



Doctoral Thesis

## Computational Design of Structured Materials

**Author(s):**

Schumacher, Christian

**Publication Date:**

2018

**Permanent Link:**

<https://doi.org/10.3929/ethz-b-000341779> →

**Rights / License:**

[In Copyright - Non-Commercial Use Permitted](#) →

This page was generated automatically upon download from the [ETH Zurich Research Collection](#). For more information please consult the [Terms of use](#).

Diss. ETH No. 25445

# Computational Design of Structured Materials

A thesis submitted to attain the degree of  
**Doctor of Sciences of ETH Zurich**  
(Dr. sc. ETH Zurich)

presented by

**Christian Schumacher**

Master of Science ETH in Computer Science

born on September 3, 1987

citizen of Luzern (LU), Switzerland

accepted on the recommendation of

**Prof. Dr. Markus Gross**, examiner

**Prof. Dr. Chiara Daraio**, co-examiner

**Prof. Dr. Bernhard Thomaszewski**, co-examiner

2018



# Abstract

Advances in manufacturing and processing techniques allow us to create objects with ever increasing complexity. This complexity can be used to create not only aesthetic forms, but also *structured materials*, whose geometries determine their deformation behavior. Systematic control over the elastic properties is possible with the right structures, but the relationship between structure and elastic properties is nontrivial.

This highlights the need for new computational design approaches that can leverage the geometric complexity provided by advanced fabrication techniques, either by providing feedback about the deformation behavior throughout the design process, or by directly translating high-level functional targets into the appropriate material structure. This thesis introduces a novel set of techniques to characterize and create structured materials with various functional targets.

We first introduce a data-driven microstructure synthesis approach. After creating a database of small-scale structures that show a clearly defined macroscopic deformation behavior when tiled, we use an interpolation method to create new structures from this database and design objects with spatially varying elastic properties.

We then propose a design method to create surfaces with decorative cutouts that integrates aesthetics and stability into a single design process. We combine a discrete element texture approach with a topology optimization method to automatically optimize the distribution of cutouts, creating a synergy of structure and function.

Finally, we develop an approach to characterize the mechanical properties of tessellation-based networks of rods. We create a compact representation of the elastic behavior of such networks and explore the space of a specific type of tessellation, isohedral tilings, to show that we can create a broad range of mechanical behaviors from combinations of simple shapes. We present the resulting mechanical characterization using an intuitive visual representation that lends itself to an efficient exploration of the underlying space of structures.



# Zusammenfassung

Technische Fortschritte in der Herstellung und Verarbeitung erlauben es uns, Objekte mit stetig steigender Komplexität herzustellen. Diese Komplexität kann nicht nur für die Gestaltung von ästhetischen Formen verwendet werden, sondern auch zur Erzeugung von *strukturierten Materialien*, welche durch ihre Geometrie das Deformationsverhalten des Materials bestimmen. Mit der richtigen Struktur ist eine gezielte Kontrolle der elastischen Eigenschaften möglich. Dieses Verhältnis zwischen Struktur und Eigenschaften ist jedoch nicht offensichtlich.

Dies zeigt den Bedarf nach neuen computergestützten Entwurfsansätzen auf, welche die geometrische Komplexität dieser neuen Herstellungsverfahren ausnützen können, entweder durch das Bereitstellen von Rückmeldungen über das Deformationsverhalten während des Entwurfsprozesses, oder direkt durch das Umwandeln von abstrakten funktionellen Vorgaben zu den entsprechenden Materialstrukturen. Diese Doktorarbeit stellt eine Reihe neuer Methoden zur Charakterisierung und Erzeugung von strukturierten Materialien mit verschiedenen funktionellen Zielen vor.

Wir führen zuerst einen datenbankgestützten Ansatz zur Synthese von Mikrostrukturen ein. Nach dem Erstellen einer Datenbank von kleinräumigen Strukturen, welche ein klar definiertes makroskopisches Verhalten aufzeigen, wenn man sie aneinanderreicht, benutzen wir ein Interpolationsverfahren, um mit Hilfe dieser Datenbank neue Strukturen zu erzeugen und Objekte mit räumlich-ändernden elastischen Eigenschaften auszufüllen.

Wir stellen anschliessend eine Entwurfsmethode zur Erstellung von Oberflächen mit dekorativen Ausschnitten vor, welche Ästhetik und Stabilität in einem einzelnen Prozess vereint. Wir kombinieren den Ansatz der Texturen aus diskreten Elementen mit einer Methode der Topologieoptimierung zur automatischen Optimierung der Verteilung der Ausschnitte und kreieren damit eine Synergie von Struktur und Funktion.

Schliesslich entwickeln wir einen Ansatz zur Charakterisierung der mechanischen Eigenschaften von Stab-Netzwerken, welche auf Parkettierungen basieren. Wir erzeugen eine kompakte Repräsentation des elastischen Verhaltens solcher Netzwerke und untersuchen den Raum einer spezifischen Parkettierung, der isohedralen Parkettierung, um zu zeigen, dass wir durch Kombinationen von einfa-

chen Formen eine breite Auswahl an mechanischen Verhalten erzeugen können. Wir präsentieren die resultierende mechanische Charakterisierung mit Hilfe einer intuitiven visuellen Darstellung, welche sich für die effiziente Erkundung des zugrundeliegenden Raums der Strukturen eignet.

# Acknowledgments

First, I would like to thank my advisor, Prof. Markus Gross, for offering me the opportunity to be part of the amazing research environment formed by the Computer Graphics Lab and Disney Research. His guidance and vision were essential for my doctorate.

My deepest gratitude goes to Prof. Bernhard Thomaszewski, whose support and dedication have guided me through my doctorate. He always motivated me to keep pushing when things did not work out at first.

I will always be grateful to Prof. Bernd Bickel, who guided me at the beginning of my doctorate. He introduced me to all the ins and outs of academic research, and inspired me to tackle difficult problems.

I would also like to thank all the collaborators with whom I had the pleasure to work: Prof. Steve Marschner, whose unique insights and ability to make sense of a complicated project were invaluable to this thesis. Prof. Chiara Daraio and Dr. Jan Rys, who made me appreciate the many different facets of all research communities, and whose enthusiasm has always been contagious.

I also want to extend my gratitude to all former and present members of Disney Research, the Computer Graphics Lab, and the Interactive Geometry Lab. Working in such a great environment made even the long days and weekends enjoyable.

Finally, I could not be where I am without the unconditional support of my family and friends. I will forever be grateful.



# Contents

<b>Abstract</b>	<b>iii</b>
<b>Zusammenfassung</b>	<b>v</b>
<b>Acknowledgements</b>	<b>vii</b>
<b>Contents</b>	<b>ix</b>
<b>Introduction</b>	<b>1</b>
1.1 Challenges . . . . .	3
1.2 Contributions . . . . .	4
1.3 Thesis Outline . . . . .	5
1.4 Publications . . . . .	5
<b>Related Work</b>	<b>7</b>
2.1 Physically-based Simulation of Deformable Objects . . . . .	7
2.1.1 Shells and Rods . . . . .	7
2.1.2 Extended Finite Element Method . . . . .	8
2.1.3 Homogenization . . . . .	8
2.1.4 Analysis and Visualization of Material Models . . . . .	9
2.1.5 Material Modeling and Capture . . . . .	9
2.2 Structured Materials . . . . .	10
2.2.1 Mechanical Metamaterials . . . . .	10
2.3 Computational Design . . . . .	11
2.3.1 Mass Property Optimization . . . . .	11
2.3.2 Structural Optimization . . . . .	12
2.3.3 Deformation Control . . . . .	13
2.3.4 Microstructure Optimization and Inverse Homogenization . . . . .	14
2.3.5 Appearance Control . . . . .	15
2.3.6 Surface Design . . . . .	15
2.4 Textures & Tilings . . . . .	15
2.4.1 Texture Synthesis . . . . .	16
2.4.2 Polygonal Patterns and Other Tilings . . . . .	16

## Contents

<b>Fundamentals</b>	<b>19</b>
3.1 A Very Brief Introduction to Elasticity . . . . .	19
3.1.1 Deformations and Strains . . . . .	20
3.1.2 Elastic Energy . . . . .	20
3.1.3 Stress . . . . .	21
3.2 Linear Elasticity . . . . .	21
3.3 Elasticity and the Finite Element Method . . . . .	22
3.3.1 Simulation . . . . .	23
3.3.2 Removing Rigid Body Motion . . . . .	24
3.3.3 Periodic Boundary Conditions . . . . .	25
3.4 Implicit Function Theorem and the Adjoint Method . . . . .	26
<b>Controlling Elasticity through Microstructures</b>	<b>29</b>
4.1 Overview . . . . .	30
4.2 Background . . . . .	32
4.2.1 Numerical Coarsening . . . . .	32
4.3 Microstructure Optimization . . . . .	32
4.3.1 Problem Formulation . . . . .	33
4.3.2 Regularization . . . . .	34
4.3.3 Numerical Methods . . . . .	38
4.4 Metamaterial Spaces . . . . .	39
4.4.1 Microstructure Representation . . . . .	40
4.4.2 Interpolation . . . . .	42
4.4.3 Generating Metamaterial Families . . . . .	43
4.5 Structure Synthesis . . . . .	45
4.6 Results . . . . .	48
4.6.1 Metamaterial Space Construction . . . . .	48
4.6.2 Validation . . . . .	49
4.6.3 Application Examples . . . . .	53
4.7 Conclusion . . . . .	55
<b>Designing Structurally-Sound Surfaces with Decorative Patterns</b>	<b>57</b>
5.1 Overview . . . . .	59
5.2 Stencil Patterns . . . . .	59
5.2.1 Stencil Representation & Projection . . . . .	59
5.2.2 Pattern Objectives . . . . .	60
5.2.3 Initialization . . . . .	62
5.2.4 User Control . . . . .	63
5.3 Structural Optimization . . . . .	64
5.3.1 Simulation . . . . .	64
5.3.2 Extension to Stenciled Shells . . . . .	65
5.3.3 Optimization . . . . .	66

5.3.4	Numerical Solution . . . . .	68
5.4	Results . . . . .	68
5.4.1	Validation . . . . .	68
5.4.2	Physical Results . . . . .	69
5.5	Conclusion . . . . .	73
5.5.1	Limitations and Future Work . . . . .	73
<b>Mechanical Characterization of Structured Sheet Materials</b>		<b>75</b>
6.1	Structured Sheet Mechanics . . . . .	76
6.1.1	Phenomenology . . . . .	77
6.1.2	Mesosopic Model . . . . .	78
6.1.3	Macroscopic Model . . . . .	79
6.1.4	Homogenization . . . . .	79
6.1.5	Representing Direction-Dependent Elasticity . . . . .	83
6.2	Tilings and Symmetries . . . . .	84
6.2.1	Isohedral Tilings . . . . .	85
6.2.2	Material Symmetries . . . . .	86
6.2.3	Shared Symmetries . . . . .	87
6.3	Results . . . . .	87
6.3.1	Mechanical Testing . . . . .	89
6.3.2	Space of Structures . . . . .	90
6.3.3	Analysis of Individual Structures . . . . .	93
6.3.4	Structure Optimization . . . . .	96
6.4	Conclusions . . . . .	98
6.4.1	Limitations and Future Work . . . . .	99
<b>Conclusion</b>		<b>101</b>
7.1	Future Work . . . . .	103
<b>Appendix</b>		<b>105</b>
A.1	Microstructure Optimization . . . . .	105
A.1.1	Numerical Coarsening . . . . .	105
A.1.2	Compression Test Data . . . . .	107
A.2	Structured Sheet Materials . . . . .	108
A.2.1	Curved Periodic Boundary Conditions . . . . .	108
A.2.2	Computing the Macroscopic In-Plane Strain and Stress . . . . .	109
A.2.3	Direction-Dependent Elasticity Measures . . . . .	109
A.2.4	Membrane . . . . .	110
A.2.5	Test Results . . . . .	111
<b>References</b>		<b>115</b>



# Introduction

*Whether it be the sweeping eagle in his flight, or the open apple-blossom, the toiling work-horse, the blithe swan, the branching oak, the winding stream at its base, the drifting clouds, over all the coursing sun, form ever follows function, and this is the law.*

– Louis Sullivan [1896]

*Form follows function*: the design principle coined by American architect Louis Sullivan states that, simply put, the shape of an object should originate from its purpose and function. He compared the principle to the multitude of examples found in nature—the *sweeping eagle* or *open apple-blossom*—where the shape of a part seems to be made for its function. Of course, rather than this being a manual design process, it is the result of natural selection, telling us that some forms are simply better than others at fulfilling a certain function. But the fact that form and function are closely related goes even deeper than the typical applications of this design principle. The form of an object is not limited to its macroscopic shape, but encompasses all possible scales down to the microscopic and atomic structure. And in nature, there are just as many examples for the importance of small-scale structures as there are for large-scale ones: our bones are light, yet strong due to their porous interior structure. Butterfly wings show an extraordinary range of colors thanks to the interaction of light with their microstructure. And the atomic structure is what distinguishes a diamond from a pencil. In this light, maybe Sullivan's contemporary, Frank Lloyd Wright, put it slightly better: "*Form and function are one*" [Wright, 1953].

## Introduction

The fact that we as designers—in the broadest sense—are not limited to constructing shapes that adhere to the tight connection between form and function makes this design principle even more important. Throughout history, we have been learning about this connection, creating new materials with different compositions and new forms that better fulfill their function. **But while structures at large scales have become easy to control, creating *structured materials* with specific properties derived from their small-scale geometry has been a very difficult task.** Often, stochastic processes can create fascinating new materials like metal foams [Ashby et al., 2000], but the exact geometry of their structure is hard to control. New and improved fabrication methods, especially from the field of additive manufacturing, promise to change that: with the ability to create complex geometries, we can control material properties in a targeted way. These *metamaterials*—materials that obtain their properties from an architected structure rather than just the material from which they are made—promise a wealth of new design possibilities, where a variety of material properties can be tailored to specific applications. This thesis focuses on *mechanical metamaterials*, and specifically the *deformation behavior* of ideally elastic materials.

