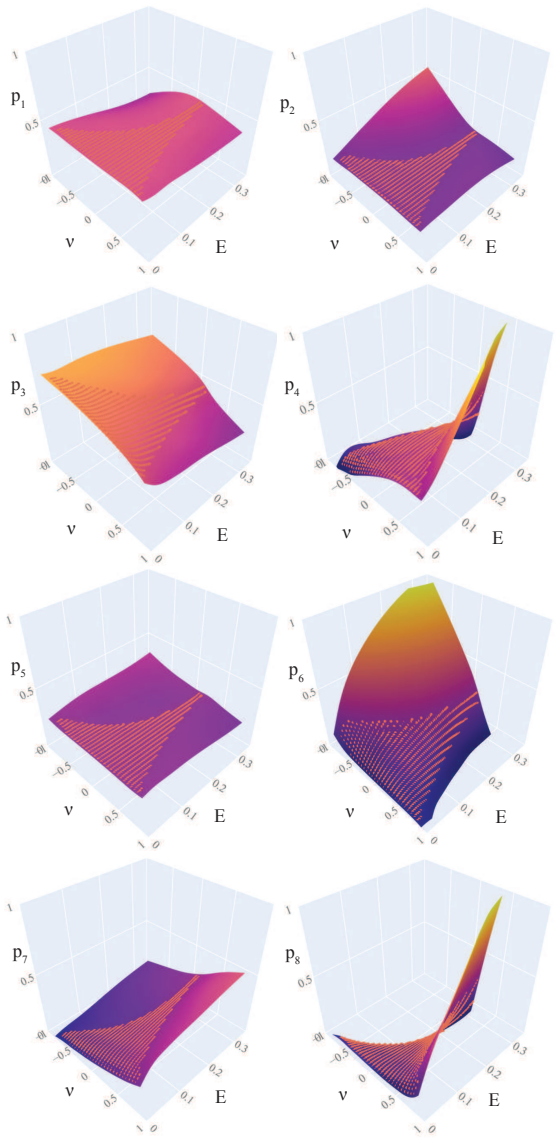

Fig. 17. Plots of parameters $p_1 \dots p_8$ as functions of (E, ν) for $\alpha = \pi/2$. The color of the surface corresponds to parameter values.

Discussion of alternative approaches. The closest works on 2D microstructures are [Ostanin et al. 2018] and [Martínez et al. 2019], which introduce low-parametric families of structures. However, both approaches are limited to regular grids in a profound way: they rely on symmetries of triangular/hexagonal grids to produce isotropic effective properties, a requirement that is lost on deformed lattices with fewer symmetries.

Regular quad/hex grids are used by topology optimization-based methods [Schumacher et al. 2015; Zhu et al. 2017], and in the family of structured proposed in [Panetta et al. 2017, 2015]. While these methods could be adapted to rhombic tiles by applying a bilinear map to their structures, this construction introduces large errors (see Section 8 for more details).

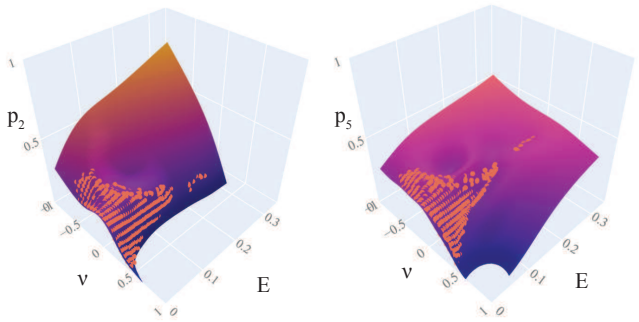

Fig. 18. Plots of parameters for $\alpha = \pi/4$. The color of the surface corresponds to the parameter values.

Our construction is based on the optimization techniques introduced in [Panetta et al. 2017, 2015], by addressing two of their limitations: (1) we lift the regular lattice cell assumption (for the 2D planar case), directly targeting rhombic tiles and (2) we address the lack of smoothness of $\mathcal{P}$ (cf. Figures 17 and 11).

7 QUAD MESH OPTIMIZATION

We tessellate the polygonal domain Ω into (nearly-)rhombic elements by starting with a quadrangulation of Ω and optimizing the shape of its cells.

Quad mesh generation. Our algorithm can be applied to any planar quadrilateral mesh: for the examples in the paper, we used a combination of meshes automatically generated using the *Integer Grid Maps* algorithm [Bommes et al. 2013] (Figure 34, disk), *Pattern-Based Quadrangulation* [Takayama et al. 2014] (Figure 34, bar), and manually designed using Blender (Figure 34, pliers, ghost). Generally, denser meshes are easier to optimize for our algorithm since they have more degrees of freedom.

Quad to rhombi projection. Our algorithm uses an alternating optimization inspired by ShapeUp [Bouaziz et al. 2012]: we compute a target rhombic shape for each quad (see inset), and then optimize the vertex positions to match the target shape of each element in least-squares sense. To obtain the ideal rhombic shape R for a quad F, we fit a linear transformation T of a canonical unit square Q (a parallelogram) to F, after translating its barycenter to the origin:

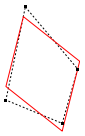

$$\min_T \sum_i \|f_i - Tq_i\|^2, \tag{12}$$

where f_i and q_i are the vertices of F and Q respectively. Letting θ_T be the smaller angle of the parallelogram, we define the unit target rhombic shape R_u as the rhombus with unit length edges and minimal angle $max(\theta_T, \theta_m)$, where θ_m is the smallest angle covered by our parameter sweep (45 degrees). Capping the minimal angle pushes the optimization toward rhombic shapes within the reachable range of our parametric family. The scaled target rhombic

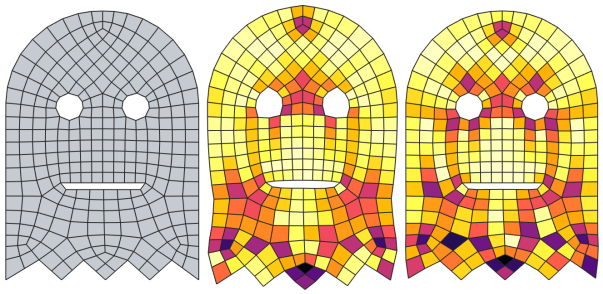

Fig. 19. A quadrilateral mesh (left) is optimized with (right) and without (middle) the boundary term. The shape of the eyes and the mouth is lost in the middle image, while it is preserved on the right. The color for each quad maps to α (the minimum angle of the cell).

shape R is computed by translating (by a vector t) and scaling (by a diagonal matrix A) R_u to best fit the quad F:

$$\min_{A,t} \sum_i \|f_i - Ar_i - t\|^2, \tag{13}$$

where r_i and f_i are the vectors of vertices of R_u and F respectively.

Exponential rhombic distortion energy. Equipped with a target reference for each quad, we define a distortion energy by penalizing the deviation of each element from its ideal rhombic shape using the symmetric Dirichlet energy [Smith and Schaefer 2015]:

$$\mathcal{D}(J_F) = \|J_Q\|_F^2 + \|J_Q^{-1}\|_F^2,$$

where J_F is the Jacobian of the bilinear map mapping F to its target rhombic shape R. Ideally, we would like to minimize the L_∞ norm of the distortion integrated over the entire domain Ω, which is however a challenging problem to solve numerically. We opt for an approximation, using a continuation approach [Rabinovich et al. 2017] minimizing the following exponential energy for increasing values of s:

$$E_{\mathcal{D}}(x) = \sum_{Q \in \Omega} A_Q e^{s\mathcal{D}(J_Q)} \tag{14}$$

where A_Q is the area of the quad Q.

