

Polar Interpolants for Thin-Shell Microstructure Homogenization

Antoine Chan-Lock
Universidad Rey Juan Carlos
Madrid, Spain
antoine.chanlock@gmail.com

Miguel A. Otaduy
Universidad Rey Juan Carlos
Madrid, Spain
miguel.otaduy@urjc.es




Figure 1: The images compare a twist deformation (far from the training data) of a real-world thin-shell micro-texture (left), a micro-scale simulation (center), and a macro-scale simulation using our homogenized material model (right).

ABSTRACT

This paper introduces a new formulation for material homogenization of thin-shell microstructures. It addresses important challenges that limit the quality of previous approaches: methods that fit the energy response neglect visual impact, methods that fit the stress response are not conservative, and all of them are limited to a low-dimensional interplay between deformation modes. The new formulation is rooted on the following design principles: the material energy functions are conservative by definition, they are formulated on the high-dimensional membrane and bending domain to capture the complex interplay of the different deformation modes, the material function domain is maximally aligned with the training data, and the material parameters and the optimization are formulated on stress instead of energy for better correlation with visual impact. The key novelty of our formulation is a new type of high-order RBF interpolant for polar coordinates, which allows us to fulfill all the design principles. We design a material function using this novel interpolant, as well as an overall homogenization workflow. Our results demonstrate very accurate fitting of diverse microstructure behaviors, both quantitatively and qualitatively superior to previous work.

CCS CONCEPTS

• Computing methodologies → Physical simulation.

KEYWORDS

micro-textures, homogenization

ACM Reference Format:

Antoine Chan-Lock and Miguel A. Otaduy. 2024. Polar Interpolants for Thin-Shell Microstructure Homogenization. In *SIGGRAPH Asia 2024 Conference Papers (SA Conference Papers '24)*, December 3–6, 2024, Tokyo, Japan. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3680528.3687693>

1 INTRODUCTION

Real-world materials exhibit complex nonlinear and anisotropic behaviors that are difficult or impossible to reproduce with standard constitutive material models. For this reason, we have seen many attempts at modeling these complex behaviors using strain-space interpolation techniques [Miguel et al. 2012; Wang et al. 2011; Xu et al. 2015]. Microstructures are particularly challenging materials. By designing the micro-scale material distribution, one can achieve macro-scale behaviors that even surpass the nonlinearity and anisotropy of natural materials [Bickel et al. 2010; Konaković-Luković et al. 2018; Panetta et al. 2015; Schumacher et al. 2015].

In this work, we study the material homogenization of thin-shell microstructures, i.e., the design of constitutive material models that match the macro-scale membrane (in-plane) and bending (out-of-plane) response of thin-shell microstructures. Such thin-shell microstructures pose important challenges to material homogenization: the continuum deformation space is six-dimensional, anisotropy arises along arbitrary directions, and there may be a complex interplay between membrane and bending deformations. For instance, the microstructure shown in Fig. 1 releases energy in two complex ways when stretched: it exhibits auxetic behavior (i.e., it widens when stretched), and it produces characteristic folds (See also Fig. 9).

When designing interpolation material models for thin shells, we set the following desiderata: they should model the membrane-bending interplay, they should be conservative, and they should optimize the strain-stress response, as this shows good correlation with visual impact. Despite the large number of previous methods, to the best of our knowledge none of them meets all these criteria. We provide a detailed discussion in Section 2, but in summary,

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

SA Conference Papers '24, December 3–6, 2024, Tokyo, Japan

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1131-2/24/12.

<https://doi.org/10.1145/3680528.3687693>




Figure 2: Conservativeness of the mechanical model is critical for simulation robustness. The same microstructure is simulated with energy-based line-search optimization (left), suitable for conservative models, and with stress-norm line-search optimization (right), the common choice for non-conservative models. The latter fails to converge.

previous methods either fit the strain-stress response and are not conservative [Wang et al. 2011], or they fit the energy response to be conservative but then fail to account for visual impact [Miguel et al. 2016; Sperl et al. 2020]. We present *polar interpolants*, a novel approach for material homogenization of thin-shell microstructures that meets all desiderata.

The first ingredient of our approach is to represent both membrane and bending strain in terms of magnitude and direction, as detailed in Section 3. We refer to this representation as *eigen-strain*. This strain representation is not new, and it has several advantages: it is easier to produce deformation data that is well aligned with the function domain, and fitting quality improves because the function domain is less curved. However, the eigen-strain representation complicates the design of a conservative material model that optimizes the strain-stress relationship. We are not aware of any suitable interpolation-based material that handles the periodicity of the eigen-strain representation.

The second ingredient of our approach is the design of high-order polar interpolants, i.e., interpolation functions that fit high-order derivatives on periodic domains. Our polar interpolants are an extension of the high-order interpolants of Chan-Lock et al. [2022] to periodic domains. In Section 4 we detail their formulation and we show how we design interpolation-based thin-shell energy models as a function of eigen-strain.

The third and last ingredient of our approach is the workflow for material homogenization, which we describe in Section 5. We present the generation of training data and the progressive refinement of the material model.

With all ingredients together, we achieve high-quality fitting of complex thin-shell microstructures, both qualitatively and quantitatively superior to previous work [Sperl et al. 2020]. Even though our training data is made of uniform deformations with uniaxial curvature only, we show good match of non-uniform deformations with biaxial curvature, which extend well beyond the training conditions. We also show qualitative matching of real-world 3D-printed microstructures.

All the data and code of our work can be found in the following URL <https://github.com/antoine-chan-lock/PolarInterpolants>.

2 RELATED WORK

The most popular approach to model a material’s mechanical response is to use a constitutive law with few parameters [Kim 2020;




Figure 3: The same microstructure (left) is stretched 10% (middle) and bend (right). Energy is $100\times$ higher under stretch, which indicates little correlation between energy and visual deformation.