The effect of small-scale structures in a material has long been studied for the analysis and modeling of existing materials [Gibson and Ashby, 1997]. The precise control of such structures opens new possibilities in engineering, where we can use the existing knowledge base on structured materials to formulate design principles for tailored metamaterials. **But besides creating metamaterials based on existing structures, material properties can also be architected in new ways, by introducing spatially varying patterns, tailored instabilities that create application-specific responses [Florijn et al., 2014], or mechanism-like structures that lend themselves to reconfiguration [Overvelde et al., 2017].** The ability to create an architected behavior also makes metamaterials ideal for pushing the boundaries of material properties, for example through ultrastiff, ultralight metamaterials that show an extreme stiffness to weight ratio [Schaedler et al., 2011; Zheng et al., 2014], materials with a negative Poisson’s ratio [Grima and Evans, 2000], or materials with a negative thermal expansion coefficient [Qu et al., 2017].

The expressiveness of mechanical metamaterials is also connected to the length scales at which structures can be fabricated. Most widely available additive manufacturing technologies target features at the millimeter scale and above, and can not yet take full advantage of structured materials due to this limitation in resolution. However, advanced lithography approaches are able to create structures at the micrometer scale and below, which allows us to control geometry over multiple length scales. The introduc-

tion of structures at different length scale creates hierarchical metamaterials that have been shown to have material characteristics that outperform non-hierarchical metamaterials [Meza et al., 2015]. Structures at the micrometer length scale are also crucial in biomechanical applications: localized material properties, which can be tailored through structure, influence cell growth, migration, and specification [Hahn et al., 2006]. Metamaterials, in combination with new fabrication capabilities, offer opportunities for improvements in many fields, but require the right set of tools and design principles in order to be useful.

And the same holds true also beyond the laboratory and industry setting. Even though the fabrication technologies available for personal use are not as sophisticated, cheaper devices and fabrication services offer an appealing way for *personal fabrication*. These possibilities create a demand for design tools that can integrate functional descriptions and adapt to complex targets, and still be accessible to non-expert users.

### 1.1 Challenges

The increased complexity from new fabrication methods opens new possibilities, but also challenges existing design approaches. While it might have been feasible to design objects with large-scale structures using a traditional iterative approach—a structure is fabricated or possibly simulated, and a designer improves the shape based on the result—designing small-scale structures with a much larger complexity in the same way is simply infeasible. Likewise, the precise fabrication processes that have become more readily available to non-experts cannot be fully taken advantage of because the connection between form and function is not intuitive, and traditional design tools are not able to translate functional requirements into a geometric description.

Automated approaches, such as topology or shape optimization, that create large-scale structures that best approximate a given functional specification have been successful in augmenting traditional design approaches. However, even such approaches often fail if they are naively applied to the massive design space that modern fabrication approaches offer: we cannot optimize or even simulate the complete structure of an object down to the micrometer or nanometer scale in any meaningful way. Rather, we have to find ways to create complex small-scale structures without the need for a global high-resolution simulation and optimization. This also includes carefully selecting the material properties that we need to model: introducing

complex nonlinear and anisotropic effects into a simulation or optimization greatly increases the computational complexity of a problem, and should be justified by tangible improvement in the results.

Further problems arise if we want to create structures that, besides functional targets, also include aesthetic considerations: when the structure of an object will be exposed, it often needs to fulfill visual criteria to be acceptable. However, aesthetics is not a property that is easily quantifiable, and generally requires informed user feedback. A carefully crafted design space is crucial to provide an environment where such an informed feedback is possible, and where even an automatic design optimization might be feasible.

## 1.2 Contributions

This thesis proposes three methods that tackle the control of the mechanical behavior of objects through their structure. The goal of these methods is to either create materials with specific elastic properties, or to prevent deformations that would lead to failure while maintaining a visually appealing appearance. It makes the following main contributions.

- *A method to create microstructures with specific elastic properties and efficiently synthesize assemblies of spatially varying structures.* We optimize the time-intensive task of creating microstructures for a heterogeneous distribution of material properties by using a data-driven approach to synthesize new structures on demand. We analyze the influence of the spatial variation and propose a method to select the optimal structure from a pool of candidates with similar elastic properties.
- *A design approach to create structurally-sound surfaces with decorative patterns.* We integrate simulation and optimization into the design process of surfaces with artistic cutouts, where the aesthetic target of a design has to be balanced with its structural integrity in a synergistic way. This integration allows for a visualization of critical regions of the design, and together with a discrete element texture objective that captures its visual aspect, can be used to automatically improve the stability of the result. We propose a fast simulation approach inspired by material interpolation schemes from topology optimization to provide interactivity and user control.
- *A method to characterize the mechanical properties of structured sheets materials.* We introduce a homogenization method that captures the

salient membrane and bending properties of structures created from tessellations of the plane. Such tessellations represent a large design space that combines visually appealing geometry with interesting mechanical properties, though the relation between these properties is highly unintuitive. We distill the material properties of these structures into a compact visual representation that allows for an intuitive exploration of the space of tilings.

### 1.3 Thesis Outline

Following the introduction in Chapter 1, this thesis is structured as follows:

Chapter 2 provides an overview of the current state of research in computational design, structural optimization, fabrication-oriented algorithms, elasticity simulation, and homogenization, and explains how this thesis fits into the related work.

Chapter 3 introduces the basics of elasticity theory and the simulation of deformable objects. It provides a basis for the methods presented in this thesis.

Chapter 4 introduces a novel method to create microstructures with specific elastic properties, and shows how a data-driven approach can be used to efficiently create microstructures with spatially varying properties. This chapter is based on [Schumacher et al., 2015].

Chapter 5 describes an interactive method to create and optimize structurally-sound shell structures with artistic cutouts. This chapter is based on [Schumacher et al., 2016].

Chapter 6 introduces a method to characterize the deformation behavior of rod networks based on tessellations of the plane. This chapter is based on [Schumacher et al., 2018].

Finally, Chapter 7 summarizes the thesis, and provides an outlook for future research directions.

### 1.4 Publications

In the context of this thesis, the following peer-reviewed publications have been accepted:

## Introduction

- Christian Schumacher, Bernd Bickel, Jan Rys, Steve Marschner, Chiara Daraio, and Markus Gross. Microstructures to Control Elasticity in 3D Printing. *ACM Transactions on Graphics*, 34(4):136:1–136:13, 2015.
- Christian Schumacher, Bernhard Thomaszewski, and Markus Gross. Stenciling: Designing Structurally-Sound Surfaces with Decorative Patterns. *Computer Graphics Forum*, 35(5):101–110, 2016.
- Christian Schumacher, Steve Marschner, Markus Gross, and Bernhard Thomaszewski. Mechanical Characterization of Structured Sheet Materials. *ACM Transactions on Graphics*, 37(4):148:1–148:15, 2018.

The contents of these papers are included in this thesis.

## Related Work

This chapter provides an overview of the research efforts in computational design, structured materials, and related fields.

### 2.1 Physically-based Simulation of Deformable Objects

A central part of computational design methods is the prediction of how an object behaves in the physical world. As such, physically-based simulations play an important role in many of these methods. Especially in the context of stability and elasticity, modeling the deformation behavior of a material properly is crucial. Such *deformable models* have originally been introduced to computer graphics in the context of animation [Nealen et al., 2006], where the focus has been on realism and simplicity, but not necessarily accuracy. With the stricter requirements in computational design approaches, the *finite element method* (FEM) has become a popular approach for simulations. An excellent introduction can be found in Sifakis and Barbič [2012].

#### 2.1.1 Shells and Rods

Apart from volumetric simulations that can model arbitrary geometries, specialized models have been developed for shells and rods, whose thin features would otherwise require a prohibitively high resolution of the underlying mesh. The simulation of thin shells has been extensively used in computer graphics to animate cloth, but has also started to find applications in computational design tools in recent years, where it is used to

## Related Work

predict and optimize different types of surfaces. Grinspun et al. [2003] introduced a widely used shell model with hinge-based bending elements, which has been further investigated and extended [Gingold et al., 2004; Garg et al., 2007]. We employ a modified version of this shell model in Chapter 5.

Specialized rod simulation approaches have also become popular in the simulation of one-dimensional structures and networks of beams. Bergou et al. [2008; 2010] introduced popular models with space-parallel and time-parallel frames that have also been used as the basis to computational design methods. We use rod simulations extensively in Chapter 6.

### 2.1.2 Extended Finite Element Method

An emerging simulation approach in computer graphics is the *extended finite elements method* (XFEM), where a regular FEM simulation is augmented with additional basis functions, creating *enriched* elements that can recreate more complex deformations such as cracks. For shell simulations, Kaufmann et al. [2009] introduced such an approach that uses enrichment textures to augment a coarse simulation with detailed cuts.

### 2.1.3 Homogenization

Inspired by the seminal work of Hashin and Shtrikman [1963], homogenization theory was developed to efficiently simulate inhomogeneous materials with fine structures, allowing for microscopic behavior to be averaged into a coarser representation with equivalent behavior at the macroscopic scale [Hassani and Hinton, 1998; Michel et al., 1999; Cioranescu and Donato, 1999; Miehe and Koch, 2002]. A similar idea has already been used in laminates, where non-dominant dimensions are analytically contracted to create a simpler model [Reddy, 2004]. For general shells, homogenization approaches use numeric simulation [Coenen et al., 2010], and similar methods have also been applied to beam structures [Cartraud and Messenger, 2006; Buannic and Cartraud, 2001].

To improve the simulation of complex geometries in the context of computer graphics, Nesme et al. [2009] encoded the material stiffness within coarse elements using shape functions after a fine-level static analysis, and Kharevych et al. [2009] turned the heterogeneous elastic properties represented by a fine mesh into anisotropic elastic properties of a coarse mesh

## 2.1 Physically-based Simulation of Deformable Objects

that effectively captures the same physical behavior. Chen et al. [2018] introduced a coarsening approach for complex inhomogeneous and nonlinear materials by tailoring the shape functions of a coarse mesh.

Chen et al. used a database of material combinations to create a data-driven finite element method [2015], and introduced a specialized homogenization approach for dynamic simulations [2017]. Geers et al. [2007] described a nonlinear homogenization approach for thin sheets that uses a full solid simulation to compute the material response on the microscopic level without fitting a material model.

We apply homogenization approaches in Chapter 4 and Chapter 6 to describe the macroscopic elastic behavior of structured materials.

### 2.1.4 Analysis and Visualization of Material Models

While an abstract or numerical representation of a material model is sufficient to simulate the elastic response of an object, it does not lend itself to the analysis and visualization of the material behavior. To allow for a more intuitive understanding especially of anisotropic materials, Marmier et al. [2010] and Gaillac et al. [2016] introduced simple visualization approaches that create graphical representations of important material properties.

François et al. [2017] use directional Young's modulus plots to visualize the material properties of triangular lattices, based on a parameterized analytical model. We use such a visualization as the basis of our characterization method presented in Chapter 6.

### 2.1.5 Material Modeling and Capture

As important as it is to use the appropriate simulation approach for a given situation, it would be useless without the right material parameters. For the purpose of animation, such material properties can be designed virtually with the help of intuitive methods [Xu et al., 2015b; Li and Barbič, 2015]. However, for simulations that reproduce real objects, capturing and accurately modeling their material properties is essential. Apart from using standardized tests to determine a set of material parameters, methods that can capture deformation properties of complex nonlinear and anisotropic materials have become popular. Bickel et al. [2009] presented a data-driven material model for soft tissue. Becker and Teschner [2007] introduced a method to fit spatially-varying material parameters from measurements of real-world objects. For the purpose of characterizing the properties of cloth, Wang et

## Related Work

al. [2011] and Miguel et al. [2012] presented systems to capture deformations of different types of cloth in various configurations and determine the corresponding material models.

## 2.2 Structured Materials

While computer graphics has only recently developed an interest in structured materials, the analysis of such structures has been an important topic in mechanical engineering and materials science for a long time [Gibson and Ashby, 1997]. Such structures, whether they are the result of stochastic processes in naturally occurring materials or part of specifically designed materials, are often modeled as random or regular structures, such as foams, fiber networks, or lattices [Ashby, 2005; Picu, 2011], or as *mechanism-like* structures such as truss networks [Hutchinson and Fleck, 2006] that model the behavior of more idealized materials and are used to explore the theoretical boundaries of material properties.

### 2.2.1 Mechanical Metamaterials

*Metamaterials* describe structured materials whose geometry has been tailored to have desired material properties. *Mechanical metamaterials* specifically create an architected deformation behavior through their structure [Bertoldi et al., 2017]. Inspiration for such metamaterials can come from structures found in nature [Wegst et al., 2015], or from the analysis and modeling of structured materials [Gibson and Ashby, 1997; Fleck et al., 2010]. But with increased control over the geometry of structures, new design approaches also look into other areas, with origami- and kirigami-inspired designs [Xu et al., 2017], combinatorial assemblies of structures to tailor materials with spatially varying material responses [Coulais et al., 2016], or structures that use mechanical instabilities for tailored material responses [Mullin et al., 2007; Kochmann and Bertoldi, 2017].

Metamaterials can also explore the theoretical boundaries of specific material properties, with the goal of creating unusual material behavior. Auxetic materials, which exhibit a negative Poisson's ratio—expanding laterally when stretched [Lakes, 1987]—have become a topic of interest due to their beneficial behavior for, e.g., shock absorption, with designs often based on *mechanism-like structures* [Grima and Evans, 2000; Rafsanjani and Pasini,

2016]. Similarly, materials with negative thermal expansion have been designed using embedded multi-material structures [Qu et al., 2017]. The boundaries of metamaterials have also been pushed by the ability to 3D print structures below the micrometer scale [Lee et al., 2012], which can be leveraged to construct ultrastiff, ultralight materials [Schaedler et al., 2011; Zheng et al., 2014]. This precision also allows for the creation of hierarchical metamaterials that incorporate structures at different length scales, showing promising results for ultralight materials with unexpectedly high resilience [Meza et al., 2015].

## 2.3 Computational Design

*Computational Design* has become an important approach for the next generation of design tools. With an increase in design complexity and possibilities in functional specifications, computational design can augment or replace traditional geometry-based design approaches with fabrication-oriented methods. As such, computational design approaches often combine simulation and optimization methods with traditional user interaction.

In computer graphics, we are currently witnessing an increasing interest in fabrication-oriented design for reproducing 3D physical artifacts from virtual representations, tackling different problems: from controlling the appearance, elasticity, or stability of fabricated objects, to solving very specific design problems tied to unique fabrication techniques. Chen et al. [2013] presented an abstraction mechanism for translating functional specifications to fabricable 3D prints, and Vidimčič et al. [2013] introduced a programmable pipeline for the procedural evaluation of geometric detail and material composition, allowing models to be specified easily and efficiently.

