

Computational Homogenization for Inverse Design of Surface-based Inflatables

YINGYING REN, EPFL, Switzerland

JULIAN PANETTA, UC Davis, USA

SEIICHI SUZUKI, EPFL, Switzerland

UDAY KUSUPATI, EPFL, Switzerland

FLORIN ISVORANU, EPFL, Switzerland

MARK PAULY, EPFL, Switzerland

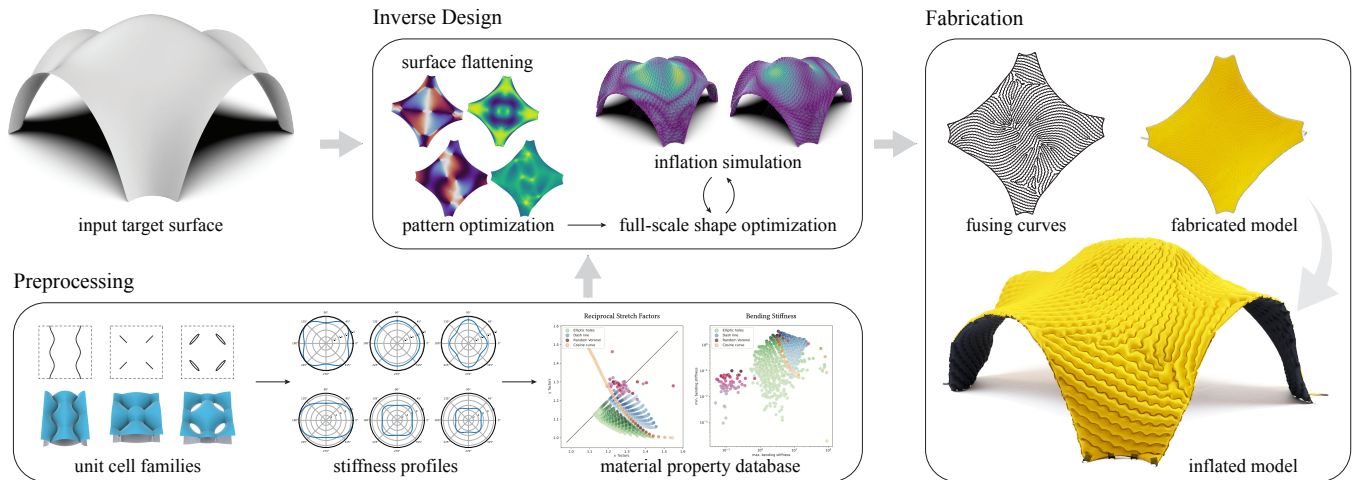


Fig. 1. Overview of our computational pipeline. Using a novel numerical homogenization approach, we process a large set of parameterized unit cell patches to evaluate their in-plane contraction as well as stretching and bending stiffnesses. Our inverse design optimization tracks the inflated equilibrium state and blends suitably chosen unit cells with good structural performance to best approximate a given target surface.

Surface-based inflatables are composed of two thin layers of nearly inextensible sheet material joined together along carefully selected fusing curves. During inflation, pressure forces separate the two sheets to maximize the enclosed volume. The fusing curves restrict this expansion, leading to a spatially varying in-plane contraction and hence metric frustration. The inflated structure settles into a 3D equilibrium that balances pressure forces with the internal elastic forces of the sheets.

We present a computational framework for analyzing and designing surface-based inflatable structures with arbitrary fusing patterns. Our approach employs numerical homogenization to characterize the behavior of parametric families of periodic inflatable patch geometries, which can then be combined to tessellate the sheet with smoothly varying patterns. We

Authors' addresses: Y. Ren, EPFL, yingying.ren@epfl.ch; J. Panetta, UC Davis, julian.panetta@gmail.com; S. Suzuki, EPFL, seiichi.suzuki@epfl.ch; U. Kusupati, EPFL, uday.kusupati@epfl.ch; F. Isvoranu, EPFL, florin.isvoranu@epfl.ch; M. Pauly, EPFL, mark.pauly@epfl.ch.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM 0730-0301/2024/7-ART87

<https://doi.org/10.1145/3658125>

propose a novel parametrization of the underlying deformation space that allows accurate, efficient, and systematic analysis of the stretching and bending behavior of inflated patches with potentially open boundaries.

We apply our homogenization algorithm to create a database of geometrically diverse fusing patterns spanning a wide range of material properties and deformation characteristics. This database is employed in an inverse design algorithm that solves for fusing curves to best approximate a given input target surface. Local patches are selected and blended to form a global network of curves based on a geometric flattening algorithm. These fusing curves are then further optimized to minimize the distance of the deployed structure to target surface. We show that this approach offers greater flexibility to approximate given target geometries compared to previous work while significantly improving structural performance.

CCS Concepts: • **Applied computing** → **Computer-aided design**; *Computer-aided manufacturing*; • **Computing methodologies** → **Physical simulation**; *Mesh geometry models*.

Additional Key Words and Phrases: deployable structure, physics-based simulation, homogenization, metamaterial, numerical optimization, computational design, fabrication

ACM Reference Format:

Yingying Ren, Julian Panetta, Seiichi Suzuki, Uday Kusupati, Florin Isvoranu, and Mark Pauly. 2024. Computational Homogenization for Inverse Design of Surface-based Inflatables. *ACM Trans. Graph.* 43, 4, Article 87 (July 2024), 18 pages. <https://doi.org/10.1145/3658125>

1 INTRODUCTION

Shape-morphing structures can transition between multiple geometric configurations. Unlike traditional kinematic mechanisms, shape-morphing structures are usually compliant, lightweight, and bending-active, which facilitates a range of applications in diverse fields such as space technology, consumer products, medical devices, soft robotics, and architecture [Rivas-Adrover 2015].

In this paper, we focus on inflatable shape-morphing structures that are fabricated by joining two sheets of thin membrane, potentially perforated with holes, along a large number of fusing curves. We call these structures *surface-based inflatables* (Figure 2). When inflated, the unfused regions of the membrane will separate to maximize the enclosed volume, which induces an in-plane contraction. By prescribing suitably designed fusing and cutting patterns, we can achieve spatially varying contraction and hence metric frustration that pushes the structure into a curved 3D shape.

Surface-based inflatables are characterized by the inflated thickness being small relative to the overall sheet extent so that the inflated structure can be considered as a thin shell. This kind of inflatable structure has been studied in prior work with approaches tailored to certain specific classes of fusing curves, e.g. [Gao et al. 2020; Ou et al. 2016; Panetta et al. 2021; Siéfert et al. 2020]. Restricting to a limited class of fusing curves not only reduces aesthetic diversity but also has implications for the range of target shapes that can be approximated (due to the limited gamut of achievable contraction behaviors) and the structural performance of the deployed inflatable. For example, the inverse design framework introduced by Panetta et al. [2021] generates networks of parallel tubes that suffer from an extreme anisotropy of bending and stretching stiffnesses: in the case of straight tubes, the inflated structure is effectively a collection of stiff beams connected by hinges of negligible stiffness that permit the sheet to easily roll and collapse transversely. This can lead to a catastrophic loss of structural stability, as documented in [Gao et al. 2020], and presents difficulties for shape approximation since tubes will resist the bending needed to follow directions of high normal curvature. Ensuring stiffness of the deployed structure is crucial for many applications, e.g. in architecture, and together with improving surface approximation power is a key motivation for our investigation of more general fusing curve geometries as well as patterns featuring holes.

Since searching over all possible topologies and geometries of a full-scale inflatable is computationally intractable at high resolution, we propose a two-scale approach as illustrated in Figure 1. In a precomputation phase, we define several families of procedurally generated periodic patterns that we use to instantiate a large number of periodic cell geometries. We then apply a novel, efficient numerical homogenization approach to analyze their effective in-plane contractions and material properties under inflation. These results are organized in a database that can be queried for cells with desired contraction values and stiffness profiles and provides a smoothly interpolated pattern-to-property map for gradient-based optimization. With these offline computations complete, we can efficiently solve the inverse design problem in two stages. We first apply a flattening optimization algorithm to the design surface that determines both the uninflated shape of the structure and the spatially




Fig. 2. Surface-based inflatable unit cell. Two thin, planar sheets are joined together along a pattern of fusing curves (top right). Certain regions can be cut out to create a permeable cell (bottom left). Using periodic boundary conditions, we can simulate inflation, which causes the cell to bulge vertically and contract in plane (bottom right). Our framework characterizes the geometric and physical properties of unit cells with arbitrary fusing and cutting patterns.

varying geometric parameters whose corresponding homogenized properties best encode the desired target shape. We then blend these patterns into a global network of fusing curves and holes that is post-optimized to improve target-surface fitting. Crucially, our analysis yields an analytical formula for the bending stiffnesses of the patterns that can be incorporated directly into the flattening to improve the extrinsic curvature approximation.

Contributions. The main goal of our work is to build a computational framework that can accurately and efficiently analyze the deformation behavior and material properties of inflatables with arbitrary fusing and cutting patterns. To achieve this goal, we make the following core technical contributions:

- We develop a numerical homogenization technique for periodic inflatable patches to calculate the in-plane contraction and stiffnesses of arbitrary patterns, including those with open volume.
- We propose a highly efficient approach for characterizing the small-strain bending stiffnesses of thin plates that admits a concise analytical formula for the exact stiffness along any single bending axis.
- We present a database of patterns that exhibit a large range of material properties and deformation behaviors.
- We introduce a complete inverse design pipeline that incorporates information from the pattern database to design inflatables that satisfy multiple objectives, including shape approximation, structural soundness, and pattern aesthetics.

We fabricate several physical prototypes to validate our predictive model and illustrate potential applications of our new class of surface-based inflatables. Please find the material database of fusing patterns and the full implementation of our homogenization and inverse design framework at <https://go.epfl.ch/Inflatables>.

2 RELATED WORK

Active research in computer graphics, material science and mechanical engineering has contributed to the emergence of a wide range of shape morphing structures [Oliver et al. 2016]. A popular subset are deployable structures that transform from a planar state to a curved 3D surface, which offers significant advantages in terms of fabrication. Examples include hydrogels [Gladman et al. 2016], auxetic materials [Chen et al. 2021; Konakovic-Lukovic et al. 2018], rod or tiles on pre-stretched fabrics [Guseinov et al. 2017; Jourdan et al. 2020, 2022; Pérez et al. 2017], and deployable linkage systems [Becker et al. 2023; Panetta et al. 2019; Pillwein et al. 2020; Ren et al. 2022].

Inflatable structures are lightweight, compliant, and easy to deploy, and hence have found use in a variety of application domains, such as soft robotics [Andrade-Silva and Marthelot 2023], HCI [Niiyama et al. 2020], emergency shelters [Melancon et al. 2021], space technology [Schenk et al. 2014], and medical devices [Hsiao et al. 2019]. Prior research has also explored combining inflatable structures with other elastic elements [Jones et al. 2023], rigid elements [Yang et al. 2023], and sensors [Tang et al. 2022]. We classify pure inflatable structures into three categories based on their dimensions: curve-based, surface-based, and volumetric inflatables (Figure 3).

Curve-based inflatables deploy into tubular approximations of space curves that can be combined with other sensors or actuators [Baines et al. 2023; Luo et al. 2022] and can be highly effective as soft grippers [Andrade-Silva and Marthelot 2023]. One strategy for shaping the inflation of such tubular structures is by altering the membrane’s material properties, *e.g.*, by embedding stiff kirigami structures [Jin et al. 2020], printing the membrane using an optimized arrangement of two different materials [Tanaka et al. 2023], or with a 3D-printed reinforcement designed with topology optimization [Maestre et al. 2023]. These works inspire future directions for using the surface-based inflatables we address in this paper in similar soft robotics applications.

Volumetric inflatables approximate 3D shapes with a single-layer membrane, pressuring the shape’s interior. The pressure acts to invert any inward bumps, preventing approximation of surfaces with negative mean curvature unless additional constraints such as interior supports or strings are added. Skouras et al. [2014] proposed a system for designing collections of flat panels that stitch together and inflate into a target volumetric shape. Konakovic-Lukovic et al. [2018] construct auxetic kirigami patterns that can constrain balloons to inflate into target shapes. Melancon et al. [2021] introduce an origami-inspired structure that deploys via pneumatic inflation, locking into a second stable equilibrium.

Surface-based inflatables differ in that only a thin layer around the target surface is pressurized, enabling the approximation of non-water-tight shapes and surfaces with inward bumps. Ou et al. [2016] presented a simulation-based design interface for a variety of parameterized inflatable patterns that can function as hinge mechanisms to induce sheets to curl or twist. They also demonstrated a variety of fabrication methods and materials for realizing designs, including manual sealing and heat pressing. Experiments in [Zhang et al. 2023b] have further demonstrated this material system’s capacity for multistability. Siéfert et al. [2020] introduced a class of



Fig. 3. Representative examples of existing works on inflatables: (a) curve inflatables [Baines et al. 2023], (b) surface-based inflatables [Panetta et al. 2021], and (c) volumetric inflatables [Skouras et al. 2014]. Copyrights for the three images are held by the authors of the cited papers.

inflatables using networks of parallel air channels whose widths, spacings, and orientations can be varied to produce different Gaussian curvatures. Gao et al. [2020] generalized this design space using zigzag curves to achieve biaxial contractions and improved stiffness. Similar structures have been employed by Boutemy et al. [2023] as formwork for concrete shell construction. A common limitation of these works is the absence of an inverse design method for finding optimal fusing patterns to approximate a given target shape. Panetta et al. [2021] addressed this challenge in the specific case of parallel tube patterns, which suffer significant drawbacks as discussed in the introduction. We show that the larger set of pattern geometries unlocked by our contributions can improve shape approximation, boost structural stiffness, and offer new aesthetics.