Li and Barbić 2015; Smith et al. 2018; Wen and Barbić 2023]. Consequently, many works have addressed the characterization of thin-shell materials (e.g., cloth) using low-dimensional parametric models. Bhat et al. [2003] pioneered the estimation of cloth material parameters from video. In recent years, efforts are aimed at simplifying the characterization methodology, in particular by leveraging large data sets for learning-based material estimation from limited data [Feng et al. 2022; Rodriguez-Pardo et al. 2023].

Using low-dimensional parametric models works well for authoring VFX or when an average fit to the mechanical response is sufficient. However, research on microstructures soon showed that these low-dimensional parametric models cannot capture the richness of microstructure behaviors, and researchers moved to strain-dependent material parameterizations [Bickel et al. 2010].

When choosing a strain-dependent material parameterization, one of the properties often desired is conservativeness, as it can guarantee simulation stability and robustness via line-search energy optimization [Gast et al. 2015; Li et al. 2020]. In the absence of conservativeness, a common choice for line search is the stress norm, but this is not robust for our microstructure simulations, as shown in Fig. 2. Sperl et al. [2020] achieved conservativeness by modeling the material as an energy function and fitting measured energy values. However, we have observed that, for thin-shell microstructures, energy shows little correlation with visually apparent deformation. Fig. 3 compares the same microstructure under stretch and bending. Due to the large difference in stiffness, the energy is $100\times$ higher in the stretch deformation. A least-squares fitting of energy values would dramatically blur the response under bending.

Miguel et al. [2016] fitted energy functions to deformation examples. This approach combines conservativeness with a visually meaningful fit; nevertheless, the optimization process is challenging and prone to local minima. The authors partially circumvented this issue by formulating additive energy models in an incremental fashion. In our numerical homogenization approach, training data can be produced under much more controlled conditions, and the optimization process is simplified. Xu et al. [2015] designed conservative materials following the Valanis-Landel assumption, with energy terms dependent on principal strains. Our eigen-strain representation also uses principal strains, but it includes principal directions to form a polar-coordinate representation for anisotropy.

In contrast to fitting energy values, fitting stress values allows one to account for the mechanical response on each deformation mode separately, hence it leads to better correlation with visual deformation. Several authors parameterized the strain-stress response by interpolation functions, e.g., Lamé parameters [Bickel et al. 2009], the stiffness tensor [Wang et al. 2011], or stress values

directly [Wang et al. 2020]. Unfortunately, none of these methods is conservative. We would like to fit stress values by parameterizing an energy function, to achieve both good correlation with visual deformation and conservativeness. However, naïve fitting of stress values starting from an energy spline or RBF leaves us short of parameter degrees of freedom, i.e., with thin-shell stress being a six-dimensional derivative of energy, stress is only well defined by a six-tuple of energy values. Chan-Lock et al. [2022] addressed this challenge with high-order RBF interpolants, but they used it only for three-dimensional in-plane strain. We extend high-order RBF interpolants to more complex domains, showing accurate fitting of both in-plane and out-of-plane deformation.

Neural networks have also been used for strain-dependent material parameterization. The approach has a long history in computational mechanics [Colasante and Gosling 2016; Ghaboussi et al. 1991; Le et al. 2015; Shen et al. 2005; Xu et al. 2021]. Some of these works even model conservative materials [Linka et al. 2021; Masi et al. 2021]. The recent work of Li et al. [2023] is particularly interesting for metamaterial design, as it models the mechanical response of material families with a single neural network. Neural networks offer an interesting approach to material parameterization due to their flexibility, but they suffer from parameter explosion and they typically lack explainability. Our approach, based on RBF interpolants, provides closer insight into the material behavior. One interesting approach for future work would be to use our high-order polar interpolants as building block for neural material design.

To conclude this discussion, it is worth noting that multiple works on microstructure characterization or design have used various forms of homogenization, as a way to condense their complex material response [Li et al. 2023; Schumacher et al. 2018; Sperl et al. 2022; Zhang et al. 2023]. Our work introduces a new material parameterization, aimed at improving the accuracy of previous homogenization methods, and can serve as a building block for material design techniques.

3 THIN-SHELL EIGEN-STRAIN

In this section, we describe our choices of membrane and bending deformation metrics. Both use a representation based on eigen-analysis, with principal directions of deformation, which has notable advantages for the design of interpolation-based material functions. We start the section summarizing the choice of fundamental deformation metrics. Then, we describe the eigen-representation, and we discuss its advantages and challenges.

3.1 Membrane and Bending Deformation

To describe membrane (in-plane) deformation, we start from the well-known quadratic Green strain. With deformation gradient F , Green strain is defined as:

$$E = \frac{1}{2} (F^T F - I). \quad (1)$$

To describe bending (out-of-plane) deformation, we start from the triangle-averaged shape operator [Grinspun et al. 2006]:

$$\Lambda = \sum_i \frac{\xi_i}{2 A l_i} t_i \otimes t_i, \quad (2)$$




Figure 4: The distribution of training data is more regular and better suited for RBF interpolation on the eigen-strain domain $(\lambda_E, \mu_E, \theta_E)$ (right) than on the Green-strain domain E (left).

with A the triangle area, $\{\xi_i\}$ edge-based signed dihedral angles, $\{l_i\}$ edge lengths, and $\{t_i\}$ outward-pointing edge vectors. Given edge vectors $\{e_i\}$ and triangle normal n , we compute $l_i = |e_i|$ and $t_i = e_i \times n$. All values are measured in the deformed setting.

3.2 Eigen-Representation of Strain

We use an eigen-representation of strain for both membrane and bending deformations. Let us denote either strain as $S \in \{E, \Lambda\}$. Given its eigen-decomposition, we denote the eigen-values as $\lambda, \mu, \|\lambda\| > \|\mu\|$. We characterize the strain S as a 3-vector (λ, μ, θ) , where λ and μ are, by definition, the principal strains, and θ is the angle of the eigen-vector of λ , i.e., the principal strain direction.

