

Taking a Moment to Characterize the Bending Response of Thin Sheet Materials

P. Xie^{1,2} , J. S. Montes Maestre² , S. Coros²  and B. Thomaszewski² 

¹ISTA, Austria

²ETH Zürich, Switzerland

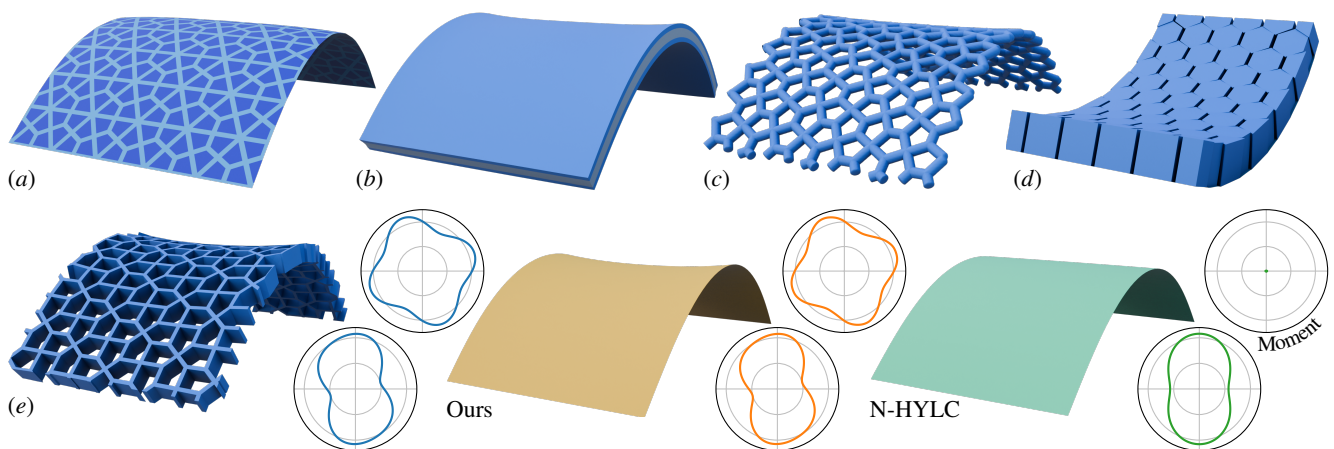


Figure 1: Using both energy and bending moment data for homogenization, our method is able to accurately model the bending response of a wide range of sheet materials, including bi-material shells (a), thick sandwich plates (b), rod network materials (c), scaled sheets (d), and thick structured plates (e). As illustrated in these uniaxial bending examples, our approach properly captures the material’s moment response in the principal (bottom plots) and orthogonal directions (top plots). An existing method (N-HYLC by [FWYX24]) that uses only energy data for homogenization fails to model this moment-curvature coupling (bottom right).

Abstract

Structured sheet materials such as 3D-printed rod networks, multi-material thin shells, and multi-layer laminates exhibit diverse mechanical behaviors. To avoid the computational burden of native-scale simulations, data-driven homogenization offers a promising alternative. This process involves probing a representative patch of material—a unit cell—with a set of stretching and bending tests subject to periodic boundary conditions. Because macro-scale bending moments are not directly available from native-scale simulations, existing methods exclusively rely on elastic energy data. Unfortunately, using only elastic energy from uniaxial tests is not sufficient for capturing the full moment-curvature relationship, and imposing biaxial curvature states would necessarily break periodicity. We present a moment-based homogenization method that infers curvature coupling using only uniaxial bending tests. Our method computes macro-scale bending moments from native-scale simulations for a wide range of mechanical models. To this end, we translate internal deformations into elastic stresses and then integrate these stresses through the thickness and across the unit-cell patch. We use the resulting homogenized bending moments along with energy data to fit neural bending energy density functions. We demonstrate our method on a diverse set of materials, including multi-material shells, rod networks, and multi-layer sheets. Our results show improved accuracy compared to existing approaches and realistic double-curvature behavior when applied to larger samples.

CCS Concepts

• **Computing methodologies** → **Physical simulation**; **Shape modeling**;

1. Introduction

3D-printed rod networks, multi-layer laminates, and multi-material thin shells are examples of structured sheet materials whose micro-scale composition gives rise to complex macromechanical behaviors. Modeling these materials in an accurate yet computationally efficient way is an important problem in material design and visual computing.

Simulating structured materials on their native scale is often computationally prohibitive for real-world problem sizes. A common strategy to reduce complexity while capturing the macromechanical behavior is numerical homogenization. In this class of approaches, a representative material element, subject to periodic boundary conditions, is probed with different states of membrane and bending deformations. With a sufficient number of uni- and bi-axial stretching tests, the often nonlinear in-plane behavior of the material can be determined. In principle, an analogous strategy can be followed to characterize the out-of-plane behavior using uniaxial and biaxial bending tests. However, double curvature states cannot be obtained without violating periodicity conditions—a direct consequence of Gauss' *Theorema Egregium*. Existing methods deal with this problem by either discarding biaxial curvature tests, or by abandoning periodicity conditions. Consequently, the response to double curvature will either be globally inaccurate or biased towards specific sample sizes.

In this work, we present a new data-driven approach for homogenizing the bending response of structured sheet materials using only uniaxial curvature tests. Our key observation is that uniaxial bending is sufficient for inferring curvature coupling if the material's full stress response, including bending moments, is taken into account: even for cylindrical bending, the material will tend towards a double curvature state if it is energetically favorable. This preference will manifest as non-zero bending moments in the direction orthogonal to the imposed axis of curvature, see Fig. 1 (*bottom row*) for an example. To extract curvature coupling information from uniaxial bending tests, we must determine bending moments throughout the representative material patch on which deformations are imposed. Since bending moments are typically not directly available from native-scale simulations, we first convert the model's internal deformations to elastic stresses in the intrinsic coordinate system of the sheet. The integral of these stresses through the thickness then yields the point-wise bending moments. By performing uniaxial bending tests with different magnitudes and directions, we obtain energy- and moment-curvature data, which we use to fit a neural energy density function.

Our moment-based homogenization approach can be applied to a large range of sheet materials and native-scale simulation models. We show examples of discrete shells with heterogeneous material properties, multilayer sheets modeled with solid shell elements, thin networks made from discrete elastic rods, and scaled sheet materials that exhibit contact-induced stiffening during bending; see Fig. 1 for an overview. Our quantitative analysis shows that our new moment-based homogenization approach outperforms existing baselines in terms of accuracy with respect to native-scale simulation data. We furthermore illustrate the qualitative differences when using our homogenized model for simulating finite-sized material samples with double curvature.

2. Related Work

Capturing Real-World Materials. Creating digital representations of real-world deformable objects has been a focus of graphics research for decades [JP99, BBO*09, PRWH*18]. Due to their relevance for clothing and textile simulation, thin sheet materials have received special attention. Traditional capture pipelines often deform physical sheets under controlled boundary conditions and utilize optical or mechanical setups to measure the resulting deformation, which is then used to fit various material models [WOR11, MBT*12]. However, these techniques are usually limited to uniaxial curvature settings and thus overlook a material's resistance to double curvature—an important aspect of material behavior that our work aims to capture.

Mechanical Metamaterials. Metamaterials with strategically designed microstructures can produce a vast range of macromechanical properties. As a physical analogue to texture synthesis, microstructure design has also attracted attention from the visual computing community. The spectrum of studied architectures ranges from 3D lattices [MDL16, MSD17] and voxels [SBR*15], to multilayer stacks [BBO*10] and 2D structured sheets [MSS*19, LCT23, ZNCT25]. While the elasticity of these structures already leads to considerable nonlinearities, internal contacts arising during deformation complicate their response even further [MMDH*24, TTCB25]. Beyond designing periodic unit cells, another stream of research optimizes material layouts globally, such as stripe-shaped bi-material distributions [MMDH*23] or pneumatic channels that follow desired orientation fields [PIC*21]. Regardless of the specific architecture, computational design and utilization of these metamaterials fundamentally rely on accurate material characterization. Rather than designing new types of metamaterials, we focus on the problem of finding macromechanical descriptions for structured sheet materials.