Geometric abstractions for inverse design. At the core of the design algorithm introduced by Panetta et al. [2021] for parallel-tube inflatables is a coarsened formulation of the design optimization as a surface flattening problem that is based on a geometric abstraction of the tubes’ physics. Errors introduced by this approximation are then corrected by a native-scale physics-based shape optimization. Similar formulations have been applied for other physical systems involving either anisotropic contraction [Aharoni et al. 2018; Jourdan et al. 2023, 2022] or uniform expansion [Wang et al. 2022]. The bespoke geometric abstractions used in these works are generally defined in an *ad hoc* way, relying on intuitions that do not generalize to the more complex patterns considered in our work. Our primary technical contributions are related to defining rigorous abstractions for arbitrary periodic pattern geometries using numerical homogenization of the relevant physics and building from them a more general coarsened design optimization.

Homogenization. Many works in graphics and computational fabrication have employed homogenization to accelerate elasticity simulations [Chan-Lock et al. 2022; Chen et al. 2015; Kharevych et al. 2009; Sperl et al. 2020] and rigid-body simulation [Tang et al. 2023], and to design volumetric [Chen et al. 2018; Panetta et al. 2017, 2015; Schumacher et al. 2015], sheet [Chen et al. 2021; Li et al. 2022; Martínez et al. 2019; Schumacher et al. 2018; Tozoni et al. 2020; Zhang et al. 2023a] and ribbon [Rodriguez et al. 2022] metamaterials. To our knowledge, we are the first to develop a periodic homogenization technique for analyzing the contraction and the bending and stretching stiffnesses of inflated pattern tilings. Our method for calculating homogenized bending stiffnesses is most closely related to those of [Schumacher et al. 2018; Sperl et al. 2020], but we avoid a more expensive constrained optimization problem by introducing

a highly efficient formulation with a novel parametrization of the microscopic deformation state. Furthermore, we derive a minimal set of equations, involving analytically computed derivative terms, that can be solved to obtain a highly compact, exact expression for the directional uniaxial small-strain bending stiffness profile. Also related is [Tang et al. 2023], which analyzes bending *limits* (not stiffnesses) of chainmail fabrics composed of interlocking quasi-rigid elements by bending finite patches onto parabolic approximations of constant-curvature midsurfaces.

Stripe patterns for material design. The stripe patterns parametrization algorithm [Knöppel et al. 2015; Ray et al. 2006] has been employed as a dehomogenization technique in the design of structured materials, e.g. to convert the output of a homogenized compliance minimization algorithm to a network of orthogonal beams [Allaire et al. 2019] or to convert a homogenized shape-morphing structure into a pattern of tool paths for fusing or filament deposition [Jourd'an et al. 2023; Panetta et al. 2021]. Our dehomogenization step uses *two* orthogonal stripe patterns as texture coordinates to map our unit cell patterns onto the flattened sheet domain. Montes Maestre et al. [2023] recently proposed a single-scale optimization of bi-material structures within the space of stripe patterns, using X-FEM to avoid the need to explicitly resolve moving material interfaces with remeshing. We will discuss the relation of this work to our method in more detail in Section 8.3.

3 OVERVIEW

Figure 1 gives a visual overview of our approach. In a preprocessing phase, we build a database of inflatable unit cells with a variety of different fusing curve patterns.

We first review in Section 4 the simulation model based on previous work to compute the equilibrium shape of a pressurized inflatable. We then present in Section 5 our novel thin-shell homogenization technique that allows accurately, efficiently, and systematically analyzing the geometric and mechanical behaviors of inflatable unit cells. Conceptually, for each cell we analyze the properties of an infinite planar tiling that is simulated by enforcing a special form of periodic boundary conditions on the cell domain.

In Section 6, we present analytical formulas for computing the small-strain stretching and single-axis bending stiffnesses of inflated unit cells. We show that the complete uniaxial bending stiffness behavior of a thin plate is determined by exactly five scalar coefficients, from which the directions of weakest and strongest bending resistance can be determined.

We apply our homogenization framework to a diverse set of fusing curve patterns to build a database of unit cells. We show in Section 7 how suitably chosen parametric patches allow spanning a large gamut of macroscopic material behavior.

The database is then utilized in our global inverse design algorithm described in Section 8. This method takes a desired target shape as input and computes the planar fusing curves of the inflatable sheets. By incorporating a novel bending regularization formulation, the inverse optimization automatically selects patterns whose directional stiffnesses are well-suited for approximating the target surface.




Fig. 4. Tiling a unit cell. The model in the bottom left has been simulated with periodic boundary conditions as used in homogenization. The simulated and fabricated tilings with closed boundaries indicate how the geometric and physical properties of the center patch of these grids converge to the periodic result as the grid extends.

We present results in Section 9 and evaluate our approach on a variety of simulated and fabricated designs. In particular, we illustrate the improvements in shape approximation and the superior structural properties of our designs compared to the parallel tube patterns of [Panetta et al. 2021]. We conclude with a discussion of limitations of our approach and several avenues for future work.

4 SIMULATION MODEL

To compute the equilibrium shape of a pressurized inflatable, we adopt the formulation of [Panetta et al. 2021], which we briefly summarize here. The input to the simulation is a sheet domain $\mathcal{D} \subset \mathbb{R}^2$ with embedded fusing curves $\gamma \subset \mathcal{D}$ that define a partitioning of \mathcal{D} into fused regions \mathcal{D}_f and air channels \mathcal{D}_a . For improved stiffness, our work employs patterns with zero-width fusing, meaning $\mathcal{D}_f = \gamma$, while the parallel-tube inflatables of [Panetta et al. 2021] rely on finite-area fused regions of varying width. The sheet domain is then fused with a second copy of itself along $\partial\mathcal{D}$ to obtain a watertight surface $\hat{\Omega}$ with top and bottom copies of \mathcal{D}_f denoted as $\hat{\Omega}_f^+$ and $\hat{\Omega}_f^-$, respectively.

Simulating the equilibrium under applied pressure p is formulated as finding the deformation map $\Phi : \hat{\Omega} \rightarrow \mathbb{R}^3$ that minimizes the total potential energy:

$$\min_{\Phi} E_{\text{elastic}}[\Phi] - p \text{Vol}(\Omega) \quad \text{s.t.} \quad \Phi(\hat{\Omega}_f^+) = \Phi(\hat{\Omega}_f^-), \quad (1)$$

where $p \text{Vol}(\Omega)$ measures the work done by the pressure forces throughout the inflation into deformed shape $\Omega = \Phi(\hat{\Omega})$. The functional E_{elastic} measures the membrane energy stored by the deformation. As in [Panetta et al. 2021; Skouras et al. 2014], we omit a bending energy functional due to the thin sheets' negligible bending resistance and employ a neo-Hookean membrane energy density relaxed by tension field theory (neglecting compressive stresses that are released by wrinkling) [Steigmann 1990].

This simulation problem is discretized into a finite-dimensional optimization by meshing \mathcal{D} with triangles conforming to the poly-line approximations of γ and employing linear finite elements. For efficiency, the equality constraints in (1) are enforced by assigning the same deformation variables to vertices in $\hat{\Omega}_f^+$ and $\hat{\Omega}_f^-$, leading to an unconstrained optimization.

This simulation can be directly applied to study the inflation of small tilings of our pattern geometries as shown in Figure 4. However, the resulting behavior could change as the tiling is extended due to boundary effects, and it is not obvious how to define effective contraction factors and stiffnesses from the resulting simulations. These difficulties are addressed by our homogenization approach introduced in the next section, which provides a natural definition of the effective material properties in the limit of an infinite tiling.

5 PERIODIC HOMOGENIZATION

Our coarsened design optimization relies on an understanding of the metric change induced by inflating a pattern as well as the bending stiffness around the inflated state. Periodic homogenization [Nakshatrala et al. 2013; Schumacher et al. 2018; Sperl et al. 2020; Zhang et al. 2023a] is a rigorous approach for defining a periodic metamaterial's properties by studying the behavior of an infinite tiling of its pattern unit cell Y . The conceptually infinite simulation is made tractable by an assumption that the tiling's translational symmetry is preserved by the deformation¹. This enables us to represent the full tiling's deformation by introducing a compact parametrization of its admissible deformations and simulating the physics within just a single unit. Throughout the rest of this section, $\hat{\Omega}$ will refer to the inflatable patch contained within Y . We note that the period cell boundary $\partial Y \cap \hat{\Omega}$ is typically unfused, so $\hat{\Omega}$ is not necessarily watertight; we address this issue in Section 5.2.

5.1 Admissible Deformations of a Periodic Tiling

In the absence of bending, translational symmetry forces the deformed position of each material point $\mathbf{X} \in \hat{\Omega}$ to take the form:

$$\Phi(\mathbf{X}) = \bar{\mathbf{F}}\mathbf{X} + \mathbf{u}(\mathbf{X}), \quad \mathbf{u} \text{ periodic on } \partial Y, \quad (2)$$

where macroscopic deformation gradient $\bar{\mathbf{F}} \in \mathbb{R}^{3 \times 2}$ represents the average stretching experienced by each cell in the tiling, and $\mathbf{u} : \hat{\Omega} \rightarrow \mathbb{R}^3$ is a periodic *fluctuation displacement* allowing the sheet to relax into a lowest-energy configuration consistent with that average stretch. Periodicity of \mathbf{u} is necessary for the deformed patch $\Phi(\hat{\Omega})$ to stitch seamlessly with its deformed neighbors and ensures that the microscopic deformation gradient $\nabla\Phi$ indeed averages to $\bar{\mathbf{F}}$ over Y . Rigid rotation of this tiling (belonging to the nullspace of the simulation energy in (1)) can be eliminated with the constraint

$$\bar{\mathbf{F}} = \begin{bmatrix} s_1 & s_3 \\ s_3 & s_2 \\ 0 & 0 \end{bmatrix},$$

¹The nonconvexity of (1) admits the possibility of symmetry-breaking buckling modes whose periods span an ensemble of unit cells [Miehe et al. 2002]. Determining the size of these ensembles is a fundamental challenge. In this work, we expand our unit cells Y to contain 2×2 repetitions of the translational symmetry unit for pattern families that we empirically determine have buckling modes with that period.




Fig. 5. Bending the midsurface. Left and center: schematic illustrations of the variables used to describe cylindrical bending of the unit cell midsurface. Right: an example inflatable patch with a bent equilibrium state along with the macroscopic midsurface.

which can be interpreted as replacing $\bar{\mathbf{F}}$ with the symmetric part of its polar decomposition (defined with appropriate padding) and forces the tiling grid to remain parallel to the $x_3 = 0$ coordinate plane. In the following, we denote $\bar{\mathbf{S}}$ as the symmetric nonzero block of $\bar{\mathbf{F}}$. Finally, rigid translation can be eliminated by pinning $\mathbf{u}(\mathbf{X}_0) = \mathbf{u}_c$ to a constant for an arbitrary material point \mathbf{X}_0 .

5.1.1 Bending the midsurface. The deformed state described by (2) resembles a thin, flat plate whose midsurface we can shift to $x_3 = 0$ by adjusting \mathbf{u}_c . We now consider bending of this midsurface. For all cells to remain congruent to one another, the midplane must be bent into a surface of constant curvature. However, really only singly-curved surfaces are admissible since introducing nonzero Gaussian curvature requires spatially varying metric distortion that would violate the tiling's translational symmetry. Due to this complication, like in [Sperl et al. 2020], we restrict our analysis to single-axis bending, where the midplane wraps onto a cylinder with curvature $\bar{\kappa}$ and unit axis vector $\hat{\mathbf{a}}$ and principal curvature direction $\hat{\mathbf{a}}^\perp$. We note that characterizing single-axis bending behavior is most relevant to our design problem since the stretching inherent in biaxial bending will be resisted by our patterns' robust stretching stiffnesses, penalizing deviations away from the target surface's metric encoded in their contractions. Unlike Sperl et al. [2020], who consider bending only along two fixed orthogonal directions, we permit arbitrary bending axes $\hat{\mathbf{a}}$ in simulation and derive a minimal set of equations that can be efficiently solved to fully characterize the small-strain uniaxial bending stiffnesses of planar equilibria.

