

Elastic Textures for Additive Fabrication

Julian Panetta¹ *

Qingnan Zhou¹ *

Luigi Malomo²

Nico Pietroni²

Paolo Cignoni²

Denis Zorin¹

¹New York University

²Visual Computing Lab, CNR-ISTI




Figure 1: Six basic elastic textures are used to obtain a large range of homogenized isotropic material properties. A $3 \times 3 \times 1$ tiling of each pattern is shown, along with rendered (left) and fabricated (right) cell geometry below. The naming convention is explained in Section 4.

Abstract

We introduce *elastic textures*: a set of parametric, tileable, printable, cubic patterns achieving a broad range of isotropic elastic material properties: the softest pattern is over a thousand times softer than the stiffest, and the Poisson's ratios range from below zero to nearly 0.5. Using a combinatorial search over topologies followed by shape optimization, we explore a wide space of truss-like, symmetric 3D patterns to obtain a small family. This pattern family can be printed without internal support structure on a single-material 3D printer and can be used to fabricate objects with prescribed mechanical behavior. The family can be extended easily to create anisotropic patterns with target orthotropic properties. We demonstrate that our elastic textures are able to achieve a user-supplied varying material property distribution. We also present a material optimization algorithm to choose material properties at each point within an object to best fit a target deformation under a prescribed scenario. We show that, by fabricating these spatially varying materials with elastic textures, the desired behavior is achieved.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—[Geometric algorithms, languages, and systems];

Keywords: additive fabrication, microstructures, deformable objects, homogenization, shape optimization, goal-based material design

* Joint first authors

ACM Reference Format

Panetta, J., Zhou, Q., Malomo, L., Pietroni, N., Cignoni, P., Zorin, D. 2015. Elastic Textures for Additive Fabrication. *ACM Trans. Graph.* 34, 4, Article 135 (August 2015), 12 pages. DOI = 10.1145/2766937 <http://doi.acm.org/10.1145/2766937>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA. Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM 978-1-4503-3331-3/15/08 ... \$15.00. DOI: <http://dx.doi.org/10.1145/2766937>

1 Introduction

Rapid advances in the accessibility of additive fabrication has a significant impact on how manufacturable geometric models are constructed. A key distinctive feature of common additive fabrication technologies is that the cost and time of production is practically uncorrelated with structural complexity: in fact, a complex structure using less material may be both cheaper and faster to produce.

Complex structures, aside from potentially reducing costs, open up many new possibilities, in particular for manufacturing deformable objects. By varying a small-scale structure, one can adjust a variety of material properties, from elasticity to permeability. Importantly, these properties can be varied nearly continuously over the object, something that is not commonly done in traditional processes. As it was observed in prior work, this opens up many new possibilities for object behavior.

Small-scale structures present a set of new design challenges: in all but the simplest cases, these are hard or impossible to design by hand to meet specific goals. At the same time, computational optimization of fine-scale variable structure over a whole object, even of moderate size, can easily result in numerically difficult topology and shape optimization problems with millions of variables.

In this paper, we describe *elastic volumetric textures*, a library of tileable parameterized 3D small-scale structures that can be used to control the elastic material properties of an object. Applying such textures to a hex mesh with target material properties specified per element is similar to using dithering to achieve a continuous variation of brightness or color.

In a sense, almost *all* material properties owe themselves to small-scale structures at the molecular or crystal level, and a large body of work in nanoscience aims to control material properties precisely by structure design. These works must accommodate constraints imposed by the specific properties of the elements and molecules used, the need for self-assembly, and other considerations.

Our focus is on larger-scale structures, which can be manufactured using existing 3D printing technology. With feature sizes at the scale of $10\mu\text{m}$ - $100\mu\text{m}$, these are well described by conventional elasticity theory. While this type of structure was also extensively

studied, typically this was in the context of a specific problem, such as optimizing strength for a given material volume fraction. Our goal is to maximize the range of effective material properties that can be obtained using a single material by varying the structure.

We consider variable-thickness *truss-like structures*—i.e. structures composed of connected bars—as these cover a considerable range of properties on the one hand, and on the other hand, allow us to work with a relatively small number of parameters. We present a method for building a dictionary of structures that cover a large space of material properties. These structures are *tileable*, which makes it possible to vary material properties across an object, and *printable*.

We demonstrate that elastic volumetric textures allow one to control the deformation behavior of objects, either by painting material properties directly or by a two-stage *shape optimization* procedure, involving solving for variable continuous properties then approximating them using our texture dictionary. We validate our results by measuring samples for different choices of parameters and topologies and by demonstrating the deformation behavior of objects fabricated with spatially varying structures.

2 Related Work

Microstructure design and optimization. There is a huge literature on theoretical studies of effective moduli of composites (our periodic structures are an extreme example of a composite combining a material with void). Recent monographs include [Cioranescu and Donato 1999; Milton 2002; Torquato 2002]. Much of the literature focuses on identifying microstructures with *extremal* effective behavior, i.e., with effective elasticity properties at the boundary of the achievable zone for a given class of composites [Allaire 2002; Cherkaev 2000; Milton 2002]. Many classes of extremal structures were described (see, e.g., [Cadman et al. 2013]), however most of these classes—e.g. sequentially laminated microstructures [Avelaneda 1987] and microstructures based on inclusions [Grabovsky and Kohn 1995; Liu et al. 2007]—are either difficult or impossible to manufacture at this time. Interchangeable composites and other structures were found that maximize simultaneously, e.g., the bulk modulus and permeability [Guest and Prévost 2006] or electrical conductance [Torquato et al. 2002; Torquato et al. 2003; Torquato and Donev 2004], but these designs are of limited use for tailoring elastic behavior.

The closest work to ours is [Sigmund 1995], which constructs truss microstructures with prescribed elasticity tensors. It starts with a full “ground structure” containing about 2000 candidate members, then optimizes the members’ thicknesses *but not offsets* to obtain a microstructure period cell whose homogenized properties (computed using a truss model) match the desired properties. Neither tileability of structures for different parameters nor printability can be guaranteed. We discuss the differences in greater detail at the end of Section 4. Further exploration of periodic structures of this type was done more recently in [Chu et al. 2010], comparing different methods for optimizing these structures.

A number of microstructures were obtained using various types of *topology optimization*, which was originally designed for global structure optimization. In the case of microstructure design, these methods look for a periodic structure minimizing, e.g., compliance for a fixed total volume fraction. The result is normally a single-scale structure, with scale controlled by the resolution of the simulation grid or other types of regularization. Important methods proposed for solving these problems include solid isotropic material with penalization (SIMP) and rational approximation of material properties (RAMP) [Bendsøe 1989; Bendsøe and Sigmund 2003;

Nakasone and Silva 2010]. [Radman et al. 2013] demonstrated design of isotropic materials maximizing bulk modulus.

Topology optimization offers more flexibility in the choice of structure, but it requires a relatively expensive optimization for each specific problem. The ability to undergo topological transitions under continuous parameter changes is both a strength, as it allows exploration of a broader space of structures, and a weakness, as it considerably complicates design of parametric families satisfying printability and tileability constraints, which motivates our approach.

Microstructure fabrication. Several groups focusing on additive fabrication have recently obtained encouraging results. In particular, materials previously thought to be unmanufacturable were produced and behave as expected. Notably, the work of Hollister and collaborators [Lin et al. 2004a; Lin et al. 2004b; Hollister 2005; Kang 2010] in the context of bone scaffold design and fusion cage design demonstrated the use of optimized microstructures. The possibility of manufacturing auxetic (negative Poisson’s ratio) materials was demonstrated in [Greaves et al. 2011], and in [Schwerdtfeger et al. 2011; Bückmann et al. 2012; Andreassen et al. 2014].

The idea of fabricating tileable structures with varying properties also appears in [Hiller and Lipson 2009] in which the authors discuss “digital materials,” as composed of a set of discrete voxels with predefined shapes that can be connected. Similarly, a building-block based approach was also used in the context of bio-printing [Mironov et al. 2009], where the authors use spheroids of living materials with evolving and controllable composition, varying material and biological properties in time.