There are two main advantages of using the eigen-representation of strain as the domain for interpolation-based material design. The first advantage is that the representation is naturally well suited for capturing anisotropic material response, i.e., direction-dependent differences in material stiffness. The second advantage is that the homogenization training data fits more compactly the strain domain. Fig. 4 compares the in-plane training strains of our homogenization process in eigen-domain $(\lambda_E, \mu_E, \theta_E)$ vs. regular Green-strain domain (E_{xx}, E_{yy}, E_{xy}) . The eigen-representation produces a more regular and less curved distribution of the training data, and this leads to higher accuracy in material estimation.

We are not the first to choose an eigen-representation of strain, but we find the treatment of this representation in prior work too simplistic. Wang et al. [2011] used the directions of principal stretch and principal curvature to interpolate in-plane stiffness tensors and bending stiffness separately. However, their model is not conservative and they ignore force Jacobian terms due to changes in stiffness. Therefore, they never cared to compute derivatives wrt the direction of strain. Sperl et al. [2020] used the direction of principal curvature to interpolate bending energies, but they proposed this only for runtime evaluation, not during the material fitting process. In Section 6 we demonstrate the improved accuracy of our model.

The eigen-representation of strain comes with one challenge though: the periodic nature of the principal strain direction θ . The design of interpolation-based materials cannot be addressed from a Euclidean perspective, and it requires instead special treatment of strain distances on a periodic domain. In the next section, we present our key novel development, polar high-order interpolants,

which enable stress fitting of conservative material functions on periodic domains.

4 POLAR INTERPOLANTS

In our design of interpolation-based thin-shell energy functions, we wish to follow the methodology of high-order RBF interpolants. However, RBFs require special distance functions on periodic domains; therefore, it is necessary to extend the fundamental theory of high-order RBF interpolants to such periodic domains.

We start the section with a short review of high-order RBF interpolants in regular Cartesian coordinates. Then, we derive our new polar high-order RBF interpolants. With these ingredients, we address the design of interpolation-based thin-shell energies for membrane and bending deformation. Our narrative combines motivation, formalization of variables and functions, and discussion. In the supplementary document, on the other hand, we include a more implementation-oriented description of the energy model.

4.1 Cartesian High-Order RBF Interpolants

The development of matrix-valued RBFs [Fuselier 2008; Narcowich and Ward 1994] considers interpolants of the type:

$$\psi = w^T \nabla \Phi(r), \quad (3)$$

where w is a vector of RBF weights, Φ is an RBF, and $r = \|u\|$ with $u = \Delta x = x - x_i$ the regular difference vector for Cartesian coordinates (between the point of evaluation of the RBF x and the RBF center x_i). This definition of RBFs yields a vector function $v = \nabla \psi$. The key difference with scalar RBFs is in the vector-valued weights, which provide the necessary degrees of freedom for interpolating vector fields, while starting from scalar interpolants, i.e., they are conservative by definition.

Chan-Lock et al. [2022] showed that the scalar interpolant (3) can be recast using RBFs directly, instead of RBF gradients as:

$$\psi = w^T u \phi(r). \quad (4)$$

Their demonstration is based on redefining the RBF $\phi(r) = \frac{\Phi'(r)}{r}$. They also extended the methodology beyond vector-field interpolation, and they defined high-order interpolants for higher-order derivatives of conservative fields (e.g., Hessians). We do not use this aspect of their work, as we found that fitting Hessians did not increase accuracy with respect to the parameter count, probably because we fit energies in 5D vs. their 3D work.

4.2 Polar High-Order RBF Interpolants

We aim to extend the method of high-order RBF interpolants to polar coordinates, i.e., coordinate sets that combine linear coordinates and periodic angular coordinates. We start by defining a new difference vector, which we refer to as *polar difference vector*:

$$u = \begin{cases} \Delta x, & \text{if } x \text{ is linear.} \\ \sin(\Delta x), & \text{if } x \text{ is angular.} \end{cases} \quad (5)$$

For angular coordinates, we use a difference function that is smooth and evaluates to 0 at 0 and π . This ensures continuity and smoothness of the difference function for periodic angular coordinates.

Starting from the regular matrix-valued RBF interpolant (3), we evaluate the RBF gradient with our new polar difference vector:

$$\psi = w^T \Phi'(r) \nabla u \frac{u}{r}. \quad (6)$$

Same as Chan-Lock et al., we redefine the RBF $\phi(r) = \frac{\Phi'(r)}{r}$. Then, we obtain the following polar high-order RBF interpolant:

$$\psi = w^T \nabla u u \phi(r), \quad (7)$$

with $\text{diag}(\nabla u) = \begin{cases} 1, & \text{if } x \text{ is linear.} \\ \cos(\Delta x), & \text{if } x \text{ is angular.} \end{cases}$

The resulting polar high-order interpolant resembles strongly the Cartesian high-order interpolant (4), with two key differences for angular coordinates: the new difference function $\sin(\Delta x)$, and a weight $\cos(\Delta x)$. These subtle but key changes allow us to use the power of high-order RBF interpolants together with the material design accuracy of eigen-representations of strain.

4.3 RBF Energy Model

The high-order polar RBF interpolants are the main ingredient to design a general interpolation-based energy for thin-shell membrane and bending deformation. Nevertheless, we need to address two additional questions. First, we need to define the specific interpolation domain for the fitted energy function, and we do this with awareness of the available training data. Second, we need to define the energy behavior under isotropic strain ($\lambda = \mu$), where the eigen-strain is not differentiable because θ is undefined, and we do this by blending anisotropic and isotropic energy evaluations.