Picking a point \mathbf{c} on the $x_3 = 0$ plane to remain fixed during the wrapping, we can obtain the cylindrically-wrapped coordinates $\hat{\mathbf{x}}$ of an arbitrary point $\mathbf{x} = \mathbf{c} + \mathbf{v}$ also on the plane by traversing the geodesic corresponding to tangent vector \mathbf{v} : $\hat{\mathbf{x}} = \exp_c(\mathbf{v})$ (employing a *geodesic polar coordinate system* for the cylinder [Do Carmo 2016; Melv er and Reimers 2012]). Through this traversal, the midplane normal rotates by an angle of $\bar{\kappa}(\hat{\mathbf{a}}^\perp \cdot \mathbf{v})$ around axis $\hat{\mathbf{a}}$. Therefore, the normal at $\hat{\mathbf{x}}$ is $R(\boldsymbol{\omega})\mathbf{e}_3$, where $R(\boldsymbol{\omega}) = e^{[\boldsymbol{\omega}]_\times}$ is the rotation matrix corresponding to angle-axis vector $\boldsymbol{\omega} := \bar{\kappa}\hat{\mathbf{a}}(\hat{\mathbf{a}}^\perp \cdot \mathbf{v})$. The following explicit formula for $\hat{\mathbf{x}}$ is derived in the supplement by integration:

$$\hat{\mathbf{x}} = \mathbf{x} + \xi(\boldsymbol{\omega})\mathbf{v}, \quad (3)$$

$$\xi(\boldsymbol{\omega}) := \left(\frac{1 - \cos(\|\boldsymbol{\omega}\|)}{\|\boldsymbol{\omega}\|^2} \right) [\boldsymbol{\omega}]_\times + \left(\frac{\|\boldsymbol{\omega}\| - \sin(\|\boldsymbol{\omega}\|)}{\|\boldsymbol{\omega}\|^3} \right) [\boldsymbol{\omega}]_\times^2.$$

Here $[\boldsymbol{\omega}]_{\times}$ denotes the skew-symmetric cross product matrix. We further show in the supplement that *under single-axis bending*, different selections of \mathbf{c} produce wrappings differing only by a rigid transformation (and thus identical homogenized properties). Consequently, \mathbf{c} can be picked arbitrarily.

5.1.2 Double curvature. Our formulas generalize to doubly curved mid-surfaces by setting $\boldsymbol{\omega} = \nabla\boldsymbol{v}$ with $\nabla\boldsymbol{v} = \bar{\kappa}_1 \hat{\mathbf{a}} \otimes \hat{\mathbf{a}}^{\perp} - \bar{\kappa}_2 \hat{\mathbf{a}}^{\perp} \otimes \hat{\mathbf{a}}$. This new expression reduces to the original when $\bar{\kappa}_2 = 0$, and we observe that the upper-left two-by-two block of

$$[\mathbf{e}_3]_{\times} \nabla\boldsymbol{\omega} = -\bar{\kappa}_1 \hat{\mathbf{a}}^{\perp} \otimes \hat{\mathbf{a}}^{\perp} - \bar{\kappa}_2 \hat{\mathbf{a}} \otimes \hat{\mathbf{a}}$$

is the second fundamental form (and thus the macroscopic bending strain $\bar{\epsilon}_b$) of the surface evaluated at \mathbf{c} . Unfortunately, when $\bar{\kappa}_1 \bar{\kappa}_2 \neq 0$, the resulting deformation necessarily violates periodicity, introducing distortion that depends on the choice of \mathbf{c} (see the supplement for an illustration).

Our work effectively constrains $\bar{\epsilon}_b$ to states of single-curvature that we parametrize using the macroscopic variables $(\bar{\kappa}, \bar{\alpha})$. Angle $\bar{\alpha}$ represents the principal curvature direction $\hat{\mathbf{a}}^{\perp} = [\cos(\bar{\alpha}), \sin(\bar{\alpha}), 0]^{\top}$. The singularity in this parametrization at $\bar{\kappa} = 0$ turns out to be important for the bending stiffness analysis of Section 6.

5.1.3 Bending the full space. Formula (3) for $\hat{\mathbf{x}}$ works only for points on the $x_3 = 0$ plane. To extend the bending map to the entire space, we employ the following Kirchhoff-Love-type deformation:

$$\hat{\mathbf{x}}(\mathbf{x}) = \mathbf{x}_{2D} + \xi(\boldsymbol{\omega})(\mathbf{x}_{2D} - \mathbf{c}) + h\mathbf{e}^{[\boldsymbol{\omega}]_{\times}} \mathbf{e}_3, \quad (4)$$

where $\mathbf{x}_{2D} = (I - \mathbf{e}_3 \otimes \mathbf{e}_3)\mathbf{x}$ and $h = x_3 = u_3$. In other words, points at height h above the $x_3 = 0$ plane remain at offset h from the bent plane along the rotated normal $\mathbf{e}^{[\boldsymbol{\omega}]_{\times}} \mathbf{e}_3$. We emphasize that this formula does not actually constrain the microscopic deformation to obey Kirchhoff-Love kinematic assumptions as the deformed position $\hat{\mathbf{x}}$ corresponding to an interior point \mathbf{X} still can move freely to minimize energy by adjusting all components of $\mathbf{u}(\mathbf{X})$. Rather, it provides a computationally convenient way to *define* the midsurface offset in terms of closest point projections. The approach in [Sperl et al. 2020] adds the fluctuation displacement after midsurface bending, while we add it before, which we find conceptually simpler and easier to implement (avoiding the need for parallel transport).

Now, suppose that the translational pin constraint value \mathbf{u}_c was picked inappropriately so that the true midplane was at

$$\bar{u}_3 := \frac{1}{|\hat{\Omega}|} \int_{\hat{\Omega}} u_3(\mathbf{X}) \, dA \neq 0$$

before bending. As detailed in the supplementary material, we can re-parameterize the macroscopic state $(\bar{S}, \bar{\kappa}, \bar{\alpha}, \mathbf{u})$ to enforce $\bar{u}_3 = 0$ *without changing the represented deformation*, moving within a nullspace. This re-parametrization strategy enables us to replace the expensive equality constraint formulation of [Sperl et al. 2020] (which requires solving a saddle point problem) with an inexpensive pin constraint. While our approach is not well suited for imposing specific macroscopic strains $(\bar{S}, \bar{\kappa}, \bar{\alpha})$ as needed in that work, it is highly efficient for computing the minimum energy state under an imposed pressure and detecting non-planar equilibria. It also enables calculation of small-strain bending stiffnesses as detailed in Section 6.

5.1.4 The cell problem. Having established our parametrization of the microscopic deformation, we can now solve for the macroscopic contraction \bar{S} and bending state $(\bar{\kappa}, \bar{\alpha})$ by minimizing the total potential energy within the unit cell:

$$\begin{aligned} \min_{\bar{S}, \bar{\kappa}, \bar{\alpha}, \mathbf{u}} E[\hat{\mathbf{x}}(\bar{F}\mathbf{X} + \mathbf{u})] \\ \text{s.t. } \mathbf{u}(\hat{\Omega}_f^+) = \mathbf{u}(\hat{\Omega}_f^-), \mathbf{u} \text{ periodic, } \mathbf{u}(\mathbf{X}_0) = \mathbf{u}_c. \end{aligned} \quad (5)$$

Here, $E[\Phi] = E_{\text{elastic}}[\Phi] - p \text{Vol}(\Phi(\hat{\Omega}))$ denotes the total potential energy of (1), and $\hat{\mathbf{x}}(\cdot)$ depends implicitly on $\bar{\kappa}$ and $\bar{\alpha}$ via the formulas of the previous subsections.

5.2 Periodic Volume Calculation

One issue remains in our homogenization formulation (5): how to unambiguously define the volume enclosed by a single inflatable patch of the tiling when it contains openings at the cell boundary. We prove here that *any* way of delineating the volume attributed to Y by inserting end cap surfaces that are compatible with the periodic tiling yields the same definition of $\text{Vol}(\Phi(\hat{\Omega}))$ that can be computed solely in terms of unclosed surface $\Phi(\hat{\Omega})$. Note that once the invariance to end cap surface geometry is established, we can simply close the holes by an arbitrary periodic triangulation of the open boundary loops (inserting no new vertices), which we found simplified the implementation.

Proof: Label each pair of identified boundaries of Y as ∂Y_i^{\pm} ; for the rectangular unit cells studied in this work, $i \in \{1, 2\}$, but our formulation applies to other tiling cells (e.g., hexagonal). Now consider filling the open holes of $\Phi(\hat{\Omega} \cap \partial Y)$ with arbitrary surfaces denoted by Γ_i^{\pm} to obtain the closed-off surface $\bar{\Omega}$. For $\bar{\Omega}$ to tile properly, we require that $\Gamma_i^+ = R_i \Gamma_i^- + \mathbf{t}_i$, where R_i and \mathbf{t}_i are the rotation and translation relating the two identified boundaries that have explicit formulas in terms of the macroscopic state $(\bar{S}, \bar{\kappa}, \bar{\alpha})$ given in the supplement. This implies that the outward-pointing normals of these inserted surfaces satisfy $\mathbf{n}|_{\partial \Gamma_i^+} = -R_i \mathbf{n}|_{\partial \Gamma_i^-}$.

Using a standard surface integral formula for the volume enclosed by $\bar{\Omega}$ yields:


$$3 \text{Vol}(\bar{\Omega}) = \int_{\bar{\Omega}} \mathbf{n} \cdot \mathbf{x} \, dA = \int_{\Omega} \mathbf{n} \cdot \mathbf{x} \, dA + \sum_i \left(\int_{\Gamma_i^+ \cup \Gamma_i^-} \mathbf{n} \cdot \mathbf{x} \, dA \right).$$

The contribution from each pair of inserted cap surfaces depends only on \mathbf{t}_i and the *integral* of the cap's normal:

$$\int_{\Gamma_i^+ \cup \Gamma_i^-} \mathbf{n} \cdot \mathbf{x} \, dA = \int_{\Gamma_i^-} \mathbf{n} \cdot \mathbf{x} - (R_i \mathbf{n}) \cdot (R_i \mathbf{x} + \mathbf{t}_i) \, dA = \mathbf{t}_i \cdot \int_{\Gamma_i^+} \mathbf{n} \, dA.$$

It follows from Stoke's theorem (see supplement) that this integral depends only on the hole boundaries,

$$\mathbf{t}_i \cdot \int_{\Gamma_i^+} \mathbf{n} \, dA = \frac{\mathbf{t}_i}{2} \cdot \int_{\partial \Gamma_i^+} \mathbf{x} \times d\mathbf{s},$$






Fig. 6. Representative unit cells for the different pattern families. From top to bottom: fusing curves, inflated patch, bending stiffness profile, and stretching stiffness profile. On the left, the parallel tubes pattern shows its characteristic highly anisotropic stiffness profiles. Note that the peak stiffnesses of this pattern are roughly an order of magnitude higher than the other patterns'. The central column shows the three parametric families we apply in our inverse design pipeline. On the right, we show a pattern with 3-fold rotational symmetry, which leads to fully isotropic stiffness profiles. Due to its very low stiffness, we do not use this pattern for inverse design.

and not how the hole is filled. Hence, the inflated cell's volume can be computed without reference to the cap surfaces as:

$$\text{Vol}(\bar{\Omega}) = \frac{1}{3} \int_{\Omega} \mathbf{n} \cdot \mathbf{x} \, dA + \frac{1}{6} \sum_i \left(\mathbf{t}_i \cdot \int_{\Omega \cap \partial Y_i^+} \mathbf{x} \times d\mathbf{s} \right). \quad (6)$$

6 STIFFNESS ANALYSIS

We compute homogenized stretching and bending stiffnesses of an inflated tiling around its equilibrium state using analytical formulas obtained via sensitivity analysis, the general idea of which has been introduced in past works [Chan-Lock et al. 2022; Chen et al. 2021;

Zhang et al. 2023a]. However, our method differs from the previous ones on calculating bending stiffnesses [Schumacher et al. 2018; Sperl et al. 2020], which fit constitutive models to a large number of sampled deformations at finite strains. We propose a novel approach for obtaining the *exact* infinitesimal uniaxial bending stiffnesses that uses a *minimal* number of solves. We will first introduce notations and the established technique for computing stiffness (Section 6.1, Section 6.2). Then we will discuss the challenges in characterizing bending stiffnesses and present our solution (Section 6.3).




Fig. 7. The five basis functions for single-axis bending stiffness $k_b(\bar{\alpha})$.

6.1 Generic Stiffness Calculation via Sensitivity Analysis

We start by denoting the full state vector defining the microscopic deformation as $\mathbf{y} = (\bar{S}, \bar{\kappa}, \bar{\alpha}, \mathbf{u})$ and \mathbf{y}^* as the equilibrium state solving (5). Calculating stiffnesses amounts to computing second derivatives of the averaged energy density stored in the inflated tiling with respect to an “independent” macro-state variable $\bar{\mathbf{y}}$ while letting the remaining state variables $\tilde{\mathbf{y}}$ relax into the lowest-energy configuration. To this end, we define energy density as a function of $\bar{\mathbf{y}}$,

$$E^*(\bar{\mathbf{y}}) := \frac{1}{\det \bar{S}^* |\mathcal{D}|} \min_{\tilde{\mathbf{y}}} E(\bar{\mathbf{y}}, \tilde{\mathbf{y}}),$$

committing a slight abuse of notation to write the energy functional from (5) now as a function of the decomposed state. For example, to compute stretching stiffness, we take $\bar{\mathbf{y}}$ to be \bar{S} and $\tilde{\mathbf{y}}$ to be $(\bar{\kappa}, \bar{\alpha}, \mathbf{u})$. We note that the energy density is defined with respect to the inflated cell patch area (accounting for the constant $\det \bar{S}^*$ factor from \mathbf{y}^*) because we think of the inflated equilibrium as defining the rest-state of the thin-shell metamaterial whose stiffnesses we seek.