Compliant mechanisms. The material optimization method that we present solves a similar problem to that of compliant mechanism design. [Bendsøe and Sigmund 2003] reviews several existing approaches to designing mechanisms that maximize mechanical advantage/output deflection or tune an output displacement to a particular path. These approaches have little control over the resulting structure’s macroscopic shape, whereas tuning deformation behavior using our microstructure approach creates a “mechanism” that still looks like the input shape.

Fabrication and computer graphics. A broad variety of fabrication-related work has been done in the computer graphics community. Several techniques have been proposed to design paper craft objects [Mitani and Suzuki 2004], plush objects [Mori and Igarashi 2007], and objects made of interlocking planar slices [Cignoni et al. 2014; Schwartzburg and Pauly 2013; Hildebrand et al. 2012]. Other techniques use geometric techniques to change surface appearance by synthesizing surface microgeometry [Weyrich et al. 2009] or changing the shape to generate custom target caustics [Schwartzburg et al. 2014].

Another close work to ours, [Bickel et al. 2010], introduces an optimization process to find the best combination of stacked layers to satisfy an input deformation, enabling fabrication of objects with complex heterogeneous materials using multi-material 3D printers. Our work can be viewed as complementary, focusing on the design of structures that can be, e.g., used as a part of deformation behavior design; our material optimization method provides an alternative to the method in that paper. In [Skouras et al. 2013], multi-material printing and discrete material optimization is used in a similar way on complex characters to achieve desired deformations with actuation. Our elastic textures can be viewed as a tool for solving this type of problem. Our structures also can be employed in systems like [Chen et al. 2013] and [Vidimče et al. 2013].

Homogenization. A central tool in our work, homogenization was used in graphics for reducing complexity of physical models in [Kharevych et al. 2009], finding the constitutive parameters of a

low resolution discretization that best approximates the behavior of the original higher complexity material. The periodic homogenization method that we use is based on the one described in [Allaire 2002].

3 Overview and Main Results

In this section, we describe our overall approach, visualized in Figure 2, and a specific set of patterns that we have obtained.

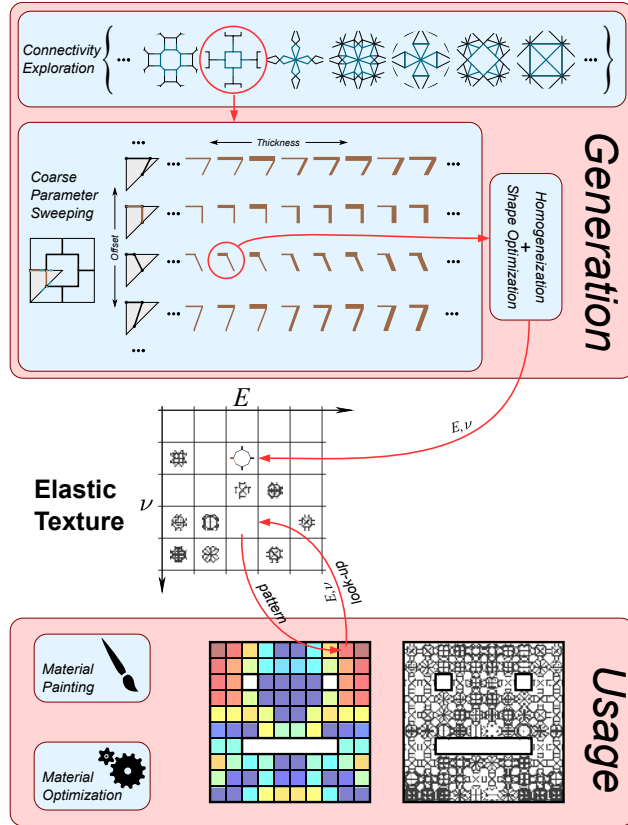


Figure 2: Overview of elastic texture generation and use.

Problem. The general problem we solve can be formulated as follows: for each tensor C from a given range of elasticity tensors, and a base isotropic material with Young's modulus E^b and Poisson's ratio ν^b , find a structure made out of the base material in a unit cubic cell, such that if the cell is infinitely tiled in space, the resulting homogeneous material has elasticity tensor C .

As discussed in the introduction, we aim to construct a family of patterns that are printable and tileable to enable creation of variable material properties.

Printability is heavily dependent on the choice of technology. We focus on printability criteria related to stereolithography, the most accurate 3D printing method available at this time, but our approach can be easily modified to handle other technologies.

As the printing process proceeds layer-by-layer, we assume that the structure is defined with respect to a fixed coordinate system X, Y , and Z aligned with the printer, with Z being vertical. The (idealized) printability criteria that we use are:

1. There are no enclosed voids.

2. For any point of the structure, the extent covered by the structure in the X, Y , and Z directions from the point are above a printability threshold d_{\min} .
3. Every point of the pattern is supported: for every XY slice, all connected components of the slice have at least one point connected to lower points in the structure by a segment contained in the structure. While this condition does not prevent long horizontal bars supported at single points, which can be difficult to print, we have found it sufficient in practice for pattern sizes up to 10mm and a d_{\min} of 0.3mm.

We also make our primary goal to generate periodic structures with isotropic homogenized properties. Such patterns have the elasticity tensor C defined by two parameters, Young's modulus E and Poisson's ratio ν , and its inverse, compliance tensor S , has the (Voigt notation) form

$$S = \frac{1}{E} \begin{pmatrix} 1 & -\nu & -\nu & 0 & 0 & 0 \\ -\nu & 1 & -\nu & 0 & 0 & 0 \\ -\nu & -\nu & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 2(1+\nu) & 0 & 0 \\ 0 & 0 & 0 & 0 & 2(1+\nu) & 0 \\ 0 & 0 & 0 & 0 & 0 & 2(1+\nu) \end{pmatrix}. \quad (1)$$

Expressing in terms of the shear modulus, $G = E/(2(1+\nu))$, the last three diagonal terms of S are simply $1/G$.

While for many tasks anisotropic materials are either sufficient or preferable, periodic structures with isotropic homogenized properties are easiest to use, as cell orientation is decoupled from material properties. In addition, once an isotropic starting point is obtained, it is easy to obtain a controlled anisotropic behavior.

Searching the space of all possible structures in a cell, even at a finite resolution, is an impossible task. Instead, we choose a space of structures with a limited but sufficiently large set of parameters, that can be optimized to achieve specific material properties.

Truss-like structures. We focus on truss-like structures (*patterns*) as shown in Figure 1, consisting of bars of different thicknesses connecting a set of nodes in the cell. Unlike real truss structures, the connections between bars are not pin joints, and flexural rigidity at the nodes plays a major role.

This particular space of structures is motivated by several considerations. First, the space is known to contain both very stiff and very weak patterns, providing a broad range of behaviors. Second, tileability and printability requirements yield specific geometric conditions, expressed mostly as constraints on the structure's geometry. For example, the requirement of no enclosed voids is automatically satisfied if the frame structure has no self-intersections; the bound on extents can be obtained by bounding the thickness from below; and the support condition is easily formulated as a constraint on node positions. These conditions are detailed in Section 4.

Symmetry considerations, as well as restrictions on the number and placement of nodes, yield a space of patterns parametrized by their set of edges connecting some subset of the 15 candidate nodes we define on a tetrahedron (their *topology*), thicknesses of these edges, and offsets of the nodes from their default positions. This space is still very large, and we explore it using both topology and geometry searches, described in Section 4. These rely fundamentally on the homogenization and shape optimization procedures described in Sections 5 and 6.

Resulting family. The search procedure's final result is shown in Figure 3; the six pattern topologies themselves are illustrated in Figure 1. The complex boundaries of the (E, ν) regions arise from the

multiple types of geometric constraints enforcing printability. Figure 4 shows some patterns with topology “ $(E1,E2)(E1,E4)(E2,E4)$.”

The family of topologies covers a large range of Young’s moduli, with a largest-to-smallest Young’s modulus ratio of 1800, and a sizable range of Poisson’s ratios, -0.16 to 0.48 .

We note that the range of negative Poisson’s ratios is somewhat limited, while on the high end, we are able to achieve ratios close to the theoretical maximum. This observation is consistent with [Sigmund 1995]: while it is relatively easy to obtain more extreme negative Poisson’s ratios for patterns with cubic symmetry but with shear modulus too low for isotropy, the isotropy requirements restrict the range. Printability constraints restrict it further.