### 2.3.1 Mass Property Optimization

The mass distribution of an object is a simple property that is hard to intuitively control for a user, but relatively easy to compute. Prévost et al. [2013] were the first to introduce a straightforward method to optimize the center of mass of an object with the goal to make it stand, and later also included movable masses to create reconfigurable mass properties [Prévost et al., 2016]. Bächer et al. [2014] extended the computation of mass properties to the moment of inertia, allowing them to optimize the rotational properties needed to create spinning tops with complex geometries. Musial-

## *Related Work*

ski et al. [2015; 2016] introduced a novel discretization of the interior surface to simplify the computation and optimization of mass properties.

### **2.3.2 Structural Optimization**

Finding designs with good structural properties has long been a central topic in mechanical engineering. Topology optimization [Bendsøe and Sigmund, 2003; Deaton and Grandhi, 2014] has been used to find structures that provide minimal compliance, given a fixed amount of material. These approaches often use material interpolation schemes that allow them to transform the discrete problem—where to put material—to a continuous one—how much material is at a given point [Bendsøe and Sigmund, 1999]. Wu et al. [2016] have recently shown that even high-resolution topology optimization problems can be solved efficiently. Other approaches have applied similar principles to the design of structural shells [Hassani et al., 2013].

In the computer graphics community, several papers have proposed efficient methods to increase the stability of 3D printed objects. Stava et al. [2012] improved the structural strength by automatic hollowing, thickening, and strut insertion. Wang et al. [2013] proposed a method for computing skin-frame structures for the purpose of reducing the material cost of the printed object. Lu et al. [2014] optimize a distribution of internal voids for stability to create honeycomb-like structures with an optimal strength-to-weight ratio. This approach can be seen as an extension of traditional infill patterns common in 3D printing [Livesu et al., 2017]. Topology optimization methods generally assume a fixed force distribution. To find forces that most likely lead to failure, Zhou et al. [2013] presented a method that determines the worst-case loads for fabricated models. Using a similar approach, Langlois et al. [2016] presented a method that determines a stochastic measure of failure points and can be used for inverse design approaches that increase the stability based on this measure.

While pure structural optimization methods generally lack the ability to control the aesthetics of the result, methods that combine form and function have been successful at bridging the gap between mechanical and aesthetic objectives. Statics aware grid shells by Pietroni et al. [2015] are an example of physics-based surface designs with an aesthetic component, though the appearance is restricted to Voronoi patterns.

Creating a more direction combination of aesthetics and stability through objects with decorative patterns, Dumas et al. [2015] used a voxel-based representation to remove material from an object by projecting a binary texture

onto a mesh. With a similar goal, Zehnder et al. [2016] introduced a method to create ornamental curve networks using a rod-based parameterization that avoids rasterization and provides a high degree of artistic control. Martinez et al. [2015] directly combine appearance and compliance objectives in a unified optimization. Their approach is based on traditional topology optimization and tries to find an aesthetically-pleasing structure within a given stability bound, using a fixed amount of material.

We present our own optimization of shells with decorative patterns, inspired by topology optimization approaches, in Chapter 5.

### 2.3.3 Deformation Control

Designing objects to show specific deformations under a given load goes one step further—we no longer want to maximize the structural strength, but aim to change the behavior of an object such that it fulfills a functional target. Skouras et al. [2013] introduced a method to create a two-material distribution that best reproduces a series of example deformations of the object. Pérez et al. [2015] presented an approach to create light-weight and cost-effective rod meshes with a desired deformation behavior and later combined two different material domains by printing rod structures directly on a prestretched cloth to create curved surfaces [Pérez et al., 2017]. Guseinov et al. [2017] use a similar approach—combining 3D-printed parts with prestretched membranes—to create shells that bend in a desired way and keep their shape. Connecting the advantages of silicone with the expressiveness of structured materials, Zehnder et al. [2017] introduced a material design approach that uses inclusions embedded in silicone to create various material behaviors.

Apart from these generally static targets, other methods have investigated motion and dynamics of objects. Hiller et al. [2012] introduced an approach to achieve high-level functionality such a locomotion of soft robots. Ma et al. [2017] use pneumatic actuation to deform objects and optimize the individual chambers to match a target pose. Megaro et al. [2017] optimize the behavior of compliant mechanisms for trajectories and robustness in order to replace traditional rigidly-articulated mechanisms. Bern et al. [2017] design plush toys with internal strings that can actuate the toy in a desired way. And Chen et al. [2016] use a numerical coarsening approach to efficiently optimize the dynamic behavior of objects.

In a purely virtual approach, Xu et al. [2015a] presented an interactive material design tool which computes a spatial distribution of material properties given user-provided displacements and forces at a set of mesh vertices.

### 2.3.4 Microstructure Optimization and Inverse Homogenization

A typical inverse design problem for structured materials is to find a specific small-scale structure or distribution of structures that show a desired macroscopic deformation behavior. While this often involves an inverse homogenization, there are also more direct approaches. Bickel et al. [2010] presented a data-driven process for designing and fabricating objects with a desired deformation behavior, based on a combination of predetermined structures whose properties have been measured in advance.

Classical inverse homogenization approaches generally avoid these measurements, instead using analytical or numerical approaches to determine the macroscopic elastic properties from a single *unit cell*—the smallest repeatable unit in the structure. The inverse design approach then looks for suitable small-scale structures by optimizing the material distribution in this unit cell. Researchers have proposed various parameterizations of the material distribution, such as networks of bending beams [Hughes et al., 2010], spherical shells patterned with an array of circular voids [Babaei et al., 2013], or rigid units [Attard and Grima, 2012]. Alternatively, the domain of a base cell can be discretized into small material voxels, and a discrete value problem has to be solved. Due to the combinatorial complexity, direct search methods are prohibitively expensive, and the problem is usually solved using a relaxed formulation with continuous material density variables [Sigmund, 2009] or advanced search heuristics [Huang et al., 2011].

Rodrigues et al. [2002] and Coelho et al. [2008] suggest methods for hierarchical topology optimization, computing a continuous material distribution on a coarse level and matching microstructures for each coarse cell. While in their approach each microstructure cell can be optimized independently, each of them still needs to be computed based on a costly optimization scheme. For functionally graded materials with microstructures, Zhou et al. [2008] guarantee the matching of boundaries either by prescribing connectors or by incorporating a complete row of cells that form a gradient during a single optimization.

Researchers in computer graphics have been especially interested in creating efficient inverse homogenization approaches that can create large tilings of microstructures with spatially varying properties. Panetta et al. [2015] introduced a microstructure optimization approach based on a truss-like parameterization that can create a database of small-scale structures with a large range of elastic properties. Later, they extended this approach to include a stress-based worst-case objective that also minimizes the structure's likelihood to break [Panetta et al., 2017]. Zhu et al. [2017] followed

the more traditional voxel-based optimization approach, applying it to the synthesis of dual-material microstructure, but also optimize for the macroscopic distribution of material parameters in a two-scale approach. Martínez et al. [2016] employed an efficient stochastic approach to create isotropic Voronoi foams and determined the mapping between structure parameter and mechanical properties. They extended their foam-based approach to orthotropic foams [Martínez et al., 2017], and also included fabrication constraints [Martínez et al., 2018]

We propose a data-driven microstructure optimization approach similar to [Panetta et al., 2015], but using a voxel-based discretization, in Chapter 4.

### 2.3.5 Appearance Control

Recent work also investigated the reproduction of appearance through a combination of material and structure, creating a macroscopic appearance that is governed by microscopic details. These approaches include modulating the surface structure to achieve desired reflection properties [Weyrich et al., 2009; Lan et al., 2013; Rouiller et al., 2013], interleaving different colored materials on the surface [Reiner et al., 2014], or volumetric combination of multiple materials [Hašan et al., 2010; Dong et al., 2010] to control subsurface scattering behavior.

### 2.3.6 Surface Design

Several researchers have proposed approaches to design different forms of physical surfaces. Skouras et al. [2012] optimize the shape of rubber balloons to match a target shape once inflated, and later applied a similar approach to the design of inflatable, non-stretchable structures [Skouras et al., 2014]. Other methods include design approaches for interlocking structures [Skouras et al., 2015], plush toys [Mori and Igarashi, 2007], bead-work [Igarashi et al., 2012], wire-mesh models [Garg et al., 2014], and thermoformed surfaces with textures [Schüller et al., 2016].

## 2.4 Textures & Tilings

Computer graphics has a long history of creating decorative patterns on digital surfaces. These patterns are often used to add additional detail on otherwise simple geometries, or to create interesting geometric patterns on

## Related Work

such surfaces. When we take the step from purely digital models to physical models, such patterns need to not only be visually appealing, but also fulfill fabrication constraints, such as being structurally sound. At the same time, geometric patterns also offer an extensive space of possible mechanical properties that arise from their spatial configuration, as we will demonstrate in Chapter 6.

### 2.4.1 Texture Synthesis

Example-based texture synthesis methods use a small input exemplar to create a seamless output in a larger domain; see for instance Wei et al. [2009]. In contrast to traditional pixel-based or voxel-based textures, discrete example-based textures [Ijiri et al., 2008; Landes et al., 2013] offer more flexibility to modify a texture without destroying its appearance, and are well suited for applications in digital fabrication. Texture synthesis approaches that use energy-based formulations to define their neighborhood matches, for example the methods introduced by Ma et al. [2011] and Roveri et al. [2015], are especially suited to be included in an optimization-based computational design approach.

We use discrete element textures [Ma et al., 2011] to quantify the aesthetics of a distribution of cutouts on a shell in Chapter 5.

### 2.4.2 Polygonal Patterns and Other Tilings

Computer Graphics has traditionally been interested in procedural pattern generation for decorating digital surfaces, both with texture and with geometric structures. Besides architectural applications of meshing and tiling methods [Tang et al., 2014; Jiang et al., 2015], some of these approaches take inspirations from physical pattern creation in quilting [Zhou et al., 2006] or weaving [Akleman et al., 2009]. Others use tilings to create efficient samplings [Ostromoukhov, 2007]. In a combination of geometry and fabrication, Konaković et al. [2016] used the unique properties of auxetic mechanical materials to design surfaces that can be fabricated from a single flat piece of material.

*Space-filling tilings* of the plane offer an interesting parameterization for visually appealing structures, and have been studied extensively in the mathematics community [Grünbaum and Shephard, 1986]. Kaplan and Salesin [2000] used a special class of space-filling tilings, isohedral tilings

with curved edges, and in subsequent work investigated the synthesis of Islamic star patterns [Kaplan and Salesin, 2004]. We take inspiration from the aesthetics and the continuous parameterization of these tilings, and investigate the space of mechanical properties that they create in Chapter 6.

## *Related Work*

# Fundamentals

The simulation of deformable structures is at the core of all methods presented in this thesis. This chapter will give a brief introduction to the theory of elasticity and the finite element method, and its application to the simulation of deformable structures.

We will also present a short overview and derivation of the implicit function theorem and the adjoint method, which provide the basis for many modern computational design approaches, including two methods presented in this thesis.

## 3.1 A Very Brief Introduction to Elasticity

The general principle that governs elasticity is the tendency of an elastic object to return to a *natural state*, or *rest state*. This phenomenon can be intuitively described in terms of energy: the rest state represents the state of least energy, and any deformation away from this state introduces more energy into the object. This additional energy generates an *elastic response*—an internal force that pulls the object to its rest state—and the interaction of this response with external forces and constraints is what we simulate in a deformable object.

Elasticity excludes any effects that change the rest state of an object, either due to excessive stresses (plasticity) or creep (viscosity). While backed by a wealth of theory of their own, these *inelastic* phenomena are often undesirable, and computational design approaches generally try to prevent inelastic deformations.

### 3.1.1 Deformations and Strains

Models in elasticity generally work with a notion of *deformation*—a change in shape or size. While there are many ways to describe a deformation, a description that is useful in the context of elasticity should capture one of its central properties: elasticity is invariant to translations and rotations, or in other words, rigid motions will not induce an elastic response. We generally refer to such a measure of deformation as a *strain*.

We start with the definition of a strain by looking at how we can describe a deformation. We consider the relationship between the position  $\mathbf{X}$  of a point of the object in its rest state, and its position  $\mathbf{x}(\mathbf{X})$  in a deformed state. In order to describe a deformation, we want to know how the points close to this position change: if they experience a different displacement, a deformation will be induced. We can quantify this with the derivative of the deformed position, and for that reason define the *deformation gradient* as

$$\mathbf{F} = \frac{\partial \mathbf{x}}{\partial \mathbf{X}}. \quad (3.1)$$

This deformation gradient measures the local deformation, but while it is invariant to translations, it will change when rotated. We can remove the rotational part of the deformation gradient by multiplying it with itself, creating the *right Cauchy-Green deformation tensor*  $\mathbf{C} = \mathbf{F}^T \mathbf{F}$ .

We can now use the right Cauchy-Green deformation tensor to define a strain measure, the *Green strain*, as

$$\mathbf{E} = \frac{1}{2}(\mathbf{C} - \mathbf{I}). \quad (3.2)$$

Based on this strain measure, we can now define constitutive equations for the elastic response.

### 3.1.2 Elastic Energy

The behavior of elastic objects can conveniently be described through an *elastic energy*. Such an energy is defined with respect to a given deformation, and describes the amount of work needed to reach this deformation. The elastic energy of an object is generally described through an *elastic energy density*  $\Psi(\mathbf{E})$ , a function of some strain measure  $\mathbf{E}$ —in our case generally the Green strain or the Cauchy strain for a linear material (see below)—which is integrated over the whole object domain  $\Omega$  to arrive at the elastic energy  $W$ :

$$W = \int_{\Omega} \Psi(\mathbf{E}(\mathbf{x})) d\mathbf{X} \quad (3.3)$$

The actual form of the elastic energy is material-dependent, and generally parameterized by a set of *material parameters*. The St. Venant-Kirchhoff material model is a simple material model that uses the Green strain in its definition. It uses two material parameters,  $\mu$  and  $\lambda$ , and has the form

$$\Phi_{\text{StVK}}(\mathbf{E}) = \frac{\mu}{2} \text{tr}(\mathbf{E}^t \mathbf{E}) + \lambda \text{tr}(\mathbf{E})^2. \quad (3.4)$$

### 3.1.3 Stress

The tendency of an elastic object to return to its rest state gives rise to *internal forces*. The *stress* is a measure of these forces inside a continuum. As such, it is related to the elastic energy—the internal forces push the object towards a lower energy state. Just as there are different ways to measure strain, there are also different ways to define a stress. For example, the *second Piola-Kirchhoff stress tensor*  $\mathbf{P}$  of an elastic energy is defined as the derivative of the elastic energy with respect to the Green strain:

$$\mathbf{P} = \frac{\partial W}{\partial \mathbf{E}}. \quad (3.5)$$

## 3.2 Linear Elasticity

We often use a linearized version of the elasticity equations for efficiency reason. If we can keep the strain measure as well as the material model and the quadrature rule linear, we can generally simplify the simulation of an elastic object to the solution of a single system of linear equations. While such a model only captures a subset of all deformation effects, namely only effects that involve small deformations, it is often sufficient for applications that only require an approximate solution.