The tangent stiffness with respect to $\bar{\mathbf{y}}$ around the inflated equilibrium $\bar{\mathbf{y}}^*$ is defined as the second derivative [Chen et al. 2021]:

$$\frac{d^2 E^*}{d\bar{\mathbf{y}}^2}(\bar{\mathbf{y}}^*) = \frac{\partial^2 E}{\partial \bar{\mathbf{y}}^2} - \frac{\partial^2 E}{\partial \tilde{\mathbf{y}} \partial \bar{\mathbf{y}}} \left(\frac{\partial^2 E}{\partial \tilde{\mathbf{y}}^2} \right)^{-1} \frac{\partial^2 E}{\partial \tilde{\mathbf{y}} \partial \bar{\mathbf{y}}}, \quad (7)$$

where all partial derivative terms on the right are evaluated at $\bar{\mathbf{y}}^*$. The rightmost term is evaluated using a Cholesky factorization of the same sparse Hessian matrix factorized during the Newton solve of (5), but with rows and columns corresponding to $\tilde{\mathbf{y}}$ removed. Care must be taken as detailed in the supplement to maintain $\bar{u}_3 = 0$ during sensitivity analysis.

6.2 Stretching Stiffness

The complete small-strain stretching stiffness behavior is captured by the fourth-order tangent homogenized membrane elasticity tensor C^s , which is defined by taking $\bar{\mathbf{y}} = \bar{S}$ in (7). For example, the uniaxial stretching stiffness k_s (i.e., the directional Young’s modulus [Schumacher et al. 2018]) along any requested direction specified




Fig. 8. Selected unit cells of the Random Voronoi family, from left to right: highest maximal bending stiffness, highest maximal stretching stiffness, highest minimal bending stiffness, and highest minimal stretching stiffness.

by unit vector $\mathbf{d} \in \mathbb{R}^2$ can be computed inexpensively as:

$$k_s(\mathbf{d}) = \frac{1}{(\mathbf{d} \otimes \mathbf{d}) : (C^s)^{-1} : (\mathbf{d} \otimes \mathbf{d})}.$$

6.3 Bending Stiffness

Ideally, the same approach could be used to characterize bending stiffnesses as for stretching: compute the tangent bending stiffness tensor $C^b := \frac{d^2 E^*}{d\bar{\mathbf{e}}_b^2}(\bar{\mathbf{e}}_b^*)$ by taking the second fundamental form $\bar{\mathbf{e}}_b$ as the independent macro-state variable $\bar{\mathbf{y}}$. Unfortunately, as discussed in Section 5.1.2, imposing a full-rank second fundamental form is incompatible with the periodic boundary conditions used in homogenization, and we therefore restricted ourselves to analyzing single-axis-bending configurations parametrized by $(\bar{\kappa}, \bar{\alpha})$. This precludes calculating the necessary derivatives directly. The challenge of this section is to nevertheless completely characterize the directional single-axis bending stiffness of planar tilings exactly and with a minimal set of computations, all while exclusively probing states of single curvature.

Around bent equilibria with $\bar{\kappa}^* \neq 0$, stiffnesses with respect to perturbing curvature magnitude and bending axis can be obtained simply via (7) with $\bar{\mathbf{y}} = (\bar{\kappa}, \bar{\alpha})$. However, for the planar equilibria relevant to our coarsened design optimization (which relies on actuated metric distortion rather than bending to program surface curvature), this direct approach fails to fully characterize bending stiffness. In fact, entries $\frac{\partial^2 E^*}{\partial \bar{\alpha}^2}$ and $\frac{\partial^2 E^*}{\partial \bar{\kappa} \partial \bar{\alpha}}$ of the resulting stiffness tensor vanish when $\bar{\kappa}^* = 0$: at this coordinate singularity, $\bar{\mathbf{e}}_b$ (and thus E^*) is constant with respect to $\bar{\alpha}$, and, by optimality, $\frac{\partial E^*}{\partial \bar{\kappa}}(0, \bar{\alpha}) = 0$




Fig. 9. Properties of unit cell pattern families. From left to right: in-plane reciprocal stretch factors in principal directions, minimal and maximal bending stiffness, and minimal and maximal stretching stiffness. Colors represent minimal bending stiffness for each patch, where darker colors indicate higher values.

independent of $\bar{\alpha}$. The single nonzero entry $\frac{\partial^2 E^*}{\partial \bar{\kappa}^2}$ reveals stiffness only along the current, arbitrary bending axis $\bar{\alpha}^*$. While $\bar{\alpha}^*$ can be modified to probe stiffnesses along different bending directions (without altering the planar equilibrium represented), this would require evaluating (7) for each such direction, making high-resolution profiles expensive to compute. This difficulty is fundamental to the rank-one constraint on $\bar{\epsilon}_b$ and is not eased by choosing, e.g., scaled-axis vector variable $\bar{\kappa}\hat{\mathbf{a}}$ as $\bar{\mathbf{y}}$.

Our solution is to assume the existence of C^b , determine exactly which of its components are probed by single-axis bending experiments on planar equilibria, and derive equations for characterizing them in the $(\bar{\kappa}, \bar{\alpha})$ parametrization. We first evaluate the uniaxial bending stiffness k_b using the (unknown) tensor C^b :

$$k_b(\bar{\alpha}) = \frac{\partial^2 E^*}{\partial \bar{\kappa}^2}(0, \bar{\alpha}) = (\mathbf{d} \otimes \mathbf{d}) : C^b : (\mathbf{d} \otimes \mathbf{d}), \quad \mathbf{d} = [\cos(\bar{\alpha}), \sin(\bar{\alpha})]^\top,$$

where principal curvature direction \mathbf{d} is the unpadding, 2D version of $\hat{\mathbf{a}}^\perp$. This simplifies to the following expression after accounting for major and minor symmetries of C^b :

$$k_b(\bar{\alpha}) = \sum_{i=0}^4 q_i \cos(\bar{\alpha})^{4-i} \sin(\bar{\alpha})^i, \quad (8)$$

with $\mathbf{q} = \{C_{1111}^b, 4C_{1112}^b, 2C_{1122}^b + 4C_{1212}^b, 4C_{2212}^b, C_{2222}^b\}$. Thus, five of the six independent components of C^b participate in uniaxial bending, and effects of moduli C_{1122}^b and C_{1212}^b cannot be disentangled. The basis functions associated with each q_i are visualized in Figure 7.

The five unknown coefficients in \mathbf{q} can be determined uniquely by solving an interpolation problem, characterizing uniaxial stiffness along five sampled directions $\bar{\alpha}_i$ and solving a five-by-five linear system. We emphasize that our approach is *exact*, as opposed to the least-squares approximations used in past work (e.g., fitting a linear model to many finite-strain samples [Schumacher et al. 2018]); these samples form a minimal set of constraints defining the coefficients \mathbf{q} that *precisely* reconstruct the ground-truth uniaxial

bending stiffness profile $\frac{\partial^2 E^*}{\partial \bar{\kappa}^2}(0, \bar{\alpha})$ known to take form (8). We chose directions $\bar{\alpha}_i = \frac{\pi i}{5}$, which empirically appear to yield a near-optimal condition number for the system (~ 4.9), ensuring an accurate and robust fit. Each directional stiffness sample $k_b(\bar{\alpha}_i)$ is obtained by *updating* $\bar{\mathbf{y}}^*$ with $\bar{\alpha}^* = \bar{\alpha}_i$ and then applying (7) with $\bar{\mathbf{y}} = (\bar{\kappa}, \bar{\alpha})$. For efficiency, we only actually compute the nonzero stiffness entry of interest, $\frac{\partial^2 E^*}{\partial \bar{\kappa}^2}$. While the Hessian blocks in (7) that involve $\bar{\mathbf{y}}$ must be recomputed for each of these stiffness samples, the Hessian factorization can be reused, making the cost of bending stiffness characterization comparable to one Newton iteration.

The parametrization singularity also has implications for capturing buckling instabilities during simulation. A structure with a flat rest state must break symmetry to buckle into a non-planar equilibrium, and this happens by following a direction of negative curvature in the potential energy landscape away from an unstable equilibrium. When homogenizing inflatables, however, there exist unstable planar equilibria whose physical directions of negative curvature are invisible under certain settings of $\bar{\alpha}^*$. This manifests itself in our framework as the solver returning a purportedly stable planar equilibrium whose $k_b(\bar{\alpha})$ profile computed with our approach reveals negative stiffnesses. In this case we can set $\bar{\alpha}^*$ to one of the negative stiffness directions and restart the optimization, at which point our solver escapes along the newly revealed buckling mode.

7 HOMOGENIZED FUSED PATTERNS

Our homogenization framework allows computing the effective material properties of a unit cell patch from its fusing curves and holes. Specifically, we compute the two principal stretches and characterize the bending and membrane stiffness profiles. To build our unit cell database, we aim to create families of patterns with large ranges of scale factors and varied stiffness profiles. We experimented with the small set of parameterized families illustrated in Figure 6, but emphasize that our homogenization approach can be applied to arbitrary patterns that could be adapted to a specific application.

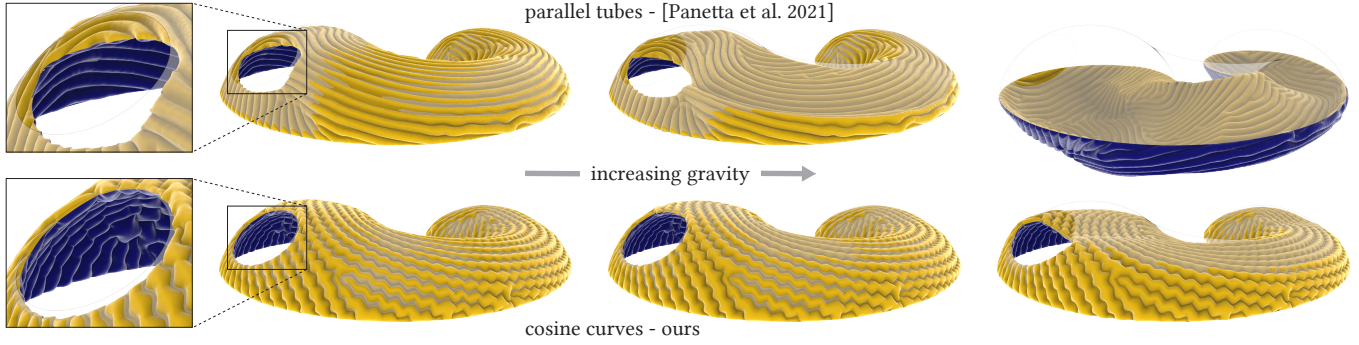


Fig. 10. Impact on the inflated shape of a fictitious gravity force. On the left, no external forces are applied; following the arrow, the gravity force increases. The simulation indicates a complete collapse and inversion of the parallel tube model while the cosine pattern model maintains its shape well.

To build our patch database, we opted for square unit cells with reflectional symmetry across both axes as required for our dehomogenization implementation based on two orthogonal stripe patterns. We enforce this symmetry by designing the fusing curve pattern in one quadrant and generating the other three by mirroring.

Motivated by the zigzag curves proposed in [Gao et al. 2020], our first pattern, the *cosine pattern*, uses a continuous cosine curve of varying amplitude. Note that the parallel tubes pattern of [Panetta et al. 2021] is a special case of this pattern as highlighted in Figure 6. When used for inverse design, the cosine pattern creates continuous fusing curves that often lead to long narrow tubes in the inflatable. This can potentially cause difficulties for inflation as documented in [Panetta et al. 2021]. As a remedy, the authors proposed a heuristic for removing small segments from the curves to promote better airflow and to guarantee that all air channels form one single inflatable volume.

A simpler and more direct approach is to use disconnected curves within each unit cell, which automatically ensures a single connected volume. The *dash pattern* is defined by a single oriented line segment, parameterized by orientation angle $\alpha \in [0, \pi/4]$ and length $l \in [0.2s, s]$, where s is the width of the unit cell square.

To illustrate that our approach is not limited to simple cell domains, we propose the *elliptical hole pattern* that is created by cutting out an elliptical region positioned at the quadrant center. This family is parameterized by three parameters, one for the orientation of the ellipse and two for lengths of its principal axes. Besides interesting structural properties, this pattern also has a unique visual quality and offers functional benefits by generating perforated inflatables with increased permeability.

We also experimented with stochastic patterns. The *Voronoi pattern* is constructed by placing five randomly sampled points into the domain and computing their periodic Voronoi diagram. Each Voronoi edge is then contracted by a factor of three to create a connected inflatable region. Figure 8 shows some representatives of this family with extreme properties.