Homogenization. Characterizing material behaviors at macro-scale is typically achieved through *homogenization*, a process that finds parameters of a given homogeneous material that best approximates the native-scale behavior [KMOD09]. Under the assumption of small displacements and linear elasticity, this task amounts to determining a single stiffness tensor that describes the stress-strain relationship around the rest state [PZM*15, SBR*15]. Accounting for large deformations and corresponding nonlinearities naturally increases the complexity of the homogenization pipeline [ZBJ*23].

For thin sheet materials, their stretching and bending behaviors are typically characterized independently. Schumacher et al. [SMGT18] fit linear membrane and bending stiffness tensors to simulation data of elastic rod networks, where spherical curvature tests are used to determine the bending response but inherently conflict with periodic boundary conditions. Tang et al. [TCT23] lay aside periodicity and instead enforce finite-sized specimens onto paraboloid surfaces as boundary conditions while acknowledging the bias incurred by this approach. To avoid the geometric incompatibility, Sperl et al. [SNW20] restrict their scope to uni-axial bending tests when homogenizing yarn-level cloth. However, relying only on energy densities measured from uniaxial tests is insufficient to determine the full bending response, leaving the material's resistance to bi-axial states of curvature unconstrained. Rather than fitting bending models to a large number of simulation data, Ren

et al. [RPS*24] use sensitivity analysis to determine the uniaxial bending stiffness in different directions, noting that the curvature-coupling terms cannot be disentangled. While they use a directional integration approach to bypass this ambiguity, the process implies a specific Poisson’s ratio and therefore does not extend to general materials.

In summary, none of these existing approaches is able to accurately model the resistance of thin sheet materials to double curvature states. By leveraging macro-scale bending moments derived from native-scale simulations, our work proposes an accurate and efficient solution for this problem.

3. Method

We present a novel approach for data-driven homogenization of thin sheet materials that only requires uniaxial bending tests to completely determine the macroscopic constitutive bending model. To set the stage, we first describe our motivation for including moments in bending homogenization in Sec. 3.1, using an isotropic plate as a simple example. We then derive methods to compute bending moments for different native-scale simulation models in Sec. 3.2. Finally, we introduce our neural approach in Sec. 3.3 that uses both energy and moment data for bending homogenization.

3.1. Motivation

Determining macroscopic material properties from simulation data requires probing a representative element (unit cell) with a set of deformation states subject to periodic boundary conditions. For the membrane part, one usually applies uniaxial and biaxial deformations in different directions, extracts the elastic energy from the corresponding equilibrium states, and computes material parameters from these deformation-energy pairs through numerical fitting—see, e.g., Schumacher et al. [SMGT18]. Applying this same strategy to the bending part, however, leads to a fundamental problem: while uniaxial curvature can be imposed in any direction without violating periodicity, biaxial curvature necessarily breaks periodicity. This limitation is a direct consequence of Gauss’ *theorema egregium*, which implies that surfaces with double curvature must exhibit non-uniform in-plane deformations. Consequently, unit cells with double curvature cannot be tiled without gaps, violating one of the fundamental assumptions of homogenization theory. While uniaxial bending energy alone is insufficient to fully characterize the material behavior, we demonstrate below that the *bending moments* induced by uniaxial curvature states uniquely determine the material parameters.

The response of an isotropic, homogeneous plate to double curvature is governed by its Poisson’s ratio, which controls the interaction between bending in orthogonal directions. For a plate of thickness h , Young’s modulus E , and Poisson’s ratio ν , subjected to a uniaxial curvature κ , the bending energy density is given by

$$\Psi_B = \frac{Eh^3}{24(1-\nu^2)} \kappa^2. \tag{1}$$

Data consisting of curvature–energy pairs therefore allow the identification of an effective bending stiffness, $B = \frac{Eh^3}{12(1-\nu^2)}$. However, such data constrain only this combined parameter and do not

provide sufficient information to disentangle Poisson’s ratio from Young’s modulus.

On the other hand, Poisson’s ratio is explicitly encoded in the moment-curvature relationship of the material,

$$\begin{bmatrix} m_{xx} \\ m_{yy} \\ m_{xy} \end{bmatrix} = \frac{Eh^3}{12(1-\nu^2)} \begin{bmatrix} 1 & \nu & 0 \\ \nu & 1 & 0 \\ 0 & 0 & \frac{1-\nu}{2} \end{bmatrix} \begin{bmatrix} \kappa_{xx} \\ \kappa_{yy} \\ 2\kappa_{xy} \end{bmatrix}, \tag{2}$$

where $\mathbf{m} = (m_{xx}, m_{yy}, m_{xy})^T$ is the vectorized moment tensor in Voigt notation. It is evident from the above expression that for $\nu \neq 0$, uniaxial curvature states induce bending moments both along the axis of curvature and in the orthogonal direction. The Poisson’s ratio can thus be easily and uniquely determined from a single uniaxial curvature test. A collection of bending moments for different uniaxial curvature directions provides the necessary information to fully reconstruct anisotropic bending stiffness tensors for more general materials. In Sec. 3.3, we show how this insight can be generalized to capture nonlinear moment-curvature relationships.

3.2. Computing Bending Moments

To probe the bending response of the material, we impose uniaxial curvature states on a unit cell with periodic boundary conditions as described by Montes et al. [MMDH*24], see also App. A. We then compute energy densities and averaged bending moments for the corresponding equilibrium configurations. To this end, we first compute point-wise bending moment tensors by integrating the second Piola-Kirchhoff stress tensors \mathbf{S} through the thickness,

$$\mathbf{m}(\mathbf{x}) = \int_{-h/2}^{h/2} \mathbf{S}(\mathbf{x}) z dz. \tag{3}$$

The averaged moment tensor is then obtained by integrating the point-wise tensor across the periodic patch Ω ,

$$\mathbf{m}_{\text{avg}} = \frac{1}{|\Omega|} \int_{\Omega} \mathbf{m}(\mathbf{x}) dA. \tag{4}$$

In the following, we explain how we implement this strategy for different mechanical models and discretizations.

Volumetric Elements. We consider two types of volumetric finite elements in this work: 10-node quadratic tetrahedra and 15-node prismatic solid shell elements. For these elements, we build a continuous geometry field via quadratic basis functions \mathbf{N}_i ,

$$\mathbf{x}(\mathbf{q}) = \sum_i \mathbf{N}_i(\xi_1, \xi_2, \xi_3), \tag{5}$$

where $\mathbf{q} = (\xi_1, \xi_2, \xi_3)$ are parametric coordinates. The double integral implied by (3) and (4) can be rewritten as a triple integral in parametric coordinates over the entire volume. For prismatic solid shell elements we have,

$$\mathbf{m}_{\text{avg}} = \frac{1}{|\Omega|} \int_0^1 \int_0^{1-\xi_1} \int_{-1}^1 \mathbf{S}(\mathbf{x}(\mathbf{q})) z(\mathbf{q}) \left| \frac{\partial \mathbf{x}}{\partial \mathbf{q}} \right| d\xi_3 d\xi_2 d\xi_1, \tag{6}$$

where $\xi_1 \in [0, 1]$, $\xi_2 \in [0, 1 - \xi_1]$, $\xi_3 \in [-1, 1]$. Bending moments for quadratic tetrahedra are obtained in an analogous way. We approximate the corresponding integrals using Gauss-Legendre quadrature with four and six quadrature points for quadratic tetrahedra and solid shell elements, respectively. Following classical

plate theory where only membrane strains are relevant, we ignore stress components in the thickness direction when computing bending moments.