In defining the eigen-strain domain and the full energy model, we follow this notation. We denote as $\hat{x} = (\lambda_E, \mu_E, \theta_E, \lambda_\Lambda, \mu_\Lambda, \theta_\Lambda) \in \mathbb{R}^6$ the full thin-shell eigen-strain, as described in Section 3. However, as emphasized by Sperl et al. [2020], the training data for homogenization is limited to a 5D subspace with $\mu_\Lambda = 0$ (i.e., uniaxial/cylindrical bending, see more details in Section 5.1). Then, we define a cylindrical-bending eigen-strain domain $\tilde{x} = (\lambda_E, \mu_E, \theta_E, \lambda_\Lambda^*, \theta_\Lambda^*) \in \mathbb{R}^5$. Details on the values of the bending strain $(\lambda_\Lambda^*, \theta_\Lambda^*)$ will follow. Finally, we normalize the principal strain components $\lambda_E, \mu_E, \lambda_\Lambda^*$ based on the range of the training data, to ensure that all training strains are in the range $[-1, 1]$. The polar difference operator achieves this for angular strains, hence no normalization is necessary. As a result, we obtain the RBF interpolation domain $x \in \mathbb{R}^5$. In this normalized strain domain, we construct fundamental RBF interpolants $\{\psi_i\}$ following (7), and we combine them to produce the fitted energy model which we denote as Ψ . Finally, we define a full energy model Ψ_{full} to robustly support isotropic deformations, which details described next.

We follow the approach of Sperl et al. [2020] to address the dimensionality mismatch of 5D training data $\tilde{x}$ vs. 6D runtime strain $\hat{x}$. We evaluate each RBF interpolant ψ_i (with the same RBF center and weights) for both principal curvatures. This yields the fitted energy model:

$$\Psi = \sum_i \psi_i(\lambda_E, \mu_E, \theta_E, \lambda_\Lambda^* \leftarrow \lambda_\Lambda, \theta_\Lambda^* \leftarrow \theta_\Lambda) + \psi_i(\lambda_E, \mu_E, \theta_E, \lambda_\Lambda^* \leftarrow \mu_\Lambda, \theta_\Lambda^* \leftarrow \theta_\Lambda + \frac{\pi}{2}). \quad (8)$$

Previous interpolation-based energy models [Miguel et al. 2016; Sperl et al. 2020] typically defined interpolants in low-dimensional subspaces, and then combined multiple energy terms in an additive manner. We find that this is not necessary for our homogenization methodology, and we can design a 6D energy directly, thanks to our polar interpolation approach.

The fitted energy (8) is robust under anisotropic strain ($\lambda \neq \mu$). However, as mentioned above, it is not differentiable in the isotropic case ($\lambda = \mu$), where θ is undefined. For runtime evaluation, we handle this through localized blending of isotropic evaluations of the fitted energy (8), forcing the principal strain direction $\theta = 0$. Then, our full thin-shell energy model becomes:

$$\begin{aligned} \Psi_{\text{full}} = & (1 - \omega(\lambda_E, \mu_E)) (1 - \omega(\lambda_\Lambda, \mu_\Lambda)) \Psi(\lambda_E, \mu_E, \theta_E, \lambda_\Lambda, \mu_\Lambda, \theta_\Lambda) \\ & + \omega(\lambda_E, \mu_E) (1 - \omega(\lambda_\Lambda, \mu_\Lambda)) \Psi(\lambda_E = \mu_E, \theta_E = 0, \lambda_\Lambda, \mu_\Lambda, \theta_\Lambda) \\ & + (1 - \omega(\lambda_E, \mu_E)) \omega(\lambda_\Lambda, \mu_\Lambda) \Psi(\lambda_E, \mu_E, \theta_E, \lambda_\Lambda = \mu_\Lambda, \theta_\Lambda = 0) \\ & + \omega(\lambda_E, \mu_E) \omega(\lambda_\Lambda, \mu_\Lambda) \Psi(\lambda_E = \mu_E, \theta_E = 0, \lambda_\Lambda = \mu_\Lambda, \theta_\Lambda = 0), \end{aligned} \quad (9)$$

with a weight function $\omega(\lambda, \mu) = e^{-(\lambda-\mu)^Y}$ that localizes the isotropic evaluations. It is not necessary to design multiple energy terms, we fit just a single energy (8), and then we blend four evaluations of this energy at runtime: anisotropic stretch and bending, isotropic stretch, isotropic bending, and isotropic stretch and bending.

The final energy model Ψ_{full} does not guarantee zero stress at zero strain by construction. We could add the zero-strain stress compensation term of Chan-Lock et al. [2022] a posteriori of model fitting, but we found this to be unnecessary in practice, because the fitted stress at zero strain was negligible in all our experiments.

5 HOMOGENIZATION WORKFLOW

In this section, we describe the data generation and fitting process. Our approach to data generation is largely similar to previous work based on simulation with periodic boundary conditions, which simplifies the exploration of continuum strain values. In the description of the energy-model fitting process, we discuss data normalization and the optimization methodology, but we refer to the supplementary document for full details.

5.1 Generation of Training Data

To generate training data, we follow a controlled simulation approach. In a nutshell, we simulate microstructure deformations under homogeneous coarse strain x using periodic boundary conditions, and we compute the homogenized stress $\frac{\partial \Psi}{\partial x}$. In this way, we obtain data pairs $(x, \frac{\partial \Psi}{\partial x})$, which we use for fitting our RBF energy model, as we discuss in detail in the next subsection.

We sample the strain domain x in the following way. For stretch, we start by sampling the first principal stretch λ_E and the principal direction θ_E . For each pair, we obtain the minimum-energy orthogonal stretch, and we sample the second principal stretch μ_E in a positive and negative range around this minimum-energy configuration. For bending, as discussed by Sperl et al. [2020], it is not possible to apply arbitrary uniform coarse bending. We follow their approach, and apply only uniaxial bending, i.e., we sample the first principal curvature λ_Λ and the principal direction θ_Λ , with the other principal curvature $\mu_\Lambda = 0$. To summarize, for each microstructure,

we sample regularly a 4D domain $(\lambda_E, \theta_E, \lambda_\Lambda, \theta_\Lambda)$. Then, for each 4D sample we adapt the sampling of μ_E around the minimum energy configuration, and we leave $\mu_\Lambda = 0$.