We also studied the *3-star pattern* shown on the right in Figure 6, inspired by the kirigami-style auxetics used in [Konakovic-Lukovic et al. 2018]. With its threefold rotational symmetry, this pattern produces fully isotropic bending and stretching stiffness profiles. However, we observed that its bending resistance is too low to be

useful for inverse design. Please find inflated physical models of these patterns and validation experiments on their stiffnesses in the supplement.

Material Property Database. For each of our parametric fusing curve families, we run our homogenization framework in parallel on unit cells sampled on a dense rectangular grid in parameter space (Figure 9). Every unit cell is inflated to the same pressure, chosen to produce reasonable membrane strains across all analyzed cells.

Within a chosen family, each sampled parameter vector $\mathbf{p}_i \in \mathbb{R}^{n_p}$ is mapped by our homogenization to contraction $\bar{\mathbf{S}}^*$, stretching stiffness tensor C^s and bending stiffness coefficients \mathbf{q} . The reflectional symmetry of the patterns guarantees in absence of buckling that the principal stretch axes of $\bar{\mathbf{S}}^*$ are aligned with the pattern coordinates, meaning $s_3 = 0$ and only the two diagonal entries $\mathbf{s} = (s_1, s_2)$ need to be stored. The preserved symmetry of the inflated equilibrium means that C^s is *orthotropic*, containing only four nonzero moduli that must be recorded. Likewise, C^b has orthotropic symmetry ($C_{1112}^b = C_{2212}^b = 0$), meaning only three of the five bending stiffness coefficients are nonzero.

In summary, the effective behavior of mirror-symmetric patches can be encoded by nine nonzero coefficients. We fit tensor-product splines mapping from the regular parameter grid to each of these coefficients. These splines define smooth property functions $\mathbf{s}(\mathbf{p})$ and $\mathbf{q}(\mathbf{p})$ for the family that we use in our coarsened design optimization below. We further construct the convex hull of the principal contraction data $\mathbf{s}(\mathbf{p}_i)$ in the (s_1^{-1}, s_2^{-1}) reciprocal stretch plane for use in the local-global flattening scheme.

8 INVERSE DESIGN PIPELINE

We now turn to the problem of designing an inflatable structure that approximates an input target surface $\mathcal{M} \subset \mathbb{R}^3$ using one of our parametric families of unit cell patterns selected by the user. Our pipeline consists of three high-level phases: (i) a coarsened design problem determining a suitable rest shape of the sheet and spatially varying pattern parameters encoding the shape; (ii) a dehomogenization step converting this coarsened design into a seamlessly blended patterning of the sheet; and (iii) a physics-based shape optimization of the full-scale design.

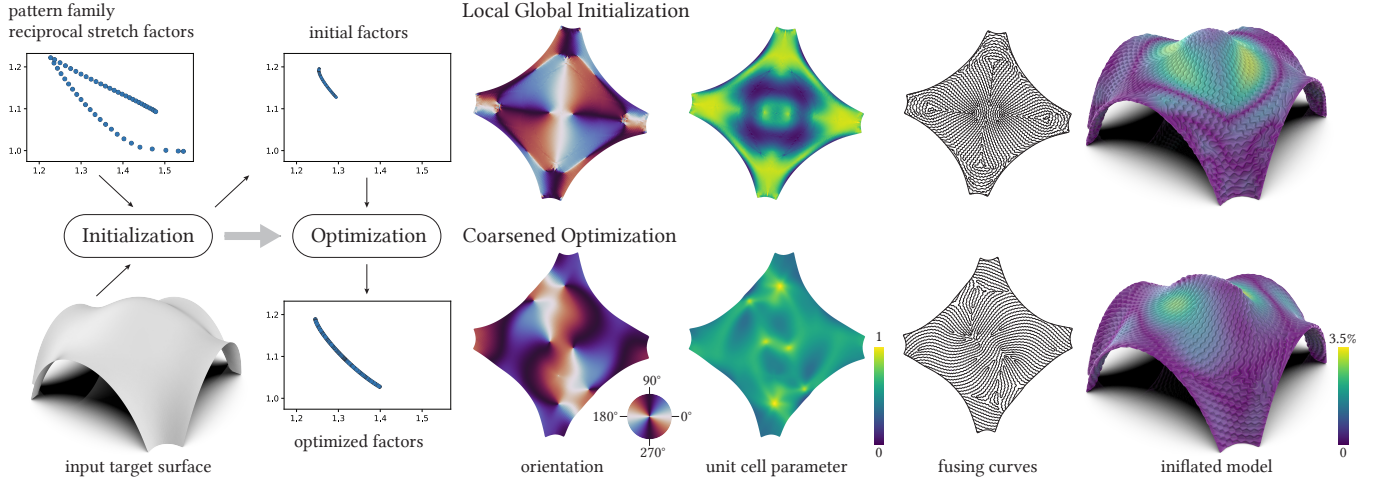


Fig. 11. Coarsened Design Optimization. For a given target surface and pattern family, the local/global flattening computes a first approximation of the scale factor distribution and corresponding cell orientation. These are converted to pattern parameters by the subsequent optimization, which also significantly improves the target fitting and globally adjusts the fusing curves pattern. Color-coding indicates the distance to the target surface relative to the bounding box diagonal.

8.1 Coarsened Design Optimization

We first solve a homogenized approximation of the design problem that generalizes the flattening algorithm specific to parallel-tubes of [Panetta et al. 2021]. To jointly find an optimal flattening $f : \mathcal{M} \rightarrow \mathbb{R}^2$ and parameter vector field $\mathbf{p} : \mathcal{M} \rightarrow \mathbb{R}^{n_p}$, we define the following weighted objective:

$$E_{\text{flat}}(f, \mathbf{p}) := E_{\text{fit}}(f, \mathbf{p}) + w_b E_{\text{bend}}(f, \mathbf{p}) + E_{\text{reg}}(f, \mathbf{p}). \quad (9)$$

The flattening's *inverse* f^{-1} serves as a homogenized model of deployment, with the singular value decomposition $\nabla f = U \Sigma V^T$ describing how principal stretch directions ($\mathbf{u}_1, \mathbf{u}_2$) in the sheet are scaled by *reciprocal* singular values ($\sigma_1^{-1}, \sigma_2^{-1}$) and lifted to corresponding vectors ($\mathbf{v}_1, \mathbf{v}_2$) in \mathcal{M} 's tangent plane. Our objective promotes mappings that are consistent with the homogenized material properties of a smoothly varying pattern.

The first term seeks to fit σ_1, σ_2 to the *reciprocal* principal stretches of the pattern field $\mathbf{s}(\mathbf{p})$:

$$E_{\text{fit}}(f, \mathbf{p}) := \frac{1}{2|\mathcal{M}|} \int_{\mathcal{M}} \left(\sigma_1 - 1/s_1(\mathbf{p}) \right)^2 + \left(\sigma_2 - 1/s_2(\mathbf{p}) \right)^2 dA. \quad (10)$$

Note that we are implicitly enforcing a hard constraint on the pattern orientation that snaps the pattern's principal contraction axes to $(\mathbf{u}_1, \mathbf{u}_2)$.

The bending regularization term penalizes patterns with stiff uniaxial bending stiffness along directions of high normal curvature of \mathcal{M} , reducing the pattern's resistance to this extrinsic curvature:

$$E_{\text{bend}}(f, \mathbf{p}) := \frac{1}{|\mathcal{M}|} \int_{\mathcal{M}} \int_0^{2\pi} \kappa_{\mathcal{M}}(\alpha)^2 k_b(\alpha; \mathbf{p}) d\alpha dA.$$

Here, $\kappa_{\mathcal{M}}(\alpha)$ denotes the normal curvature of \mathcal{M} along direction $\cos(\alpha)\mathbf{v}_1 + \sin(\alpha)\mathbf{v}_2$. By Euler's theorem for normal curvature, we have $\kappa_{\mathcal{M}}(\alpha) = \kappa_1 \cos^2(\alpha + \delta) + \kappa_2 \sin^2(\alpha + \delta)$, where κ_1, κ_2 are the two principal curvatures, and the phase offset δ represents the angle of \mathbf{v}_1 relative to the first principal curvature direction. The inner

integral can then be evaluated analytically after substituting in the bending stiffness expansion (8):

$$E_{\text{bend}}(f, \mathbf{p}) = \frac{1}{|\mathcal{M}|} \int_{\mathcal{M}} \left(\sum_{i=0}^4 q_i(\mathbf{p}) g_i(\delta, \kappa_1, \kappa_2) \right) dA. \quad (11)$$

Concise expressions for g_i obtained by symbolic integration are presented in the supplement.

Finally, the regularization terms promote the smooth variation of pattern orientations and parameters over the domain:

$$E_{\text{reg}}(f, \mathbf{p}) := \frac{w_\phi}{4} \int_{\mathcal{M}} \left\| \nabla \begin{bmatrix} \cos(2\phi) \\ \sin(2\phi) \end{bmatrix} \right\|_F^2 dA + \frac{w_s}{2} \int_{\mathcal{M}} \|\nabla \mathbf{p}\|_F^2 dA. \quad (12)$$

Here, ϕ is the angle between \mathbf{u}_1 (which coincides with the coordinate axes of our mirror-symmetric patterns) and the global x -axis of the flattened domain. The first term is unchanged from [Panetta et al. 2021], while the second term penalizes rapid variations in pattern parameters \mathbf{p} that would induce geometric distortions and worsen deviations from the homogenized properties.

We impose bounds $\mathbf{p}_{\min}, \mathbf{p}_{\max} \in \mathbb{R}^{n_p}$ *pointwise* on the parameter field \mathbf{p} to restrict to a region of parameter space yielding valid geometries with planar equilibria. For improved structural performance of the inflated equilibria, we adjust these bounds to exclude patterns with excessively low stiffnesses. Our coarsened design problem is therefore formulated as a bound-constrained minimization:

$$\min_{f, \mathbf{p}} E_{\text{flat}}(f, \mathbf{p}) \quad \text{s.t.} \quad \mathbf{p}_{\min} \leq \mathbf{p} \leq \mathbf{p}_{\max}. \quad (13)$$

8.1.1 Optimization strategy. The full objective E_{flat} is highly non-convex and requires a high-quality initialization. Starting from a least-squares conformal map [Lévy et al. 2002], we first apply a generalization of the anisotropic local-global optimization scheme from [Aharoni et al. 2018; Panetta et al. 2021]. In this scheme, all regularization terms are dropped, leaving only E_{fit} , and our optimization

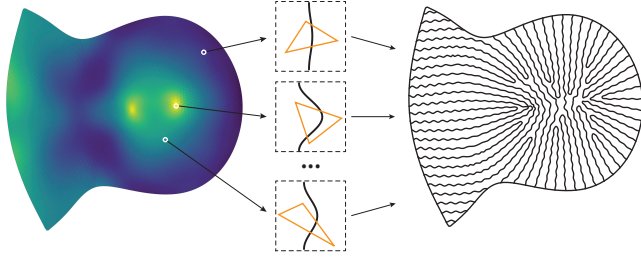


Fig. 12. Fusing curves (right) are generated by merging together per-triangle polylines defined by the pattern parameter field \mathbf{p} and a high-frequency stripe pattern texture map. On the left, we show the distribution of the cosine family’s 1D parameter field computed by our flattening optimization.

problem is re-written in the equivalent form:

$$\min_{f, \mathbf{s}} \min_{\substack{\tilde{U} \in SO(2) \\ \tilde{V} \in SO(3)}} \int_{\mathcal{M}} \left\| \nabla f - \tilde{U} \begin{bmatrix} s_1^{-1} & 0 & 0 \\ 0 & s_2^{-1} & 0 \end{bmatrix} \tilde{V}^T \right\|_F^2 dA,$$

where \mathbf{s} is now treated as an independent field (bypassing the $\mathbf{s}(\mathbf{p})$ interpolation). Holding f fixed, the minimization can be solved in closed form by taking $\tilde{U} = U$ and $\tilde{V} = V$ and projecting σ_1, σ_2 onto the convex hull of achievable reciprocal contraction factors constructed in Section 7. Then, holding $\mathbf{s}, \tilde{U}, \tilde{V}$ fixed, the optimal f is computed by solving a Poisson equation.

Finally, to solve the full problem (13), we first initialize \mathbf{p} by minimizing with the optimal f from the local-global scheme held fixed. Then we free f and run the full optimization until convergence. Please find more details on the weight-tuning scheme in the supplement. In Figure 11, we illustrate the changes in the design parameters and improvements in the inflated structures made by this full regularized optimization versus the local-global initialization.

8.1.2 Implementation details. We discretize (9) for a triangulated target surface \mathcal{M} using linear finite elements for f and a piecewise constant, per-element field for \mathbf{p} . To evaluate the principal curvatures and directions needed in E_{bend} , we average a per-vertex shape operator computed via [Panozzo et al. 2010] onto each triangle and then compute its Eigen-decomposition, obtaining piecewise constant curvature data. The regularization terms in E_{reg} measure the smoothness of piecewise constant fields and are implemented using an intrinsic Delaunay triangulation of \mathcal{M} ’s barycentric dual mesh as proposed in [Panetta et al. 2021]. We solve the discretized version of (13) using Knitro [Artelys 2019]. After optimization, we average the per-element parameter fields \mathbf{p} and orientation frames $(\mathbf{u}_0, \mathbf{u}_1)$ onto the vertices (using a circular mean for the latter) to define even smoother, piecewise linear fields used in dehomogenization.