The linearized version of the Green strain is called the *Cauchy strain*, and is defined as

$$\epsilon = \frac{1}{2}(\mathbf{F} + \mathbf{F}^T) - \mathbf{I}. \quad (3.6)$$

One effect of the linearization of the strain is that while it is still invariant to translations, it is only invariant to the *linearized* rotations, which are generally not a good approximation of the full rotation. Nonetheless, the Cauchy strain is a suitable strain measure for many applications.

In the linearized version of elasticity, the elastic energy can be written in a general form as

$$W = \frac{1}{2} \boldsymbol{\epsilon} : \mathbb{C} : \boldsymbol{\epsilon}. \quad (3.7)$$

The rank-4 tensor  $\mathbb{C}$  is commonly called the *material stiffness tensor*. Due to the symmetry of the energy and the underlying physical principles, this tensor has to be partially symmetric, with the minor symmetries  $\mathbb{C}_{ijkl} = \mathbb{C}_{jikl}$  and  $\mathbb{C}_{ijkl} = \mathbb{C}_{ijlk}$ , as well as the major symmetry  $\mathbb{C}_{ijkl} = \mathbb{C}_{klij}$ . A quantity that is connected to the material stiffness tensor is the *material compliance tensor*  $\mathbb{S} = \mathbb{C}^{-1}$ , which is its (symmetric) inverse. While we can define a material model with either one of these, the compliance tensor will be important in extracting different material properties from the material tensor.

Based on this energy formulation, we can now also define the linearized version of the stress. The *linear Cauchy stress*  $\boldsymbol{\sigma}$ , which is linearly dependent on the strain, is defined as

$$\boldsymbol{\sigma} = \mathbb{C} : \boldsymbol{\epsilon}. \quad (3.8)$$

Now we can also see other relationships emerge. For example, the linear elastic energy can also be written as a product of stress and strain,  $W = \frac{1}{2} \boldsymbol{\epsilon} : \boldsymbol{\sigma}$ . And the material compliance tensor  $\mathbb{S}$  can be used to compute the strain corresponding to a given stress,  $\boldsymbol{\epsilon} = \mathbb{S} : \boldsymbol{\sigma}$ .

### 3.3 Elasticity and the Finite Element Method

One of the most popular methods to simulate elastic objects is the *finite element method* (FEM), which discretizes the domain of the simulation with vertices and elements—often triangles and tetrahedra in two and three dimensions, respectively—and allows us to approximate the energy integral on a per-element basis. We will not go into the details of the method here, but refer to [Bathe, 2006] for an in-depth introduction. We will, however, quickly summarize the result of applying the finite element method to an elasticity problem.

### 3.3 Elasticity and the Finite Element Method

The finite element method defines a *basis function* for each element that approximates the function that we want to reproduce—in our case the deformed configuration  $\mathbf{x}$  of the elastic object. While this basis function can be arbitrarily complex, we will limit this introduction to *linear* basis functions. In the linear case, where the deformed position is linearly interpolated inside each element, the deformation gradient, and therefore also the strain, is constant. This significantly simplifies the finite element formulation of the energy, as we do not need to apply complex quadrature rules for the per-element integration, but can simply use a single quadrature point. The finite element approximation of Equation (3.3) then becomes

$$W = \sum_i V_i \Psi(\mathbf{E}_i), \quad (3.9)$$

where we take the sum over the products of the energy density for the strain  $\mathbf{E}_i$  inside the  $i$ -th element, and the element volume  $V_i$ . In the case of linear elasticity, we can replace the energy density by the linear energy formulation, and arrive at

$$W = \sum_i \frac{1}{2} V_i \boldsymbol{\epsilon}_i : \mathbf{C} : \boldsymbol{\epsilon}_i. \quad (3.10)$$

#### 3.3.1 Simulation

The simulation of a deformable object involves finding the elastic response to a set of external forces and boundary conditions. While there is also a possible dynamic component to the solution, we are only interested in the static case in this thesis, and solve for the *equilibrium* of the object where all the net forces are zero. There are multiple ways to approach this problem, and an elegant solution is to formulate it as an energy minimization problem: we balance the internal elastic energy against an external energy defined by the potential of the applied external forces. In our case, these forces are generally defined per vertex—for each vertex  $i$ , we have a force  $\mathbf{f}_i$ . The combined energy of the system is then given as

$$W_{\text{total}} = W - \sum_i (\mathbf{x}_i - \mathbf{X}_i)^T \mathbf{f}_i, \quad (3.11)$$

where  $\mathbf{x}_i$  and  $\mathbf{X}_i$  are the deformed and undeformed positions of vertex  $i$ , respectively. We can now minimize this energy using, for example, a Newton-

Raphson approach. A solution that minimizes the energy will, at vertex  $i$ , fulfill the first-order optimality condition

$$\frac{\partial W_{\text{total}}}{\partial \mathbf{x}_i} = \frac{\partial W}{\partial \mathbf{x}_i} - \mathbf{f}_i = \mathbf{0}. \quad (3.12)$$

We see that for such a solution, the external forces match the derivative of the elastic energy,  $\frac{\partial W}{\partial \mathbf{x}_i} = \mathbf{f}_i$ , which just happens to be the internal force generated by the elastic energy. We have arrived at a force equilibrium. To include all vertices in a single equation, we will often concatenate all vertex positions and forces into single vectors,  $\mathbf{x} = [\mathbf{x}_1^T \cdots \mathbf{x}_n^T]^T$  and  $\mathbf{f} = [\mathbf{f}_1^T \cdots \mathbf{f}_n^T]^T$ , and then formulate the problem as

$$\nabla_{\mathbf{x}} W(\mathbf{x}) = \mathbf{f}. \quad (3.13)$$

### 3.3.2 Removing Rigid Body Motion

Since the elastic energy  $W(\mathbf{x})$  is invariant to translation and rotation, the solution to Equation (3.13) is not unique. A common work-around to this is to constrain enough degrees of freedom to get rid of this null space. However, the choice of degrees of freedom might influence the solution in the presence of forces. Another approach to resolve the null space is to introduce constraints on the moments of the object, as presented in Zhou et al. [2013]. These constraints take the form

$$\begin{aligned} \mathbf{c}_1(\mathbf{x}) &= \sum_{i=1}^n (\mathbf{x}_i - \mathbf{X}_i) = \mathbf{0} \\ \mathbf{c}_2(\mathbf{x}) &= \sum_{i=1}^n ((\mathbf{x}_i - \mathbf{X}_i) \times (\mathbf{x}_i - \bar{\mathbf{X}})) = \mathbf{0}, \end{aligned} \quad (3.14)$$

where  $\mathbf{X}_i$  is the rest state position of vertex  $i$ , and  $\bar{\mathbf{X}}$  is the mean rest state position. For simplicity, we can combine these constraints into a single vector  $\mathbf{c}(\mathbf{x}) = [\mathbf{c}_1(\mathbf{x})^T \ \mathbf{c}_2(\mathbf{x})^T]^T$ . Intuitively, these constraints fix the mean translation and linearized rotation. For the special case of computing  $\mathbf{c}_2$  for two-dimensional elasticity, we treat positions as points on the  $z = 0$  plane and use the  $z$ -component of the cross product.

We can then combine Equation (3.13) and Equation (3.14) into a single equilibrium problem:

$$\nabla_{\mathbf{x}}W(\mathbf{x}) = \mathbf{f} \quad \text{s.t.} \quad \mathbf{c}(\mathbf{x}) = \mathbf{0} \quad (3.15)$$

In the case of linear elasticity, this problem can be expressed as a system of linear equations and solved efficiently.

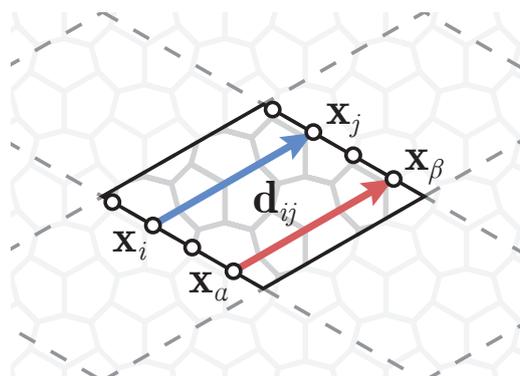
### 3.3.3 Periodic Boundary Conditions

The simulation of periodic structures, such as the ones presented in Chapter 4 and Chapter 6, requires a slight modification of the regular simulation approach. Such structures can be represented by a *unit cell* (UC) or *representative volume element* (RVE)—the smallest unit that can be tiled to build an arbitrarily large structure. When we want to determine the elastic behavior of such a tiling, we generally do not want to simulate such a complete tiling. Rather, we want to simulate a single unit cell in a way that emulates this tiling, and then reason about the tiling from this single unit cell. Such a tiling can be emulated using *periodic boundary conditions*, a set of boundary conditions that match the shape of boundaries, and effectively simulate a version of the unit cell that is surrounded by virtual copies of itself.

We incorporate these periodic boundary conditions by imposing constraints on vertices on opposite boundaries. These boundary conditions guarantee that the patch behaves as if it were surrounded by identical copies of itself, without having to fix the position of vertices, which would artificially stiffen the structure. The relation of a unit cell to its virtual neighbors is given by a simple translation. Since patches are connected at the boundary, we emulate a tiling by guaranteeing that opposite boundaries have the same shape, without restricting the actual shape of the boundary. This requirement can be formulated as a constraint on the difference between corresponding pairs of vertices on opposite boundaries [Smit et al., 1998]. For any pair of vertices  $(\mathbf{x}_i, \mathbf{x}_j)$  on opposite boundaries and a translation  $\mathbf{d}_{ij}$ , we can express one of the vertices through the other as

$$\mathbf{x}_j = \mathbf{x}_i + \mathbf{d}_{ij}. \quad (3.16)$$

The translation  $\mathbf{d}_{ij}$  describes the tileability of the patch, and is identical for every pair of vertices on the same pair of boundaries (see Figure 3.1). We keep this translation as a degree of freedom by choosing a *reference vertex pair*  $\mathbf{x}_\alpha$  and  $\mathbf{x}_\beta$  on each boundary pair, and leave their positions unconstrained. We then express the translation as  $\mathbf{d}_{ij} = \mathbf{x}_\beta - \mathbf{x}_\alpha$ , which ensures that while the distance between two opposite boundaries is free, it is the same for all



**Figure 3.1:** Periodic boundary conditions link all pairs of vertices on opposite boundaries by a translation  $\mathbf{d}_{ij}$ .

boundary vertex pairs on corresponding boundaries. These boundary conditions can be efficiently integrated into a simulation by removing the corresponding vertices from the degrees of freedom—we lose  $\mathbf{x}_j$  as a degree of freedom, since its value is determined by  $\mathbf{x}_i$ ,  $\mathbf{x}_\alpha$ , and  $\mathbf{x}_\beta$ .

We can also use the vertices  $\mathbf{x}_\alpha$  and  $\mathbf{x}_\beta$ , and their counterpart on the other pairs of boundaries, to apply macroscopic strains to the structure. Specifically, if we fix the displacement between  $\mathbf{x}_\alpha$  and  $\mathbf{x}_\beta$  along a single direction, we impose a *uniaxial strain*, and if we fix the full displacement, it will be a full strain.

### 3.4 Implicit Function Theorem and the Adjoint Method

Computational design methods are often based on a forward problem—a simulation of an elastic object, for example—and try to solve an inverse design problem with the assumption that this forward problem is adjusted as the design changes. If we assume that the forward problem is the solution of a static equilibrium elasticity simulation, and our inverse design problem can be formulated as a minimization of a function  $g(\mathbf{p}, \mathbf{x})$  of some parameters  $\mathbf{p}$  and the static equilibrium  $\mathbf{x}$ , we can formulate a general computational design problem as

$$\min_{\mathbf{p}} g(\mathbf{p}, \mathbf{x}) \quad \text{subject to} \quad \nabla_{\mathbf{x}} W(\mathbf{x}, \mathbf{p}) = \mathbf{f}. \quad (3.17)$$

The constraint  $\nabla_{\mathbf{x}} W(\mathbf{x}, \mathbf{p}) = \mathbf{f}$  ensures that the static equilibrium condition is fulfilled (see Equation (3.13)). While we could solve such a problem with

### 3.4 Implicit Function Theorem and the Adjoint Method

a constraint optimization approach, another option is to directly integrate the constraint into the function  $g$  and then use the *implicit function theorem* to compute the correct derivative. If we assume that the constraint is always satisfied, then the variable  $\mathbf{x}$  is implicitly dependent on the variable  $\mathbf{p}$ , since the equilibrium condition changes if the parameters change. If we want to compute the derivative of  $g$  with respect to the parameters  $\mathbf{p}$ , as we generally do if we want to apply a gradient-based optimization approach, we then end up with

$$\frac{dg(\mathbf{p}, \mathbf{x})}{d\mathbf{p}} = \frac{\partial g(\mathbf{p}, \mathbf{x})}{\partial \mathbf{p}} + \frac{\partial g(\mathbf{p}, \mathbf{x})}{\partial \mathbf{x}} \frac{d\mathbf{x}}{d\mathbf{p}}. \quad (3.18)$$

The implicit function theorem states that we can compute the derivative  $\frac{d\mathbf{x}}{d\mathbf{p}}$  from the constraint by differentiating it.