Boundary projection. Without additional constraints, minimizing (14) would change the shape of the boundary of Ω (Figure 19). We add a penalty term to avoid changes to the boundary, while allowing nodes to slide on it:

$$E_\partial(x, y) = \sum_{x_i \in \partial\Omega} \|x_i - y_i\|^2, \tag{15}$$

where x_i are the coordinates of a boundary vertex of a point on the boundary of the quad mesh, and $y_i = \Phi(x_i)$ is the closest point on the boundary of Ω.

Optimization. We create a rhombic mesh by minimizing the combined energy:

$$E(x) = (1 - \lambda)E_{\mathcal{D}}(x) + \lambda E_\partial(x, \Phi(x)), \tag{16}$$

with $\lambda = 10000$ a weight balancing the two terms. We use the solver proposed in [Rabinovich et al. 2017] to find vertex positions minimizing eq. (16). This solver uses a local/global optimization strategy in order to minimize eq. (16). In the local step, the target rhombi are rotated to minimize $E_{\mathcal{D}}(x)$ (eq. (4) in [Rabinovich et al. 2017]), and we compute $y_i^k = \Phi(x_i^k)$, the projection on the boundary $\partial\Omega$ for border vertices after iteration k. In the global step, the Jacobian J_Q in $E_{\mathcal{D}}$ is replaced by $J_Q^k = J_Q(x^k)$, $E_\partial(x, \Phi(x))$ is replaced by $E_\partial(x, y^k)$, and E is minimized wrt x.

8 EVALUATION AND EXAMPLES

In this section, we explore the advantages of our approach, compared to "naive" application of existing techniques, and its limitations. Specifically, we explore the following features of our method:

- the use of boundary-aligned, irregular quad meshes instead of regular grids with cut cells;
- the optimization of the cell geometry structure taking cell shape into account, compared to remapping geometry optimized for square cells;
- continuous dependence of geometric parameters on material parameters, compared to pointwise inverse homogenization used in previous work;
- mesh independence and the effects of deviation from rhombicity on the accuracy of the method.

Finally, we show how our method handles a number of test problems. Many of the comparisons are performed on models with variable material properties generated using material optimization, as in [Panetta et al. 2015]. Specifically, the deformation on one part of the boundary is given, and the material properties for interior cells are optimized to obtain a user-defined target deformation on an other part of the boundary. For example, we use a plier shape for several examples; the prescribed deformation moves the handles together, and the target deformation moves the jaws together. In the optimization, the cells of the mesh are used as finite elements with constant material properties.

Comparison with regular grids with cut cells. The simplest approach that allows one to use existing techniques with microstructures optimized for squares is to overlay a regular grid on the geometry of the shape and remove parts of the cells outside the input shape boundary. We used a regular grid with square side length equal to the average edge length of our rhombic geometry to obtain approximately the same number of cells. The material optimization was performed with the same boundary conditions and target deformations but for each mesh separately (cut regular grid vs rhombic cells), rather than, e.g., interpolating target material properties from one mesh to the other. For the purposes of FEM simulation used in the material optimization, cut cells are triangulated. In both cases we assign the geometry for each cell based on the computed material. As an additional step for the regular cell grid, we crop the final fine-scale geometry using the original object's boundary. We run the deformation simulation on full fine-scale geometry for comparison. Figure 20 shows an example with significant deformations compared to the ground truth. We note that one feature of our method, the continuous dependence of geometric parameters on material parameters remains useful even for objects that can be meshed with regular grids, and improves accuracy.

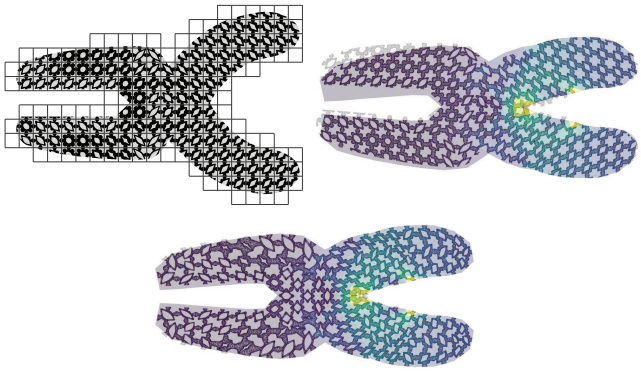

Fig. 20. Top left: example of a cropped regular grid of perfect square cells. Top right: deformation using the model produced with square grid as shown on the left. Bottom: result using our new method.

To experiment with higher resolution quad meshes, we applied one level refinement on each of the two configurations shown in Figure 20, keeping the same material distribution. The result is presented in Figure 21. For both methods, the error is reduced, although it is still considerably higher when using regular grids with cut cells. Notice, however, that (although being a valid experiment) refinement is not a very practical solution in real world scenarios since most fabrication technologies have resolution limitations.

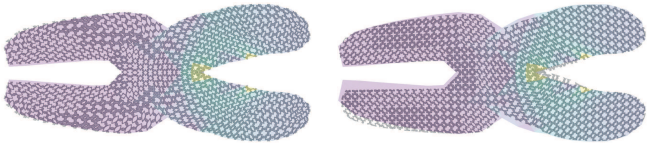

Fig. 21. Left: deformation result after one level of quad mesh refinement when using our new method. Right: result using regular grid with cut cells.

Comparison with remapping square cell structures. In principle, the structures obtained for square cells can be directly applied to rhombic cells, by simply applying a bilinear transformation to the geometry. Figure 22 shows the effect of not using structures specialized to set with specific angles: the increase in the error is substantial.

Comparison with pointwise inverse homogenization. In this experiment, the same mesh and the same variable material properties, resulting from optimization, are used with two different ways of assigning a structure to each cell (Figure 23). The first approach performs the nearest-neighbor lookup in the database of structures obtained by pointwise shape optimization [Panetta et al. 2017], the other using our smooth parametric family $\mathcal{P}(E, \nu; \alpha)$. In both cases, the structures are mapped to the arbitrary quads of the mesh by a bilinear map. We observe that the result is closer to the reference simulation using a variable-material solid. Figure 24 shows that the $E - \nu$ space coverage of two methods is close.

Mesh independence. If all cells are assigned the same effective material properties, the resulting deformation should not change significantly as we change the domain discretization.

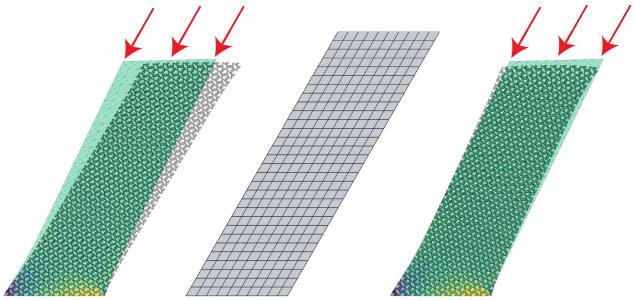

Fig. 22. Middle: non-deformed solid shape. Left: deformation of the structure using geometry optimized for $\alpha = \pi/2$ everywhere; the solid shape in the background shows the reference deformation of a single-material solid. Right: using correct cell geometry with parameters $\mathcal{P}(\alpha)$ everywhere. The microstructure has material parameters $E = 0.02$ and $\upsilon = -0.4$. Bottom part is fixed (0 Dirichlet), and on the top part we apply a force parallel to the slanted side (at an angle of $4\pi/3$).