Discrete Shells. We use discrete shells [GHDS03] to model thin sheets with multimaterial distributions. Unlike volumetric elements, neither deformation in the thickness direction nor bending moments are directly available in discrete shells. To recover this information, we assume a constant curvature state within each triangle, which is computed via the triangle-averaged shape operator described in [GGRZ06]. Together with Kirchhoff-Love assumptions, the point-wise moment of each triangle is directly determined via Eq. 2 and can be expressed compactly as $\mathbf{m}_i = \mathbf{B}_i \boldsymbol{\kappa}_i$. The homogenized moment tensor for the entire patch is then obtained by averaging element-wise moments over the reference configuration,

$$\mathbf{m}_{\text{avg}} = \frac{1}{|\Omega|} \sum_i |\Omega_i| \mathbf{B}_i \boldsymbol{\kappa}_i. \quad (7)$$

It should be noted that we assume the shell to be initially flat such that all triangles in the reference configuration lie on the neutral plane of the homogenized patch. This assumption simplifies the computation of homogenized moments, which are derived by integrating the in-plane stresses through the thickness with respect to the mid-surface. If a given triangle were offset along the thickness direction in the rest state, this displacement would need to be taken into account during integration.

Discrete Elastic Rods. Discrete elastic rods from [BWR*08, BAV*10] are highly efficient for simulating slender, curve-like materials that resist stretching, bending, and twisting. While the discrete model provides bending and twisting moments, we must convert per-rod quantities to patch moments. For simplicity, we focus on rods with circular cross sections that are initially flat and straight. In this setting, rods in pure bending will not experience any stretch along their centerline such that only bending and twisting stresses contribute to the moments. For a given rod segment with scalar curvature κ and twist ϑ , we follow Megaro et al. [MZB*17] to compute 3D strain tensors $\boldsymbol{\varepsilon} = \boldsymbol{\varepsilon}_b(\boldsymbol{\kappa}) + \boldsymbol{\varepsilon}_t(\boldsymbol{\vartheta})$ for the cross section of the rod,

$$\boldsymbol{\varepsilon}_b(\boldsymbol{\kappa}) = -\kappa \mathbf{z} \begin{bmatrix} 1 & 0 & 0 \\ 0 & -\nu & 0 \\ 0 & 0 & -\nu \end{bmatrix}, \boldsymbol{\varepsilon}_t(\boldsymbol{\vartheta}) = \frac{\vartheta}{2} \begin{bmatrix} 0 & -z & y \\ -z & 0 & 0 \\ y & 0 & 0 \end{bmatrix}. \quad (8)$$

Ignoring components in the thickness direction, the corresponding stresses are obtained using the constitutive law of the rod,

$$\boldsymbol{\sigma} = -Ez \begin{bmatrix} \boldsymbol{\kappa} & \frac{\vartheta}{2(1+\nu)} \\ \frac{\vartheta}{2(1+\nu)} & 0 \end{bmatrix}. \quad (9)$$

We then perform integration over the cross-sectional area S to obtain effective bending moments,

$$\mathbf{m} = \int_S \boldsymbol{\sigma} z dA = -E \frac{\pi r^4}{4} \begin{bmatrix} \boldsymbol{\kappa} & \frac{\vartheta}{2(1+\nu)} \\ \frac{\vartheta}{2(1+\nu)} & 0 \end{bmatrix}, \quad (10)$$

which can be directly expressed using the bending stiffness B and torsional rigidity β of the discrete elastic rod,

$$\mathbf{m} = \begin{bmatrix} -B\boldsymbol{\kappa} & -\frac{\beta\vartheta}{2} \\ -\frac{\beta\vartheta}{2} & 0 \end{bmatrix}. \quad (11)$$

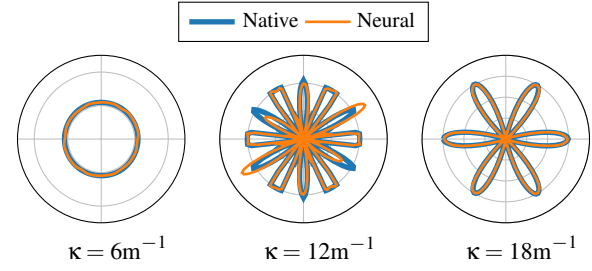


Figure 2: Bending response for a scaled sheet material. The polar plots show bending moments in the direction orthogonal to an imposed uniaxial curvature with different directions and magnitudes before contact (left), at the onset of contact (middle), and after full contact (right).

Finally, the resulting moment tensor is then transformed from the rod's cross sectional space to the material space of the patch, and the total averaged moment of the homogenized patch is

$$\mathbf{m}_{\text{avg}} = \frac{1}{|\Omega|} \sum_i \mathbf{R}_i \mathbf{m}_i \mathbf{R}_i^T, \quad (12)$$

where \mathbf{R}_i is the rotation transforming from the cross-sectional space of rod segment i to the material space of the patch. For elliptical cross sections, we provide analogous derivations in App. B.

Scaled Sheets. Scaled sheets are made from quasi-rigid scales and a thin, flexible substrate [MMDH*24]. We use discrete shells and tetrahedron elements to model the substrate and scales, respectively. Bending moments of the scaled sheet are computed by combining the moments of the substrate and the scales.

3.3. Neural Homogenization

With energy and moment data from uniaxial curvature tests in hand, we now seek a macromechanical bending model that best matches the native-scale simulations. For increased flexibility, we opt for a neural approach based on a learned energy density function. We largely follow [LCT23] and train a multi-layer perceptron (MLP) to predict energy density values Φ given curvature tensors $\boldsymbol{\kappa}$ as input. We use four layers with 32 neurons each and SiLU activation functions for input and hidden layers. The output layer uses a soft-plus activation function to ensure positive energy. Once trained, we compute moments from the network through standard back-propagation. For training, we use both energy densities Ψ_i and moments \mathbf{m}_i from simulation data and minimize the loss function

$$L_p(\boldsymbol{\kappa}) = \frac{1}{N} \sum_{i=1}^N \left(\frac{\Phi(\boldsymbol{\kappa}_i) - \Psi_i}{\Psi_i} \right)^2 + \frac{1}{N} \sum_{i=1}^N \left\| \frac{1}{\|\mathbf{m}_i\|_2} \left(\frac{\partial \Phi(\boldsymbol{\kappa}_i)}{\partial \boldsymbol{\kappa}_i} - \mathbf{m}_i \right) \right\|_2^2, \quad (13)$$

where N is the number of samples.