$$\begin{aligned} \frac{\partial W(\mathbf{x}, \mathbf{p})}{\partial \mathbf{x}} &= \mathbf{f} \\ \frac{d}{d\mathbf{p}} \left( \frac{\partial W(\mathbf{x}, \mathbf{p})}{\partial \mathbf{x}} \right) &= \frac{d}{d\mathbf{p}} \mathbf{f} \\ \frac{\partial^2 W(\mathbf{x}, \mathbf{p})}{\partial \mathbf{x}^2} \frac{d\mathbf{x}}{d\mathbf{p}} + \frac{\partial^2 W(\mathbf{x}, \mathbf{p})}{\partial \mathbf{x} \partial \mathbf{p}} &= \mathbf{0} \\ \frac{d\mathbf{x}}{d\mathbf{p}} &= - \left( \frac{\partial^2 W(\mathbf{x}, \mathbf{p})}{\partial \mathbf{x}^2} \right)^{-1} \frac{\partial^2 W(\mathbf{x}, \mathbf{p})}{\partial \mathbf{x} \partial \mathbf{p}} \end{aligned} \quad (3.19)$$

We see that  $\frac{\partial^2 W(\mathbf{x}, \mathbf{p})}{\partial \mathbf{x}^2}$  is the *Hessian* of the elastic energy, and by factorizing it, we could compute the derivative  $\frac{d\mathbf{x}}{d\mathbf{p}}$  by solving a series of systems of equation, one for every parameter in  $\mathbf{p}$ . However, there is a more efficient way: the *adjoint method*. We can insert Equation (3.19) into Equation (3.18), and get

$$\frac{dg(\mathbf{p}, \mathbf{x})}{d\mathbf{p}} = \frac{\partial g(\mathbf{p}, \mathbf{x})}{\partial \mathbf{p}} - \frac{\partial g(\mathbf{p}, \mathbf{x})}{\partial \mathbf{x}} \left( \frac{\partial^2 W(\mathbf{x}, \mathbf{p})}{\partial \mathbf{x}^2} \right)^{-1} \frac{\partial^2 W(\mathbf{x}, \mathbf{p})}{\partial \mathbf{x} \partial \mathbf{p}}. \quad (3.20)$$

We can now see that we multiply the inverse matrix by a vector on the left side and a matrix on the right side. Instead of solving a system of linear equations for a number of right-hand side, we can instead solve a single system of equations and arrive at the same result. We first solve

$$\frac{\partial^2 W(\mathbf{x}, \mathbf{p})}{\partial \mathbf{x}^2} \mathbf{y} = \frac{\partial g(\mathbf{p}, \mathbf{x})}{\partial \mathbf{x}}^T \quad (3.21)$$

## Fundamentals

for  $\mathbf{y}$ , and then compute the derivative as

$$\frac{dg(\mathbf{p}, \mathbf{x})}{d\mathbf{p}} = \frac{\partial g(\mathbf{p}, \mathbf{x})}{\partial \mathbf{p}} - \mathbf{y}^T \frac{\partial^2 W(\mathbf{x}, \mathbf{p})}{\partial \mathbf{x} \partial \mathbf{p}}. \quad (3.22)$$

# Controlling Elasticity through Microstructures

The functional design of objects often requires precise control over how the object deforms under given loads, and consequently control of its spatially varying elastic properties. However, even modern additive fabrication methods offer a limited palette of materials, and many can only use a single material at a time. Structured materials can alleviate this problem by manipulating the elastic properties by means of geometry, creating a *meta-material*—assemblies of small-scale *microstructures* that exist at a scale much smaller than the object they fill.

While many manually-designed metamaterials have been proposed, creating a material with specific properties in an automatic fashion is still a difficult and time-intensive problem. Furthermore, creating spatially varying elastic properties requires that we vary the microstructure throughout an object using a large number of such structures with specific properties. We face a complex inverse problem: to determine a discrete small-scale material distribution at the resolution of the 3D printer that yields the desired macroscopic elastic behavior. Inverse problems of this type have been explored for designing periodic structures that can be tiled to synthesize homogeneous volumes, but the methods are computation-intensive and do not scale to designing non-periodic structures for objects with spatially varying material properties.

In this chapter, we propose a data-driven approach that efficiently assembles models out of precomputed small-scale structures such that the result resembles the desired local elasticity. We first build a database of tiled structures

indexed by their elastic properties. We want these structures to cover a large and ideally continuous region in the space of possible elastic behaviors. To achieve this goal, we introduce an optimization method for sampling structures that exhibit a range of desired behaviors, but are also sufficiently similar to allow interpolation.

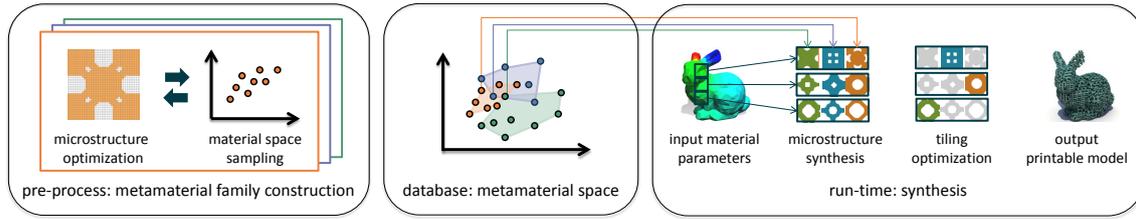
We then use this database of structures to synthesize microstructures with specific elastic properties, and tile the interior of an object to create a metamaterial. We take into account that the same elastic behavior can be reproduced by structures that differ significantly in shape. This might create multiple candidate structures with the same elastic properties, and we propose an efficient global optimization algorithm that selects an optimized tiling for spatially varying structures that takes into account the geometric compatibility between neighboring structures.

### 4.1 Overview

The goal of our system is to automatically convert an object with given spatially varying elastic properties into a 3D printable representation that requires only a single base material for fabrication, and mimics the desired elastic behavior. In this chapter, we limit ourselves to small strains and demonstrate our approach for both isotropic and anisotropic elastic materials. As outlined in Figure 4.1, our system consists of two main stages: a pre-processing step that constructs metamaterial structures covering the space of reproducible material properties, and a synthesis stage that uses those structures to generate microstructures for a given object. We start the description by defining the most important terms we use throughout the paper.

**Metamaterial Space** A *metamaterial space* is a specific organization of metamaterials. We target reproducing elastic behavior using 3D printers and represent their behavior using the  $n$  parameters of the underlying constitutive model. For example, linear isotropic materials are represented by  $n = 2$  parameters, the Young's modulus and Poisson's ratio. In addition, every metamaterial space has a mapping function, mapping the  $n$  parameters to one or several microstructures. This mapping is one-to-many because different microstructures might reproduce the same elastic behavior.

**Metamaterial Space Construction** Our first stage aims to define a function that efficiently maps a given elastic material to a small, tileable structure with the same bulk behavior. Our process starts by systematically sampling



**Figure 4.1:** An overview of our system. In a preprocessing step, we compute metamaterial families. Each family consists of multiple related microstructures that can be interpolated to smoothly vary the material properties. We store all families in a database, representing our metamaterial space. Given as input an object with specified material parameters, we synthesize locally microstructures that resemble the desired deformation behavior. As multiple, topologically different structures can have the same bulk behavior, we potentially have multiple candidates of microstructures for a single location. Using these candidates, we compute an optimized tiling, ensuring that neighboring structures are properly connected. Finally, the physical prototype is 3D printed.

the material parameter space. Our strategy is to compute a sparse set of representative structures that (i) cover a wide range in the space of elastic material parameters and that (ii) allow the interpolation of neighboring structures in parameter space. Via a weighted combination of these samples, we then reconstruct a continuous mapping from elasticity parameters to microstructures. We call such a set of interpolatable structures a *metamaterial family*. A metamaterial family defines a one-to-one mapping from material properties to structure. In practice, a family usually only covers a partial gamut of the material space. Therefore, we precompute several families until we sufficiently cover a desired range of elastic behavior. All metamaterial families together constitute our metamaterial space. Note that, as shown in Figure 4.9, the gamut of these metamaterial families might be partially overlapping, yielding a one-to-many mapping from parameters to microstructures.

**Synthesis** To synthesize the microstructure for a given object, we tile its interior. For each tile, we interpolate a microstructure from each precomputed family of structures. This provides us with a small set of candidate structures for each cell, out of which we have to select exactly one. These choices are not independent; they must be made consistently so that the structures connect well with their neighboring tiles. We suggest a carefully designed metric that quantifies the compatibility of structures and phrase this selection as a combinatorial problem, which we solve efficiently using

an optimization method based on message passing. Finally, we use the selected structures to fabricate the object using rapid prototyping.

## 4.2 Background

In order to determine and optimize for the behavior of a microstructure, we rely on a physical model. We use a standard linear elastic model (see Section 3.2), which we will later extend with properties specific to microstructure simulation and topology optimization.

We want to reiterate here that the linear material stiffness tensor,  $\mathbf{C}(\mathbf{p})$ , is a function of a small set of material parameters  $\mathbf{p}$  that depend on the material model. For example, an isotropic material has two parameters, the Young's modulus and the Poisson's ratio. For more complex materials, additional parameters such as the shear modulus and direction-dependent Young's moduli and Poisson's ratios are used.

### 4.2.1 Numerical Coarsening

Optimizing a microstructure is an inverse problem, corresponding to the forward problem of determining the coarse-scale behavior from the microstructure. This forward problem can be defined using the idea of homogenization: compute a material stiffness tensor for a homogeneous material whose elastic behavior matches that of the tiled microstructure. We use the *Numerical Coarsening* approach [Kharevych et al., 2009], which uses a set of load cases to approximate the coarse elastic behavior of a given structure. Essentially, given the deformations  $\mathbf{h}$  that these load cases induce, which are called *harmonic displacements*, the method computes a single material stiffness tensor  $\mathbf{C}(\mathbf{h})$  that describes the homogenized material behavior of a microstructure, which we will use to solve the inverse problem. We refer to Appendix A.1.1 for a more detailed summary to the Numerical Coarsening approach.

## 4.3 Microstructure Optimization

Our microstructure optimization method solves the inverse problem to the Numerical Coarsening method mentioned in the previous section, solving for a microstructure that coarsens to a given stiffness tensor.

Optimizing a microstructure requires a way to define and alter the material distribution within a cell. A common approach in topology optimization is to discretize the material distribution by subdividing the cell into a grid of material voxels [Sigmund, 2009], where each voxel is associated with a binary *activation* that describes whether the voxel is full (1) or void (0). However, optimizing the microstructure using these binary variables directly would be infeasible for moderately large grids. Instead, the problem is usually relaxed by allowing the activations to vary smoothly between 0 and 1 during the optimization, and only requiring them to converge to a binary solution at the end of the optimization. For the continuous activations, a meaningful interpolation between void and full voxels has to be defined such that the activation corresponds to a physical quantity in the simulation. A simple way to define this is by interpolating between stiffness tensors. For any voxel  $i$  ( $1 \leq i \leq m$ , with  $m$  being the number of voxels), an individual material stiffness tensor  $\mathbf{C}_i$  is defined as an interpolation between the base material stiffness tensor  $\mathbf{C}_{base}$  and air, which is assumed to have a zero material stiffness tensor:

$$\mathbf{C}_i = \alpha_i \mathbf{C}_{base}. \quad (4.1)$$

To ensure numerical stability, the minimum of  $\alpha_i$  is set to  $\alpha_{min} = 10^{-5}$ . This interpolation scheme follows the established *solid isotropic material with penalization* (SIMP) approach for an exponent of 1 [Sigmund, 2009]. Choosing a different exponent would help to converge to a binary solution in topology optimization problems with extremal objectives, where adding more material improves the objective and the maximum amount of material is fixed by a constraint. However, we do not have such an objective, and have to resort to other means to reach a binary solution. As a consequence, the exponent we choose does not influence the convergence.

The number of activations can be reduced by exploiting symmetries of the goal material. For example, for a cubic material, the response along each principal axis has to be identical, and we mirror the activations along all axes and all diagonal planes accordingly.

#### 4.3.1 Problem Formulation

We pose the problem of finding a microstructure that exhibits a large-scale behavior identical to a homogeneous material with desired material parameters  $\mathbf{p}_{goal}$  (see Section 4.2) as a least squares problem. From the parameters  $\mathbf{p}_{goal}$ , a stiffness tensor  $\mathbf{C}^{goal} = \mathbf{C}(\mathbf{p}_{goal})$  can be computed. The optimization

then modifies the vector of activations  $\alpha = [\alpha_1 \cdots \alpha_m]$  such that the homogenized stiffness tensor  $\mathbf{C}(\mathbf{h}(\alpha))$ , which is indirectly dependent on the activations through the harmonic displacements  $\mathbf{h}(\alpha)$ , matches the goal stiffness tensor as closely as possible:

$$\begin{aligned} \min_{\alpha} \quad & \|\mathbf{C}^{goal} - \mathbf{C}(\mathbf{h}(\alpha))\|_F^2 + R \\ \text{s.t.} \quad & \alpha_{min} \leq \alpha_i \leq 1 \quad 1 \leq i \leq m. \end{aligned} \quad (4.2)$$

Here,  $R$  is a combined regularization term that penalizes less desirable results. This formulation differs from most other microstructure optimization approaches that typically try to find extremal properties for a specific amount of material. It is related to the formulation in [Zhou and Li, 2008], though it does not use a volume fraction constraint. Panetta et al. [2015] also use a similar objective, but measure the error through the material compliance tensor instead of the material stiffness tensor.

### 4.3.2 Regularization

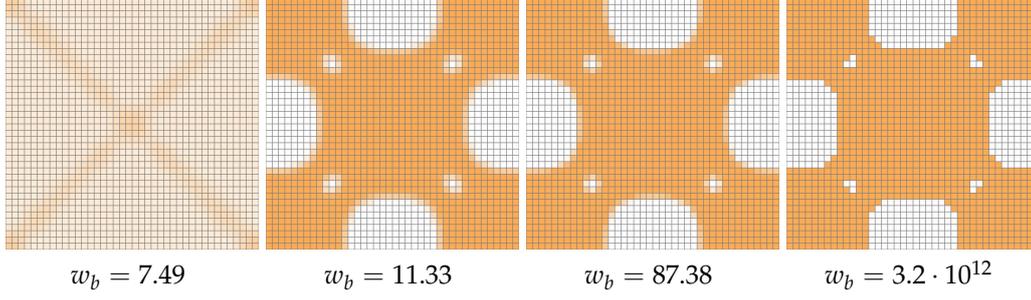
While the optimization problem (4.2) could be solved without any regularization, there is no guarantee that the result can be fabricated. In order to enforce manufacturability, we add three different regularization terms  $R_b$ ,  $R_s$  and  $R_{cb}$  with corresponding weights  $w_b$ ,  $w_s$  and  $w_{cb}$  to the objective. The combined regularization term  $R$  is defined as

$$R = w_b R_b + w_s R_s + w_{cb} R_{cb}. \quad (4.3)$$

Figure 4.2 and 4.3 show the influence of the individual regularization terms. In the following, we will elaborate on each term.