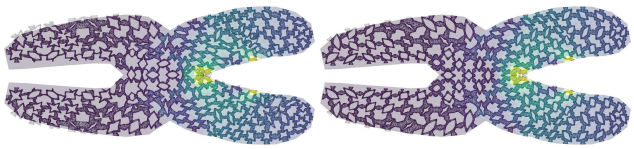

Fig. 23. Left: deformation of an object constructed using a table of structures obtained using pointwise inverse homogenization as in [Panetta et al. 2017]. Right: deformation of the object constructed using our family, with the same boundary conditions. The gray shape in the background shows the ground truth obtained by simulation on a solid quad mesh with materials assigned per quad.

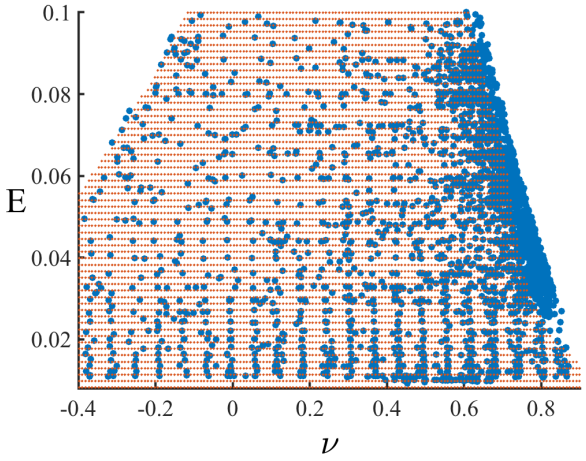

Fig. 24. Coverage in (E, ν) plane for $\alpha = \pi/2$ obtained by pointwise shape optimization (blue) in the range $E = 0.01 \ldots 0.1$ and $\nu = -0.4 \ldots 0.9$. The coverage of our family is shown in orange.

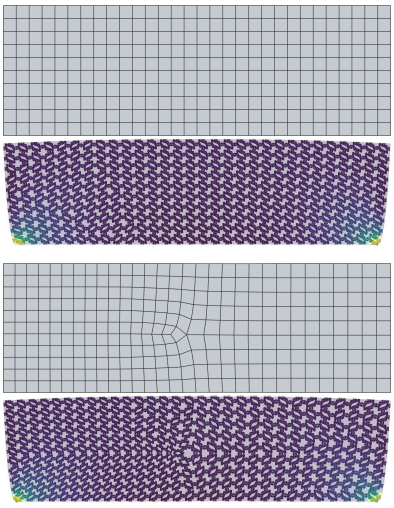

Fig. 25. Top: deformation of a shape with constant assigned materials and regular connectivity. Bottom: deformation of the same shape, with irregular mesh connectivity. The deformed pattern (simulated using FEM on a fine mesh) is compared to the simulated homogeneous mesh with equivalent material properties ($E = 0.05$, $\nu = 0.4$). In both cases, a uniform vertical load is applied on the top, while we keep a 0 Dirichlet condition at the bottom.

We performed two experiments to evaluate the effect of changing the mesh. The first experiment (Figure 25) confirms that the deformations obtained for an irregular mesh on a bar, which requires different rhombic structures for different cells, are nearly the same as for a regular mesh with identical structure, optimized for the same square cell everywhere.

In the second experiment, we study how deviation from rhombic shape affects the result, when the cell shape is intentionally distorted to be far from rhombic. We create a *twisted structure* – a compound cell subdivided into 4 non-square subcells, and compute the properties of the material composed of such compound cells, using homogenization on the whole cell. Specifically, starting from a 4×4 square quad mesh on the cell, with each quad assigned $E = 0.05$, $\nu = 0.0$, we gradually rotate the cross formed by the edges incident at the central vertex of the mesh (5 degrees at a time), and optimize the shape of the elements of the twisted mesh (keeping the cross fixed, Figure 26). Figure 27 shows how the Frobenius norm of the distance between the target elasticity tensor used to obtain the structure in each subcell and the homogenized elasticity tensor of the compound cell changes with rotation angle. In the plot, we also show the error we observe if we use the structure optimized for $\alpha = \pi/2$ everywhere.

Notice how the error stays stable when we increase the rotation angle, which is not the case when using only the structures for $\alpha = \pi/2$ in all quads. We emphasize that this experiment was intentionally conducted with a very coarse grid of subcells. The error of our method further decreases for a finer grid: if we subdivide each quad and placing the same initial structure in all sub-quads (see Figure 28). For example, for a rotation angle of 10 degrees, we had an initial error of 0.0123 with our method, while, with one level

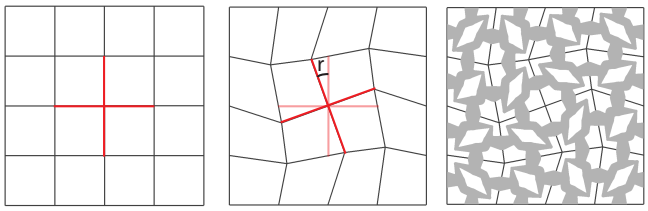

Fig. 26. An example of the twisted structure. Left: initial mesh with perfect squares. Middle: twisted mesh by a rotation r of 20 degrees. Right: final structure after rotation.

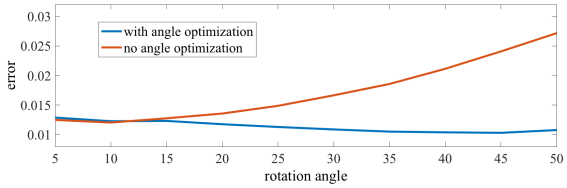

Fig. 27. Plot of twisting accuracy when the rotation angle r is increased. The plot shows the Frobenius norm of the distance between the target elasticity tensor and the actual elasticity tensor obtained by running homogenization on the compound pattern.

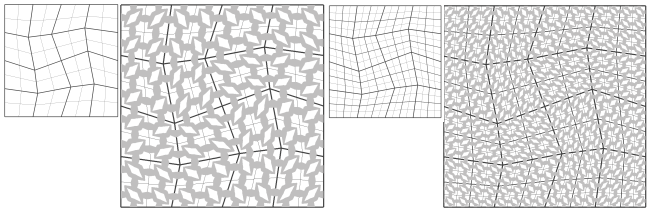

Fig. 28. Refinement of the twisted structure. Left: one level of quad refinement. Right: two levels of quad refinement.

of refinement, the error is almost halved, decreasing to 0.0069, and two levels decrease the error to 0.0040.

Non-rhombicity error. In our approach, all cells are approximated by rhombi for the purposes of computing the cell structure. To study the effect of deviation of cell shapes from rhombic, we experimented with a simple example shown in Figure 29. Using a 2×2 quad grid in a compound cell, and starting from 4 perfect squares, we move the center vertex in the direction of the top-right corner by an increasing distance d in each axis, which increases the non-rhombicity of every cell. For each d, the homogenized properties of the material consisting of compound cells is computed.

We compute then, for each value of d, the *rhombic errors* of each of the 4 quads as the sum of the distances from the quad vertices to the vertices of the closest rhombus (which can be obtained as explained in Section 7). We also compute the *target error* as the Frobenius norm of the difference between the homogenized elasticity tensor for the compound cell and the target one. Results are shown in Figure 30. The distance to target material changes smoothly with the total rhombic error. Moreover, to analyse quantitatively the quality of

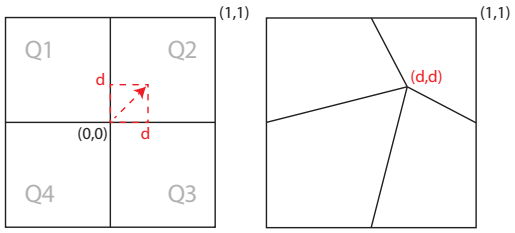

Fig. 29. Non-rhombic cell deformation: the center vertex is slowly translated from $(0, 0)$ to $(1, 1)$, increasing non-rhombicity of cells and distance from the homogenized properties of the target cell to target material properties.