Quite remarkably, four out of the six topologies can be transformed into each other by simple operations (single vertex splits, addition of cross-shaped supports connecting some nodes). The other two are also related to each other by a simple transform, but are not related to the first sequence.

We do not claim that the proposed family is in any sense optimal. It is most likely possible to extend the coverage or to cover the same domain with fewer topologies. However, the presented set is already quite useful for controlling material properties, as the examples of Section 7 show.

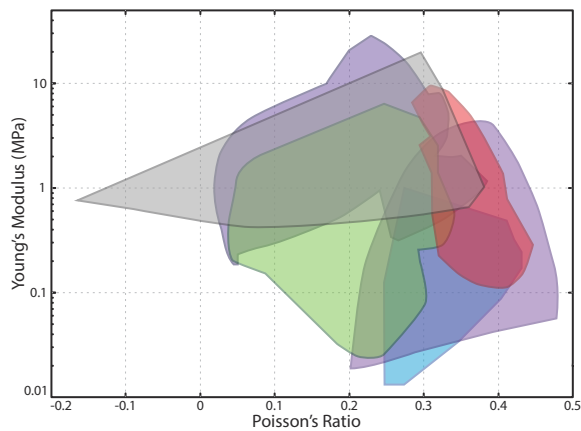


Figure 3: Region of the (E, ν) space covered by the selected set of patterns. Each topology’s coverage is shown in a different color.

Accuracy. We fabricated eight patterns with different homogenized Young’s moduli using the B9Creator SLA printer and tested their stiffness using the BOSE ElectroForce 3200 measurement system. The machine gradually compressed our samples in the Z direction between two compression plates and measured the displacement resulting from the applied force at each step. We used $6 \times 6 \times 2$ tilings of 5mm cells for this test.

We note that our measured force/displacement slopes are roughly proportional to the homogenized Young’s moduli (Figure 5a), implying that the measurements are consistent with some (unmeasured) base Young’s modulus. The curvature seen could be explained partly by friction in the compression testing setup (Figure 5b). Another significant source of error is the inaccuracy of our B9Creator, which tends to thicken thin geometries.

We used a lower-accuracy setup to measure Poisson’s ratio but still obtained reasonable agreement with homogenization (Figure 6). We compressed the microstructures in the Z direction between two lubricated metal blocks and manually measured the expansion/contraction in the X and Y directions. From these displace-

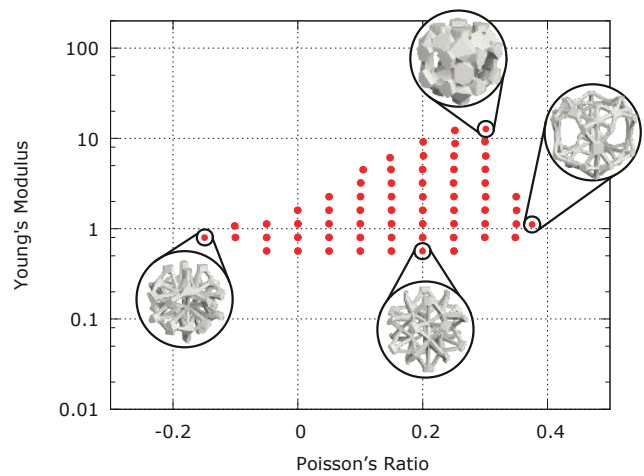


Figure 4: Samples of the (E, ν) space reached by patterns with topology “ $(E1,E2)(E1,E4)(E2,E4)$.”

ment measurements, we computed the X , Y , and Z strains and their ratios.




Figure 5: Compression test results for eight patterns with varying homogenized Young’s moduli ($6 \times 6 \times 2$ tiling of 5mm cells). (a) Slopes extracted from the measured force vs. displacement curves along with a best-fit line through the origin. (b) Moduli extracted from simulated compression tests, with and without modeling compression plate friction. Without friction, the simulated test agrees with homogenization perfectly, but friction introduces error.

We also validated the patterns’ isotropy by printing a block filled with a tiled pattern that was rotated by 45° around the Z axis and clipped (Figure 7). The measured effective Young’s moduli in the rotated orientations were in good agreement with the unrotated orientation: a compression test in the X , Y , and Z directions extracted effective Young’s moduli of 0.635MPa, 0.6293MPa, and 0.628MPa for the example shown.

Base Material. We used estimated base material properties of $E^b = 200\text{MPa}$ and $\nu^b = 0.35$ for all results. We estimated E^b using a three point bending test on rectangular bar samples, but the base Young’s modulus can also be estimated from the microstructure compression test.

Using a different base Young’s material would not qualitatively change our results, apart from making the patterns uniformly softer or stiffer. The homogenized Young’s moduli depend linearly on E^b , and the Poisson’s ratios are independent of E^b , so scaling the base modulus simply softens/stiffens every pattern by the same factor. In particular, for the displacement-based material optimization of Section 7, fabricating the solution with a different base Young’s




Figure 6: Poisson's ratios measured from $3 \times 3 \times 1$ printed tilings of 10mm cells vs. homogenized properties. The $\nu = -0.67$ sample, outside our family's range, violates isotropy and printability constraints (we added support structure manually for this experiment).




Figure 7: We extracted a 45° rotated rectangular block from a regularly tiled 10mm cell microstructure to test Young's modulus in non-axis aligned directions.

modulus maintains the same target deformation behavior (although the required force will change).

We note that our patterns are not very sensitive to moderate changes in the base Poisson's ratio. Changing from $\nu^b = 0.35$ to $\nu^b = 0.35 \pm 0.05$ results in a median relative change in Young's modulus of 0.683% (max: 1.88%) and a median absolute change in Poisson's ratio of 0.00229 (max: 0.0129) over all patterns. Since most additive fabrication materials fall within this range, we expect similar results for other printers and materials.

4 Search for Efficient Patterns

In this section, we describe the class of patterns that we consider and the main steps of the search method.

Ground class of patterns. The topology of patterns is defined by a set of edges connecting nodes in the cube cell. We generate the geometric variations by adding *offsets* to the node positions and by changing edge *thicknesses*.

Motivated by the isotropy requirement, we constrain our search to patterns with cube symmetries, which are guaranteed to have the same Young's moduli and Poisson's ratios in every axis-aligned direction. That is, the compliance tensor has the form (1) *except* the last three diagonal entries $1/G$ may not equal $2(1 + \nu)/E$. This yields an easy-to-check isotropy measure:

$$A = 2(1 + \nu)G/E, \quad (2)$$

which we use to identify isotropic patterns in our search. The symmetry will also dramatically reduce the space of pattern topologies to a tractable size after a few additional constraints are introduced. Note, however, that cube symmetry is not necessary for isotropy; other isotropic structures exist with, e.g., tet symmetry.

Consider the group of symmetries of a cube O_h , which includes reflections about three symmetry planes orthogonal to the X, Y, Z axes and the six planes orthogonal to the bisector of each pair of axes. By partitioning the cube according to these symmetry planes, we obtain 48 equal tetrahedra as in Figure 8a. O_h maps a single one of these tetrahedra to any other, so it is sufficient to define the nodes and edges of the pattern graph—as well as their offsets and thicknesses—on a single tetrahedron.




Figure 8: (a) The tetrahedral cube decomposition used to generate 3D patterns; (b) The 15 nodes defined on a tetrahedron together with their degrees of freedom.

We generate the different topological configurations by changing the connectivity between 15 nodes on a tetrahedron (see Figure 8a): vertex nodes $\{V_0, V_1, V_2, V_3\}$, edge nodes $\{E_0, E_1, E_2, E_3, E_4, E_5\}$, faces nodes $\{F_0, F_1, F_2, F_3\}$, and a single internal node, T_0 . Configurations are named by their graphs' edge sets (see Figure 1). Each node is constrained to stay on its respective simplex to preserve the topology, so vertex nodes are fixed, edge nodes have a single offset, and so on (Figure 8b).




Figure 9: Symmetry orbits are colored with yellow, red and green. Left: vertex symmetry orbits. Right: edge symmetry orbits.

Figure 9 shows an example topology colored by its node and edge orbits with respect to symmetry group O_h , and Figure 10 demonstrates the effects of the node offset and edge thickness parameters.