Using this data loss alone, the network cannot learn the correct output for arbitrary double curvature states because there are only simulation data for uniaxial curvature. Nevertheless, the paired energy-moment data completely determine the response to double

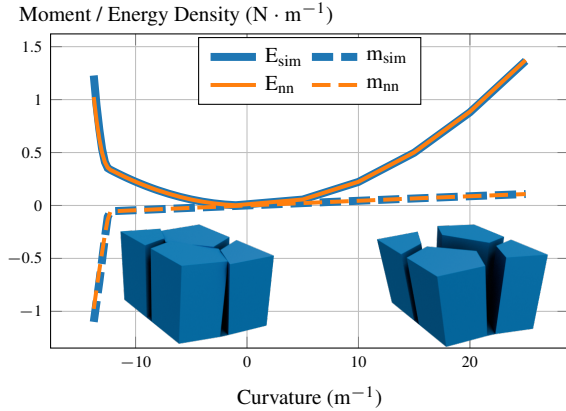


Figure 3: Comparison between moment and energy data obtained from native-scale simulations (dashed and solid blue) and our network (dashed and solid orange) for a uniaxial curvature test on a periodic scaled sheet.

curvature states for materials with locally *linear* moment-curvature relation. Moreover, we observed from our native-scale simulations that, although nonlinearities with respect to direction of curvature can be substantial (see Fig. 2), most materials display a quasi-linear moment-curvature relationship even for larger curvatures. For this reason, we introduce an additional loss term that encourages the network to learn an energy density function that is *as-quadratic-as-possible* along the directions of principal curvatures, while allowing full nonlinearity with respect to angular directions. This regularizer is conveniently expressed by penalizing third-order derivatives with respect to principal curvatures as

$$L_d(\kappa) = \frac{1}{N} \sum_{i=0}^N \left(\frac{\partial^3 \Phi(\kappa_i)}{\partial \kappa_{i1}^3} \right)^2 + \left(\frac{\partial^3 \Phi(\kappa_i)}{\partial \kappa_{i1}^2 \partial \kappa_{i2}} \right)^2 + \left(\frac{\partial^3 \Phi(\kappa_i)}{\partial \kappa_{i1} \partial \kappa_{i2}^2} \right)^2 + \left(\frac{\partial^3 \Phi(\kappa_i)}{\partial \kappa_{i2}^3} \right)^2, \quad (14)$$

where κ_{i1} and κ_{i2} are principal curvature magnitudes.

Since the material’s response to single curvature is already fully determined by uniaxial bending tests, regularization is only required in regions without any simulation data, i.e., double curvature states. Therefore, rather than regularizing the network at the same data points used for uniaxial testing, we sample new data points with varying magnitudes and directions of principal curvatures in the double curvature domain and evaluate the third-order derivative loss on these new data points in the training process.

Our method has similarities with the neural homogenization approach (N-HYLC) by Feng et al. [FWYX24], but there are several important differences. First and foremost, N-HYLC estimates the gradient of the bending energy density function via finite difference along the uniaxial deformation path. While the gradient in the orthogonal direction is directly available in the moment tensor, it cannot be estimated via finite difference due to the lack of biaxial data. Furthermore, N-HYLC evaluates the third-order derivative loss exclusively within single curvature states whereas our loss is

evaluated across arbitrary curvature configurations. As single curvature behavior is already heavily constrained by simulation data, it is the double curvature domain that remains unconstrained and demands regularization. As shown in Sec. 4, these differences lead to contrasting behaviors between our method and N-HYLC.

4. Results

We evaluate our moment-based homogenization method for structured sheet materials by comparing it to native-scale simulations and existing alternatives. We begin with two basic experiments that test the consistency between energy densities and our bending moments, and the ability of our homogenized model to accurately predict the double curvature response on finite-sized samples. We then apply our method to a diverse set of sheet materials and evaluate its quantitative accuracy on a larger-scale simulation experiment. Finally, we conduct an ablation study to demonstrate the importance of using moment data and our third-order derivative loss for training. Taken together, our results underline the importance of using moment data in order to accurately capture the nonlinear and complex bending response of structured sheet materials.

4.1. Validation

Energy-Moment Consistency. Our network learns an energy density function given energy and moment data from native-scale simulations. In order for this to work, it is crucial that the simulation data is self-consistent, i.e., it must satisfy the differential relationship $\mathbf{m} = \frac{\partial \mathcal{E}}{\partial \kappa}$. To verify that our approach for computing bending moments is indeed consistent, we train our network on native-scale simulation data from a scaled sheet, and then compare moment and energy values as a function of curvature for both simulation data and network output. As can be seen from Fig. 3, both moment and energy curves match the ground truth data almost perfectly. Together with the fact that homogenized moments are computed by differentiating through the network, this low discrepancy indicates that our bending moment computation is indeed consistent with the elastic energy of the native-scale model. Our experiments with the remaining mechanical models described in Sec. 3.2 confirm this observation.

Double Curvature Response. While double curvature is fundamentally at odds with periodicity assumptions, we generally expect to see nonzero Gaussian curvature for finite-sized material samples. To verify that our homogenized model predicts the correct orthogonal curvature, we compare native- and macro-scale simulations on finite-sized quadratic patches defined as a height field over the xy -plane, $f(\mathbf{x}, A, B, C) = x_z - Ax_x^2 - Bx_y^2 - Cx_x x_y$. Without loss of generality, we prescribe curvatures of a given magnitude along the x -axis by setting A to be a constant and minimize the elastic energy for the remaining patch variables,

$$\min_{\mathbf{x}, B, C} E(\mathbf{x}) \quad \text{s.t.} \quad f(\mathbf{x}_i, A, B, C) = 0 \quad \forall i. \quad (15)$$

where \mathbf{x}_i denote per-vertex positions. We solve the above problem for the native-scale simulation and our homogenized model and plot the optimized paraboloid parameters in Fig. 4. It is evident from these plots that our method produces double curvature states with orthogonal curvature values close to the ground-truth data.

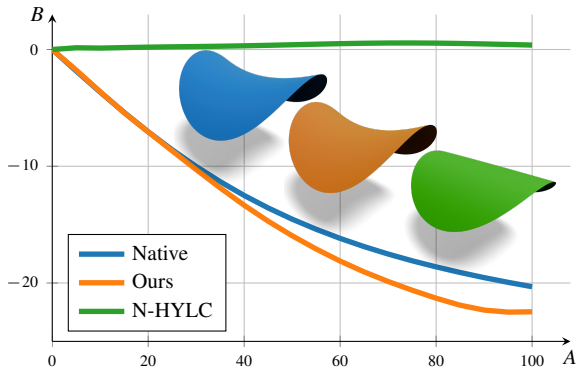


Figure 4: Curvature response of a disk-shaped patch with an isotropic homogeneous material simulated using different methods. Using the paraboloid formulation of Eq. (15), we impose curvature of varying magnitudes (A) along a given direction and plot the corresponding curvature in the orthogonal direction (B) at equilibrium. Whereas our model (orange) shows good agreement with the native scale simulation (blue), N-HYLC [FWYX24] (green) predicts a single-curvature state. The paraboloid parameter C is close to 0 for all three cases.

Varying Patch Size. Per Gauss' *Theorema Egregium*, deforming an initially planar sample into a double curvature configuration requires in-plane deformation, i.e., stretching or compression. Taking the above experiment as an example, we consider the amount of stretching at the boundary of the sample that is required to produce a given Gaussian curvature at its center. For small sample sizes, the required stretching is likewise small and, consequently, there is relatively little resistance to Gaussian curvature. As sample size increases, however, the level of stretching at the boundary required to maintain the Gaussian curvature at the center increases rapidly, and so does the sample's resistance to double curvature. In the limit of infinite sample size, nonzero Gaussian curvature becomes energetically impossible. It is worth noting that this observation is consistent with the fact that homogenization—which assumes an infinite tiling of the deformed unit cell—is incompatible with double curvature. To study this size-dependent behavior in simulation, we employ the same experimental setup as before and apply our homogenized material model to disk patches of varying diameters. As shown in Fig. 5, the native-scale simulations for larger disks exhibit smaller curvature along the transverse direction as expected. Our homogenized model faithfully reproduces this size-dependent behavior and tracks the native-scale data reasonably well.

4.2. Homogenizing Thin Sheet Materials

Our method can be used to compute macromechanical models for a wide range of thin sheet materials. In the following, we present examples for different types of materials and evaluate the accuracy of our homogenized models by comparing to native-scale simulations.

Experimental Setup. To compare the behavior of native-scale and homogenized models, we design a test case that is inspired by the real-world experimental setup described in [MGN*21].