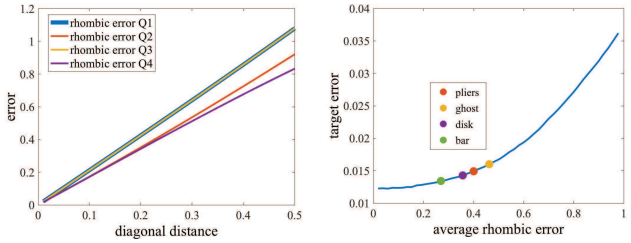

Fig. 30. Plots of accuracy in non-rhombicity experiment. Left: plot shows how rhombic error varies when d value is increased. Notice that the change is almost linear and, as expected, error for Q1 and Q3 are the same. Right: smooth variation of target error when rhombic error increases. Dots indicate corresponding error for our example models in Figure 30.

our quad meshes and to show how much damage non-rhombicity could cause in accuracy, we computed the average rhombic error for our meshes in Figure 34, normalizing the results with the squared root of the average area of quads. The resulting numbers for pliers, ghost, disk and bar models can be seen on the right side of Figure 30. Notice that these numbers would indicate a low target error.

Scalability. Using the same example shown in Figure 29 with d fixed at 0.25, we also experimented with the scalability of our solution by running the entire pipeline (quad optimization, material optimization, material to geometry mapping and final mesh construction) on different levels of refinement on the initial quad mesh. The result is shown in Figure 31. Notice that our running time scales linearly with the number of cells of the quad mesh.

Merge error. Once the geometry is defined per cell, it needs to be merged into a single geometry for the whole object. For our choice of pattern (but without using our smooth spline fit), adjacent cells share a single node and the structures of two cells may have different geometric parameters (radii) assigned to the shared node: these radii need to be averaged, affecting the properties of cells. In contrast, with our microstructures family, if the transition of materials is smooth, the change in radii will also be smooth and almost no averaging will be required. In fact, for square cells ($\alpha = \pi/2$), the radius at the boundary will always be very close to 0.3 (see Figure 17), because of the way we fix p_5 in our map construction, as described in Section 6.2.

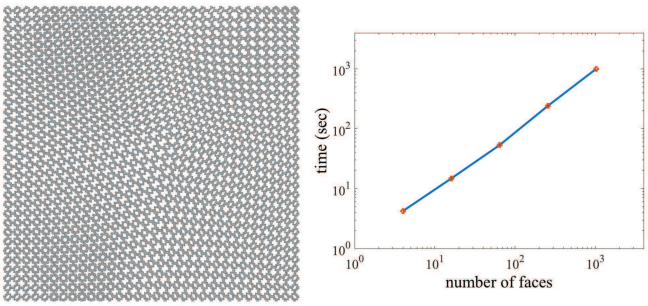

Fig. 31. Scalability experiment. Left: example of optimized structure with 4 levels of refinement on the initial mesh, which means a final number of 1024 faces. Right: Plot (in log–log scale) showing how our running time grows with the number of quad faces. For this experiment, we used a machine with 20 cores and 32GB of RAM.

Figure 32 shows an example where there is a significant mismatch between radii assigned to the shared node by two cell structures that have same material properties ($E = 0.0315$, $\nu = 0.75$ and anisotropy measurements a_1 and a_2 in the order of 10^{-5} and 10^{-4}, respectively). After the merge, the new homogenized properties can be obtained by building a new square cell as shown on the right side of Figure 32. The new material has significantly different properties, not being isotropic anymore ($a_1 = 0.028$ and $a_2 = 0.037$).

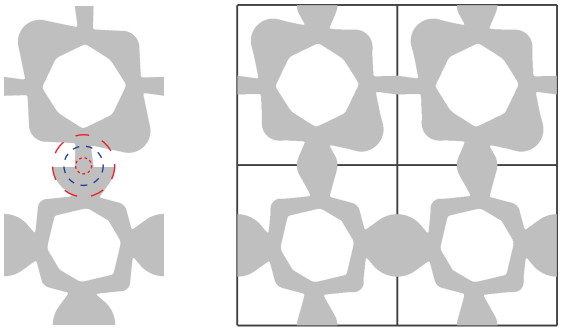

Fig. 32. Error due to merging cell structures. Left: radii of contact nodes of two structures with same material properties (in red) and resulting radius after merge (in blue). Right: base homogenization cell to verify the properties of merged structure (which are significantly different than original).

Interpolation accuracy. Finally, we show that our spline approximation is quite accurate. Figure 33 shows the absolute differences $|H_\nu(\mathcal{P}(E_j, \nu_j)) - \nu_j|$ and $|H_E(\mathcal{P}(E_j, \nu_j)) - E_j|$, at points (E_j, ν_j) not present in the originally sampled set: these are differences between actual homogenized material properties of the geometric structures obtained using our map $\mathcal{P}$, and the values at which $\mathcal{P}$ was evaluated. Smoothness of the map $\mathcal{P}$ is essential for the interpolation of the relatively sparse set of values to be accurate.

Examples of manufactured structures. Finally, we applied our complete pipeline for a set of 2D examples similar to [Panetta et al.

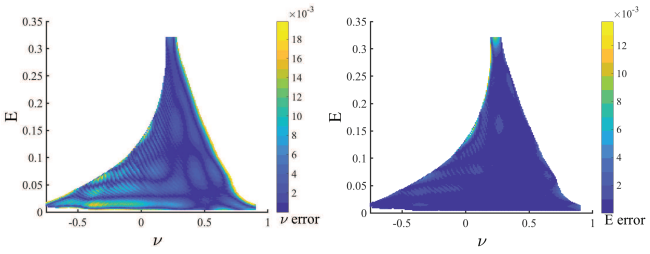

Fig. 33. Accuracy of $\mathcal{P}$ away from sample points, measured as the difference between directly computed and interpolated E and ν.

2015], with a crucial difference being that we use irregular, boundary-conforming meshes optimized for rhombic shape of cells. The per-cell material property distribution is obtained using the optimization described in [Panetta et al. 2015]: we prescribe fixed Dirichlet boundary conditions on a part of the boundary, and target positions of another (target) part, optimizing the material distribution to obtain the desired deformation.

As we worked with 2D elasticity, we used 0.5" sheets of closed-cell foam with small (< 0.2 mm) foam cell size to fabricate our examples using a laser cutter. We also experimented with other materials, as 1/8" acrylic sheets, which allow for higher resolution when compared to foam (see pliers example in Figure 34). These examples have primarily illustrative purpose. We expect more practical applications to be accomplished by integration into new CAD design pipelines e.g, nTop Platform [nTopology 2020], which already supports a range of small-scale parametric structures.

9 CONCLUSIONS AND LIMITATIONS

We have demonstrated that it is possible to construct a family of geometric structures for rhombic shapes with a large range of angles, parameterized directly by material parameters of these shapes. The map from the material parameters to the geometric parameters is smooth and represented in a compact form as a set of 3D splines. Within explicitly defined bounds, any structure in the family is isotropic. While the material-to-geometry map depends on the base material properties, the transformation between different materials can be achieved with a simple analytic formula; for this reason, the proposed family can be viewed as independent of the base material properties.

We demonstrate that using this family improves the accuracy with which target material properties can be approximated using cellular structures manufactured from a single material.