The space of possible connectivities, even after accounting for symmetries, is far too large to explore completely (on the order of 10^{32} configurations). We enforce the following constraints to reduce the space of patterns:

- **Connected:** the tiled pattern is a single connected component.
- **No coinciding edges:** no edge is contained within another. E.g., if graph edge (V_0, V_1) is chosen and E_0 is the midpoint node of the corresponding tet edge, graph edge (V_0, E_0) is forbidden since it overlaps the first for any offset.




Figure 10: The results of varying the thickness (top) and offset (bottom) parameters of a particular pattern topology.

- **No dangling edges:** every node has valence greater than 1.
- **Number of edges:** at most 3 graph edges per tetrahedron.
- **Max node valence:** node valences do not exceed 7.

Valences are computed on the graph after periodic tiling of the cube cell. The first two criteria reduce the space to 16221 topologies, and the remaining three to 1205 topologies.

Printability. For truss-like patterns, printability is affected by two main factors: the pattern graph structure and the edge thicknesses.

The first printability criterion can be defined on the nodes by considering their offset positions. We say that a node n_1 has supporting node n_2 , if these are connected by an edge and n_1 is strictly above n_2 . We say that n_1 and n_2 are at the same level if they have equal Z coordinates. A pattern is printable only if every connected set of nodes at one level has at least one supporting node.

Printability can be tested by a simple algorithm: we first mark as supported all nodes with a supporting node (considering periodicity). Then we propagate the front of supported nodes to neighbors at the same level. When this breadth first search terminates, the pattern is printable if and only if all nodes are marked as supported. The procedure is illustrated in Figure 12. We also note that this constraint can be expressed algebraically as a set of inequality constraints on the offset variables, which can be enforced by an optimization solver.




Figure 12: 2D examples of the printability detection algorithm. Vertices with supporting nodes are marked (green), then a breadth-first search extends the supported vertex front to horizontal neighbors. The remaining unmarked nodes are unsupported (red). Two cases are shown: unprintable (top) and printable (bottom).

Tileability. The tileability requirement means that all pattern topologies should belong to the same *family*, meaning topologies with the same set of nodes and edges appearing on the faces of the cube cell (Figure 11).

Searching the space of topologies. The goal of our search is to identify a family of pattern topologies that covers as much as possible of the (E, ν) space while satisfying the printability and tileability requirements.

The initial space consists of all pattern topologies satisfying the constraints on graph connectivity mentioned previously. We proceed in the following steps:

1. *Coarse geometry sweep.* Geometric variations are generated for each pattern topology by trying thicknesses of 0.3mm and 0.7mm and node offsets corresponding to barycentric coordinates of 0.2, 0.35, 0.5, 0.65, and 0.8 on the associated tetrahedron simplex. The resulting *printable* patterns are meshed, self-intersecting meshes are discarded, and the remaining patterns' effective elasticity parameters are computed using periodic homogenization (Section 5).
2. *Isotropy filtering.* We select a subset of the patterns closest to isotropic (we use a heuristic bound of $0.8 < A < 1.2$), which we consider promising candidates as starting points for optimizing pattern parameters to precisely match a range of isotropic elasticity tensors.
3. *Topological family selection.* At this point, we have a rough map of the area in (E, ν) space covered by our set of patterns. We obtain a rough estimate of each topology's coverage by taking the convex hull associated with its nearly isotropic geometric configurations. We manually pick the single family whose pattern topologies cover the largest region of (E, ν) based on these estimates.
4. *Selection of a minimal covering set of topologies.* For the selected family, we run a finer sweep of offsets and thicknesses, again filtering for printability, to compute a more precise estimate of the boundary of the (E, ν) domain that each pattern topology can cover. Among all topologies in the family, we selected 6 such that the union of their coverage areas contains most of the domain covered by the family.
5. *Lookup map construction.* Finally, using the shape optimization machinery of Section 6 and the initial nearly-isotropic points for each of the 6 topologies chosen, we optimize each patterns' parameters to reach a grid of isotropic elasticity tensors evenly spaced in $(\log(E), \nu)$.

Our procedure is similar to [Sigmund 1995], but with several key differences. First, [Sigmund 1995] uses a simplified truss model, whereas our method directly homogenizes and optimizes the printable geometry. Second, by using a full topology (including all possible edges between nodes in a ground structure) and permitting zero edge thickness in optimization, the work avoids the topol-




Figure 11: Two pattern topologies from each of three different families, shown with the families' interfaces (nodes on the cube cell faces).

ogy enumeration stage. As a side-effect, it cannot accommodate the lower thickness bounds or support criteria needed for printability. While a mixed-integer formulation like [Mela and Koski 2013] could allow enforcement of d_{min} by introducing separate binary variables to disable members, this would involve a difficult mixed-integer nonlinear programming problem in our tensor-fitting setting. Finally, by introducing offset variables, our novel shape optimization approach enables much finer control of the elasticity tensors as the design is not limited to the discrete node positions of a ground structure.

5 From Patterns to Material Properties

Our goal is to find the homogenized elasticity tensor C^H , describing the effective properties of the microstructure when it is fabricated at a small enough scale and periodically repeated to fill the space. This elasticity tensor is almost never the spatial average of elasticity tensors (for example, if a cell is almost completely filled with material but is disconnected from other cells, it has zero Young's modulus). We first define more precisely what a homogenized elasticity tensor is and then explain how to compute it.

Defining the homogenized tensor. Consider heterogeneous object Ω^ϵ filled with a periodic microstructure, as shown schematically in Figure 13. Parameter ϵ determines the size of cell Y relative to the object Ω^ϵ and permits asymptotic analysis as $\epsilon \rightarrow 0$.




Figure 13: (Schematic) Periodic tiling of a domain Ω with base cell Y having geometry ω and length scale ϵ .

The elastic response of an object under macroscopic external load \mathbf{f} is governed by the linear elastostatic equation

$$-\nabla \cdot [C : \varepsilon(\mathbf{u}^\epsilon)] = \mathbf{f} \text{ in } \Omega^\epsilon, \quad (3)$$

complemented by appropriate boundary conditions. Here, C is the periodically varying elasticity tensor, and $C : \varepsilon$ denotes its double contraction with strain ($C_{ijkl}\varepsilon_{kl}$) to compute stress. Considering a sequence of problems indexed by ϵ and letting \mathbf{u} denote the limit of \mathbf{u}^ϵ as $\epsilon \rightarrow 0$, the homogenized elasticity tensor C^H is defined as the tensor that satisfies

$$-\nabla \cdot [C^H : \varepsilon(\mathbf{u})] = \mathbf{f} \text{ in } \Omega, \quad (4)$$

with same boundary conditions. Vectors \mathbf{u} and \mathbf{u}^ϵ denote the displacement, and $\varepsilon(\mathbf{u}) := \frac{1}{2}(\nabla \mathbf{u} + (\nabla \mathbf{u})^T)$ is the Cauchy strain tensor.

Expressions for the homogenized tensor. The standard derivation of the homogenized elasticity tensor based on a two-scale asymptotic expansion is provided in the additional material. Here we give an intuitive motivation for the periodic homogenization equations.

\mathbf{u}^ϵ has a high frequency periodic component that is averaged out to obtain \mathbf{u} as period $\epsilon \rightarrow 0$. So $\mathbf{u}(\mathbf{x})$ can be thought of as the average displacement over the infinitesimal base cell Y at point \mathbf{x} . Likewise, $\varepsilon(\mathbf{u})$ is the average strain in the cell. For the object to be in equilibrium, (4) should represent an average force balance over the cell, meaning $C^H : \varepsilon(\mathbf{u})$ should be the average stress tensor. That gives the following intuitive interpretation of C^H : it maps the average strain applied at a point to the average stress resulting within the microstructure geometry.

Thus, applying C^H to $\varepsilon(\mathbf{u})$ is equivalent to simulating the microstructure's deformation under that average strain and averaging its stress. We formulate this simulation inside a single base cell Y by assuming that the displacement consists of a linear term (with constant strain $\varepsilon(\mathbf{u})$) plus a Y -periodic "microscopic fluctuation" term, \mathbf{w} (with zero average strain by periodicity). This assumption of periodicity is reasonable because, by the translational symmetry of an infinite tiling, every cell deforms identically. Now we simply solve for fluctuation \mathbf{w} putting the microstructure in equilibrium:

$$-\nabla \cdot (C(\mathbf{y}) : [\varepsilon(\mathbf{w}(\mathbf{y})) + \varepsilon(\mathbf{u})]) = 0 \text{ in } Y,$$

where \mathbf{y} is the microscopic variable (the coordinate in Y). Then the average stress is $C^H : \varepsilon(\mathbf{u}) = \frac{1}{|Y|} \int_Y C : [\varepsilon(\mathbf{w}) + \varepsilon(\mathbf{u})] \, d\mathbf{y}$.