**Enforcing Binary Values** While the simulation allows the activations  $\alpha$  to vary freely between  $\alpha_{min}$  and 1, these configurations do not correspond to valid physical objects. For fabrication, the activations have to be binary (excluding the small offset  $\alpha_{min}$  to ensure numerical stability). In order to reach such a solution, the regularization term  $R_b$  acts as a penalty for activations that are not equal to either  $\alpha_{min}$  or 1:

$$R_b = \sum_{i=1}^m (\alpha_i - \alpha_{min})(1 - \alpha_i). \quad (4.4)$$



**Figure 4.2:** Optimization at different stages (with different penalty weights to force the activations towards  $\alpha_{min}$  or 1).

We gradually increase the weight  $w_b$  during the optimization, transitioning from a continuous to a discrete solution. For each value of the weight, a full optimization is run until convergence is reached. If the solution is not binary, the weight is increased and the optimization resumed. Figure 4.2 shows different stages of the optimization, with various weights  $w_b$ .

**Smoothness** The size of a single microstructure cell in a fabricated object is largely defined by two factors: the resolution of the 3D printer, and the size of the smallest detail in the structure. Smaller cells provide a better approximation of a continuous material, and since the printer resolution is assumed to be fixed, structures without small details are generally preferred.

The regularization  $R_s$  views the activations as an approximation of a material distribution field, and uses a second-order finite difference approach to penalize deviations from smoothness. For this purpose, any component of  $\alpha$  is assumed to have two indices in 2D, such that  $\alpha_{i,j}$  corresponds to the voxel  $(i, j)$ . The regularization in 2D then has the form

$$R_s = \sum_{i,j} (\alpha_{i-1,j} + \alpha_{i+1,j} + \alpha_{i,j-1} + \alpha_{i,j+1} - 4\alpha_{i,j})^2. \quad (4.5)$$

In 3D, we instead use six neighbors to compute the second-order finite difference approximation. Assuming each component of  $\alpha$  is associated with three indices, such that  $\alpha_{i,j,k}$  corresponds to the voxel  $(i, j, k)$ , the regularization has the form

$$R_s = \sum_{i,j,k} (\alpha_{i-1,j,k} + \alpha_{i+1,j,k} + \alpha_{i,j-1,k} + \alpha_{i,j+1,k} + \alpha_{i,j,k-1} + \alpha_{i,j,k+1} - 6\alpha_{i,j,k})^2 \quad (4.6)$$

**Checkerboard Patterns** An artifact that frequently appears in topology optimization is elements that are connected by a single vertex, called *checkerboard patterns* [Sigmund and Petersson, 1998]. To avoid such structures, the regularization  $R_{cb}$  penalizes configurations that contain checkerboard patterns, as illustrated in Figure 4.3. In 2D, this regularization is based on  $2 \times 2$  patches of voxels and has the form

$$R_{cb} = \sum_{i,j} (1 - \alpha_{i,j})(\alpha_{i+1,j} - \alpha_{min})(\alpha_{i,j+1} - \alpha_{min})(1 - \alpha_{i+1,j+1}) + (\alpha_{i,j} - \alpha_{min})(1 - \alpha_{i+1,j})(1 - \alpha_{i,j+1})(\alpha_{i+1,j+1} - \alpha_{min}). \quad (4.7)$$

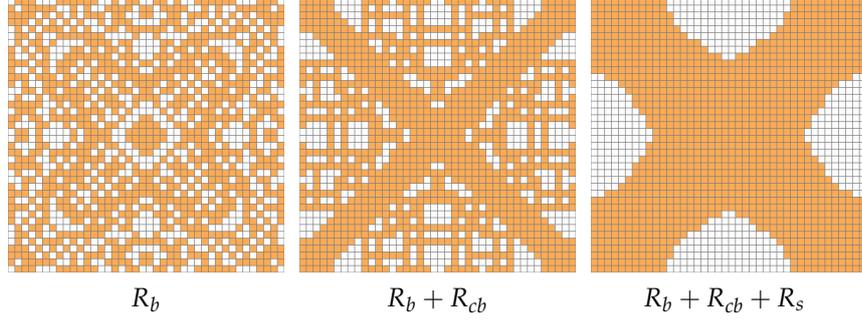
In the case of binary activations,  $R_{cb}$  is only non-zero if the structure contains a checkerboard pattern. In the continuous case, the regularization also acts as an additional regularizer that pushes the activations towards  $\alpha_{min}$  or 1.

In 3D, the number of different local checkerboard patterns increases. Checkerboard patterns now include structures that are connected by a single vertex or a single edge. To cover these two cases, we will split the regularization term in two components,  $R_{cb,v}$  and  $R_{cb,e}$ , respectively.

To check for structures connected by a single vertex,  $R_{cb,v}$  has to cover patches of  $2 \times 2 \times 2$  voxels. For these patches, there are four configurations that are undesirable and will not be covered by  $R_{cb,e}$ . These are the only configurations for which a binary solution should lead to a regularization value larger than 0. This condition can be formulated as

$$R_{cb,e} = \sum_{i,j,k} (\alpha_{i,j,k} - \alpha_{min})(1 - \alpha_{i+1,j,k})(1 - \alpha_{i,j+1,k})(1 - \alpha_{i,j,k+1}) \\ (1 - \alpha_{i+1,j+1,k})(1 - \alpha_{i+1,j,k+1})(1 - \alpha_{i,j+1,k+1})(\alpha_{i+1,j+1,k+1} - \alpha_{min}) \\ + (1 - \alpha_{i,j,k})(\alpha_{i+1,j,k} - \alpha_{min})(1 - \alpha_{i,j+1,k})(1 - \alpha_{i,j,k+1}) \\ (1 - \alpha_{i+1,j+1,k})(1 - \alpha_{i+1,j,k+1})(\alpha_{i,j+1,k+1} - \alpha_{min})(1 - \alpha_{i+1,j+1,k+1}) \\ + (1 - \alpha_{i,j,k})(1 - \alpha_{i+1,j,k})(\alpha_{i,j+1,k} - \alpha_{min})(1 - \alpha_{i,j,k+1}) \\ (1 - \alpha_{i+1,j+1,k})(\alpha_{i+1,j,k+1} - \alpha_{min})(1 - \alpha_{i,j+1,k+1})(1 - \alpha_{i+1,j+1,k+1}) \\ + (1 - \alpha_{i,j,k})(1 - \alpha_{i+1,j,k})(1 - \alpha_{i,j+1,k})(\alpha_{i,j,k+1} - \alpha_{min}) \\ (\alpha_{i+1,j+1,k} - \alpha_{min})(1 - \alpha_{i+1,j,k+1})(1 - \alpha_{i,j+1,k+1})(1 - \alpha_{i+1,j+1,k+1}) \quad (4.8)$$

Structures that are connected by a single edge can be detected by looking at patches of  $2 \times 2 \times 1$  voxels, similar to the regularization in 2D. Considering all the orientations of this patch, this condition can be written as



**Figure 4.3:** *Influence of different regularizations: Optimization results where only binary values are enforced (left), with additional anti-checkerboard regularization (middle) and with smoothness regularization (right). The objective value for the last two results is similar, while the first result has a worse objective value.*

$$\begin{aligned}
 R_{cb,v} = & \sum_{i,j,k} (1 - \alpha_{i,j,k})(\alpha_{i+1,j,k} - \alpha_{min})(\alpha_{i,j+1,k} - \alpha_{min})(1 - \alpha_{i+1,j+1,k}) \\
 & + (\alpha_{i,j,k} - \alpha_{min})(1 - \alpha_{i+1,j,k})(1 - \alpha_{i,j+1,k})(\alpha_{i+1,j+1,k} - \alpha_{min}) \\
 & + (1 - \alpha_{i,j,k})(\alpha_{i+1,j,k} - \alpha_{min})(\alpha_{i,j,k+1} - \alpha_{min})(1 - \alpha_{i+1,j,k+1}) \\
 & + (\alpha_{i,j,k} - \alpha_{min})(1 - \alpha_{i+1,j,k})(1 - \alpha_{i,j,k+1})(\alpha_{i+1,j,k+1} - \alpha_{min}) \\
 & + (1 - \alpha_{i,j,k})(\alpha_{i,j+1,k} - \alpha_{min})(\alpha_{i,j,k+1} - \alpha_{min})(1 - \alpha_{i,j+1,k+1}) \\
 & + (\alpha_{i,j,k} - \alpha_{min})(1 - \alpha_{i,j+1,k})(1 - \alpha_{i,j,k+1})(\alpha_{i,j+1,k+1} - \alpha_{min}).
 \end{aligned} \tag{4.9}$$

**Regularization Weights** The performance of our microstructure optimization depends on the choice of weights, and how they are updated during the optimization. We list here the initial values for these weights for the 2D and 3D optimization, as well as the update that is applied once the optimization converged for the current weights:

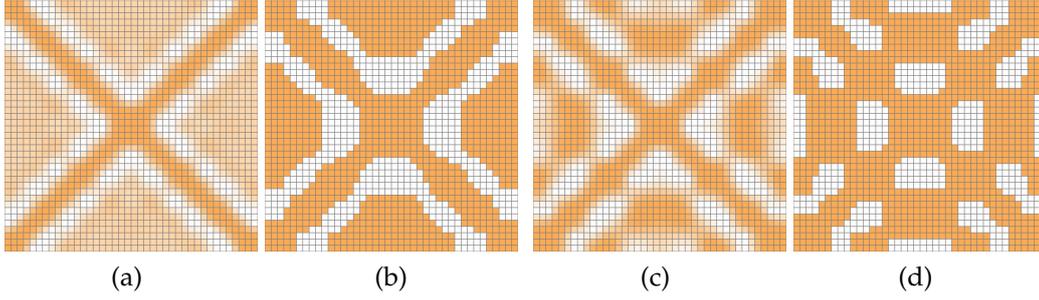
$$\begin{array}{lll}
 \text{2D} & w_b^0 = 0 & w_b^{t+1} = 1.3w_b^t + 0.1 \\
 & w_s^0 = 2 & w_s^{t+1} = 1.1w_s^t + 0.2 \\
 & w_{cb}^0 = 0 & w_{cb}^{t+1} = 5w_b^{t+1} \\
 \\ 
 \text{3D} & w_b^0 = 0 & w_b^{t+1} = 1.3w_b^t + 0.5 \\
 & w_s^0 = 5 & w_s^{t+1} = 1.1w_s^t + 1 \\
 & w_{cb}^0 = 0 & w_{cb}^{t+1} = 5w_b^{t+1}
 \end{array}$$

**Connectivity** An additional fabrication requirement is connectivity. The optimization will generally not favor binary solutions in the absence of regularization term  $R_b$ . As the influence of this term grows with increasing  $w_b$ , the previously intermediate activations will be pushed to  $\alpha_{min}$  or 1, and the structure might become disconnected (see Figure 4.4). To prevent the optimization from converging to such a solution, disconnected components are detected after every iteration. To account for the continuous nature of the activations, every activation below a threshold of 0.1 is considered inactive during the detection. If a disconnected component has been found, we compute the cost of connecting the component as the smallest change in activations that builds a connection, assuming that we set the activations to a value of 0.6. If this cost is smaller than the change in activations necessary to remove the disconnected component, we create the connection, and remove the disconnected component otherwise. The final result is then guaranteed to be connected.

In 3D, an additional fabrication constraint is necessary. While 3D printing can handle complex structures, most approaches rely on *support material* to create overhanging structures. This support material has to be removed after printing. This means that every void voxel in the structure has to be connected to the boundary of the cell. To this end, we use the same approach we used to connect components, but switch the role of void and full voxels. In practice, we did not observe any convergence problems due to these constraints.

### 4.3.3 Numerical Methods

The optimization problem (4.2) is solved with L-BFGS-B [Byrd et al., 1995], which enforces the boundary constraints. Additionally, the indirect relationship between the coarsened stiffness tensor  $\mathbf{C}(\mathbf{h}(\alpha))$  and the activations  $\alpha$



**Figure 4.4:** *Influence of the connectivity enforcement: Without using any connectivity enforcement during the optimization (a), the final structure might consist of several disconnected components (b). Enforcing connectivity with our scheme locally adjusts the activations (c) such that the final result is guaranteed to be fully connected (d).*

through the harmonic displacements  $\mathbf{h}(\alpha)$  defined in Equation (A.1) has to be taken into account when computing the derivatives of the microstructure optimization problem introduced in the paper. When the chain rule is applied to this problem, the Jacobian of the harmonic displacements with respect to the activations,  $\frac{d\mathbf{h}_{ab}}{d\alpha}$ , emerges. Since these quantities are implicitly linked by the solution of an elasticity problem, given the boundary tractions defined in (A.1), we use the implicit function theorem to compute the Jacobian (see Section 3.4). For this, we take the derivatives of both the minimization condition  $\nabla_{\mathbf{x}} U_{el} = 0$  and the constraints  $\mathbf{c} = 0$  with respect to  $\alpha$ :

$$\begin{aligned} \frac{\partial^2 U_{el}}{\partial \mathbf{h}_{ab} \partial \alpha} + \frac{\partial^2 U_{el}}{\partial \mathbf{h}_{ab}^2} \frac{d\mathbf{h}_{ab}}{d\alpha} &= 0 \\ \frac{\partial \mathbf{c}}{\partial \mathbf{h}_{ab}} \frac{d\mathbf{h}_{ab}}{d\alpha} &= 0. \end{aligned} \quad (4.10)$$

Solving this system of equations for the desired Jacobian  $\frac{d\mathbf{h}_{ab}}{d\alpha}$  requires only a single sparse matrix decomposition.

## 4.4 Metamaterial Spaces

The optimization method from the previous section is able to produce microstructures for a variety of material parameters, but may take a long time to generate a desired structure. Moreover, if the desired object contains spatially varying parameters, several optimizations need to be performed to

generate the required microstructures, making this approach infeasible in practice. To avoid this problem, we use a data-driven approach to assemble a structure with a desired behavior by interpolation from a pre-computed metamaterial family.

A metamaterial family is a collection of microstructures, each labeled with its corresponding coarse-scale material parameters, which are a point in the space of possible material properties (parameter space); we think of the microstructure as being “located at” that point. In this section, we first describe how the structures are represented and how their locations in parameter space are determined. We then propose a technique to interpolate between these structures, providing a way to efficiently compute a corresponding microstructure for any point in the parameter space. Finally, we discuss how to use the microstructure optimization approach described in the previous section to create metamaterial families by computing sets of related microstructures that cover a wide range in parameter space.