This work is a first step towards explicitly defined material-to-geometry maps of this type. We note that as an intermediate step in our construction we have obtained a map from arbitrary orthotropic, not just isotropic, properties to geometric parameters. Orthotropic materials may be useful in many contexts, and our work can be extended to this. Previous work [Milton et al. 2017; Ostanin et al. 2018] demonstrates that much broader, nearly optimal, range of parameters can be obtained by using more complex structure topology; expanding the range further is another possible future direction of work.

This work forms a foundation for developing similar parametric families in 3D, since, although requiring more parameters and computational time, we do not foresee any fundamental problems in extending the construction of cell geometry families to 3D structures. As a proof-of-concept, we produced a smooth family of microstructures for cubic cells in 3D, following exactly the process described in Section 6.2, with 2D homogenization replaced with 3D. We used a cell geometry with cubic symmetry, guaranteeing orthotropy with same Young's modulus in all directions, which means having a three dimensional material space ($m = 3$). Initially, 9 parameters defining the geometry were considered ($p = 9$). This number was reduced to 7 and then to 5 with parameter elimination step ($p' = 5$). Then, we used PCA for the final reduction from 5 to 3. Figure 35 shows the initial and also the final coverage with our 3-parameter space, from which a nonlinear map for isotropic structures can be extracted. We note that the second component of the method, decomposing shapes into hexahedral cells of suitable shape, e.g., close to rhombohedra, requires much greater adaptation, although hex meshing methods can be used as a starting point.

ACKNOWLEDGMENTS

This work was supported in part through the NYU IT High Performance Computing resources, services, and staff expertise. This work was partially supported by the NSF CAREER award 1652515, the NSF grants IIS-1320635, DMS-1436591, DMS-1821334, OAC-1835712, OIA-1937043, CHS-1908767, CHS-1901091, a gift from Adobe Research, a gift from nTopology, and a gift from Advanced Micro Devices, Inc.

REFERENCES

Niels Aage, Erik Andreassen, Boyan S. Lazarov, and Ole Sigmund. 2017. Giga-Voxel Computational Morphogenesis for Structural Design. *Nature* 550, 7674 (10 2017), 84–86. https://doi.org/10.1038/nature23911

Noam Aigerman and Yaron Lipman. 2013. Injective and Bounded Distortion Mappings in 3D. *ACM Trans. Graph.* 32, 4, Article 106 (July 2013), 14 pages. https://doi.org/10.1145/2461912.2461931

Noam Aigerman, Roi Poranne, and Yaron Lipman. 2014. Lifted Bijections for Low Distortion Surface Mappings. *ACM Trans. Graph.* 33, 4 (2014), 69:1–69:12.

Grégoire Allaire. 2002. *Shape optimization by the homogenization method.* Vol. 146. Springer.

Erik Andreassen, Boyan S. Lazarov, and Ole Sigmund. 2014. Design of manufacturable 3D extremal elastic microstructure. *Mechanics of Materials* 69, 1 (2014), 1–10. https://doi.org/10.1016/j.mechmat.2013.09.018

Martin Philip Bendsøe. 1989. Optimal shape design as a material distribution problem. *Structural optimization* 1, 4 (1989), 193–202.

Martin Philip Bendsøe and Ole Sigmund. 2003. *Topology optimization: theory, methods and applications.* Springer.

Bernd Bickel, Moritz Bächer, Miguel A. Otaduy, Hyunho Richard Lee, Hanspeter Pfister, Markus Gross, and Wojciech Matusik. 2010. Design and Fabrication of Materials with Desired Deformation Behavior. *ACM Trans. Graph.* 29, 4, Article 63 (July 2010), 10 pages. https://doi.org/10.1145/1778765.1778800

David Bommes, Marcel Campen, Hans-Christian Ebke, Pierre Alliez, and Leif Kobbelt. 2013. Integer-grid Maps for Reliable Quad Meshing. *ACM Trans. Graph.* 32, 4, Article 98 (July 2013), 12 pages. https://doi.org/10.1145/2461912.2462014

Sofien Bouaziz, Mario Deuss, Yuliy Schwartzburg, Thibaut Weise, and Mark Pauly. 2012. Shape-Up: Shaping Discrete Geometry with Projections. *Comput. Graph. Forum* 31, 5 (Aug. 2012), 1657–1667. https://doi.org/10.1111/j.1467-8659.2012.03171.x

Sofien Bouaziz, Sebastian Martin, Tiantian Liu, Ladislav Kavan, and Mark Pauly. 2014. Projective Dynamics: Fusing Constraint Projections for Fast Simulation. *ACM Trans. Graph.* 33, 4, Article 154 (July 2014), 11 pages. https://doi.org/10.1145/2601097.2601116