We can extract the components of C^H by applying it to the 6 canonical symmetric rank 2 basis tensors, $e^{kl} := \frac{1}{2}(\mathbf{e}_k \otimes \mathbf{e}_l + \mathbf{e}_l \otimes \mathbf{e}_k)$. Each application amounts to solving the *cell problem*:

$$-\nabla \cdot (C^{\text{base}} : [\varepsilon(\mathbf{w}^{kl}) + e^{kl}]) = 0 \text{ in } \omega, \quad (5a)$$

$$[C^{\text{base}} : \varepsilon(\mathbf{w}^{kl})] \hat{\mathbf{n}} = -[C^{\text{base}} : e^{kl}] \hat{\mathbf{n}} \text{ on } \partial\omega \setminus \partial Y, \quad (5b)$$

$$\mathbf{w}^{kl}(\mathbf{y}) \text{ } Y\text{-periodic}, \quad (5c)$$

$$\int_\omega \mathbf{w}^{kl}(\mathbf{y}) \, d\mathbf{y} = \mathbf{0}, \quad (5d)$$

where we rephrased the microscopic force balance as a PDE over ω , since for a structure printed with base material properties C^{base} ,

$$C(\mathbf{y}) = \begin{cases} C^{\text{base}} & \text{if } \mathbf{y} \in \omega, \\ 0 & \text{otherwise.} \end{cases} \quad (6)$$

The last constraint in (5) eliminates the rigid translation degrees of freedom that still remain after enforcing Y -periodicity.

The homogenized elasticity tensor components are finally just the average over Y of the stress components corresponding to e^{kl} :

$$C_{ijkl}^H = \frac{1}{|Y|} \int_\omega C_{ijpq}^{\text{base}} [\varepsilon(\mathbf{w}^{kl}) + e^{kl}]_{pq} \, d\mathbf{y}. \quad (7)$$

It is worth noting that C^H does not depend at all on the macroscopic details (shape Ω , force term \mathbf{f} , or boundary conditions).

FEM implementation. The cell problems (5) are solved numerically by a quadratic tetrahedral FEM discretization of ω . The piecewise linear integrand in (7) is integrated with exact quadrature.

Given the wire network of the microstructure, defining its topology (Section 4), a volume mesh is generated following the STRUT algorithm given in [Hart 2008]: (i) A polygon is created around both ends of each segment. (ii) For each vertex, the convex hull of the nearby polygons and the vertex is constructed and the polygons themselves are removed from the hull. (iii) For each edge, the convex hull of its two polygons is constructed, and again the two polygons are removed from the hull. A tetrahedral volume mesh is finally created from the resulting closed surface.

The linear elasticity solver must support periodic boundary conditions, which requires the tetrahedral mesh to have an identical tessellation on the opposite periodic cell faces.

Convergence rate. Remarkably, we have observed that the homogenized coefficients remain accurate when the microstructure varies across cells (with few or no repetitions). In our experiments, the deformation behavior of even very coarse tilings closely matches the homogenized behavior (Figure 14), and we expect similar agreement in general for smoothly varying loads.




Figure 14: Deformation of an object with varying material properties per voxel, and the same object with the material in each voxel replaced with the corresponding pattern. The deformed objects are colored by max stress.

6 Optimizing Pattern Parameters

An essential step for creating a map of elastic textures is optimizing a pattern with fixed topology to match particular elasticity parameters. This is achieved using shape optimization with respect to the pattern parameters.

The optimization problem. Our goal is to minimize a functional, $J(\omega)$, measuring the difference between the homogenized elastic properties of the pattern and a target elasticity tensor C^* . We choose an objective that is suitable for designing material distributions with large deformations under moderate forces. The distance of compliance tensors, $S^H - S^*$, as opposed to elasticity tensors, is the better choice, since the strain for a constant stress is directly proportional to S^H , not C^H . In fact, minimizing the Frobenius norm of $S^H - S^*$ can be interpreted as a multi-objective least squares optimization to fit the displacements of two cubes—one filled with C^H and the other with C^* —under a set of axis-aligned stretching and shearing loads.

We choose this Frobenius norm as our functional:

$$J(\omega) = \frac{1}{2} \|S^H(\omega) - S^*\|_F^2, \quad (8)$$

which we optimize by varying the microstructure shape, ω .

In our case, the microstructure boundary $\partial\omega$ is determined by a small number of parameters \mathbf{p} , consisting of wire mesh node offsets and thicknesses. While the number of parameters is small, we still expect multiple solutions for minimizing $J(\omega)$ with respect to these parameters. A simple regularization term (staying close to the initial point of optimization) picks a unique solution. We note that instead some quantity of importance (e.g., weight) can be optimized as it is typically done (cf., [Sigmund 1995]), adding another nonlinear term to the functional.

The derivative of the boundary $\partial\omega$ with respect to a parameter p_α is a vector field $\mathbf{v}_{p_\alpha}(\mathbf{y})$ defined at points \mathbf{y} of $\partial\omega$, with $\mathbf{v}_{p_\alpha}(\mathbf{y})$ being the velocity of \mathbf{y} if parameter p_α changes at unit speed.

Using parameters \mathbf{p} as variables, the minimization problem can be written as

$$\operatorname{argmin}_{\text{admissible } \mathbf{p}} J(\mathbf{p}) \text{ where } J(\mathbf{p}) = \frac{1}{2} \|S^H(\mathbf{p}) - S^*\|_F^2. \quad (9)$$

The admissibility of parameters is determined by geometric intersection constraints and printability constraints.

The derivative of the objective function with respect to p_α can be obtained from \mathbf{v}_{p_α} using the chain rule:

$$\frac{\partial J}{\partial p_\alpha} = [S^H - S^*] : \frac{\partial S^H}{\partial p_\alpha} = [S^H - S^*] : dS^H[\mathbf{v}_{p_\alpha}], \quad (10)$$

where $dS^H[\mathbf{v}_{p_\alpha}]$ is the *shape derivative* of S^H applied to \mathbf{v}_{p_α} .

Shape derivative of elasticity tensor The derivative of the microstructure's homogenized elasticity tensor in the direction of shape perturbation \mathbf{v} is defined as the Gâteaux derivative,

$$dC^H[\mathbf{v}] := \lim_{t \rightarrow 0} \frac{C^H(\omega(t, \mathbf{v})) - C^H(\omega)}{t}, \quad (11)$$

where $\omega(t, \mathbf{v}) := \{\mathbf{x} + t\mathbf{v} : \mathbf{x} \in \omega\}$. As shown in the additional material, the homogenized elasticity tensor, (7), can be rewritten in an energy-like form,

$$C_{ijkl}^H = \frac{1}{|Y|} \int_{\omega} (e^{ij} + \varepsilon(\mathbf{w}^{ij})) : C^{\text{base}} : (e^{kl} + \varepsilon(\mathbf{w}^{kl})) \, d\mathbf{y}, \quad (12)$$

which is shown to have shape derivative:

$$dC_{ijkl}^H[\mathbf{v}] = \frac{1}{|Y|} \int_{\partial\omega} [(e^{ij} + \varepsilon(\mathbf{w}^{ij})) : C^{\text{base}} : (e^{kl} + \varepsilon(\mathbf{w}^{kl}))](\mathbf{v} \cdot \hat{\mathbf{n}}) \, dA(\mathbf{y}). \quad (13)$$

Shape derivative of compliance tensor. The compliance tensor is the symmetric rank 4 inverse of elasticity tensor, i.e. $S_{ijkl} = C_{klmn}^{-1} = \frac{1}{2}(\delta_{im}\delta_{jn} + \delta_{in}\delta_{jm})$. Differentiating and solving for dS^H :

$$dS^H[\mathbf{v}] = -S^H : dC^H[\mathbf{v}] : S^H. \quad (14)$$

Combining the results from (10), (13), and (14), one can compute $\frac{\partial J}{\partial p_\alpha}$; the shape derivative and an example velocity field \mathbf{v}_{p_α} are shown in Figure 15.