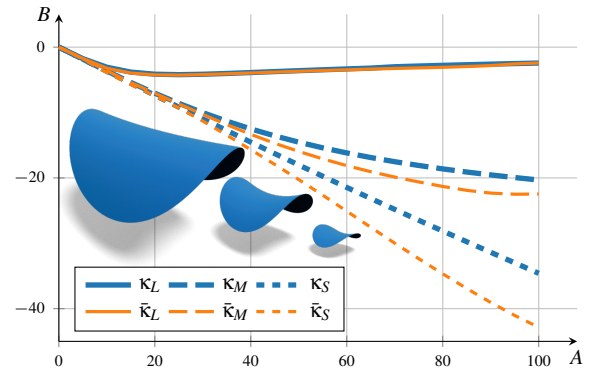


Figure 5: Optimized paraboloid parameters for disks with a diameter of 1cm (S), 2cm (M) and 4cm (L). Double curvature decreases with increasing patch size. Simulation results from our homogenized model ($\bar{\kappa}$, orange) are in good agreement with their native-scale counterparts (κ , blue).

As illustrated in the inset figure, we consider samples with in-plane dimensions of around 9 cm \times 5 cm and material-dependent thickness whose longer boundaries are displaced such as to induce compression in the sheet.

Since the boundaries are free to rotate, the sheet immediately buckles out of plane and, depending on the amount of horizontal displacement ΔL , shows curvature both along and orthogonal to the direction of compression. To quantify the deformation of the sample, we follow [MGN*21] and approximate its mid-surface as a quadratic function multiplied by a sinusoidal term,

$$z(x, y) = (A + Bx^2) \sin(Cy) + D, \quad (16)$$

where $C = \pi/L$, L is the distance between the boundaries, and A , B , and D are constants determined by fitting this function to our simulation data. The principal curvatures at the midpoint of the fitted surface, $\kappa_1 = -AC^2$ and $\kappa_2 = 2B$, are used for comparing native-scale and homogenized simulations.

Multi-material Thin Shells. Our first example is a bi-material thin sheet modeled with discrete shells. The resulting material distribution is illustrated in Fig. 6, top, with regions of stiff and soft material shown in dark and light blue, respectively. The stiffness ratio is set to 50 and Poisson's ratio is set to zero for both materials. Interestingly, even though the individual materials do not produce any Poisson effect, their composite exhibits positive Gaussian curvatures in our simulations. This indicates that the composite material has negative Poisson's ratio for out-of-plane deformation. As can be seen from the plots in Fig. 6, our homogenized model shows good agreement with the native-scale simulation in terms of principal curvatures, even for large compressions. While the orthogonal curvature κ_2 increases proportionally to the principal curvature κ_1 initially, it all but saturates at a given point. We conjecture that this effect, which is also captured by our homogenized model, is due to the fact that increasing curvature in the transverse direction would require stretching the mid-surface and thus incur a steep penalty.

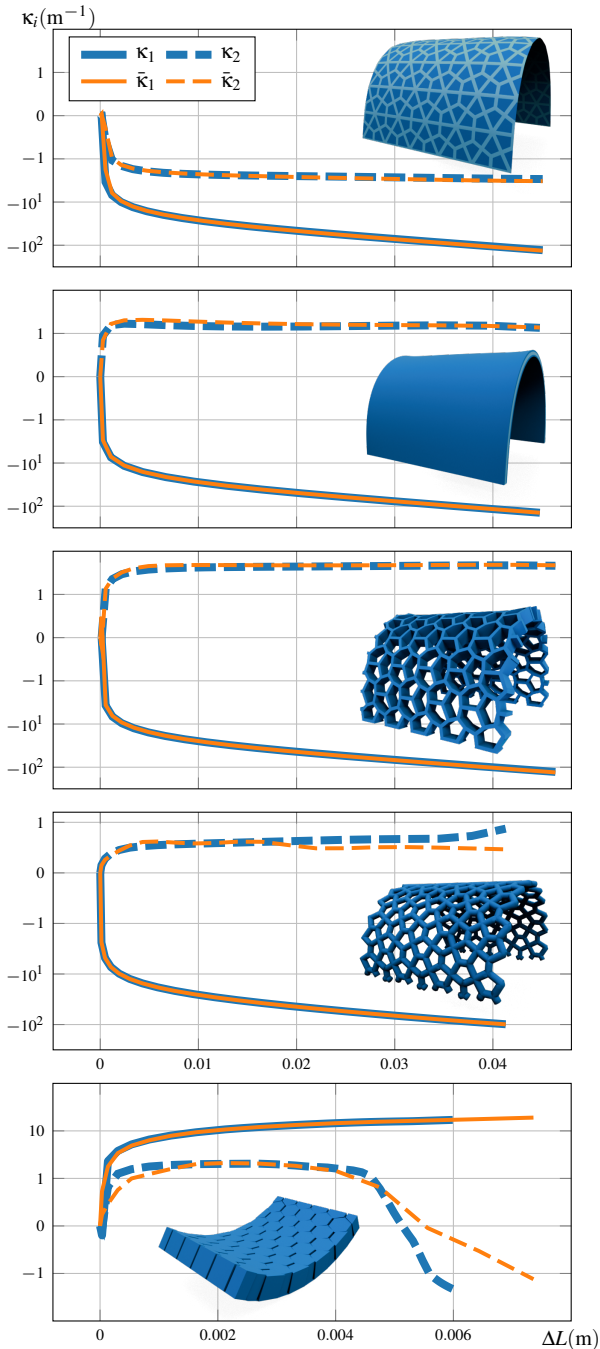


Figure 6: Comparison between native-scale and homogenized models. From top to bottom: multi-material sheet with discrete shells, multilayer sandwich sheet using quadratic solid shell elements, thick structured sheet with quadratic tetrahedra, network material using discrete elastic rods, and a flexible scaled sheet modeled with quadratic tetrahedra (scales) and discrete shells (substrate). We plot curvatures at the midpoint of the sheet vs. boundary displacement ΔL . Curvatures in the principal and orthogonal directions are denoted as κ_1, κ_2 and $\bar{\kappa}_1, \bar{\kappa}_2$ for native-scale and homogenized simulations, respectively.

Multi-layer Sheets. Our second example is a sandwich model with three layers of quadratic solid shell elements. We model the top and bottom layers with a stiff orthotropic material, whereas the middle layer uses a soft isotropic material. We set the orthotropic parameters using the formula from [LB15], with $4E_1 = E_2 = 4E_3$, $\nu = 0.25$, and principal axes rotated by $\pi/4$ and $-\pi/4$ relative to the material-space axes. For the isotropic material, we use $4E = E_3$ for Young’s modulus and 0.45 for Poisson’s ratio. As can be seen from Fig. 6, row 2, the sandwich sheet exhibits a saddle-shaped region around the midpoint of the sample, corresponding to negative Gaussian curvature. Ground truth curvatures in both directions are tracked accurately by our homogenized model.

Thick Structured Sheets. In this example, we analyze a thick sheet structured with a three-fold rotationally symmetric pattern. We use quadratic tetrahedron elements for native-scale simulation and an isotropic material with Poisson’s ratio set to 0.4. The native-scale simulation for this material (Fig. 6, row 3) shows a saddle-shaped deformation with relatively large curvature in the transverse direction compared to a homogeneous sheet with the same material. This stronger curvature response is mainly due to the lower in-plane stretching stiffness resulting from its structure. The curvature plot shows very good agreement between native-scale and our homogenized models.

Planar Rod Networks. Our fourth example is a rod network sheet modeled with discrete elastic rods. Connections between rods are modeled using rigid joints with corresponding rotation and translation degrees of freedom. See also [ZCT16, PKLI*19]. All rods are assumed to have the same circular cross section with a radius of 1mm. The Poisson’s ratio of the isotropic base material is set to 0.4. As with the two previous examples, the corresponding native-scale simulation shown in Fig. 6, row 4, likewise exhibits negative Gaussian curvature around the midpoint. The transverse curvature is somewhat less pronounced compared to the previous examples, arguably due to the smaller thickness of the sheet. Our homogenized model again shows good agreement with the native-scale simulations.