There are a few important considerations about the sampling range. For in-plane strain, we have sampled 50% compression to 100% stretch, which is well beyond the range that the material can sustain in reality without breaking. For bending strain, we define the maximum λ_Λ as the value that makes the simulation tile a perfect cylinder. As all our microstructures are homogeneous in the thickness direction, deformations are identical for positive and negative curvatures, and we sample only positive ones. For the principal stretch and bending directions, we sample the range $[0, \pi)$. Several authors have assumed 90-degree symmetry [Sperl et al. 2020; Wang et al. 2011], but we find that this is not the case for all microstructures. Fig. 5 shows the clear 90-degree asymmetry of Microstructure 1.

For the simulation of microstructures with uniform coarse deformation under periodic boundary conditions, we refer to the work of Sperl et al. [2020; 2022]. We use a simulation domain formed by exactly one tile of the microstructure pattern, and we solve the microstructure simulation with a 3D FEM discretization and a nearly incompressible Neo-Hookean hyperelastic material model [Mousanezhad et al. 2015]. To compute the coarse stress for each coarse strain, it suffices to measure microstructure forces f_k on nodes X_k of the boundary of the simulation tile [Schumacher et al. 2018; Sperl et al. 2022]: $\frac{\partial \Psi}{\partial x} = \sum_k f_k^T \frac{\partial X_k}{\partial x}$.

5.2 Fitting of the Energy Model

Our energy model is parameterized by the number of RBF kernels, their center locations in eigen-strain domain $\{x_i\}$, the RBF radii, and the vector-valued weights $\{w_i\}$. In all our experiments, we have used Gaussian RBFs. We approach the optimization of the energy model by staggering steps of kernel placement, radius refinement, and weight estimation. Prior to this, we normalize all the stress data. Recall from Section 4.3 that the eigen-strain x is also normalized based on the range of the training data.

We form a training batch with all the training data that shares the principal stretch value λ_E . Then, within each batch, we normalize each stress component to the range $[-1, 1]$, based on the range of the stress component within the batch. We form normalization batches based on λ_E because it is the dominant strain, and otherwise the samples with low λ_E (e.g., only bending but no stretch) suffer high relative stress error.

For fitting, we do not need the full energy model (9) with isotropic and anisotropic terms; it suffices to fit the fundamental energy model (8). This is possible for two reasons. One is that, in contrast to runtime simulation, there is no discontinuity issue in the principal directions. The second one is that we augment the isotropic data $\lambda = \mu$, inputting the same stress values for multiple data entries with different strain directions θ (we use 10 directions in practice).

Given all normalized training data pairs $(x, \frac{\partial \Psi}{\partial x})$, we finally fit the energy model (8). For each optimization iteration, we add an RBF center based on furthest point sampling, we optimize a common RBF radius (using COBYLA), and we fit the RBF weights. Note that we fit optimal RBF weights for each candidate RBF radius within the execution of COBYLA. We compute these weights by minimizing




Figure 5: The plot shows, for Microstructure 1, the energy density wrt bending direction under uniform uniaxial curvature. This microstructure reveals a clear 90-degree asymmetry, underlining the importance of fitting the material model on a range $[0, \pi)$ of principal directions.

the L^2 stress error at all the training samples, which amounts to a simple convex quadratic optimization problem on the RBF weights. We conclude the optimization when the RMSE of stress is below 10%, or if the optimization process stagnates. For our test materials, only 3 of them failed to reach the 10% error threshold. Please see a more detailed discussion in the next section.

6 EXPERIMENTS AND DISCUSSION

6.1 Implementation Details

Table 1 shows all 10 microstructures tested in our experiments. For all simulations, we considered a TangoGray material. We also fabricated 4 of the microstructures (ids: 1, 3, 4 and 5) using the TangoGray material, on an Object30 Prime Stratasys 3D printer. As mentioned in Section 5.1, we have modeled the micro-structure material using a nearly-incompressible hyperelastic Neo-Hookean. We adopt the energy density function and the material parameters for TangoGray from [Mousanezhad et al. 2015], in particular $\Psi = \frac{1}{2} G_0 (J^{-2/3} \text{tr}(F^T F) - 3) + \frac{1}{2} K_0 (J - 1)^2$, $G_0 = 1.7$ MPa and $K_0 = 84.43$ MPa. We have executed the training microstructure simulations using regular hexahedral meshes.

There are a few technical details worth noting. For high-quality homogenized results, we have used regular equilateral meshes, which guarantee convergence of the shape operator (2), as discussed in [Grinspun et al. 2006]. To obtain robustly the principal direction from 2D eigenvectors, we follow [Blinn 1996]. All our simulations are the result of solving for static equilibrium, and we do this with standard Newton with line search. In the supplementary document we provide gradients and Hessians of the runtime energy model wrt eigen-strain. For the derivative of eigen-strain wrt degrees of freedom, we have used symbolic differentiation, which can be found in our code. As a precaution, to prevent Newton steps from reaching outside the material’s domain, we add a stiff quadratic barrier energy per strain component at the maximum training strain. Note that no simulation actually hit the boundary at equilibrium, as our training domain is conservative.

6.2 Ablation Study

We have evaluated fitting accuracy under different model settings. In particular, we have tested two settings used by Sperl et al. [2020]: energy fitting and linear interpolation of energy models between

0° and 90° in the bending direction θ_Λ . We refer to the latter as “linear θ_Λ ”, in contrast to our method “RBF θ_Λ ”. Table 2 compares fitting error in both cases, with the same number of parameters (300). It also compares results when fitting either energy or stress. Note that energy fit with linear θ_Λ (equivalent to [Sperl et al. 2020]) suffers a stress error of more than 95%. This error goes down to 25% with our approach.

We have also validated the impact of stress normalization. Under a common model complexity of 300 RBFs for all microstructures, the averaged error per stress component ranges between 7.2% and 62.7% without normalization, and 11% and 40.3% with normalization. The lowest error is for principal stretch and the highest for principal curvature in both cases. Normalization slightly reduces accuracy on stretch, but it largely improves accuracy on bending.