Numerical computation. The integrand in (13) is cubic over each boundary element ($\varepsilon(\mathbf{w}^{ij})$ and $\mathbf{v} \cdot \hat{\mathbf{n}}$ are both linear), and we use




Figure 15: Left: a shape derivative, visualized as a steepest ascent normal velocity field for objective (8). Right: the shape velocity induced by one of the pattern's thickness parameters.

quadrature that evaluates the surface integral exactly. To evaluate the gradient for a given shape we need (a) to mesh the shape (we use the TetGen package [Si 2010]); (b) solve 6 periodic elasticity problems to obtain \mathbf{w}^{ij} , as for homogenization. The cost of a single gradient evaluation (roughly 4.75s on a single core of an Intel Xeon E-2690 v2) is dominated by the cost of periodic meshing and the elasticity solves, which take roughly equal time. We use the Ceres solver [Agarwal et al.]'s Levenberg-Marquardt implementation to minimize the objective; the convergence of the solver is quite fast (Figure 16). Typical effects of optimization are shown in Figure 17.

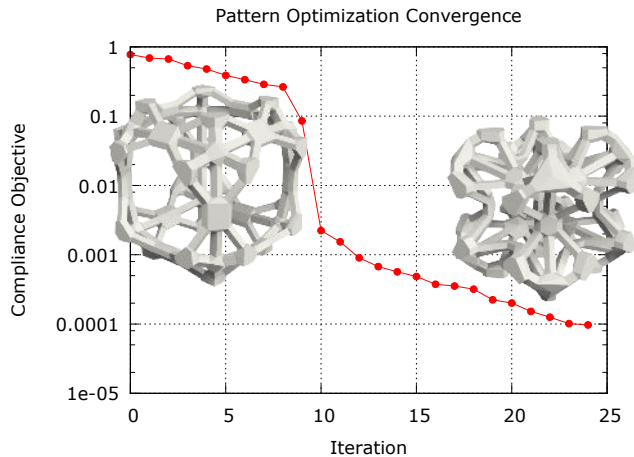


Figure 16: Convergence of a shape optimization on pattern “(E1,E2)(E1,E4)(E2,E4).” Left: optimization starting point. Right: optimized shape.

7 Applications

While our primary focus is on the design of our pattern family and the exploration of its coverage, we demonstrate the application of our elastic textures in two settings: painted material properties and specified deformation behavior. All printing was done using a B9Creator printer at 50 micron resolution with Cherry resin.

Overall workflow. The result of the preceding sections is a lookup map that, for a given (E, ν) , produces an isotropic microstructure with nearby parameters. We assume that we are given a coarse volume mesh filled with identical cube cells.

First, we assign a pair (E_i, ν_i) to each cell i , either directly or via material optimization as described below. For each cell, we retrieve

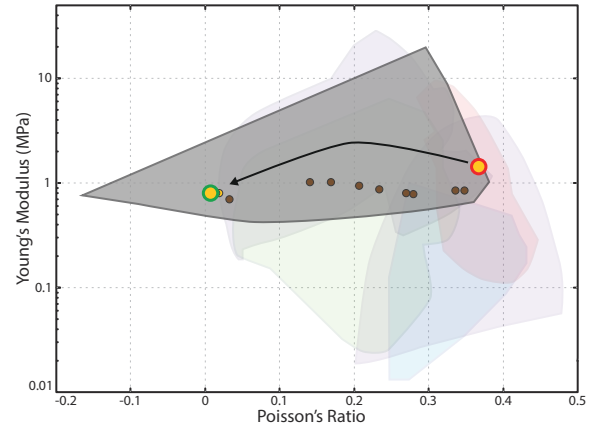


Figure 17: The path in (E, ν) space traversed by the optimization of pattern “(E1,E2)(E1,E4)(E2,E4)” shown in Figure 16. The brown points are intermediate anisotropic microstructures.

a corresponding pattern (topology id, thicknesses, and offsets) from the lookup table. The microstructures in adjacent cells are stitched together by averaging the offsets of each pair of shared face nodes so that they coincide. This might raise the lower node of the pair above some node it supports, n_s , violating printability, but printability can be restored by lowering the pair to n_s 's height.

After this step, the resulting connected wire mesh is inflated with retrieved bar thicknesses, using the process described at the end of Section 5. This results in a fine mesh that can be printed.

Material painting. The simplest approach to specifying the material properties (E, ν) is to paint them on a voxel grid. We have created an editor enabling us to paint these layer by layer. The results of fabricating several structures of this type are shown in Figure 18.

Material optimization. Manually defining material properties to achieve desired behavior may be difficult, and a more systematic approach is to solve for them. For example, consider the following problem: for given applied displacement conditions on some part of the object's boundary, we would like to get some target deformation—e.g., if a bar is compressed along the Z axis, it twists in the X-Y plane—by varying material properties. In other words, we want to find a spatially varying tensor $C_p(\mathbf{x})$, parametrized by a vector of per-cell isotropic parameters p , such that the following system has a solution:

$$\begin{aligned} -\nabla \cdot (C_p : \varepsilon(\mathbf{u})) &= 0 && \text{in } \Omega \\ \mathbf{u} &= \mathbf{u}_{in} && \text{on } \Gamma_{in} \\ \mathbf{u} &= \mathbf{u}_{trg} && \text{on } \Gamma_{trg} \\ \sigma(\mathbf{u})\hat{\mathbf{n}} &= 0 && \text{on } \Gamma_{trg} \end{aligned} \quad (15)$$

where \mathbf{u}_{in} are the applied displacements on compressed area Γ_{in} , and \mathbf{u}_{trg} are the target displacements of the surface Γ_{trg} on which no forces are applied, as indicated by the last equation. In our implementation, these target/applied conditions can be specified on a per-component basis to set up, e.g., the twisting bar example.

In general, such problems are solved using PDE constrained optimization, requiring solving an adjoint problem at each iteration. However, we found that the following local-global iteration, inspired by related “as-rigid-as-possible” (ARAP) optimization techniques in geometry, works remarkably well. For a fixed $C_p(\mathbf{x})$, we call \mathbf{u}_D the solution of the “Dirichlet problem,” with the condition of zero tractions on Γ_{trg} removed. \mathbf{u}_N is the solution of the “Neu-




Figure 18: Examples of objects with painted material properties. All are fabricated with 5mm cells.

mann problem,” in which the traction condition on Γ_{trg} is retained, but the Dirichlet condition $\mathbf{u} = \mathbf{u}_{trg}$ is removed.

We initialize the elastic tensor $C_p(\mathbf{x})$ to a constant. The iteration consists of two steps:

1. Solve the Dirichlet and Neumann problems with the current elasticity tensor, to obtain \mathbf{u}_D and \mathbf{u}_N .
2. Update C_p , minimizing the following energy:

$$\min_p \int_{\Omega} \|\varepsilon(\mathbf{u}_D) - C_p^{-1} : \sigma(\mathbf{u}_N)\|_F^2 dV \quad (16)$$

This energy can be minimized per-cell for a truly local-global method; however, in practice we find it desirable to regularize p with a Laplacian term, which requires that the “local” step be replaced by a global, but still quadratic optimization. The convergence of this method is shown in Figure 19.

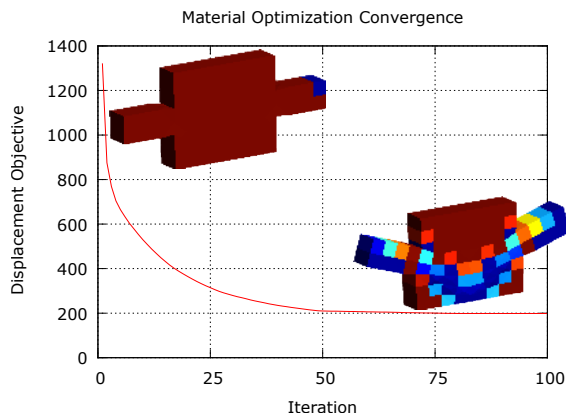


Figure 19: Convergence of material optimization.