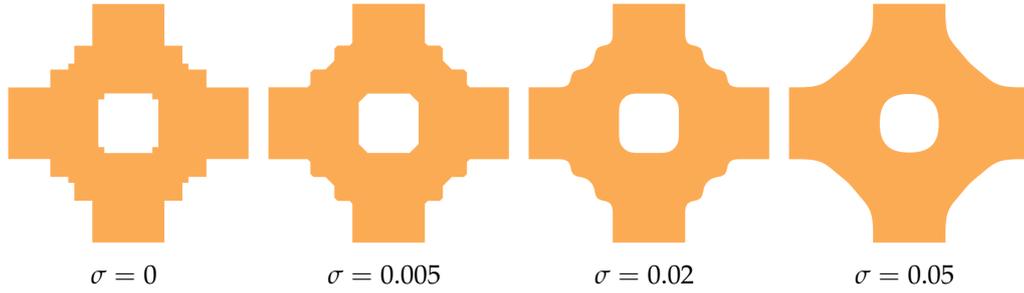
In Section 4.5 we explain how several such continuous families (collectively, a metamaterial space) are used together to assemble models spanning the entire range of achievable material properties.

### 4.4.1 Microstructure Representation

**Geometry** While the voxel-based discretization we used during the microstructure optimization is useful in that context, it poses two problems for the construction of a metamaterial space: (i) The changes to the structure are all discrete in nature, so the resulting interpolation cannot be continuous, and (ii) the resulting geometry can contain stair structures. These sharp corners can lead to localized stresses under deformation and the structure would fracture more easily.

Instead of using voxels, we use signed  $L_2$ -distance fields in  $[0, 1]^d$  to represent structures in a metamaterial space, which allows for a smooth interpolation. Additionally, we perform a Gaussian smoothing step every time we sample a microstructure from the metamaterial space, which removes unwanted discretization artifacts (Figure 4.5). To increase resolution, we scale the grid resolution by a factor of 2 and 3 compared to the original voxelization, in 2D and 3D, respectively.

**Material Parameters** Numerical Coarsening is used to compute a stiffness tensor that describes the behavior of a particular microstructure. However, for sampling and interpolation we would like to use a parameter space with



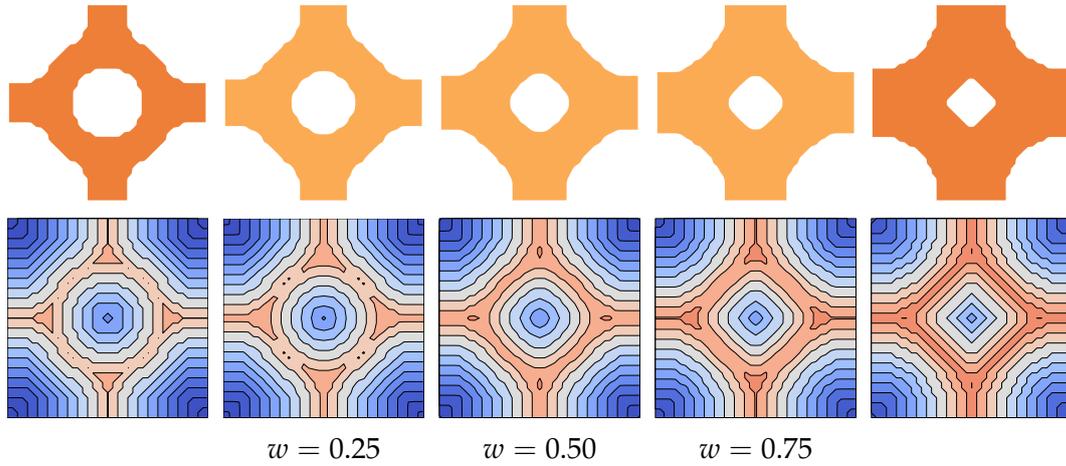
**Figure 4.5:** Results of the smoothing pass for different Gaussian spread values  $\sigma$ .

fewer degrees of freedom than this tensor has (6 in two dimensions, 21 in three dimensions). By considering only isotropic, cubic or orthotropic materials, the parameter space can be reduced to a subset of material tensors, while still offering enough freedom to show a large variety of deformation behaviors.

To approximate the parameters from a material stiffness tensor computed by Numerical Coarsening, a simple constrained least-squares approach is used. Assuming a material stiffness tensor  $\hat{\mathbf{C}}$  is given, the corresponding parameters are computed by solving

$$\begin{aligned} \min_{\mathbf{p}} \|\mathbf{C}(\mathbf{p}) - \hat{\mathbf{C}}\|_F^2 \\ \text{s.t. } p_i^{\min} \leq \mathbf{p}_i \leq p_i^{\max}. \end{aligned} \quad (4.11)$$

The function  $\mathbf{C}(\mathbf{p})$  is defined by the choice of the material model, and  $p_i^{\min}$  and  $p_i^{\max}$  are the physics-based bounds on the parameters, such as the lower limit of 0 for the Young's modulus. We additionally apply a normalization to the parameters that transforms all parameters with dimensions, such as the Young's modulus or the shear modulus, into dimensionless parameters by dividing them by the corresponding parameter of the base material. Since we use a linear material model, the resulting parameters describe the ratio of the structure's parameter to the base material's parameter, which also allows us to scale our results to materials with an arbitrary Young's modulus or shear modulus, assuming that the Poisson's ratio stays the same. Additionally, since we use this relative Young's modulus and relative shear modulus to store entries in parameter space, the distance between the parameters of two microstructures is independent of the base material's stiffness.



**Figure 4.6:** Top row: Results of the interpolation between the structures on the left and right for different weights  $w$ . Bottom row: Corresponding signed distance fields, illustrated as color plots (blue positive, red negative).

#### 4.4.2 Interpolation

The microstructures in our database describe metamaterials with certain properties; each gives a point sample of the mapping from material parameters to microstructures. Figure 4.11, 4.9, and 4.10 illustrate data points for various metamaterial families. To generate a structure for an arbitrary given set of parameters, we interpolate between points of a family, forming a weighted average over a set of microstructures with similar elastic properties. We first compute weights based on the inverse distance between the input parameters and the parameters of the metamaterial space samples, using the Wendland function with compact support [Nealen, 2004]. We chose the parameter of the Wendland function such that the weights vanish beyond a given interpolation radius, which is set to 0.1 in normalized coordinates, or the distance to the  $(M + 1)$ -nearest neighbor ( $M$  being the number of material parameters), whichever is larger. Before we interpolate, we apply the transformation  $f(x) = \text{sgn}(x) \log(|x| + \delta)$  to transform the distance fields to logarithmic space, and add a small constant  $\delta = 10^{-3}$  to keep values near the surface. In practice, we found that interpolation in logarithmic space reduces artifacts due to topology changes, e.g., holes appearing or disappearing. Given the weights and transformed distance fields, we then compute the interpolated structure by linearly interpolating the transformed distance fields. Figure 4.6 illustrates the interpolation process.

### 4.4.3 Generating Metamaterial Families

Our metamaterial space consists of several, potentially overlapping, independent metamaterial families. We start the construction of a metamaterial family from a single microstructure, which we either model by hand based on existing examples from the literature, or obtain from our microstructure optimization (Section 4.3). We then add dilated and eroded versions of this initial microstructure to achieve a large initial sampling of the metamaterial space with good interpolation properties. The dilation and erosion are performed directly on the distance field of the structure.

The metamaterial family construction then continues to do two things: (i) Generate new candidates by *evolving existing structures*, which allows us to refine regions that are already covered by samples, but with insufficient sampling density, and (ii) generate new candidates by *optimizing for new microstructures* outside of this space. We proceed by alternating between these two stages until both stages do not generate new microstructures.

**Evolving Existing Microstructures** The first stage is a heuristic search based on the existing samples in the metamaterial family. For the refinement, we create the Delaunay triangulation of all parameter points, and collect all the simplex centers. For every center, we check how well the parameters of the interpolated structure match the desired parameters, and add the interpolated structure to the metamaterial space if the deviation from the desired parameters is too large. For our experiments, we used a threshold of 0.1. In 3D, we additionally try to expand the metamaterial family in this stage. For this, we first compute the convex hull of all parameter samples. For every vertex on the convex hull, we then compute an offset point along the normal, which we use as the goal parameters for a microstructure optimization. However, instead of running the optimization, we only compute the gradient of the objective and use it to change 2% of the activations in a discretized version of the current sample, which we then add to the metamaterial family.

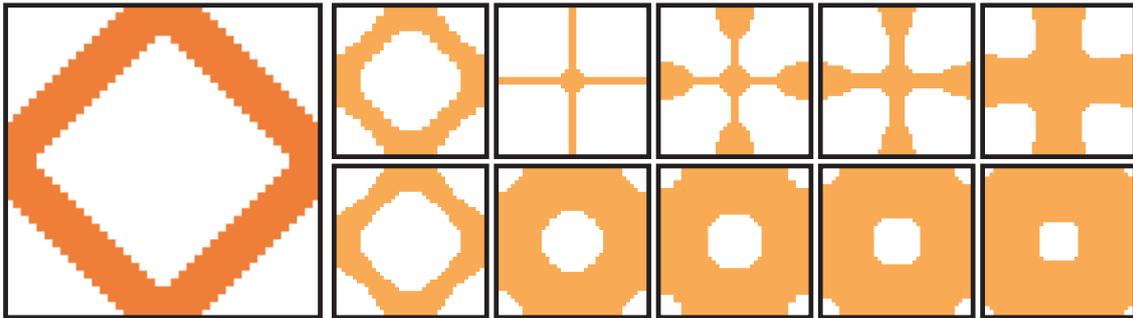
**Optimizing for New Microstructures** The second stage is based on the microstructure optimization introduced in Section 4.3. For the refinement, we again compute the Delaunay triangulation and check how well the interpolation works at the simplex centers. If it is insufficient, we run a microstructure optimization for the parameters at the simplex center with an initial guess computed from the weighted combination of all samples in the neighborhood of the simplex center. Additionally, we introduce a similarity

regularization, explained in more detail in the next paragraph. For the expansion, we again use the convex hull of all parameter samples and generate new parameter points by sampling the convex hull and offsetting the points along the normal. For each of these points, a microstructure optimization is run, constructing the initial guess in the same way as for the refinement and using the same additional regularization.

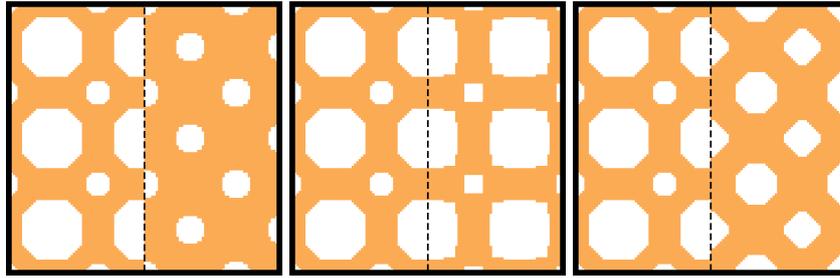
For any new microstructure optimization that is run for a given set of parameters, the result should be similar to the existing structure to improve the interpolation. To lead the optimization into the desired direction, an additional regularization is added to the optimization. This regularization penalizes the amount of change between a new structure and the structure of the  $N$  neighbors in parameter space closest to the goal parameters. To allow for small changes, this penalty uses an exponential function:

$$R_{sim} = \sum_{i=1}^N w_{sim}^i e^{\left( \frac{\frac{1}{m} \sum_{j=1}^m |\alpha_j - \alpha_j^i|}{\Delta \alpha_i} \right)^2}, \quad (4.12)$$

where  $w_{sim}^i = (N(1 + d_i))^{-1}$  is the weight for neighbor  $i$  with distance  $d_i$  to the desired parameters,  $\alpha_j^i$  is its  $j$ -th activation and  $\Delta \alpha_i = 0.1 + 2d_i$  is a threshold for the maximal desirable difference to this neighbor. Additionally, the initial activations for the optimization are set to a weighted average of the neighbors, after a smoothness step is applied. This reduces the risk of ending up in a local minimum.



**Figure 4.7:** The first 5 microstructures generated for a parameter space without similarity regularization (top) and with similarity regularization (bottom), starting from an initial structure (left).



**Figure 4.8:** *Three potential configurations of two neighboring cells. Left: The individual structures closely match the desired parameters, but the boundaries are incompatible. Middle: Opposite case. Right: Our optimization computes a trade-off between the two extremes.*

## 4.5 Structure Synthesis

Using the microstructure optimization and parameter space sampling methods, we can define several families of related structures that together span the feasible range of bulk material parameters (Figure 4.9). Synthesizing a homogeneous material volume at this stage becomes trivial: We select a family that covers the desired material behavior, compute the corresponding microstructure of the cell by interpolation as described in Section 4.4.2, and then fill the volume by repeating this cell. Note that by construction this cell is tileable. However, approximating spatially varying materials is more challenging.

Simply synthesizing microstructures for cells independently at each point in the model could lead to a mismatching boundary when multiple different cells are tiled, as illustrated in Figure 4.8. Such boundary mismatches will change the behavior of the cell, which was assumed to be in an infinite tiling of identical structures when its coarsened material parameters were computed. Therefore, both the geometry as well as the force profiles occurring at the boundaries under deformation need to be taken into account.

We propose a strategy that takes advantage of the multiple candidate structures for each cell provided by the overlapping families in our metamaterial space. To compute an optimal selection from these candidates, we propose to minimize boundary dissimilarity between each pair of neighboring structures.

For a set of cells with desired parameters  $\mathbf{p}_1, \dots, \mathbf{p}_k$  and information about the connectivity between cells, we interpolate one structure for each cell from each of the  $l$  families, resulting in structures  $S_{i,1}, \dots, S_{i,l}$  for every parameter sample  $\mathbf{p}_i$ . Finding the optimal choice of structures can then be

formulated as a labeling problem: assign a structure to each cell to minimize a given cost function.