6.3 Model Estimation Results

We have fitted our model to a relative stress RMSE threshold of 10% on all 10 microstructures. Table 1 summarizes the number of RBFs used and the final error.

For 3 of the microstructures, we could not reach an error of 10%, as the optimization process stagnated. All these 3 microstructures present auxetic behavior and a tendency to buckle. We believe that the difficulty of the optimization comes from discontinuities in the stress values, due to buckling. Nevertheless, in all cases the final fitting error was below 16%.



6.4 Controlled Validation Tests

We have tested simulations with non-uniform strain, with exactly the same boundary conditions for fine simulations at micro-structure scale, as well as coarse simulations on triangle meshes. We have used testing conditions inspired by stretch and pear-loop experiments used, respectively, for in-plane and out-of-plane cloth characterization [Sperl et al. 2022], as well as twist experiments that produce, at the same time, stretch strain and biaxial bending strain. Furthermore, we have executed these experiments for material patches cut along 4 different directions for each micro-structure: 0° , 45° , 90° , and 135° .

Fig. 7 to Fig. 12 show the final simulation results of both fine (green) and coarse (purple) simulations, for all 10 microstructures, along all 4 directions. For the stretch simulations, all cases show an excellent match. It is worth noting that our fitted models reproduce the auxetic response of the micro-structures, and even complex effects such as the longitudinal folds shown by Microstructure 4. For the bending simulations, the resulting accuracy is not as high, but the anisotropy effects are matched qualitatively, and differences across microstructures are evident. For the twist simulations, the main folds and shapes are matched, but finer wrinkles are often missed. We believe these are micro-scale effects, often due to buckling, that cannot be captured by a homogenized material.

We have compared quantitatively the bending tests vs. an energy-fit model [Sperl et al. 2020]. With our model, the average error in the pear-loop ratio height-width ratio is 15% across all materials and directions, and the maximum error is 28.5%. With the energy-fit model, the average error is 21.5% and the maximum error is 106%. This result supports the higher accuracy of bending stiffness with our model.

Table 1: The top row shows the 10 microstructures tested in our experiments, and the next row the id corresponding to each microstructure. We have fitted our homogenized material model to each microstructure, until the stress RMSE is below 10% (or if the process stagnates). The third row shows the number of RBFs of the resulting material model, and the fourth row the fitting error. Finally, the fifth row shows the number of nodes of the microstructure-level simulation meshes used for controlled validation experiments.

										
Id	1	2	3	4	5	6	7	8	9	10
RBFs	130	180	130	500	130	480	280	450	480	450
Error %	9.82	9.94	9.83	14.84	9.75	9.91	9.81	9.90	15.96	12.11
Nodes	17k	27k	22k	34k	6k	28k	57k	44k	27k	20k

Importantly, our coarse simulation results are obtained with meshes with just a fraction of the nodes of the fine simulations. Table 1 lists the mesh complexity of the fine simulations. For the coarse simulations, we have used just 400 nodes for stretch and 900 for twist, for all materials. The ratio fine vs. coarse ranges between 6 (microstructure 5, twist) and 142 (microstructure 7, stretch). For the bending experiments we only simulated a thin strip, therefore the differences are not as dramatic.

6.5 Real-World Validation Tests

We have also tested simulations vs. real-world deformations with 3D-printed microstructures. Fig. 6 shows twist deformations on 4 microstructures. We can see clear differences in the folding shapes across the 4 microstructures, and these differences are captured by our fitted materials. The bottom row of the figure shows the result of an energy-fit model. Some materials wrinkle too much, indicating that their bending stiffness is too low, which matches the result observed in the controlled validation tests.

7 CONCLUSIONS AND FUTURE WORK

We have presented a method for the homogenization of thin-shell microstructures, based on a novel type of RBF interpolant with attractive properties: it ensures model conservativeness by construction, it provides degrees of freedom to naturally fit stress data, and it ensures smoothness in polar coordinates that match well the domain of training data. We have shown the qualitative match between the resulting homogenized material models and both simulated and 3D-printed microstructures. We have also shown fitting accuracy that is superior to previous methods.

The eigen-strain parameterization enjoys important benefits, but it also comes with numerical challenges. We have introduced *polar difference* as a robust and elegant way of handling the periodicity of

the parameterization. And we propose blending of isotropic model evaluations to ensure differentiability in the isotropic case. A more efficient treatment would be desirable for the latter challenge, but we are not aware that such treatment exists.

Our current solution suffers some limitations that could motivate further work. One limitation is that the training process only explores uniaxial bending data. While the quality of our biaxial bending validations based on twist experiments is high, during training there is no guarantee that the resulting model will be accurate, as there is no biaxial training data. Unfortunately, it is not possible to obtain biaxial bending data from simulations with periodic boundary conditions. Nevertheless, we think that the current model could serve as warm start for an additional training step using non-uniform deformation data.

Another possible limitation is the high number of RBFs in the resulting model. Note that this is due to the high dimensionality of the function space, as the number of RBFs listed in Table 1 (ranging 130 to 500) is comparable to extremely coarse regular grids ($3^5 = 243$, $4^5 = 1024$). As an extension to our work, it might be sensible to balance high-dimensional and low-dimensional fitting. One could start from low-dimensional energy terms designed in an additive manner [Miguel et al. 2016; Sperl et al. 2020], and add local high-dimensional RBF corrections. If the RBF radii are small, then distant RBFs can be pruned at runtime, reducing the actual number of RBF evaluations.

The final important limitation is that our model does not account for discontinuities or multi-stability. We have confirmed that these phenomena occur in practice, due to microstructure buckling, and they hinder the fitting accuracy of our method for some materials. The nonlinear activation functions in neural networks might be a natural answer to model discontinuities. We think that imposing conservativeness is anyway a desirable property, and non-conservative effects due to, e.g., internal friction from microstructure contact, could be handled using a homogenized friction model on top of the pure conservative model [Miguel et al. 2013].