8.2 Dehomogenization

The second phase of our design algorithm takes the homogenized design, in the form of a flattened domain $\mathcal{M}_f := f(\mathcal{M})$ with smoothly graded pattern parameters \mathbf{p} and orientation frames $(\mathbf{u}_1, \mathbf{u}_2)$, and produces the corresponding sheet geometry \mathcal{D} (potentially differing from \mathcal{M}_f due to holes) and fusing curves γ needed for full-scale simulation. Rather than explicitly remeshing \mathcal{M}_f with a coarse grid




Fig. 13. This model, optimized with the elliptical holes pattern, illustrates the importance of the full-scale shape optimization in the final stage of our inverse design pipeline. The initial model after dehomogenization is significantly improved by the subsequent optimization. Colors indicate distance to the target surface relative to the bounding box diagonal.

and inserting full cells retrieved from a database as in many past microstructure works [Chen et al. 2021; Panetta et al. 2015; Tozoni et al. 2020], we opt for a texture-mapping approach that generalizes the parallel tube synthesis of [Panetta et al. 2021].

We generate two orthogonal stripe pattern parametrizations on \mathcal{M}_f aligned with the pattern orientation frame. The resulting pair of fields can be used as texture coordinates into a pattern with reflectional symmetries across both axes. In our setting, the texture being sampled varies slowly across \mathcal{M}_f according to \mathbf{p} at the same time as the texture coordinates change. For the homogenization results to be faithful, the frequencies of the texture coordinates are ideally substantially higher than the frequencies of \mathbf{p} ; this relationship is controlled by the regularization weight w_s and the stripe pattern frequency parameter. Deviations from homogenized properties are of course inevitable at finite length-scales, and will be corrected by the full-scale optimization in Section 8.3.

For patterns without zero-thickness fusing curves (*i.e.*, those for which all curves of γ enclose a finite-area fused region), texture coordinates can be used to look up a signed distance field that is then contoured at high resolution to obtain γ ; this was how the original algorithm of [Panetta et al. 2021] traced its varying-width walls. Unfortunately, for the zero-thickness line fusing curves we employ in this work for enhanced stiffness, this approach breaks down due to the difficulty of robustly extracting curves from an *unsigned* distance field.

We therefore use an explicit approach based on boolean operations with polylines as illustrated in Figure 12. We generate a high-resolution extraction grid by subdividing \mathcal{M}_f and evaluate the stripe pattern texture coordinates on it. Then, for each triangle, we use the texture coordinates to map its three corners into the pattern reflectional symmetry unit cell. We perform a boolean intersection between the mapped triangle and the pattern fusing polyline geometry corresponding to the parameters sampled at the triangle center. We append the resulting line segments to the output after transforming them by the inverse of the triangle’s texture map. Because each triangle samples slightly different textures due to changes in \mathbf{p} , these segments do not connect perfectly, and at the end we apply a merging operation with a small threshold to reconnect them, obtaining the network γ .




Fig. 14. Fabrication setup (left) and material failures (right). Damage in the fabric can occur due to melting that leads to ruptures when over-inflating.

Finally, we mesh the combined polylines of γ and $\partial\mathcal{M}_f$ using Gmsh’s MeshAdapt algorithm [Geuzaine and Remacle 2009], obtaining the triangle mesh of \mathcal{D} required for simulation. In the input to Gmsh, we mark the inner boundary loops of \mathcal{M}_f and any polylines originating from internal boundaries of the pattern unit cell as hole boundaries.

8.3 Full-scale Shape Optimization

The final phase of our inverse design pipeline directly optimizes the geometry of \mathcal{D} and γ while evaluating the quality of the inflated equilibria. This full-scale physics-based optimization allows us to correct for errors introduced by violations of the infinitesimal-tiling assumptions of periodic homogenization and by geometric distortions from the dehomogenization process. We apply the original shape optimization algorithm of [Panetta et al. 2021], which is agnostic to the source of \mathcal{D} and γ . We briefly summarize the method for completeness and to highlight one change required to support our zero-thickness fusing curves; for the full mathematical formulation, please refer to the original paper.

The variables of the discrete shape optimization are just the rest vertex positions corresponding to the fused vertices lying on $\partial\mathcal{D}$ and γ ; the remainder of the rest mesh vertex positions are determined by interpolation to maintain quality. These positions are adjusted to minimize an objective combining the shape approximation error, a barrier term to prevent degeneration of the simulation mesh, and various regularizations to maintain curve smoothness. Minimization is performed with L-BFGS-B using exact gradients computed using analytical sensitivity analysis.

The fitting term measures a distance between the inflated equilibrium shape corresponding to the current design and the target surface. However, due to the thickness of the inflated state, there is some ambiguity in determining this distance, which ideally would be computed from the inflated midsurface. This ambiguity is resolved by measuring distance only from the fused regions $\Phi^*(\hat{\Omega}_f)$, which are assumed to lie on the midsurface. Therefore, the term is formulated as

$$\int_{\Phi^*(\hat{\Omega}_f)} \|\mathbf{x} - \mathcal{P}_{\mathcal{M}}(\mathbf{x})\|^2 d\mathbf{x},$$

where $\mathcal{P}_{\mathcal{M}}(\mathbf{x})$ projects \mathbf{x} to its closest point on \mathcal{M} . Panetta et al. [2021] formulate this term as a discrete area integral, but for zero-width curves, it must be a line integral, and corresponding changes are needed to the nondimensionalization factors.

While the designs produced by the coarsened optimization and dehomogenization already yield good surface approximation quality,

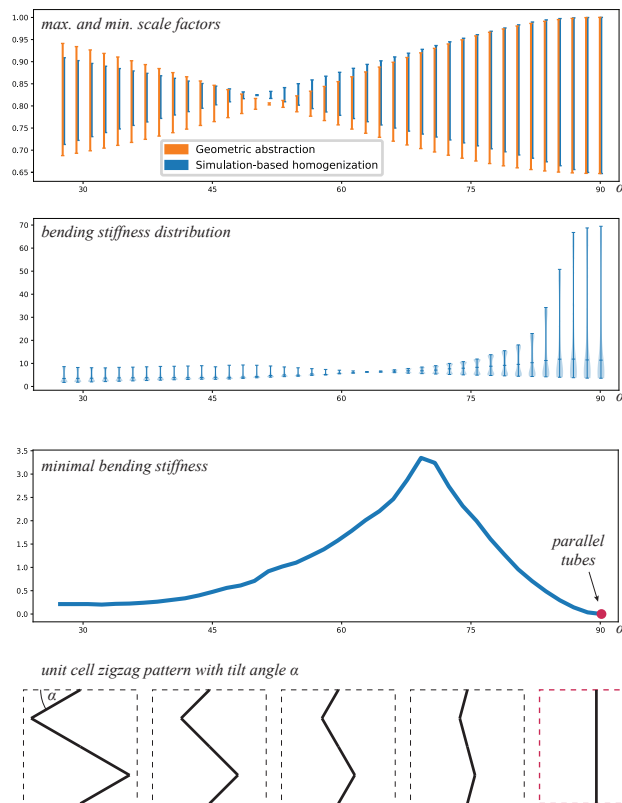


Fig. 15. Comparison of our homogenization results with the zigzag-pattern-specific geometric abstraction proposed in [Gao et al. 2020]. The ticks on each vertical line in the bending stiffness distribution plot indicate maximum, minimum, and mean values. Unit cell patterns generated by varying α are shown in the bottom row.

there is still room for this third phase to improve the design, as demonstrated in Figure 13.

Motivation for direct vertex optimization. The recent work of Montes Maestre et al. [2023] shows that it is possible to make our dehomogenization step differentiable and potentially even convert the design pipeline into a single-scale optimization over \mathbf{p} . While our coarsened optimization (Section 8.1) could rapidly obtain flattened domain shapes and high-quality parameter initializations for this style of optimization, we choose to dehomogenize once and then optimize mesh vertices for several reasons. First, it is unclear how to formulate an efficient X-FEM-style simulation for inflatable structures, especially with infinitesimal-line fusing curves. Second, while optimizing the dehomogenized structure with respect to the pattern parameter fields can in principle permit topology changes, these typically induce abrupt changes in the inflated equilibrium that are problematic for gradient-based optimization. Finally, vertex optimization enables the design to depart from the parametric pattern space where necessary to boost surface approximation.




Fig. 16. Comparison of parallel tubes (left) with our approach (right). The weak bending stiffness across the tube directions causes the free boundary roof to collapse into a sharp kink, while our approach avoids this weakness and reproduces the target shape more faithfully. Color-coding indicates distance to the target surface (overlaid with transparency) relative to the bounding box diagonal.

Table 1. Performance statistics. The first two columns report the duration of the homogenization experiments for all samples within each pattern family, which were run in parallel using 64 cores. Additionally, for each model shown in the paper, we collect the number of vertices in the input mesh and the time taken by local-global initialization, coarsened design optimization, and dehomogenization.

Pattern	Preproc. (s)	Model	Input # Vertices	Init. (s)	Coarsened (s)	Dehomogenization (s)	
						Extract Curves	Meshing
Cosine Curve	1042.36	Igloo (Fig. 12, 16)	1225	1.48	56.55	22.34	58.11
		Cashew (Fig. 10, 17, 22)	4704	5.96	651.86	187.05	30.86
		Tent (Fig. 1, 11)	3969	7.76	2739.43	164.80	81.76
Elliptic Holes	4540.10	Neck brace (Fig. 18)	2313	2.92	6736.27	224.89	98.16
		Hill (Fig. 13)	2459	3.56	2103.81	238.83	62.30
Dashline	3417.28	Squidward (Fig. 19)	3056	4.54	4205.61	94.55	16.64

9 RESULTS AND DISCUSSION

In this section, we discuss a number of different designs that we have optimized with our inverse design pipeline. Please see the accompanying video for more renderings, simulations and video footage of the inflation process, and further load-testing experiments.

Fabrication. We fabricated physical prototypes using a 976nm infrared laser to fuse two 160 μ m thin sheets of Polyamide along the optimized fusing curves. To ensure a strong fusing connection, we use sheets laminated with thermoplastic polyurethane that melts due to laser absorption. The laser welds have a constant width of about 2.3mm. The laser is mounted on a CNC bed of size 1560mm \times 2487mm and powered at 18W. With an average feed rate of 2000mm/min, we achieve fabrication times ranging from 20min

to 40min depending on model complexity. We inflate our models to a pressure of 0.5 - 1bar.

Figure 14 shows the fabrication setup and illustrates material failures as inflation pressure is increased. Overall, we are able to achieve high pressure values that lead to strong structures. Compared to the approach of [Panetta et al. 2021], which optimizes for varying-width fusing curves, our method simplifies fabrication by using constant-width fusing curves (see also Figure 16).

9.1 Evaluation

Figure 15 shows how we can use our homogenization approach to analyze the zigzag pattern proposed by Gao et al. [2020]. Their empirical model gives a reasonably accurate approximation of the in-plane contractions as a function of tilt angle α (top). The parallel

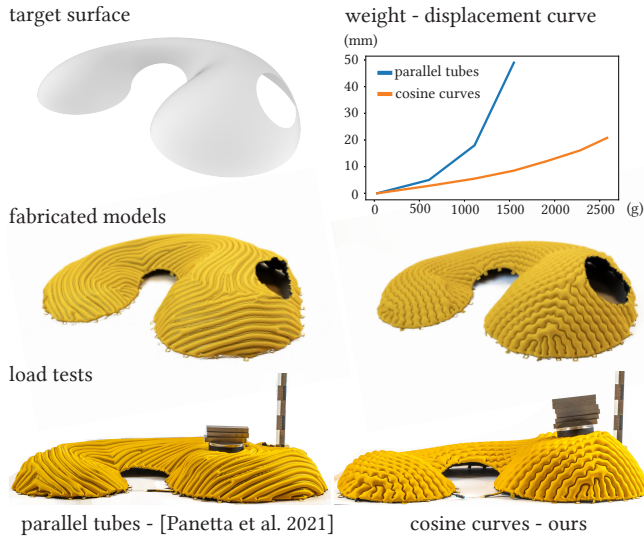


Fig. 17. Comparisons between simulation and physical prototypes for the parallel tubes of [Panetta et al. 2021] and our approach. The bottom row shows a basic load-testing experiment. See also accompanying video for more loading scenarios.

tubes of [Panetta et al. 2021] corresponding to $\alpha = 90^\circ$ have the characteristic high maximal stiffness along the tube direction but very weak minimal stiffness in the orthogonal direction as visible in the stiffness plots computed by our homogenization framework.

Figures 10, 16, and 17 show more comparisons with the method of [Panetta et al. 2021]. In their case, varying contraction rates are achieved by adapting the width of the fusing curves. This not only complicates fabrication but also introduces weaknesses when fusing curves become too wide. The main drawback of parallel tubes, however, is that the stiffness profile is highly unfavorable, as illustrated in Figure 6. Our approach offers a much richer set of cell geometries that allows decoupling the scale factors from the stiffness behavior. The optimization can thus locally select the required in-plane contraction with the most favorable stiffness profile, achieving a better approximation of the target surface.