We have used a number of simple voxelized shapes and created a variety of deformation behaviors shown in Figure 20. Finally, we have also generated a set of anisotropic samples, with controlled anisotropy ratio, one of which is shown in Figure 21.

8 Conclusions

We have presented a family of tileable and printable patterns that can be used to approximate varying isotropic material properties. The family has proved useful on a number of simple shape optimization examples: remarkably, all examples in Section 7 worked as predicted by simulation without requiring much tuning.

Limitations. There are several limitations of our pattern family. First, some parts of the (E, ν) space are poorly covered. While it




Figure 21: Compression of an anisotropic sample along the X, Y, and Z directions.

is difficult to predict which part of space is theoretically reachable, we conjecture that the space may be significantly broadened. All our simulations and constructions work in the linear regime, not taking into account, e.g., the potential for pattern buckling or other damage. Fortunately, isotropy is correlated with sufficiently high shear modulus, which makes the patterns less prone to buckling. Nevertheless, including this and other nonlinear effects in pattern design is important.

For practical use, it is difficult to restrict the tessellations of objects to equal sized cubes (though one can construct cut cells covered with relatively soft skin). A desirable solution would be to allow patterns to distort to fill arbitrary reasonably well-shaped hex cells.

Acknowledgements

We thank Professor Bob Kohn (NYU Courant Institute) for his many helpful insights. We also thank John Ricci and Yu Zhang (NYU College of Dentistry) and Oran Kennedy and Matin Lendhey (NYU School of Medicine) for their help with material testing. This work was partially supported by NSF award DMS-1436591.

References

- AGARWAL, S., MIERLE, K., AND OTHERS. Ceres solver. <http://ceres-solver.org>.
- ALLAIRE, G. 2002. *Shape optimization by the homogenization method*, vol. 146. Springer.
- ANDREASSEN, E., LAZAROV, B. S., AND SIGMUND, O. 2014. Design of manufacturable 3D extremal elastic microstructure. *Mechanics of Materials* 69, 1, 1–10.
- AVELLANEDA, M. 1987. Optimal bounds and microgeometries for elastic two-phase composites. *SIAM Journal on Applied Mathematics* 47, 6, 1216–1228.
- BENDSØE, M. P., AND SIGMUND, O. 2003. *Topology optimization: theory, methods and applications*. Springer.

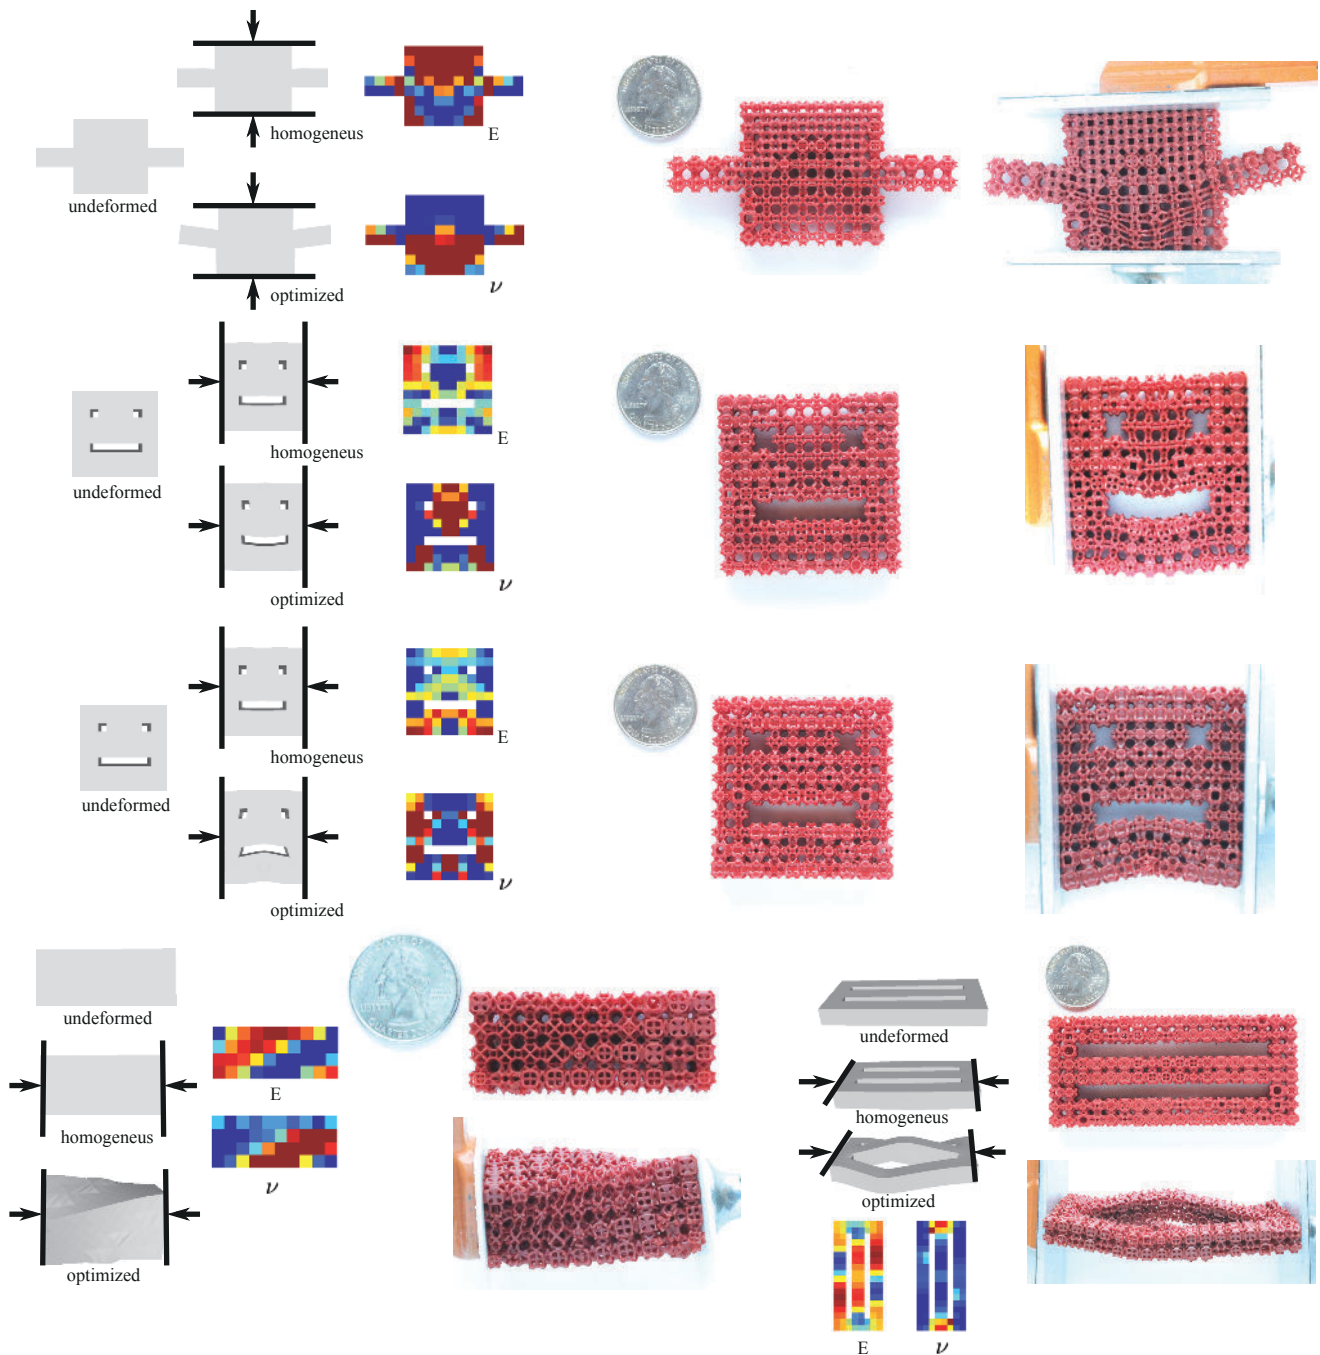


Figure 20: Examples of objects with optimized material properties. All are fabricated with 5mm cells.