To conclude, we would like to mention that our homogenized simulation renders were generated using barycentric interpolation of fine-scale geometry within the coarse triangle mesh. This can be improved with state-of-the-art solutions for strain-dependent micro-scale rendering [Montazeri et al. 2021; Sperl et al. 2021].

Table 2: Ablation study of our model (RBF θ_Λ , stress fit), evaluating both energy and stress error. All fits are computed with the same number of parameters (300).

	Linear θ_Λ		RBF θ_Λ	
	Energy Fit	Stress Fit	Energy Fit	Stress Fit
Energy error	9.19%	11.28%	4.11%	8.61%
Stress error	95.21%	31.04%	35.04%	25.17%

REFERENCES

- Kiran S. Bhat, Christopher D. Twigg, Jessica K. Hodgins, Pradeep K. Khosla, Zoran Popović, and Steven M. Seitz. 2003. Estimating cloth simulation parameters from video. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (San Diego, California) (SCA '03). Eurographics Association, Goslar, DEU, 37–51.
- Bernd Bickel, Moritz Bächer, Miguel A. Otaduy, Hyunho Richard Lee, Hanspeter Pfister, Markus Gross, and Wojciech Matusik. 2010. Design and Fabrication of Materials with Desired Deformation Behavior. *ACM Trans. Graph.* 29, 4, Article 63 (2010), 10 pages.
- Bernd Bickel, Moritz Bächer, Miguel A. Otaduy, Wojciech Matusik, Hanspeter Pfister, and Markus Gross. 2009. Capture and Modeling of Non-Linear Heterogeneous Soft Tissue. In *ACM SIGGRAPH 2009 Papers* (New Orleans, Louisiana) (SIGGRAPH '09). Association for Computing Machinery, New York, NY, USA, Article 89, 9 pages.
- J. Blinn. 1996. Consider the lowly 2 x 2 matrix. *IEEE Computer Graphics and Applications* 16, 2 (1996), 82–88. <https://doi.org/10.1109/38.486688>
- Antoine Chan-Lock, Jesús Pérez, and Miguel A. Otaduy. 2022. High-Order Elasticity Interpolants for Microstructure Simulation. *Computer Graphics Forum* 41, 8 (2022), 63–74.
- G. Colasante and P.D. Gosling. 2016. Including Shear in a Neural Network Constitutive Model for Architectural Textiles. *Procedia Engineering* 155 (2016), 103–112.
- Xudong Feng, Wencho Huang, Weiwei Xu, and Huamin Wang. 2022. Learning-Based Bending Stiffness Parameter Estimation by a Drape Tester. *ACM Trans. Graph.* 41, 6, Article 221 (2022), 16 pages.
- Edward Fuselier. 2008. Sobolev-type approximation rates for divergence-free and curl-free RBF interpolants. *Math. Comput.* 77 (2008), 1407–1423.
- Theodore F. Gast, Craig Schroeder, Alexey Stomakhin, Chenfanfu Jiang, and Joseph M. Teran. 2015. Optimization Integrator for Large Time Steps. *IEEE Transactions on Visualization and Computer Graphics* 21, 10 (2015), 1103–1115.
- J. Ghaboussi, J. H. Garrett, and X. Wu. 1991. Knowledge-Based Modeling of Material Behavior with Neural Networks. *Journal of Engineering Mechanics* 117, 1 (1991), 132–153.
- Eitan Grinspun, Yotam Gingold, Jason Reisman, and Denis Zorin. 2006. Computing discrete shape operators on general meshes. *Computer Graphics Forum (Eurographics)* 25, 3 (2006), 547–556. <https://doi.org/10.1111/j.1467-8659.2006.00974.x>
- Theodore Kim. 2020. A Finite Element Formulation of Baraff-Witkin Cloth. *Computer Graphics Forum* 39, 8 (2020), 171–179.
- Mina Konaković-Luković, Julian Panetta, Keenan Crane, and Mark Pauly. 2018. Rapid Deployment of Curved Surfaces via Programmable Auxetics. *ACM Trans. Graph.* 37, 4, Article 106 (2018), 13 pages.
- B. A. Le, J. Yvonnet, and Q.-C. He. 2015. Computational homogenization of nonlinear elastic materials using neural networks. *Internat. J. Numer. Methods Engrg.* 104, 12 (2015), 1061–1084.
- Minchen Li, Zachary Ferguson, Teseo Schneider, Timothy Langlois, Denis Zorin, Daniele Panozzo, Chenfanfu Jiang, and Danny M. Kaufman. 2020. Incremental potential contact: intersection-and inversion-free, large-deformation dynamics. *ACM Trans. Graph.* 39, 4, Article 49 (2020), 20 pages.
- Y. Li and J. Barbic. 2015. Stable Anisotropic Materials. *IEEE Transactions on Visualization & Computer Graphics* 21, 10 (2015), 1129–1137.
- Yue Li, Stelian Coros, and Bernhard Thomaszewski. 2023. Neural Metamaterial Networks for Nonlinear Material Design. *ACM Trans. Graph.* 42, 6, Article 186 (2023), 13 pages.
- Kevin Linka, Markus Hillgärtner, Kian P. Abdolazizi, Roland C. Aydin, Mikhail Itskov, and Christian J. Cyron. 2021. Constitutive artificial neural networks: A fast and general approach to predictive data-driven constitutive modeling by deep learning. *J. Comput. Phys.* 429 (2021), 110010.
- Filippo Masi, Ioannis Stefanou, Paolo Vannucci, and Victor Maffi-Berthier. 2021. Thermodynamics-based Artificial Neural Networks for constitutive modeling. *Journal of the Mechanics and Physics of Solids* 147 (2021), 104277.