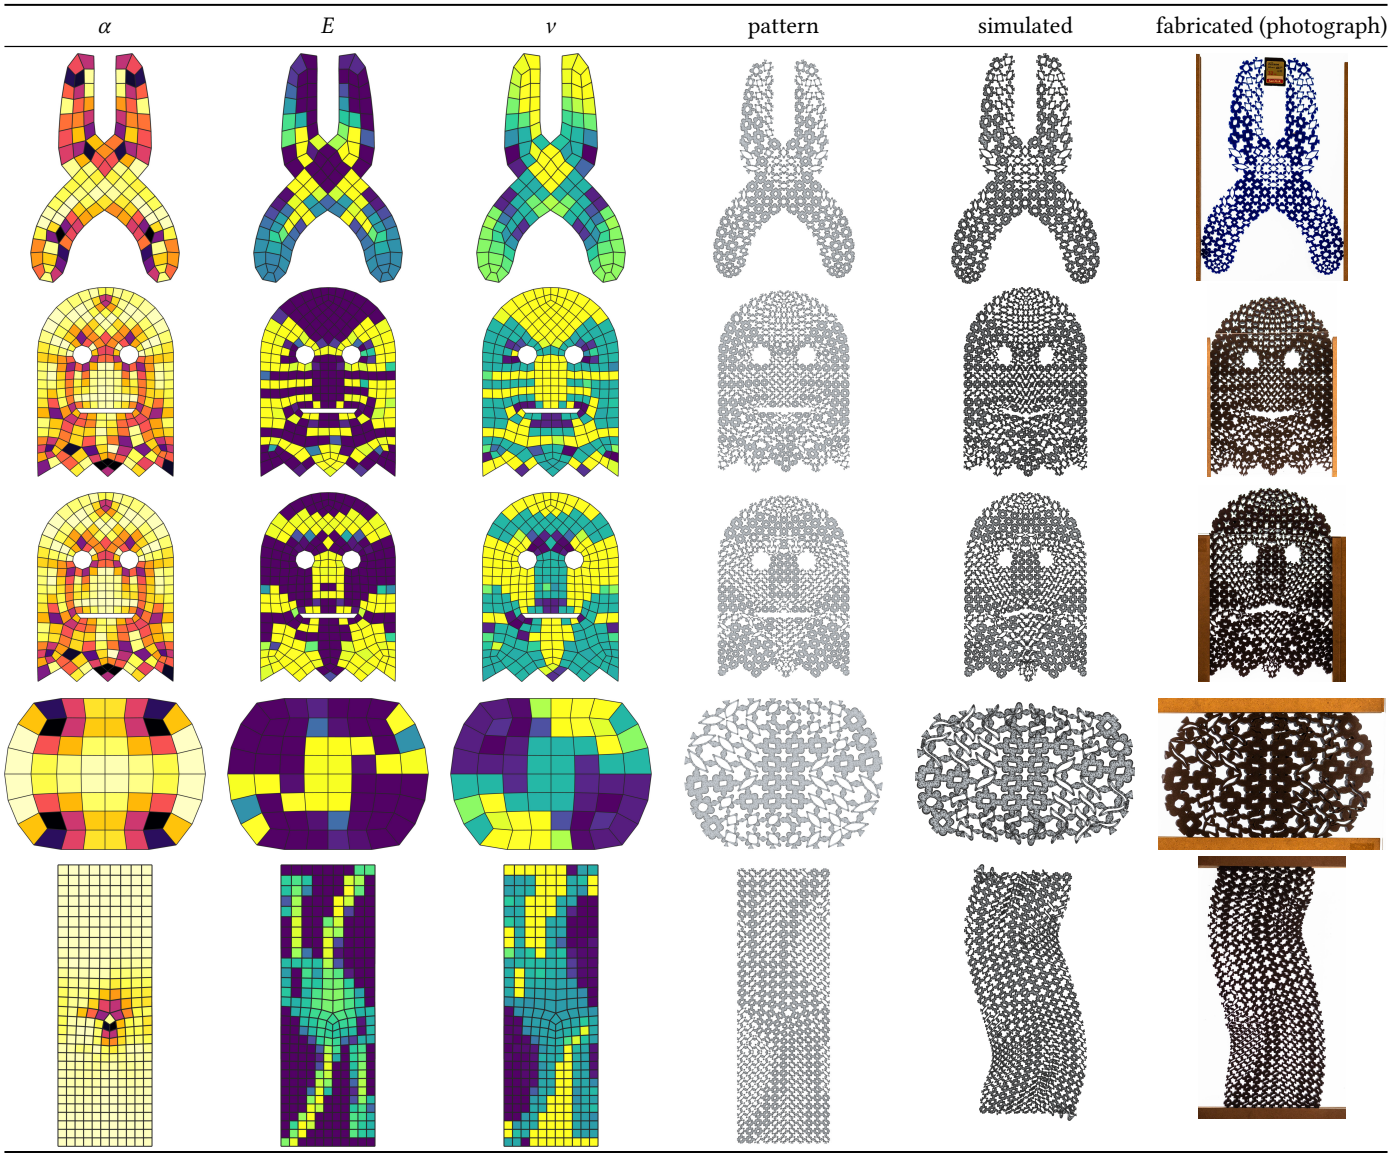

Fig. 34. Summary table of the results fabricated with our method. The plier model was laser-cut in acrylic, while the rest were laser-cut from foam. From left to right: angle per quad on the optimized mesh, optimized material distribution (Young's modulus and Poisson's ratio), final geometry at rest, deformed pattern (simulated) and photograph of deformed pattern (fabricated). Video clips of the deformation of the fabricated results are provided in the additional material. The minimum angle after optimization is 57.4 deg for the pliers, 64.5 deg for the ghost, 66.0 deg for the disk, and 67.2 deg for the bar. On these examples, our Young's moduli ranges from 0.005 to 0.25, and our Poisson's ratio ranges from -0.27 to 0.8.

M. Brewer, L. Diachin, P. Knupp, T. Leurent, and D. Melander. 2003. The Mesquite Mesh Quality Improvement Toolkit. In *Proc. of the 12th International Meshing Roundtable*. 239–250.

Tiemo Bückmann, Nicolas Stenger, Muamer Kadic, Johannes Kaschke, Andreas Frölich, Tobias Kennerknecht, Christoph Eberl, Michael Thiel, and Martin Wegener. 2012. Tailored 3D Mechanical Metamaterials Made by Dip-in Direct-Laser-Writing Optical Lithography. *Advanced Materials* 24, 20 (2012), 2710–2714.

Desai Chen, Mélina Skouras, Bo Zhu, and Wojciech Matusik. 2018. Computational discovery of extremal microstructure families. *Science advances* 4, 1 (2018), eaao7005.

Andrej Cherkaev. 2000. *Variational methods for structural optimization*. Vol. 140. Springer.

Andrei V. Cherkaev, Konstantin A. Lurie, and Graeme W. Milton. 1992. Invariant properties of the stress in plane elasticity and equivalence classes of composites. *Proceedings of the Royal Society A* 428 (1992), 519–529.

Doina Cioranescu and Patrizia Donato. 1999. *An introduction to homogenization*. Oxford University Press.

P. Degener, J. Meseth, and R. Klein. 2003. An Adaptable Surface Parameterization Method. In *Proceedings of the 12th International Meshing Roundtable*. 201–213.

Mario Deuss, Anders Holden Deleuran, Sofien Bouaziz, Bailin Deng, Daniel Piker, and Mark Pauly. 2015. *ShapeOp—A Robust and Extensible Geometric Modelling Paradigm*. Springer International Publishing, Cham, 505–515. https://doi.org/10.1007/978-3-319-24208-8_42

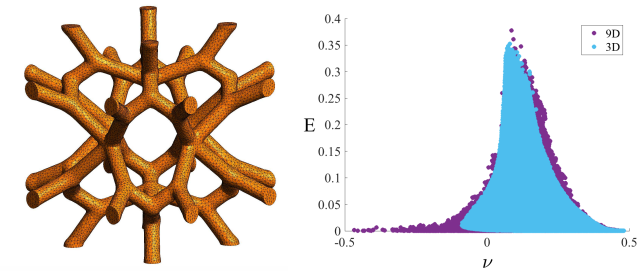

Fig. 35. Left: example of geometry in our 3D family. Right: comparison of the material properties obtained with our initial 9 parameters and the reduced set of only 3.

Xiao-Ming Fu and Yang Liu. 2016. Computing Inversion-free Mappings by Simplex Assembly. *ACM Trans. Graph.* 35, 6, Article 216 (Nov. 2016), 12 pages.

Xiao-Ming Fu, Yang Liu, and Baining Guo. 2015. Computing Locally Injective Mappings by Advanced MIPS. *ACM Trans. Graph.* 34, 4, Article 71 (July 2015), 12 pages.

Eric Garner, Helena M. A. Kolken, Charlie C. L. Wang, Amir A. Zadpoor, and Jun Wu. 2019. Compatibility in Microstructural Optimization for Additive Manufacturing. *Additive Manufacturing* 26 (03 2019), 65–75. https://doi.org/10.1016/j.addma.2018.12.007

Andrew T Gaynor and James K Guest. 2016. Topology optimization considering overhang constraints: Eliminating sacrificial support material in additive manufacturing through design. *Structural and Multidisciplinary Optimization* 54, 5 (2016), 1157–1172.

Francisca Gil-Ureta, Nico Pietroni, and Denis Zorin. 2019. Structurally Optimized Shells. *arXiv preprint arXiv:1904.12240* (2019).

G. N. Greaves, A. L. Greer, R. S. Lakes, and T. Rouxel. 2011. Poisson's ratio and modern materials. *Nature Materials* 10, 11 (2011), 823–837. https://doi.org/10.1038/nmat3134

Jeroen P. Groen and Ole Sigmund. 2017. Homogenization-Based Topology Optimization for High-Resolution Manufacturable Microstructures. *Internat. J. Numer. Methods Engrg.* 113, 8 (08 2017), 1148–1163. https://doi.org/10.1002/nme.5575

Jeroen P. Groen, Jun Wu, and Ole Sigmund. 2019. Homogenization-Based Stiffness Optimization and Projection of 2D Coated Structures With Orthotropic Infill. *Computer Methods in Applied Mechanics and Engineering* 349 (06 2019), 722–742. https://doi.org/10.1016/j.cma.2019.02.031

Jonathan Hiller and Hod Lipson. 2009. Design and analysis of digital materials for physical 3D voxel printing. *Rapid Prototyping Journal* 15, 2 (2009), 137–149.