BENDSØE, M. P. 1989. Optimal shape design as a material distribution problem. *Structural optimization* 1, 4, 193–202.

BICKEL, B., BÄCHER, M., OTADUY, M. A., LEE, H. R., PFISTER, H., GROSS, M., AND MATUSIK, W. 2010. Design and fabrication of materials with desired deformation behavior. *ACM Trans. Graph.* 29, 4 (July), 63:1–63:10.

BÜCKMANN, T., STENGER, N., KADIC, M., KASCHKE, J., FRÖLICH, A., KENNERKNECHT, T., EBERL, C., THIEL, M., AND WEGENER, M. 2012. Tailored 3d mechanical metamaterials made by dip-in direct-laser-writing optical lithography. *Advanced Materials* 24, 20, 2710–2714.

CADMAN, J. E., ZHOU, S., CHEN, Y., AND LI, Q. 2013. On design of multi-functional microstructural materials. *Journal of Materials Science* 48, 1, 51–66.

CHEN, D., LEVIN, D. I., DIDYK, P., SITHI-AMORN, P., AND MATUSIK, W. 2013. Spec2fab: a reducer-tuner model for translating specifications to 3d prints. *ACM Transactions on Graphics (TOG)* 32, 4, 135.

CHERKAEV, A. 2000. *Variational methods for structural optimization*, vol. 140. Springer.

- CHU, J., ENGELBRECHT, S., GRAF, G., AND ROSEN, D. W. 2010. A comparison of synthesis methods for cellular structures with application to additive manufacturing. *Rapid Prototyping Journal* 16, 4, 275–283.
- CIGNONI, P., PIETRONI, N., MALOMO, L., AND SCOPIGNO, R. 2014. Field-aligned mesh joinery. *ACM Trans. Graph.* 33, 1 (Feb.), 11:1–11:12.
- CIORANESCU, D., AND DONATO, P. 1999. *An introduction to homogenization*. Oxford University Press.
- GRABOVSKY, Y., AND KOHN, R. V. 1995. Microstructures minimizing the energy of a two phase elastic composite in two space dimensions. II: the Vigdergauz microstructure. *Journal of the Mechanics and Physics of Solids* 43, 6, 949–972.
- GREAVES, G. N., GREER, A. L., LAKES, R. S., AND ROUXEL, T. 2011. Poisson's ratio and modern materials. *Nature Materials* 10, 11, 823–837.
- GUEST, J. K., AND PRÉVOST, J. H. 2006. Optimizing multifunctional materials: Design of microstructures for maximized stiffness and fluid permeability. *International Journal of Solids and Structures* 43, 2223, 7028 – 7047.
- HART, G. W. 2008. Sculptural forms from hyperbolic tessellations. In *Shape Modeling and Applications, 2008. SMI 2008. IEEE International Conference on*, IEEE, 155–161.
- HILDEBRAND, K., BICKEL, B., AND ALEXA, M. 2012. Crdbrd: Shape fabrication by sliding planar slices. *Comp. Graph. Forum* 31, 2pt3 (May), 583–592.
- HILLER, J., AND LIPSON, H. 2009. Design and analysis of digital materials for physical 3d voxel printing. *Rapid Prototyping Journal* 15, 2, 137–149.
- HOLLISTER, S. J. 2005. Porous scaffold design for tissue engineering. *Nature Materials* 4, 7, 518–524.
- KANG, H. S. 2010. *Hierarchical design and simulation of tissue engineering scaffold mechanical, mass transport, and degradation properties*. PhD thesis, The University of Michigan.
- KHAREVYCH, L., MULLEN, P., OWHADI, H., AND DESBRUN, M. 2009. Numerical coarsening of inhomogeneous elastic materials. *ACM Trans. Graph.* 28, 3 (July), 51:1–51:8.
- LIN, C. Y., KIKUCHI, N., AND HOLLISTER, S. J. 2004. A novel method for biomaterial scaffold internal architecture design to match bone elastic properties with desired porosity. *Journal of Biomechanics* 37, 5, 623–636.
- LIN, C.-Y., HSIAO, C.-C., CHEN, P.-Q., AND HOLLISTER, S. J. 2004. Interbody fusion cage design using integrated global layout and local microstructure topology optimization. *Spine* 29, 16, 1747–1754. PMID: 15303018.
- LIU, L., JAMES, R. D., AND LEO, P. H. 2007. Periodic inclusion-matrix microstructures with constant field inclusions. *Metallurgical and Materials Transactions A* 38, 4, 781–787.
- MELA, K., AND KOSKI, J. 2013. Distributed loads in truss topology optimization. In *Proceedings of the 10th world congress on structural and multidisciplinary optimization, Orlando*.
- MILTON, G. W. 2002. *The theory of composites*. Cambridge University Press.
- MIRONOV, V., VISCONTI, R. P., KASYANOV, V., FORGACS, G., DRAKE, C. J., AND MARKWALD, R. R. 2009. Organ printing: tissue spheroids as building blocks. *Biomaterials* 30, 12, 2164–2174.
- MITANI, J., AND SUZUKI, H. 2004. Making papercraft toys from meshes using strip-based approximate unfolding. In *ACM SIGGRAPH 2004 Papers*, ACM, New York, NY, USA, SIGGRAPH '04, ACM, 259–263.
- MORI, Y., AND IGARASHI, T. 2007. Plushie: An interactive design system for plush toys. In *ACM SIGGRAPH 2007 Papers*, ACM, New York, NY, USA, SIGGRAPH '07, ACM.
- NAKASONE, P., AND SILVA, E. 2010. Dynamic design of piezoelectric laminated sensors and actuators using topology optimization. *Journal of Intelligent Material Systems and Structures* 21, 16, 1627–1652.
- RADMAN, A., HUANG, X., AND XIE, Y. 2013. Topological optimization for the design of microstructures of isotropic cellular materials. *Engineering Optimization* 45, 11, 1331–1348.
- SCHWARTZBURG, Y., AND PAULY, M. 2013. Fabrication-aware design with intersecting planar pieces. *Comput. Graph. Forum* 32, 2, 317–326.
- SCHWARTZBURG, Y., TESTUZ, R., TAGLIASACCHI, A., AND PAULY, M. 2014. High-contrast computational caustic design. *ACM Trans. Graph.* 33, 4 (July), 74:1–74:11.
- SCHWERDTFEGER, J., WEIN, F., LEUGERING, G., SINGER, R. F., KRNER, C., STINGL, M., AND SCHURY, F. 2011. Design of auxetic structures via mathematical optimization. *Advanced Materials* 23, 22, 2650–2654.
- SI, H. 2010. A quality tetrahedral mesh generator and a 3D Delaunay triangulator. URL <http://tetgen.berlios.de>.
- SIGMUND, O. 1995. Tailoring materials with prescribed elastic properties. *Mechanics of Materials* 20, 4, 351–368.
- SKOURAS, M., THOMASZEWSKI, B., COROS, S., BICKEL, B., AND GROSS, M. 2013. Computational design of actuated deformable characters. *ACM Transactions on Graphics (TOG)* 32, 4, 82.
- TORQUATO, S., AND DONEV, A. 2004. Minimal surfaces and multifunctionality. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* 460, 2047, 1849–1856.
- TORQUATO, S., HYUN, S., AND DONEV, A. 2002. Multifunctional composites: optimizing microstructures for simultaneous transport of heat and electricity. *Physical review letters* 89, 26, 266601.
- TORQUATO, S., HYUN, S., AND DONEV, A. 2003. Optimal design of manufacturable three-dimensional composites with multifunctional characteristics. *Journal of Applied Physics* 94, 9, 5748–5755.
- TORQUATO, S. 2002. *Random heterogeneous materials: microstructure and macroscopic properties*, vol. 16. Springer.
- VIDIČE, K., WANG, S.-P., RAGAN-KELLEY, J., AND MATUSIK, W. 2013. Openfab: A programmable pipeline for multi-material fabrication. *ACM Transactions on Graphics (TOG)* 32, 4, 136.
- WEYRICH, T., PEERS, P., MATUSIK, W., AND RUSINKIEWICZ, S. 2009. Fabricating microgeometry for custom surface reflectance. *ACM Trans. on Graphics (Proc. SIGGRAPH)* 28, 3, 32:1–32:6.