Scaled Sheet Materials. Our approach can also be applied to materials whose behavior is governed by internal contacts. Here we investigate a scaled sheet material as an example, consisting of quasi-rigid scales embedded in a soft substrate. We use discrete shell elements for the substrate, whereas the scales are modeled using quadratic tetrahedra. It is interesting to note that, despite the relatively small deformations that we observe in the scales, we found that, compared to their linear counterparts, the additional accuracy offered by quadratic elements is vital for properly capturing the bending moments, as shown in App. C. We focus on the moment-curvature relationship in the *closing* direction for which, after an initial phase, contact between the scales leads to a distinct stiffening of the material’s curvature response. As shown in Fig. 6, row 5, this nonlinear moment-curvature relationships translates into a complex curvature evolution in our compression experiment: the scaled sheet initially exhibits positive Gaussian curvature ($\kappa_1, \kappa_2 > 0$) but changes sign once contact among scales has been established throughout the patch. As can be seen from the curvature plots, our homogenized model again closely tracks the principal curvature from the native-scale simulations. The somewhat




Figure 7: Simulation results for compression experiments with different sample sizes and tiling patterns. Native-scale simulations are shown in blue in the first row. The second row displays results from our moment-based homogenization, where vertices are colored according to their nearest distance to the midsurface of the corresponding native-scale model.

larger deviations in the transverse curvature are likely explained by the fact that, as contact tightens, scales are pushed outwards which induces stretching in the κ_2 direction—an unusual coupling effect between membrane and bending deformations that our model does not capture. Nevertheless, even if curvatures are not matched perfectly in this compression experiment, our neural model still accurately captures the moment-curvature response.

This can be seen from the polar plots shown in Fig. 2, which indicate the material’s moment response in the orthogonal direction for uniaxial bending. It is interesting to note that the response is isotropic before contact, but exhibits large anisotropy at the onset of contact and afterwards. Our neural model tracks the native-scale behavior with very good accuracy.

Width-dependent Behavior. In addition to the material’s structure and corresponding mechanical properties, Gaussian curvature at equilibrium also strongly depends on the width of the sample. To study this width-dependent behavior, we create thick structured sheets with small and large widths ($S : 3.12 \text{ cm}$ and $L : 7.28 \text{ cm}$) and likewise adjust the meshes used for the homogenized model. We then perform compression experiments on these patches and plot the Gaussian curvature K in Fig. 8. The Gaussian curvature produced by our homogenized model decreases with increasing patch width, which aligns well with the native-scale simulation results.

Varying-scale Compression. We further validate our method on patches of varying sizes. For all patches we applied a compression of 27% and fixed the central tile at the boundary to serve as the boundary condition. Patch sizes range from $6 \text{ cm} \times 6 \text{ cm}$ to $14 \text{ cm} \times 14 \text{ cm}$. In contrast to the experiment adapted from [MGN*21], the remaining boundaries were left free to move, which yields a more natural configuration by imposing fewer constraints. Since we encode curvature coupling within the bending model, our homogenized results exhibit expected Gaussian curvature, as shown in Fig. 7. As the patch size grows, the homogenized models align better with the native scale ones.

4.3. Details and Statistics

Implementation Details. For each of the materials we consider, we train individual neural networks to learn bending and stretching energies. Using the same network architecture as for the bending part, the neural membrane energy is trained with uniaxial and

biaxial stretching data. We use discrete shell elements and constant strain triangles for bending and stretching energies, respectively, and solve the resulting minimization problem using Newton’s method. Our simulation models are implemented in C++ using Eigen [GJ*10] for linear algebra operations, IPC Toolkit [F*20] for contact handling, and libigl [JP*18] for geometry processing. We use *tactile* to generate isohedral tilings, Clipper2 to thicken the edges and Gmsh [GR09] for geometric operations.

Training Details. To generate training data, we uniformly sample bending directions from 0 to π and impose curvatures (in m^{-1}) ranging from -25 to 25 for scaled sheets and -150 to 150 for other materials. We use 8000 samples for scaled sheets and 750 for all other materials. We train the network using the Adam optimizer with a learning rate of 10^{-3} for 1000 epochs and a learning rate of 10^{-4} for 500 epochs. 95% of the simulation data are used for training and the remainder for validation. We set the batch size to 128 for scaled sheets and to 16 for other models. The weights for data loss and third-order derivative loss are set to 1 and 100, respectively. The training time is about 15 minutes on a MacBook Air M2 using only CPUs. The maximum relative L_2 errors in the validation set for energies and moments are listed in Table 1.

Table 1: Validation errors for different simulation models.

Material	Energy Error	Moment Error
Bi-material Shell	0.024%	0.16%
Multi-layer Sheet	0.119%	0.37%
Thick Structured Sheet	0.004%	0.01%
Rod Network	0.164%	0.53%
Scaled Sheet	0.234%	2.77%

Problem Sizes & Performance. To evaluate the computational performance of our homogenized models, we compare timings and problem sizes for native-scale and homogenized simulations for the compression experiment described in Sec. 4.2. As can be seen from Table 2, computation times for our homogenized models are about an order of magnitude smaller than for the native-scale models in general. For different material properties, our homogenized models show almost constant cost under the same mesh resolution. For scaled sheets, more solver iterations are required to reach equilibrium states due to the high nonlinearity around the onset of contact, leading to longer computation times for both native-scale and homogenized models.




Figure 8: Gaussian curvature of thick structured sheets (blue) with varying widths. As the width of the patch increases, the absolute value of its Gaussian curvature decreases. Our homogenized model (orange) reproduces this width-dependent behavior.

Table 2: Statistics for the compression experiment.

Example	# Nodes (Native / Ours)	Time [s] (Native / Ours)
Bi-material Shell	31773 / 1681	176.1 / 15.0
Multi-layer Sheet	25933 / 1681	182.9 / 11.3
Thick Structured Sheet	27000 / 1681	131.3 / 14.6
Rod Network	3948 / 1681	739.7 / 27.1
Scaled Sheet	28074 / 1681	4074.2 / 103.2

4.4. Ablation Study

To demonstrate the importance of using bending moments and third order derivative losses during training, we again consider a compression experiment but use a simple isotropic thin sheet model with discrete shell elements as the native-scale model. We investigate the evolution of the second principal curvature as a function of the first principal curvature for different options and approaches.

Third-Order-Derivative Loss. In order to predict the bending response in doubly curved states from only single curvature data, we use a third-order-derivative loss (14) that encourages the network to learn an as-quadratic-as-possible energy density function. As can be seen from Fig 9, yellow curve, without this loss, the network’s prediction becomes unstable and the accuracy deteriorates rapidly with increasing curvature.

Bending Moments. Using only uniaxial curvature tests, bending moments—especially those in the orthogonal direction—indicate the material’s response to double curvature states. Without moment data, the network does not have enough information to determine the curvature coupling and fails to predict curvature in the direction orthogonal to compression; see Fig 9, dark orange curve.

Comparison to N-HYLC. Due to the limitations discussed in Sec. 3, N-HYLC tends to wrongly predict single curvature states when the true material response exhibits nonzero Gaussian curvature—see Figs. 1, 4, 9 for examples. While this may be a reasonable approximation for materials and applications where




Figure 9: Plots of first vs. second principal curvature for different methods and options on the compression experiment. Similar to [FWYX24] (green), without moment data, our method fails to capture curvature coupling (dark orange). Omitting the third-order derivative loss, the accuracy of our network deteriorates for larger principal curvatures (yellow). Using both moment data and derivative regularization, our method (orange) closely tracks the native-scale data (blue).

stretching is predominant, it is inadequate for settings where bending and stretching can occur both combined and individually.