Scott J. Hollister. 2005. Porous scaffold design for tissue engineering. *Nature Materials* 4, 7 (2005), 518–524. https://doi.org/10.1038/nmat1421

K. Hormann and G. Greiner. 2000. MIPS: An Efficient Global Parametrization Method. In *Curve and Surface Design: Saint-Malo 1999*. 153–162.

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* (Tokyo, Japan) *(UIST '16)*. ACM, New York, NY, USA, 529–539. https://doi.org/10.1145/2984511.2984540

Alexandra Ion, David Lindlbauer, Philipp Herholz, Marc Alexa, and Patrick Baudisch. 2019. Understanding Metamaterial Mechanisms. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) *(CHI '19)*. ACM, New York, NY, USA, Article 647, 14 pages. https://doi.org/10.1145/3290605.3300877

Caigui Jiang, Chengcheng Tang, Amir Vaxman, Peter Wonka, and Helmut Pottmann. 2015. Polyhedral Patterns. *ACM Trans. Graphics* 34, 6 (2015). Proc. SIGGRAPH Asia.

Hee Suk Kang. 2010. *Hierarchical design and simulation of tissue engineering scaffold mechanical, mass transport, and degradation properties.* Ph.D. Dissertation. The University of Michigan.

Mina Konaković-Luković, Julian Panetta, Keenan Crane, and Mark Pauly. 2018. Rapid deployment of curved surfaces via programmable auxetics. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 106.

Shahar Z. Kovalsky, Noam Aigerman, Ronen Basri, and Yaron Lipman. 2015. Large-scale bounded distortion mappings. *ACM Transactions on Graphics (proceedings of ACM SIGGRAPH Asia)* 34, 6 (2015).

Shahar Z. Kovalsky, Meirav Galun, and Yaron Lipman. 2016. Accelerated Quadratic Proxy for Geometric Optimization. *ACM Trans. Graph.* 35, 4, Article 134 (July 2016), 11 pages.

Matthijs Langelaar. 2016. An additive manufacturing filter for topology optimization of print-ready designs. *Structural and Multidisciplinary Optimization* (2016), 1–13.

Chia-Ying Lin, Chun-Ching Hsiao, Po-Quan Chen, and Scott J Hollister. 2004a. Interbody fusion cage design using integrated global layout and local microstructure topology optimization. *Spine* 29, 16 (2004), 1747–1754. PMID: 15303018.

Cheng Yu Lin, Noboru Kikuchi, and Scott J. Hollister. 2004b. A novel method for biomaterial scaffold internal architecture design to match bone elastic properties with desired porosity. *Journal of Biomechanics* 37, 5 (2004), 623–636. https://doi.org/10.1016/j.jbiomech.2003.09.029

Yaron Lipman. 2012. Bounded Distortion Mapping Spaces for Triangular Meshes. *ACM Trans. Graph.* 31, 4, Article 108 (July 2012), 13 pages. https://doi.org/10.1145/2185520.2185604

Haixiang Liu, Yuanming Hu, Bo Zhu, Wojciech Matusik, and Eftychios Sifakis. 2018. Narrow-Band Topology Optimization on a Sparsely Populated Grid. *ACM Trans. Graph.* 37, 6, Article 251 (12 2018), 14 pages. https://doi.org/10.1145/3272127.3275012

Marco Livesu, Alla Sheffer, Nicholas Vining, and Marco Tarini. 2015. Practical Hex-Mesh Optimization via Edge-Cone Rectification. *Transactions on Graphics (Proc. SIGGRAPH 2015)* 34, 4 (2015).

Augustus Love. 1944. *A treatise on the mathematical theory of elasticity.* Dover publications New York. 151–155 pages.

Luigi Malomo, Jesús Pérez, Emmanuel Iarussi, Nico Pietroni, Eder Miguel, Paolo Cignoni, and Bernd Bickel. 2018. FlexMaps: Computational Design of Flat Flexible Shells for Shaping 3D Objects. *ACM Trans. Graph.* 37, 6, Article 241 (12 2018), 14 pages. https://doi.org/10.1145/3272127.3275076

Jonàs Martínez, Jérémie Dumas, and Sylvain Lefebvre. 2016. Procedural Voronoi foams for additive manufacturing. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 44.

Jonàs Martínez, Mélina Skouras, Christian Schumacher, Samuel Hornus, Sylvain Lefebvre, and Bernhard Thomaszewski. 2019. Star-Shaped Metrics for Mechanical Metamaterial Design. *ACM Transactions on Graphics* 38, 4 (07 2019), 13. https://doi.org/10.1145/3306346.3322989 Special issue, SIGGRAPH 2019.

Jonàs Martínez, Haichuan Song, Jérémie Dumas, and Sylvain Lefebvre. 2017. Orthotropic K-Nearest Foams for Additive Manufacturing. *ACM Transactions on Graphics* 36, 4 (07 2017), 1–12. https://doi.org/10.1145/3072959.3073638

Graeme W Milton. 2002. *The theory of composites.* Cambridge University Press.

Graeme W. Milton, Mark Briane, and Davit Harutyunyan. 2017. On the possible effective elasticity tensors of 2-dimensional and 3-dimensional printed materials. *Mathematics and Mechancis of Complex Systems* 5, 1 (2017), 381–391.

PH Nakasone and ECN Silva. 2010. Dynamic design of piezoelectric laminated sensors and actuators using topology optimization. *Journal of Intelligent Material Systems and Structures* 21, 16 (2010), 1627–1652.

nTopology. 2020. nTop Platform. https://ntopology.com/ Engineering software built for digital manufacturing.

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.

Julian Panetta, Abtin Rahimian, and Denis Zorin. 2017. Worst-Case Stress Relief for Microstructures. *ACM Transactions on Graphics* 36, 4 (07 2017), 1–16. 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 Transactions on Graphics (TOG)* 34, 4 (2015), 135.

Roi Poranne and Yaron Lipman. 2014. Provably Good Planar Mappings. *ACM Trans. Graph.* 33, 4 (2014), 76:1–76:11.

Roi Poranne, Elena Ovreiu, and Craig Gotsman. 2013. Interactive Planarization and Optimization of 3D Meshes. *Computer Graphics Forum* (2013). https://doi.org/10.1111/cgf.12005

Xiaoping Qian. 2016. Undercut and overhang angle control in topology optimization: a density gradient based integral approach. *Internat. J. Numer. Methods Engrg.* (2016).

Michael Rabinovich, Roi Poranne, Daniele Panozzo, and Olga Sorkine-Hornung. 2017. Scalable Locally Injective Mappings. *ACM Trans. Graph.* 36, 2, Article 16 (2017), 16 pages.

Pedro V. Sander, John Snyder, Steven J. Gortler, and Hugues Hoppe. 2001. Texture Mapping Progressive Meshes. In *ACM SIGGRAPH*. 409–416.

Christian Schüller, Ladislav Kavan, Daniele Panozzo, and Olga Sorkine-Hornung. 2013. Locally Injective Mappings. In *Symposium on Geometry Processing*. 125–135.

Christian Schumacher, Bernd Bickel, Jan Rys, Steve Marschner, Chiara Daraio, and Markus Gross. 2015. Microstructures to control elasticity in 3D printing. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 136.