In Figure 10, we evaluate how simulated gravity forces affect the inflated equilibrium shape. For this model, the boundary touching the ground is fixed. Both the approach of [Panetta et al. 2021] and our method yield good approximations of the target surface when no forces are applied, except for the free-boundary opening, where the parallel tubes pattern has difficulties approximating the target surface. Under external loading, the parallel tubes model quickly collapses due to its exceedingly low minimal bending stiffness. We further validate the superior structural performance of our design through physical load-testing experiments, as shown in Figure 17.

Figure 18 illustrates how our method can seamlessly handle cell patterns with internal holes. This is particularly relevant for use cases that require permeability, for example to reduce wind loads in architectural designs or to improve comfort for applications in fashion or personalized protective equipment.

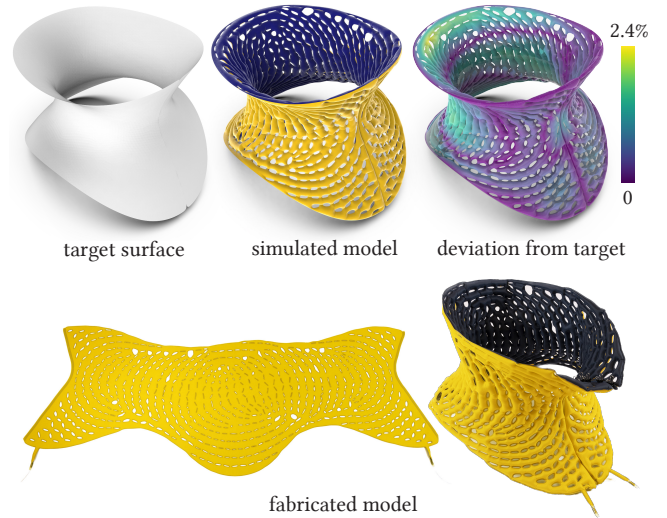


Fig. 18. Using the elliptical holes pattern family, we construct a permeable neck brace that can stabilize an injured patient’s head, illustrating a potential medical application. The fabricated sheet has been re-connected along the congruent boundary seam lines prior to inflation.




Fig. 19. The dash line pattern used on a surface with positive and negative Gaussian curvature. “Teardrops” have been added to the open ends of the line segments to reduce the risk of decollation at higher inflation pressure.

In Figure 19, we use the dashed line pattern on a double-bump surface. A small artifact is visible at the front where the dashed lines were too small to be robustly fabricated and thus were pruned, leading to a thicker inflated bump.

In Table 1, we report performance statistics for homogenization and for designing the models shown in the paper. Figure 20 shows the deviation between the scans of our physical inflated models and the corresponding input design surface. We have scanned our models using the polycam photogrammetry application and aligned the scans using iterative closest points with manual initial alignment.

9.2 Limitations

A fundamental limitation of surfaced-based inflatables, shared with other deployable systems based on metric frustration, is the limited




Fig. 20. Deviation of physical models from input surfaces. We scanned the fabricated models and color-coded the results using the distance to the closest point on the design surface relative to the bounding box diagonal.

range of in-plane contraction that is directly linked to the achievable curvature distribution. If surface flattening does not yield a result within the range of realizable stretch factors, cone singularities could be added, as commonly done in surface parameterization algorithms. However, this would require the introduction of additional cuts into the planar domain. The deployed surface would then need to be reconnected after inflation.

Our flattening algorithm currently does not guarantee that the computed mapping is free of global self-intersections. Existing strategies to ensure an injective mapping, such as [Du et al. 2021], could potentially be integrated into our approach. However, this might lead to higher distortion and thus a larger range of required contraction factors.

Overall, the fabricated models match the design inputs very well. Small discrepancies remain that could be due to several simplifications we make in our model. For example, we assume the sheet material to be isotropic, yet we observed that the woven fabric of the Polyamide sheets has anisotropic stretching resistance. Similarly, we currently do not model the effect that laser welding has on the material behavior of the fusing curves.

Our simulation only considers the final inflated shape at a specific pressure, ignoring the dynamics of inflation. Currently, a model might require assistance during inflation to ensure, for example, that a bump appears in the right orientation or that channels are not blocked due to self-collisions. We furthermore neglect collisions in our homogenization and design algorithms, though incorporating an existing technique like Incremental Potential Contact [Li et al. 2020] would not pose a fundamental challenge.

Although our homogenization framework should work with general periodic tiles, our dehomogenization framework currently relies on square unit cells with reflectional symmetry. Extending our design framework toward more flexible cell domains could follow the recent approach of Tozoni et al. [2020], who introduced rhombic unit cells for elastic metamaterials. We are also limited to using the single pattern family selected by the user for a given design, ruling out blends between different families that could produce interesting structural properties and aesthetic effects. Although our design pipeline can theoretically employ complex patterns with many parameters, our current approach of constructing tensor-product




Fig. 21. A simple parametric family of unit cells having controllable non-planar equilibrium states.

spline material property interpolants on regular-grid discretizations of parameter space would suffer from the curse of dimensionality.

Our homogenization framework can analyze unit cells with non-planar equilibrium states (Figure 21). However, our design framework currently does not utilize these cells since it is unclear how to incorporate the extrinsic curvature induced by the unit cells in addition to metric changes.

We observed delamination problems on some models, in particular with the dashed line pattern, due to stress concentrations at the fusing curve ends. We currently address this issue by adding a small “teardrop” at the end of each curve segment (see Figure 19) in a post-processing step, which sacrifices design optimality. More systematic studies are required, however, to find the optimal trade-off between the robustness of the fusing weld and the required geometric modifications to the pattern.

9.3 Future Work

Addressing the limitations discussed above provides interesting avenues for future research. We also summarize a few observations that we believe warrant further investigations.

Figure 4 shows a unit cell that exhibits a two-scale contraction. In addition to the in-plane contraction resulting from pressure forces transversely separating the sheets, the cell also folds. Currently, our approach does not exploit this kind of secondary mechanism. In future work, we want to study connections to origami-style folding to further enhance the capabilities of surface-based inflatables.

The potential of pneumatic structures in soft robotics has been widely investigated [Su et al. 2022]. We believe our systematic approach for analyzing the mechanical properties will also be beneficial in this domain, especially when coupled with a dynamic analysis of shape morphing and force transfer.

We also plan to explore the potential of our approach in applications where load-bearing capacity is essential, such as light-weight architectural rooftops (Figure 22) and concrete casting. Specifically, we plan to incorporate specific loading scenarios directly into the optimization. In this context, interesting questions also arise as to what extent combining multiple layers of inflated surface sheets enables a richer geometric space and further improved structural performance.




Fig. 22. A design study highlights the potential for surface-based inflatables for applications in architecture.

10 CONCLUSION

Surface-based inflatables are an attractive material system for creating complex 3D surfaces. With their simple fabrication, even at large scales, a high strength-to-weight ratio, and a versatile geometric and aesthetic palette, surface-based inflatables offer a rich space for design in a variety of application domains.

We have demonstrated that incorporating stiffness analysis into the inverse design pipeline allows us to build structures with significantly improved load bearing capacities that adhere closely to the desired target shape. Our work presents the first systematic study of inflatables using a novel numerical homogenization method. We believe that our approach can be applied to other material systems that involve surface bending, potentially yielding new insights for building more effective inverse design pipelines for digital fabrication.

ACKNOWLEDGMENTS

We are grateful to Nicolas Boumal for his insightful discussions regarding singularities and to Filip Goč for helping with the fabrication of physical prototypes. We also thank the anonymous reviewers for their valuable feedback. This research was supported by the Swiss National Science Foundation (Grant FNS 514543 / CF 1156).

REFERENCES

- Hillel Aharoni, Yu Xia, Xinyue Zhang, and Shu Kamien, Randall D. and Yang. 2018. Universal inverse design of surfaces with thin nematic elastomer sheets. *Proceedings of the National Academy of Sciences of the United States of America* 115, 28 (2018).
- Grégoire Allaire, Perle Geoffroy-Donders, and Olivier Pantz. 2019. Topology optimization of modulated and oriented periodic microstructures by the homogenization method. *Computers & Mathematics with Applications* 78, 7 (2019), 2197–2229. <https://doi.org/10.1016/j.camwa.2018.08.007> Simulation for Additive Manufacturing.
- Ignacio Andrade-Silva and Joel Marthelot. 2023. Fabric-Based Star Soft Robotic Gripper. *Advanced Intelligent Systems* 5, 8 (2023), 2200435. <https://doi.org/10.1002/aisy.202200435>
- Artelys. 2019. Artelys Knitro - Nonlinear optimization solver. <https://www.artelys.com/en/optimization-tools/knitro>
- Robert Baines, Sree Kalyan Patiballa, Benjamin Gorissen, Katia Bertoldi, and Rebecca Kramer-Bottiglio. 2023. Programming 3D Curves with Discretely Constrained

- Cylindrical Inflatables. *Advanced Materials* 35, 26 (2023), 2300535. <https://doi.org/10.1002/adma.202300535>
- Quentin Becker, Seiichi Suzuki, Yingying Ren, Davide Pellis, Julian Panetta, and Mark Pauly. 2023. C-Shells: Deployable Gridshells with Curved Beams. *ACM Trans. Graph.* 42, 6, Article 173 (dec 2023), 17 pages. <https://doi.org/10.1145/3618366>
- Camille Boutemy, Arthur Lebé, Mélina Skouras, Marc Mimram, and Olivier Baverel. 2023. Reusable Inflatable Formwork for Complex Shape Concrete Shells. In *Towards Radical Regeneration*. Springer International Publishing, Cham, 198–210.
- Antoine Chan-Lock, Jesús Pérez, and Miguel A Otaduy. 2022. High-Order Elasticity Interpolants for Microstructure Simulation. In *Computer Graphics Forum*, Vol. 41. Wiley Online Library, 63–74.
- Desai Chen, David IW Levin, Shinjiro Sueda, and Wojciech Matusik. 2015. Data-driven finite elements for geometry and material design. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 1–10.
- Desai Chen, Mélina Skouras, Bo Zhu, and Wojciech Matusik. 2018. Computational discovery of extremal microstructure families. *Science advances* 4, 1 (2018).
- Tian Chen, Julian Panetta, Max Schnaubelt, and Mark Pauly. 2021. Bistable auxetic surface structures. *ACM Transactions on Graphics* 40, 4 (Aug. 2021), 1–9. <https://doi.org/10.1145/3450626.3459940>
- Manfredo P Do Carmo. 2016. *Differential geometry of curves and surfaces: revised and updated second edition*. Courier Dover Publications.
- Xingyi Du, Danny M. Kaufman, Qingnan Zhou, Shahar Z. Kovalsky, Yajie Yan, Noam Aigerman, and Tao Ju. 2021. Optimizing global injectivity for constrained parameterization. *ACM Trans. Graph.* 40, 6, Article 260 (2021). <https://doi.org/10.1145/3478513.3480556>
- Tian Gao, Emmanuel Siéfert, Antonio DeSimone, and Benoît Roman. 2020. Shape Programming by Modulating Actuation over Hierarchical Length Scales. *Advanced Materials* 32, 47 (2020), 2004515. <https://doi.org/10.1002/adma.202004515>
- Christophe Geuzaine and Jean-Francois Remacle. 2009. Gmsh: A 3-D finite element mesh generator with built-in pre- and post-processing facilities. *Internat. J. Numer. Methods Engrg.* 79, 11 (2009), 1309–1331. <https://doi.org/10.1002/nme.2579> <https://onlinelibrary.wiley.com/doi/pdf/10.1002/nme.2579>
- A. Sydney Gladman, Elisabetta A Matsumoto, Ralph G Nuzzo, Lakshminarayanan Mahadevan, and Jennifer A Lewis. 2016. Biomimetic 4D printing. *Nature materials* 15, 4 (2016), 413–418.
- Ruslan Guseinov, Eder Miguel, and Bernd Bickel. 2017. CurveUps: Shaping Objects from Flat Plates with Tension-Actuated Curvature. *ACM Trans. Graph.* 36, 4, Article 64 (7 2017), 12 pages. <https://doi.org/10.1145/3072959.3073709>
- Jen-Hsuan Hsiao, Jen-Yuan Chang, and Chao-Min Cheng. 2019. Soft medical robotics: clinical and biomedical applications, challenges, and future directions. *Advanced Robotics* 33, 21 (2019), 1099 – 1111. <https://doi.org/10.1080/01691864.2019.1679251>
- Lishuai Jin, Antonio Elia Forte, Bolei Deng, Ahmad Rafsanjani, and Katia Bertoldi. 2020. Kirigami-inspired inflatables with programmable shapes. *Advanced Materials* 32, 33 (2020).
- Trevor J. Jones, Thomas Dupuis, Etienne Jambon-Puillet, Joel Marthelot, and P.-T. Brun. 2023. Soft Deployable Structures via Core-Shell Inflatables. *Phys. Rev. Lett.* 130 (Mar