- E. Miguel, D. Bradley, B. Thomaszewski, B. Bickel, W. Matusik, M. A. Otaduy, and S. Marschner. 2012. Data-Driven Estimation of Cloth Simulation Models. *Comput. Graph. Forum* 31, 2pt2 (2012), 519–528.
- Eder Miguel, David Miraut, and Miguel A. Otaduy. 2016. Modeling and Estimation of Energy-Based Hyperelastic Objects. *Computer Graphics Forum* 35, 2 (2016), 385–396.
- Eder Miguel, Rasmus Tamstorf, Derek Bradley, Sara C. Schvartzman, Bernhard Thomaszewski, Bernd Bickel, Wojciech Matusik, Steve Marschner, and Miguel A. Otaduy. 2013. Modeling and estimation of internal friction in cloth. *ACM Trans. Graph.* 32, 6, Article 212 (2013), 10 pages.
- Zahra Montazeri, Chang Xiao, Yun Fei, Changxi Zheng, and Shuang Zhao. 2021. Mechanics-Aware Modeling of Cloth Appearance. *IEEE Transactions on Visualization and Computer Graphics* 27, 1 (2021), 137–150.
- Davood Mousanezhad, Sahab Babae, Hamid Ebrahimi, Ranajay Ghosh, A.M.S Hamouda, Katia Bertoldi, and Ashkan Vaziri. 2015. Hierarchical honeycomb auxetic metamaterials. *Scientific reports* 5 (12 2015), 18306. <https://doi.org/10.1038/srep18306>
- Francis J. Narcowich and Joseph D. Ward. 1994. Generalized Hermite Interpolation Via Matrix-Valued Conditionally Positive Definite Functions. *Math. Comp.* 63, 208 (1994), 661–687.
- 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 (2015), 12 pages.
- Carlos Rodriguez-Pardo, Melania Prieto-Martin, Dan Casas, and Elena Garces. 2023. How Will It Drape Like? Capturing Fabric Mechanics from Depth Images. *Computer Graphics Forum* 42, 2 (2023), 149–160.
- 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 (2015), 136:1–136:13.
- Christian Schumacher, Steve Marschner, Markus Gross, and Bernhard Thomaszewski. 2018. Mechanical Characterization of Structured Sheet Materials. *ACM Trans. Graph.* 37, 4, Article 148 (2018), 15 pages.
- Y. Shen, K. Chandrashekhara, W.F. Breig, and L.R. Oliver. 2005. Finite element analysis of V-ribbed belts using neural network based hyperelastic material model. *International Journal of Non-Linear Mechanics* 40, 6 (2005), 875–890.
- Breannan Smith, Fernando De Goes, and Theodore Kim. 2018. Stable Neo-Hookean Flesh Simulation. *ACM Trans. Graph.* 37, 2, Article 12 (mar 2018), 15 pages.
- Georg Sperl, Rahul Narain, and Chris Wojtan. 2020. Homogenized Yarn-Level Cloth. *ACM Transactions on Graphics (TOG)* 39, 4, Article 48 (2020), 16 pages.
- Georg Sperl, Rahul Narain, and Chris Wojtan. 2021. Mechanics-aware deformation of yarn pattern geometry. *ACM Trans. Graph.* 40, 4, Article 168 (2021), 11 pages.
- Georg Sperl, Rosa M. Sánchez-Banderas, Manwen Li, Chris Wojtan, and Miguel A. Otaduy. 2022. Estimation of Yarn-Level Simulation Models for Production Fabrics. *ACM Transactions on Graphics (TOG)* 41, 4, Article 65 (2022), 15 pages.
- Bin Wang, Yuanmin Deng, Paul Kry, Uri Ascher, Hui Huang, and Baoquan Chen. 2020. Learning Elastic Constitutive Material and Damping Models. *Computer Graphics Forum* 39, 7 (2020), 81–91.
- Huamin Wang, James F. O'Brien, and Ravi Ramamoorthi. 2011. Data-Driven Elastic Models for Cloth: Modeling and Measurement. In *ACM SIGGRAPH 2011 Papers* (Vancouver, British Columbia, Canada) (SIGGRAPH '11). Association for Computing Machinery, New York, NY, USA, Article 71, 12 pages.
- Jiahao Wen and Jernej Barbic. 2023. Kirchhoff-Love Shells with Arbitrary Hyperelastic Materials. *ACM Trans. Graph.* 42, 6, Article 174 (2023), 15 pages.
- Hongyi Xu, Funshing Sin, Yufeng Zhu, and Jernej Barbic. 2015. Nonlinear Material Design Using Principal Stretches. *ACM Trans. Graph.* 34, 4, Article 75 (2015), 11 pages.
- Kailai Xu, Daniel Z. Huang, and Eric Darve. 2021. Learning constitutive relations using symmetric positive definite neural networks. *J. Comput. Phys.* 428 (2021), 110072.
- Zhan Zhang, Christopher Brandt, Jean Jouve, Yue Wang, Tian Chen, Mark Pauly, and Julian Panetta. 2023. Computational Design of Flexible Planar Microstructures. *ACM Trans. Graph.* 42, 6, Article 185 (2023), 16 pages.




Figure 6: Same as in Fig. 1, we compare twist deformations of real-world microstructures (top row), micro-scale simulations (second row), and macro-scale simulations computed with our homogenized material models (third row). But we also add simulation results with an energy-fit model (bottom row). From left to right, the images show microstructures 1, 3, 4 and 5.




Figure 7: Micro. 1 50% stretch, 90° twist, pear loop (0° , 45° , 90° , 135°)




Figure 10: Micro. 5 50% stretch, 90° twist, pear loop (0° , 45° , 90° , 135°)




Figure 8: Micro. 3 50% stretch, 90° twist, pear loop (0° , 45° , 90° , 135°)




Figure 11: Micro. 6 50% stretch, 90° twist, pear loop (0° , 45° , 90° , 135°)




Figure 9: Micro. 4 50% stretch, 90° twist, pear loop (0° , 45° , 90° , 135°)




Figure 12: Micro. 10 50% stretch, 90° twist, pear loop (0° , 45° , 90° , 135°)