Christian Schumacher, Steve Marschner, Markus Cross, and Bernhard Thomaszewski. 2018. Mechanical Characterization of Structured Sheet Materials. *ACM Trans. Graph.* 37, 4, Article 148 (07 2018), 15 pages. https://doi.org/10.1145/3197517.3201278

J. Schwerdtfeger, F. Wein, G. Leugering, R. F. Singer, C. Körner, M. Stingl, and F. Schury. 2011. Design of Auxetic Structures via Mathematical Optimization. *Advanced Materials* 23, 22 (2011), 2650–2654. https://doi.org/10.1002/adma.201004090

Jonathan Richard Shewchuk. 1996. Triangle: Engineering a 2D quality mesh generator and Delaunay triangulator. In *Workshop on Applied Computational Geometry*. Springer, 203–222.

Anna Shtengel, Roi Poranne, Olga Sorkine-Hornung, Shahar Z. Kovalsky, and Yaron Lipman. 2017. Geometric Optimization via Composite Majorization. *ACM Trans. Graph.* 36, 4, Article 38 (July 2017), 11 pages. https://doi.org/10.1145/3072959.3073618

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), 82.

Jason Smith and Scott Schaefer. 2015. Bijective Parameterization with Free Boundaries. *ACM Trans. Graph.* 34, 4, Article 70 (July 2015), 9 pages.

Olga Sorkine and Marc Alexa. 2007. As-Rigid-As-Possible Surface Modeling. In *Proceedings of EUROGRAPHICS/ACM SIGGRAPH Symposium on Geometry Processing*. 109–116.

Olga Sorkine, Daniel Cohen-Or, Rony Goldenthal, and Dani Lischinski. 2002. Bounded-distortion Piecewise Mesh Parameterization. In *Proceedings of the Conference on Visualization*. 355–362.

Kenshi Takayama, Daniele Panozzo, and Olga Sorkine-Hornung. 2014. Pattern-based Quadrangulation for N-sided Patches. In *Proceedings of the Symposium on Geometry Processing* (Cardiff, United Kingdom) *(SGP '14)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 177–184. https://doi.org/10.1111/cgf.12443

Chengcheng Tang, Xiang Sun, Alexandra Gomes, Johannes Wallner, and Helmut Pottmann. 2014. Form-finding with Polyhedral Meshes Made Simple. *ACM Trans. Graphics* 33, 4 (2014). https://doi.org/10.1145/2601097.2601213 Proc. SIGGRAPPH.

Salvatore Torquato. 2002. *Random heterogeneous materials: microstructure and macroscopic properties*. Vol. 16. Springer.

Jun Wu, Niels Aage, Rudiger Westermann, and Ole Sigmund. 2018. Infill Optimization for Additive Manufacturing—Approaching Bone-Like Porous Structures. *IEEE Transactions on Visualization and Computer Graphics* 24, 2 (02 2018), 1127–1140. https://doi.org/10.1109/tvcg.2017.2655523

Jun Wu, Christian Dick, and Rudiger Westermann. 2016. A System for High-Resolution Topology Optimization. *IEEE Transactions on Visualization and Computer Graphics* 22, 3 (03 2016), 1195–1208. https://doi.org/10.1109/tvcg.2015.2502588

Bo Zhu, Mélina Skouras, Desai Chen, and Wojciech Matusik. 2017. Two-Scale Topology Optimization With Microstructures. *ACM Trans. Graph.* 36, 4, Article 120b (07 2017). https://doi.org/10.1145/3072959.3095815

Y. Zhu, R. Bridson, and D. M. Kaufman. 2017. Blended Cured Quasi-Newton for Geometric Optimization. *ArXiv e-prints* (April 2017). arXiv:math.OC/1705.00039

A MICROSTRUCTURE GEOMETRY REPRESENTATION

In this appendix we summarize how the geometry of the microstructure is computed from the geometry parameters p_i. We use the approach of [Panetta et al. 2017] in 2D.

Inflation Graph. Microstructure is described by a graph of edges, with positions and radii and blending ratios associated with vertices. The geometry itself is a level set of a signed distance function constructed by combining distance functions associated with a primitive assigned to each edge.

Let $G = (\mathcal{V}, \mathcal{E})$ be the graph, where $\mathcal{V} = \{v_i\}_{i=1}^n$ represents the vertex indices and $\mathcal{E} = \{e_j\}_{j=1}^m$ are the edges. Recall that $F : Q \to R$ is the map from a unit square to the rhombus of interest. The shape parameters are $\{\mathbf{q}_i, r_i, b_i\}_i$ where:

- $\mathbf{q}_i = F\widetilde{\mathbf{q}}_i$ are the position of the vertices of G in R, and $\widetilde{\mathbf{q}}$ are positions in Q,
- r_i are per-vertex radii,
- b_i are per-vertex blending parameters.

In the family described in this paper, there are 12 nodes, but due to diagonal symmetries, only 4 have independent parameters, these can be chosen as in Figure 3. Moreover, symmetries also require nodes 3 and 4 to stay on diagonals (i.e., their coordinates have equal in magnitude values), and for node 1 to remain fixed. In our family we also do not vary the blending parameters keeping them at value 0.01.

Note that the positions $\{\widetilde{\mathbf{q}}_i\}_i$ correspond to the reference unit square, while $\{\mathbf{q}_i\}_i$ correspond to the rhombus.

Edge Geometry. A disk of radius r_i is placed at each vertex; then two outer common tangents are constructed for the disks centered at vertices v_i and v_j at the endpoints of an edge.

We define a signed distance function for each primitive described above, associated with an edge, and combine them together for a final shape.

Joints Geometry. The next step in the construction of the SDF for the graph G, is to determine the smooth blended geometry at a joint. A joint is defined as a pair $(v_i, \{e_k | k = 1 \dots N_i\})$, with a central vertex v_i and the list of its incident edges e_k.

Given the SDF distance d_k from query point $\mathbf{x}$ to each incident edge primitive (as defined above), the SDF to the blended joint is defined as a smooth minimum over the y_k with a blending parameter ρ_i:

$$KS(d_1, \dots, d_{N_i}, \rho_i) = -\frac{1}{\rho_i} \ln\Big(\sum_k \exp(-\rho_i d_k)\Big) \tag{17}$$

The blending parameter ρ_i is obtained as by multiplying the vertex blending parameter b_i with two spatially-varying modulation terms. One term depends on the convex hull of the joint, to avoid bulging. Another term depends on the minimum incident angle of the joint (between each pair of incident edges). The idea of the latter is to smooth more joints with small angles. [Panetta et al. 2017] explains how these quantities are computed.

Combining the Joints. The last element to consider when building the SDF for the graph as a whole, is to avoid sharp creases that can arise when combining to joints together. Here, the idea is to blend the SDF of the two closest joints using the smooth-min function Equation (17), with a blending factor that depends on the difference between the "hard-union" distance and the "smooth" SDF to those joints. More details can be found in [Panetta et al. 2017].

B SWEEP PARAMETERS

Displacements. In the initial 8-dimensional sweep, Node 2 displacement parameters (p_1, p_2) vary in the range $(0.1, 0.5)$, as $p_1 + p_2 > 0.5$ would take the node out of the triangle. Node 3 and node 4 displacements p_3, and p_4 vary in the range (0.1, 0.9), to span most of the allowed range (we avoid placing nodes too close to the boundary as this is likely to result in difficult-to-manufacture structures). We use 5 samples for all 4 offset parameters.

Radii. All radii $p_5, \dots p_8$ are allowed to vary in the $0.05 \dots 0.3$ range, and the number of samples was 6 for p_5, p_7, p_8, and 4 for p_6. The upper bound was chosen somewhat lower than the values needed to produce a solid cell, as we focused on the range of Young's moduli $0.01 \dots 0.3$ relative to the base material, and using radii about 0.3 produces E close to 0.33.