- 2023), 128201. Issue 12. <https://doi.org/10.1103/PhysRevLett.130.128201>
- David Jourdan, Pierre-Alexandre Hugron, Camille Schreck, Jonàs Martínez, and Sylvain Lefebvre. 2023. Shrink & Morph: 3D-printed self-shaping shells actuated by a shape memory effect. *ACM Trans. Graph.* 42, 6, Article 187 (2023). <https://doi.org/10.1145/3618386>
- David Jourdan, Mélina Skouras, Etienne Vouga, and Adrien Bousseau. 2020. Printing-on-fabric meta-material for self-shaping architectural models. *Advances in Architectural Geometry* (2020).
- David Jourdan, Mélina Skouras, Etienne Vouga, and Adrien Bousseau. 2022. Computational Design of Self-Actuated Surfaces by Printing Plastic Ribbons on Stretched Fabric. In *Computer Graphics Forum*, Vol. 41. Wiley Online Library, 493–506.
- Lily Kharevych, Patrick Mullen, Houman Owahdi, and Mathieu Desbrun. 2009. Numerical coarsening of inhomogeneous elastic materials. *ACM Transactions on graphics (TOG)* 28, 3 (2009), 1–8.
- Felix Knöppel, Keenan Crane, Ulrich Pinkall, and Peter Schröder. 2015. Stripe Patterns on Surfaces. *ACM Trans. Graph.* 34 (2015), Issue 4.
- Mina Konakovic-Lukovic, Julian Panetta, Keenan Crane, and Mark Pauly. 2018. Rapid Deployment of Curved Surfaces via Programmable Auxetics. *Acm Transactions On Graphics* 37, 4 (2018), 106. <https://doi.org/10.1145/3197517.3201373>
- Bruno Lévy, Sylvain Petitjean, Nicolas Ray, and Jérôme Maillot. 2002. Least Squares Conformal Maps for Automatic Texture Atlas Generation. *ACM Trans. Graph.* 21, 3 (July 2002), 362–371. <https://doi.org/10.1145/566654.566590>
- Minchen Li, Zachary Ferguson, Teso Schneider, Timothy Langlois, Denis Zorin, Daniele Panozzo, Chenfanfu Jiang, and Danny M. Kaufman. 2020. Incremental Potential Contact: Intersection- and Inversion-free Large Deformation Dynamics. *ACM Trans. Graph. (SIGGRAPH)* 39, 4, Article 49 (2020).
- Yue Li, Juan Montes, Bernhard Thomaszewski, and Stelian Coros. 2022. Programmable Digital Weaves. *IEEE Robotics and Automation Letters* 7, 2 (2022), 2891–2896.
- Yiyue Luo, Kui Wu, Andrew Spielberg, Michael Foshey, Daniela Rus, Tomás Palacios, and Wojciech Matusik. 2022. Digital Fabrication of Pneumatic Actuators with Integrated Sensing by Machine Knitting. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. Article 175. <https://doi.org/10.1145/3491102.3517577>
- Juan Montes Maestre, Ronan Hinchet, Stelian Coros, and Bernhard Thomaszewski. 2023. ToRoS: A Topology Optimization Approach for Designing Robotic Skins. *ACM Trans. Graph.* 42, 6, Article 194 (2023). <https://doi.org/10.1145/3618382>
- Jonàs Martínez, Mélina Skouras, Christian Schumacher, Samuel Hornus, Sylvain Lefebvre, and Bernhard Thomaszewski. 2019. Star-Shaped Metrics for Mechanical Metamaterial Design. *ACM Trans. Graph.* 38, 4, Article 82 (2019). <https://doi.org/10.1145/3306346.3322989>
- David Melancon, Benjamin Gorissen, Carlos J. García-Mora, Chuck Hoberman, and Katia Bertoldi. 2021. Multistable inflatable origami structures at the metre scale. *Nature* 592, 7855 (April 2021), 545–550. <https://doi.org/10.1038/s41586-021-03407-4>
- Eivind Lyche Melvær and Martin Reimers. 2012. Geodesic polar coordinates on polygonal meshes. In *Computer Graphics Forum*, Vol. 31. Wiley Online Library, 2423–2435.
- Christian Miede, Jörg Schröder, and Martin Becker. 2002. Computational homogenization analysis in finite elasticity: material and structural instabilities on the micro- and macro-scales of periodic composites and their interaction. *Computer Methods in Applied Mechanics and Engineering* 191, 44 (2002), 4971–5005.
- Juan Sebastian Montes Maestre, Yinwei Du, Ronan Hinchet, Stelian Coros, and Bernhard Thomaszewski. 2023. Differentiable Stripe Patterns for Inverse Design of Structured Surfaces. *ACM Transactions on Graphics* 42, 4 (2023). <https://doi.org/10.1145/3592114> Cited by: 0; All Open Access, Green Open Access.
- Praveen Babu Nakshatrala, Daniel A Tortorelli, and KB Nakshatrala. 2013. Nonlinear structural design using multiscale topology optimization. Part I: Static formulation. *Computer Methods in Applied Mechanics and Engineering* 261 (2013), 167–176.
- Ryuma Niiyama, Hiroki Sato, Kazzmasa Tsujimura, Koya Narumi, Young Ah Seong, Ryosuke Yamamura, Yasuaki Kakehi, and Yoshihiro Kawahara. 2020. Poimo: Portable and inflatable mobility devices customizable for personal physical characteristics. *UIST 2020 - Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology* (2020), 912–923. <https://doi.org/10.1145/3379337.3415894>
- Kate Oliver, Annela Seddon, and Richard S. Trask. 2016. Morphing in nature and beyond: a review of natural and synthetic shape-changing materials and mechanisms. *Journal of Materials Science* 51, 24 (2016), 10663–10689. <https://doi.org/10.1007/s10853-016-0295-8>
- Jifei Ou, Mélina Skouras, Nikolaos Vlavianos, Felix Heibeck, Chin-Yi Cheng, Jannik Peters, and Hiroshi Ishii. 2016. aeroMorph-heat-sealing inflatable shape-change materials for interaction design. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. 121–132.
- Julian Panetta, Florin Isvoranu, Tian Chen, Emmanuel Siéfert, Benoît Roman, and Mark Pauly. 2021. Computational inverse design of surface-based inflatables. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 1–14.
- Julian Panetta, Mina Konaković-Luković, Florin Isvoranu, Etienne Bouleau, and Mark Pauly. 2019. X-Shells: A New Class of Deployable Beam Structures. *ACM Trans. Graph.* 38, 4, Article 83 (2019), 15 pages. <https://doi.org/10.1145/3306346.3323040>
- Julian Panetta, Abtin Rahimian, and Denis Zorin. 2017. Worst-case Stress Relief for Microstructures. *ACM Trans. Graph.* 36, 4, Article 122 (July 2017), 16 pages. <https://doi.org/10.1145/3072959.3073649>
- Julian Panetta, Qingnan Zhou, Luigi Malomo, Nico Pietroni, Paolo Cignoni, and Denis Zorin. 2015. Elastic Textures for Additive Fabrication. *ACM Trans. Graph.* 34, 4, Article 135 (July 2015), 135:1–135:12 pages.
- Daniele Panozzo, E Puppo, and L Rocca. 2010. Efficient multi-scale curvature and crease estimation. *Proceedings of Computer Graphics, Computer Vision and Mathematics* 1, 6 (2010).
- Jesús Pérez, Miguel A Otaduy, and Bernhard Thomaszewski. 2017. Computational design and automated fabrication of kirchhoff-plateau surfaces. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 1–12.
- Stefan Pillwein, Kurt Leimer, Michael Birsak, and Przemyslaw Musialski. 2020. On Elastic Geodesic Grids and Their Planar to Spatial Deployment. *ACM Trans. Graph.* 39, 4, Article 125 (July 2020), 12 pages. <https://doi.org/10.1145/3386569.3392490>
- Nicolas Ray, Wan Chiu Li, Bruno Lévy, Alla Sheffer, and Pierre Alliez. 2006. Periodic Global Parameterization. *ACM Trans. Graph.* 25, 4 (Oct. 2006), 1460–1485. <https://doi.org/10.1145/1183287.1183297>
- Yingying Ren, Uday Kusupati, Julian Panetta, Florin Isvoranu, Davide Pellis, Tian Chen, and Mark Pauly. 2022. Umbrella meshes: elastic mechanisms for freeform shape deployment. *ACM Transactions on Graphics (TOG)* 41, 4 (2022), 1–15.
- E. Rivas-Adrover. 2015. *Deployable Structures*. Laurence King Publishing.
- Emmanuel Rodriguez, Georges-Pierre Bonneau, Stefanie Hahmann, and Mélina Skouras. 2022. Computational Design of Laser-Cut Bending-Active Structures. *Computer-Aided Design* 151 (2022), 103335. <https://doi.org/10.1016/j.cad.2022.103335>
- Mark Schenk, Andrew D. Viquerat, Keith A. Seffen, and Simon D. Guest. 2014. Review of inflatable booms for deployable space structures: Packing and rigidization. *Journal of Spacecraft and Rockets* 51, 3 (2014), 762–778. <https://doi.org/10.2514/1.A.32598>
- Christian Schumacher, Bernd Bickel, Jan Rys, Steve Marschner, Chiara Daraio, and Markus Gross. 2015. Microstructures to Control Elasticity in 3D Printing. *ACM Trans. Graph.* 34, 4, Article 136 (July 2015), 136:1–136:13 pages.
- Christian Schumacher, Steve Marschner, Markus Gross, and Bernhard Thomaszewski. 2018. Mechanical Characterization of Structured Sheet Materials. *ACM Trans. Graph.* 37, 4, Article 148 (jul 2018). <https://doi.org/10.1145/3197517.3201278>
- Emmanuel Siéfert, Etienne Reysat, José Bico, and Benoît Roman. 2020. Programming stiff inflatable shells from planar patterned fabrics. *Soft Matter* 16 (2020), 7898–7903. Issue 34. <https://doi.org/10.1039/D0SM01041C>
- Mélina Skouras, Bernhard Thomaszewski, Peter Kaufmann, Akash Garg, Bernd Bickel, Eitan Grinspun, and Markus Gross. 2014. Designing inflatable structures. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 1–10.
- Georg Sperl, Rahul Narain, and Chris Wojtan. 2020. Homogenized Yarn-Level Cloth. *ACM Transactions on Graphics (TOG)* 39, 4 (2020).
- D. J. Steigmann. 1990. Tension-Field Theory. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences* 429, 1876 (1990), 141–173. <http://www.jstor.org/stable/51778>
- Hang Su, Xu Hou, Xin Zhang, Wen Qi, Shuting Cai, Xiaoming Xiong, and Jing Guo. 2022. Pneumatic Soft Robots: Challenges and Benefits. *Actuators* 11, 3 (2022). <https://doi.org/10.3390/act11030092>
- Masato Tanaka, S. Macrae Montgomery, Liang Yue, Yaochi Wei, Yuyang Song, Tsuyoshi Nomura, and H. Jerry Qi. 2023. Turing pattern-based design and fabrication of inflatable shape-morphing structures. *Science Advances* 9, 6 (2023). <https://doi.org/10.1126/sciadv.ade4381>
- Pengbin Tang, Stelian Coros, and Bernhard Thomaszewski. 2023. Beyond Chainmail: Computational Modeling of Discrete Interlocking Materials. *ACM Trans. Graph.* 42, 4, Article 84 (jul 2023), 12 pages. <https://doi.org/10.1145/3592112>
- Yichao Tang, Mingtong Li, Tianlu Wang, Xiaoguang Dong, Wenqi Hu, and Metin Sitti. 2022. Wireless Miniature Magnetic Phase-Change Soft Actuators. *Advanced Materials* 34, 40 (2022), 2204185.
- Davi Colli Tozoni, Jérémie Dumas, Zhongshi Jiang, Julian Panetta, Daniele Panozzo, and Denis Zorin. 2020. A low-parametric rhombic microstructure family for irregular lattices. *ACM Transactions on Graphics (TOG)* 39, 4 (2020), 101–1.
- Yue Wang, Yingying Ren, and Tian Chen. 2022. From kirigami to hydrogels: a tutorial on designing conformally transformable surfaces. *Journal of Applied Mechanics* (2022). <https://doi.org/10.1115/1.4056350>
- Xudong Yang, Zongzheng Wang, Bojian Zhang, Tianyu Chen, Changhong Linghu, Kunlin Wu, Guohui Wang, Hailu Wang, and Yifan Wang. 2023. Self-Sensing Robotic Structures from Architected Particle Assemblies. *Advanced Intelligent Systems* 5, 1 (2023).
- Yiwei Zhang, Tomoya Tendo, and Tomohiro Tachi. 2023b. Modular design of multistable pneumatic structures from a flat pattern of air pouches. *Journal of the International Association for Shell and Spatial Structures* 64, 4 (2023), 298–305.
- Zhan Zhang, Christopher Brandt, Jean Jouve, Yue Wang, Tian Chen, Mark Pauly, and Julian Panetta. 2023a. Computational Design of Flexible Planar Microstructures. *ACM Trans. Graph.* 42, 6, Article 185 (2023). <https://doi.org/10.1145/3618396>