5. Conclusions

We presented a novel approach for homogenizing the bending response of structured sheet materials. Unlike previous methods, our approach takes into account bending moments computed from native-scale simulations to predict the material’s macromechanical response to double curvature states using only uniaxial bending tests. We demonstrated the versatility of our method by applying it to a wide range of sheet materials and mechanical models. Our quantitative results attest to our method’s excellent accuracy with respect to native-scale simulation data.

5.1. Limitations & Future Work

We compute bending moments by considering membrane deformations and the corresponding variation in stresses through the thickness. However, we do not directly consider deformations that combine bending with net stretching of the material’s mid-surface. Nevertheless, incorporating this coupling can be important for modeling the large-curvature response of structured materials with internal contacts (see, e.g., Fig. 6, row 5) and is thus a worthwhile direction for future work. As another current limitation, we have only considered purely elastic materials in our work. However, there are many interesting applications, including sheet metal forming, that involve large elastoplastic deformations. Extending our method to the plastic range is another exciting direction for future research.

Acknowledgements

We thank Yue Li and Logan Numerow for insightful discussions. This work was funded by the Swiss National Science Foundation through SNF project grant 200021_236469.

References

- [BAV*10] BERGOU M., AUDOLY B., VOUGA E., WARDETZKY M., GRINSPUN E.: Discrete viscous threads. *ACM Trans. Graph.* 29, 4 (2010). 4
- [BBO*09] BICKEL B., BÄCHER M., OTADUY M. A., MATUSIK W., PFISTER H., GROSS M.: Capture and modeling of non-linear heterogeneous soft tissue. *ACM transactions on graphics (TOG)* 28, 3 (2009), 1–9. 2
- [BBO*10] BICKEL B., BÄCHER M., OTADUY M. A., LEE H. R., PFISTER H., GROSS M., MATUSIK W.: Design and fabrication of materials with desired deformation behavior. *ACM Transactions on Graphics (TOG)* 29, 4 (2010), 1–10. 2
- [BWR*08] BERGOU M., WARDETZKY M., ROBINSON S., AUDOLY B., GRINSPUN E.: Discrete elastic rods. In *ACM SIGGRAPH 2008 Papers* (New York, NY, USA, 2008), Association for Computing Machinery. 4
- [F*20] FERGUSON Z., ET AL.: IPC Toolkit, 2020. URL: <https://github.com/ipc-sim/ipc-toolkit>. 8
- [FWYX24] FENG X., WANG H., YANG Y., XU W.: Neural-assisted homogenization of yarn-level cloth. In *ACM SIGGRAPH 2024 Conference Papers* (New York, NY, USA, 2024), Association for Computing Machinery. 1, 5, 6, 9
- [GGRZ06] GRINSPUN E., GINGOLD Y., REISMAN J., ZORIN D.: Computing discrete shape operators on general meshes. *Computer Graphics Forum* 25, 3 (2006), 547–556. 4
- [GHDS03] GRINSPUN E., HIRANI A. N., DESBRUN M., SCHRÖDER P.: Discrete shells. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (Goslar, DEU, 2003), Eurographics Association, p. 62–67. 4
- [GJ*10] GUENNEBAUD G., JACOB B., ET AL.: Eigen v3. <http://eigen.tuxfamily.org>, 2010. 8
- [GR09] GEUZAIN C., REMACLE J.-F.: Gmsh: A 3-d finite element mesh generator with built-in pre- and post-processing facilities. *International Journal for Numerical Methods in Engineering* 79, 11 (2009), 1309–1331. 8
- [JP99] JAMES D. L., PAI D. K.: Artdefo: accurate real time deformable objects. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques* (USA, 1999), ACM Press/Addison-Wesley Publishing Co., p. 65–72. 2
- [JP*18] JACOBSON A., PANOZZO D., ET AL.: libigl: A simple C++ geometry processing library, 2018. <https://libigl.github.io/>. 8
- [KMOD09] KHAREVYCH L., MULLEN P., OWHADI H., DESBRUN M.: Numerical coarsening of inhomogeneous elastic materials. *ACM Transactions on graphics (TOG)* 28, 3 (2009), 1–8. 2
- [LB15] LI Y., BARBIČ J.: Stable orthotropic materials. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (Goslar, DEU, 2015), Eurographics Association, p. 41–46. 7
- [LCT23] LI Y., COROS S., THOMASZEWSKI B.: Neural metamaterial networks for nonlinear material design. *ACM Trans. Graph.* 42, 6 (2023). 2, 4
- [MBT*12] MIGUEL E., BRADLEY D., THOMASZEWSKI B., BICKEL B., MATUSIK W., OTADUY M. A., MARSCHNER S.: Data-driven estimation of cloth simulation models. *Computer Graphics Forum* 31, 2pt2 (2012), 519–528. 2
- [MDL16] MARTÍNEZ J., DUMAS J., LEFEBVRE S.: Procedural voronoi foams for additive manufacturing. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 1–12. 2
- [MGN*21] MIRZAALI M. J., GHORBANI A., NAKATANI K., NOURIGOUSHKI M., TÜMER N., CALLENS S. J. P., JANBAZ S., ACCARDO A., BICO J., HABIBI M., ZADPOOR A. A.: Curvature induced by deflection in thick meta-plates. *Advanced Materials* 33, 30 (2021), 2008082. 6, 8
- [MMDH*23] MONTES MAESTRE J. S., DU Y., HINCHET R., COROS S., THOMASZEWSKI B.: Differentiable stripe patterns for inverse design of structured surfaces. *ACM Trans. Graph.* 42, 4 (2023). 2
- [MMDH*24] MONTES MAESTRE J. S., DU Y., HINCHET R., COROS S., THOMASZEWSKI B.: Flexscale: Modeling and characterization of flexible scaled sheets. *ACM Trans. Graph.* 43, 4 (2024). 2, 3, 4, 11
- [MSDL17] MARTÍNEZ J., SONG H., DUMAS J., LEFEBVRE S.: Orthotropic k-nearest foams for additive manufacturing. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 1–12. 2
- [MSS*19] MARTÍNEZ J., SKOURAS M., SCHUMACHER C., HORNUS S., LEFEBVRE S., THOMASZEWSKI B.: Star-shaped metrics for mechanical metamaterial design. *ACM Transactions on Graphics (TOG)* 38, 4 (2019), 1–13. 2
- [MZB*17] MEGARO V., ZEHNDER J., BÄCHER M., COROS S., GROSS M., THOMASZEWSKI B.: A computational design tool for compliant mechanisms. *ACM Trans. Graph.* 36, 4 (2017). 4
- [PIC*21] PANETTA J., ISVORANU F., CHEN T., SIÉFERT E., ROMAN B., PAULY M.: Computational inverse design of surface-based inflatables. *ACM Trans. Graph.* 40, 4 (2021). 2
- [PKLI*19] PANETTA J., KONAKOVIĆ-LUKOVIĆ M., ISVORANU F., BOULEAU E., PAULY M.: X-shells: a new class of deployable beam structures. *ACM Trans. Graph.* 38, 4 (2019). 7
- [PRWH*18] PAI D. K., ROTHWELL A., WYDER-HODGE P., WICK A., FAN Y., LARIONOV E., HARRISON D., NEOG D. R., SHING C.: The human touch: measuring contact with real human soft tissues. *ACM Trans. Graph.* 37, 4 (2018). 2
- [PZM*15] PANETTA J., ZHOU Q., MALOMO L., PIETRONI N., CIGNONI P., ZORIN D.: Elastic textures for additive fabrication. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 1–12. 2
- [RPS*24] REN Y., PANETTA J., SUZUKI S., KUSUPATI U., ISVORANU F., PAULY M.: Computational homogenization for inverse design of surface-based inflatables. *ACM Trans. Graph.* 43, 4 (July 2024). 3
- [SBR*15] SCHUMACHER C., BICKEL B., RYS J., MARSCHNER S., DARAIO C., GROSS M.: Microstructures to control elasticity in 3d printing. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 1–13. 2
- [SMGT18] SCHUMACHER C., MARSCHNER S., GROSS M., THOMASZEWSKI B.: Mechanical characterization of structured sheet materials. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–15. 2, 3
- [SNW20] SPERL G., NARAIN R., WOJTAN C.: Homogenized yarn-level cloth. *ACM Trans. Graph.* 39, 4 (2020), 48. 2
- [TCT23] TANG P., COROS S., THOMASZEWSKI B.: Beyond chainmail: Computational modeling of discrete interlocking materials. *ACM Trans. Graph.* 42, 4 (2023). 2
- [TTCB25] TANG P., THOMASZEWSKI B., COROS S., BICKEL B.: Inverse design of discrete interlocking materials with desired mechanical behavior. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers* (New York, NY, USA, 2025), SIGGRAPH Conference Papers '25, Association for Computing Machinery. 2
- [WOR11] WANG H., O'BRIEN J. F., RAMAMOORTHI R.: Data-driven elastic models for cloth: modeling and measurement. *ACM transactions on graphics (TOG)* 30, 4 (2011), 1–12. 2
- [ZBJ*23] ZHANG Z., BRANDT C., JOUVE J., WANG Y., CHEN T., PAULY M., PANETTA J.: Computational design of flexible planar microstructures. *ACM Trans. Graph.* 42, 6 (2023). 2
- [ZCT16] ZEHNDER J., COROS S., THOMASZEWSKI B.: Designing structurally-sound ornamental curve networks. *ACM Trans. Graph.* 35, 4 (2016). 7
- [ZNCT25] ZHOU H., NUMEROW L., COROS S., THOMASZEWSKI B.: Closed-form construction of voronoi diagrams with star-shaped metrics. *ACM Trans. Graph.* 44, 6 (2025). 2

Appendix A: Imposing Uniaxial Curvature

We parametrize the sheet using its mid-surface ϕ and normal \mathbf{n} as

$$\mathbf{x} = \phi(w_1, w_2) + w_3 \mathbf{n}(w_1, w_2), \quad (17)$$

where $\mathbf{w} = (w_1, w_2, w_3)^T$ denotes a given material point in local coordinates and \mathbf{x} is the corresponding world-space position. When imposing cylindrical deformation along the y -axis, the mapping is defined as

$$\phi(w_1, w_2) = \left(\frac{1}{\kappa} \sin(\kappa w_1), w_2, \frac{1}{\kappa} (\cos(\kappa w_1) - 1) \right)^T, \quad (18)$$

where κ is the prescribed curvature. Cylindrical deformations in arbitrary directions can be achieved by applying a corresponding rotation and its inverse before and after the mapping, respectively.

Enforcing Boundary Conditions. We focus on periodic materials, allowing us to study their macroscopic mechanical behavior by simulating a unit-cell patch at the micro-scale. To ensure that the material represented by the periodic patch undergoes the prescribed macroscopic deformation, local fluctuations at the micro-scale must vanish on average. We follow the first-displace-then-map approach of [MMDH*24] and average local fluctuations before cylindrical mapping. This enables us to enforce periodic boundary conditions similar to the planar case and simplifies implementation.

Our periodic boundary conditions are enforced in mid-surface coordinates using variable substitution,

$$\begin{aligned} \begin{bmatrix} w_1^i \\ w_2^i \end{bmatrix} &= \begin{bmatrix} w_1^j \\ w_2^j \end{bmatrix} + \mathbf{F} \begin{bmatrix} \bar{w}_1^i - \bar{w}_1^j \\ \bar{w}_2^i - \bar{w}_2^j \end{bmatrix}, \\ w_3^i &= w_3^j, \end{aligned} \quad (19)$$

where the superscripts i, j represent periodic components on opposite boundaries, $\mathbf{F} \in \mathbb{R}^{2 \times 2}$ is the deformation gradient of the periodic patch, and quantities decorated with a bar denote the reference configuration. Rotations are eliminated by constraining \mathbf{F} to be symmetric.

Special care must be taken for discrete shell and rod models, as bending stencils extend over multiple elements and thus require proper wrapping at the boundaries. To this end, we introduce ghost points that are copied from their periodic counterparts and then displaced based on the macroscopic deformation. For discrete elastic rods, periodicity must additionally be enforced on material frames, which we achieve by imposing corresponding constraints on reference frames and twist angles.

Appendix B: Moment of Rods for Elliptical Cross Sections

The computation of bending moments for rods with elliptical cross sections is analogous to the circular case, except that we must now integrate over an elliptical domain and consider curvature as a 2D vector $\vec{\kappa}$.

Following our derivations for the circular case, it can be generalized that $\vec{m} = \mathbf{B} \vec{\kappa}$ represents the bending moment vector in the material frames \vec{d}_1, \vec{d}_2 , where $\mathbf{B} = \frac{EA}{4} \begin{bmatrix} a^2 & 0 \\ 0 & b^2 \end{bmatrix}$ is the bending




Figure 10: Moments computed by integrating stresses and finite differencing energies for scaled sheets modeled with linear and quadratic tetrahedra.

stiffness of the discrete elastic rod, and a and b are the major and minor radii of the elliptical cross section, respectively. Since we focus on rods that are initially flat and straight, the moment from bending stresses must lie in the rod's direction, which allows us to focus on its magnitude. Again, following classical plate theory, we ignore stresses in the thickness direction by ignoring the in-plane component of the moment vector. Assuming that \vec{d}_1 is planar to the homogenized shell and \vec{d}_2 is in the thickness direction, $-\vec{m} \cdot \vec{d}_2$ is the moment from bending stresses in the rod's material space.

The Cauchy strain from twisting in an elliptical cross section is

$$\epsilon_t = \frac{\vartheta}{a^2 + b^2} \begin{bmatrix} 0 & -a^2 z & b^2 y \\ -a^2 z & 0 & 0 \\ b^2 y & 0 & 0 \end{bmatrix}. \quad (20)$$

Applying the constitutive laws, we obtain the twisting stress tensor $\sigma_t = \frac{E}{1+\nu} \epsilon_t = 2G \epsilon_t$. Ignoring components in the thickness direction and integrating the remaining terms over the ellipse, we obtain the twisting moment tensor as

$$\mathbf{m}_t = \frac{G\pi a^3 b^3}{2(a^2 + b^2)} \begin{bmatrix} 0 & -1 \\ -1 & 0 \end{bmatrix}. \quad (21)$$

Using the torsional rigidity $\beta = \frac{\pi a^3 b^3 G}{a^2 + b^2}$, we can also directly express the moment tensor for discrete elastic rods as in (11).

Appendix C: Quadratic Elements for Scales

When imposing cylindrical curvature onto scaled sheets, we expect the substrate to undergo large deformations while the scales, due to their high stiffness, stay almost undeformed. This observation may suggest that a coarse discretization of the scales with linear elements would be sufficient. Although we observe that scales' deformations remain small during bending experiments, modeling them with linear elements fails to produce proper bending moments, as shown in Fig. 10. We conjecture that the reason for the poor behavior of linear elements is their tendency to produce shear locking. Although small in magnitude, for stiff materials these strain deviations translate into large errors in stresses, likely explaining the inaccuracies in bending moments that we observe for linear elements.