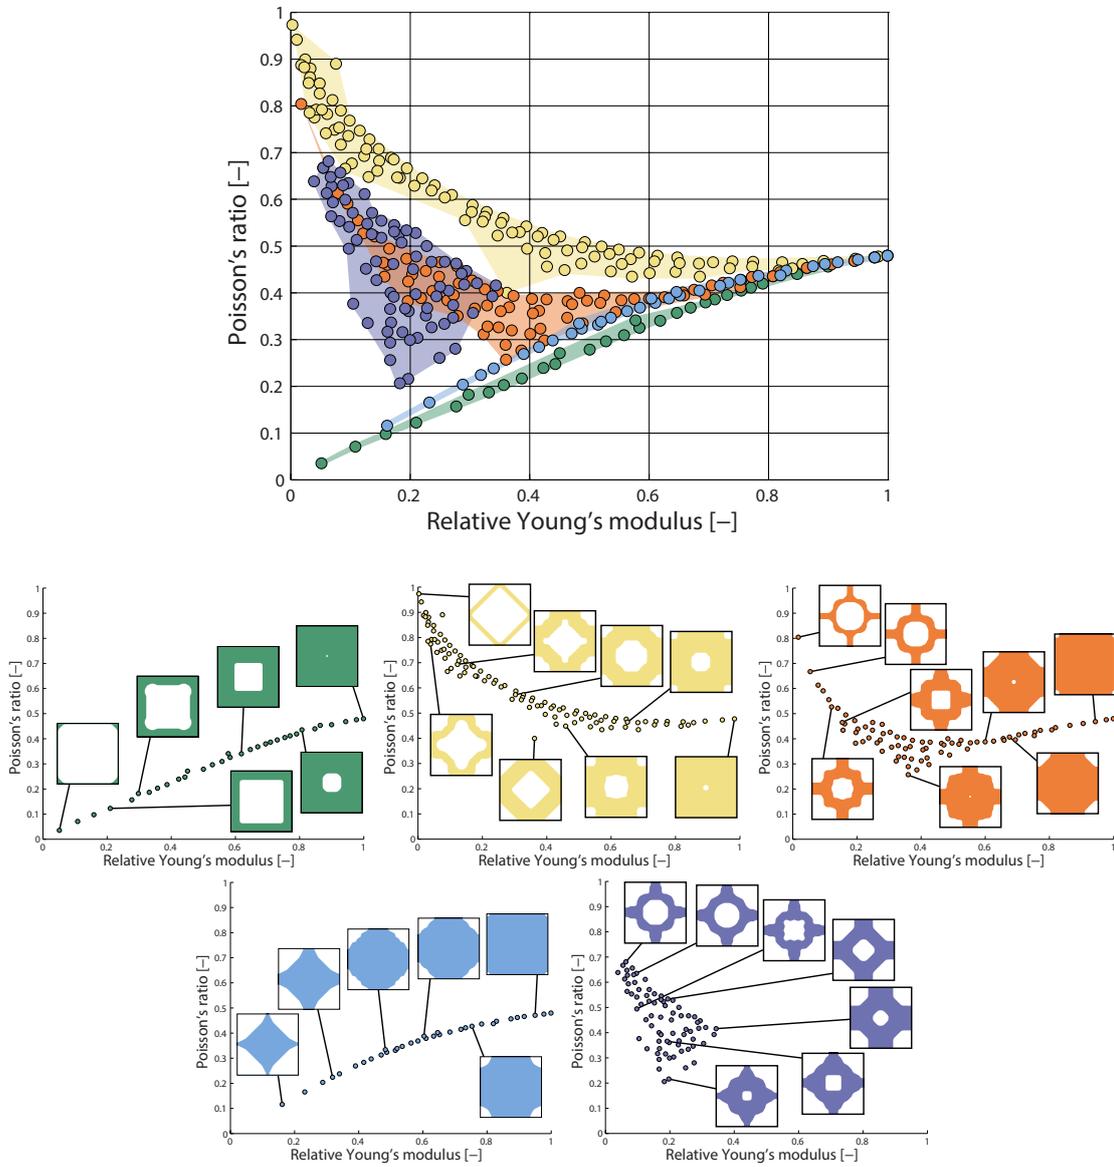
We propose a cost function that combines two types of costs. The labeling cost  $T_{i,j}^L = e^{d_{i,j}}$  describes how well a given structure  $S_{i,j}$  matches the desired parameters, and is based on the distance  $d_{i,j}$  between  $\mathbf{p}_i$  and the material parameters of the structure as computed by Numerical Coarsening.

The second cost  $T_{(i,j),(r,s)}^D = e^{g_{(i,j),(r,s)}}$  is based on the dissimilarity  $g_{(i,j),(r,s)}$  of the boundaries of two neighboring structures  $S_{i,j}$  and  $S_{r,s}$ . In most cases, it is sufficient to use the percentage of the boundary on which the two structures agree about the presence or absence of material. But since the problem of matching boundaries is linked to force discrepancies along the boundary during deformation, this dissimilarity can be improved by also considering the forces acting across the boundary. For this, we impose a unit strain on the boundary of each cell, such that they are stretched perpendicular to the boundary between the two cells. We then integrate the force magnitude as well as the force difference magnitude over the boundary, and set the dissimilarity  $g_{(i,j),(r,s)}$  to the ratio of force difference magnitude to mean force magnitude. We compare the two approaches to compute the boundary dissimilarity in Section 4.6.2.

Finding the globally optimal solution to this optimization problem is NP-hard. However, efficient algorithms exist that can find an approximate solution. We employ an iterative method using message passing based on the *alternating direction method of multipliers* (ADMM) as described in Derbinsky et al. [2013].

For the resulting structures, the distance fields can then be combined. To improve connectivity between cells, the smoothing pass is performed on the combined distance field instead of each cell individually. The final structure is reconstructed from the combined distance field using marching cubes.

**Connectivity** Our synthesis method does not guarantee connectivity. While we did not encounter cases of disconnected cells in our result, they can be detected easily and fixed by introducing additional connections between disconnected structures, at the expense of the accuracy of the approximated elastic properties.



**Figure 4.9:** *Top: Data points for five different metamaterial families for a cubic material in 2D. The values for the shear modulus are omitted. Bottom: The individual metamaterial families, including a visualization of some of the structures.*

## 4.6 Results

### 4.6.1 Metamaterial Space Construction

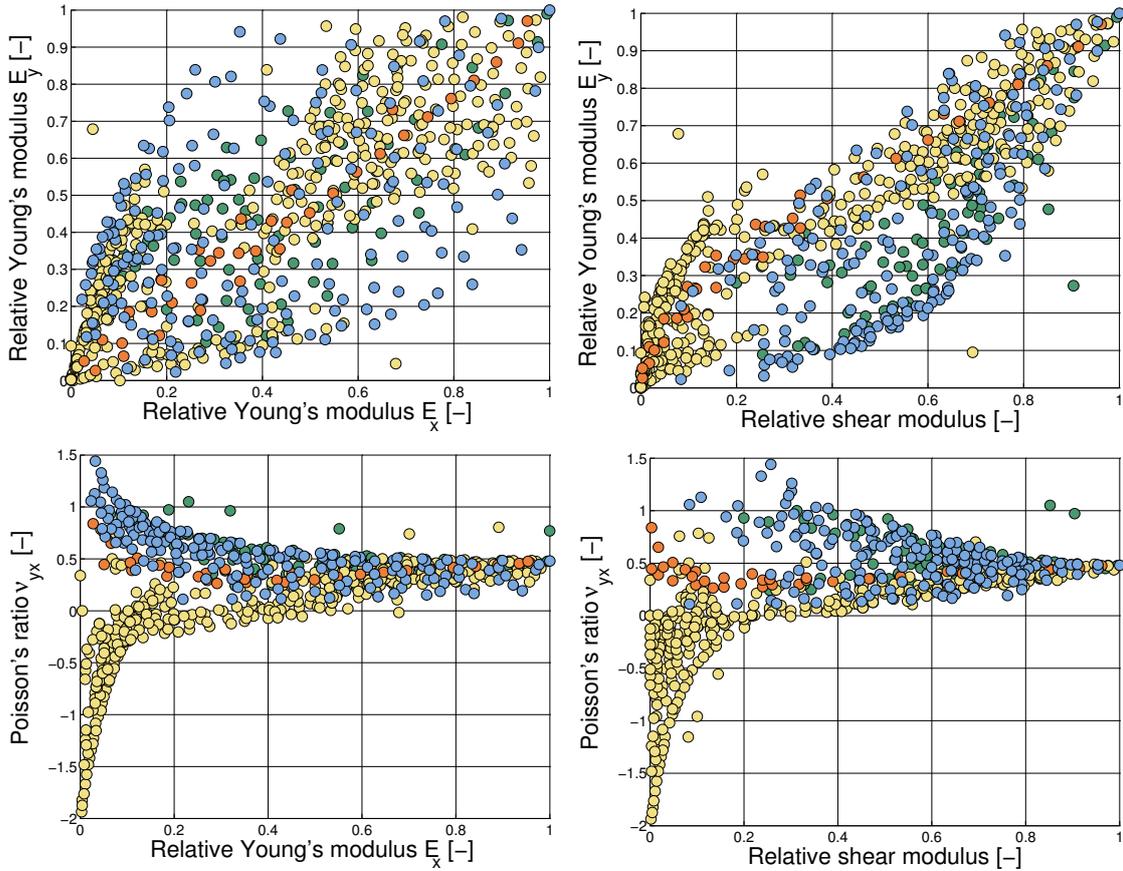
We tested our method on three different material classes. In 2D, we generated metamaterial spaces for cubic materials (3 parameters) and orthotropic materials (4 parameters), using a resolution of  $40^2$  voxels for the microstructure optimization. Due to the inherently anisotropic nature of square and cubic microstructures, we found that a cubic material space is better suited even for cases where one is only interested in the Young's modulus and Poisson's ratio defining an isotropic material.

Figure 4.9 shows the Young's modulus and Poisson's ratio of multiple metamaterial families for a cubic material. While a single family may span a wide range of parameters, this example shows that combining multiple families can significantly expand this range. Figure 4.10 shows the orthotropic metamaterial families, projected into four different combinations of the parameter axes.

We also used our method to compute a metamaterial space for cubic materials in 3D (Figure 4.11), using a  $16^3$  voxels for the microstructure optimization.

**Timings** For the microstructure optimization in 2D with a resolution of  $40^2$ , a single optimization step takes around 200 to 800 ms to compute. The optimization usually converges in fewer than 500 iterations, resulting in a total computation time of several minutes per optimization. An optimization in 3D with a resolution of  $16^3$  runs with 8 to 30 seconds per iteration, and usually also converges in fewer than 500 iterations, for a total computation time of around 2 hours per structure. Note that the metamaterial space construction is a preprocessing step that is only run once to build the database, and it can be easily parallelized.

For the structure synthesis, we tested the runtime of the optimization by combining three metamaterial spaces for different numbers of cells. For 400 cells, the optimization takes about 0.5 seconds, for 2500 cells about 2 seconds and for 10000 cells on the order of 40 seconds. Synthesizing the distance fields and running the optimization took less than 10 seconds in all of our examples.

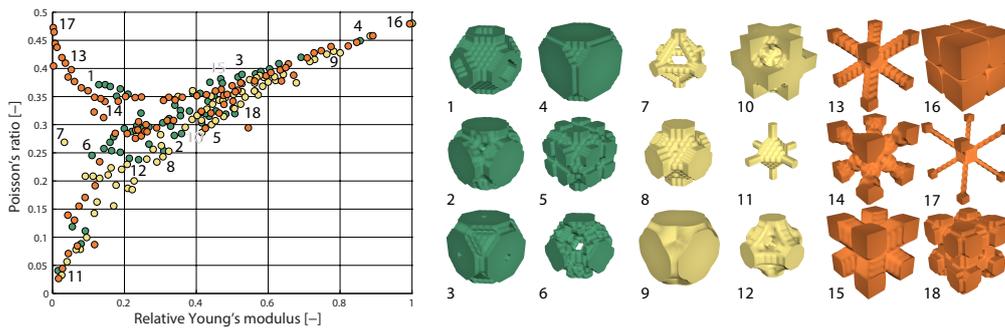


**Figure 4.10:** Data points for four different metamaterial families for an orthotropic material in 2D. We show the relative Young's moduli  $E_x$  and  $E_y$  the  $x$ -direction and  $y$ -direction, the relative shear modulus and the Poisson's ratio  $\nu_{yx}$  that describes the contraction in  $x$ -direction for an extension applied to the  $y$ -direction.

### 4.6.2 Validation

Our method relies on the ability to compute the material behavior of a microstructure from its design. To validate the results obtained by Numerical Coarsening, we tested several of our structures in a tensile test, and compared the results to our prediction. Since we are dealing with linear elasticity, the Numerical Coarsening and the optimization itself are independent of the Young's modulus of the base material. The result can be adapted to any material with the same Poisson's ratio by a simple scaling, meaning that the ratio of the computed Young's modulus of the microstructure and the Young's modulus of the base material is constant.

Our samples were fabricated by selective laser sintering of an elastic thermoplastic polyurethane (TPU 92A-1).

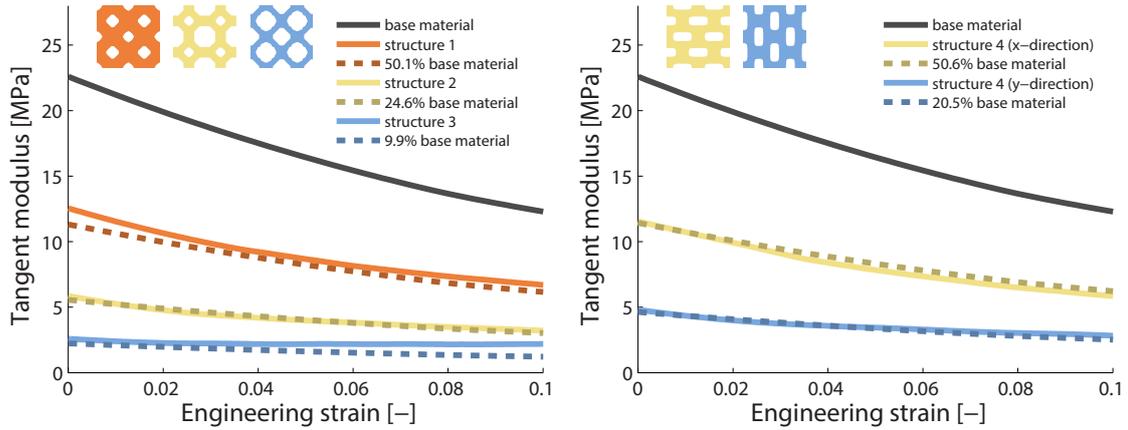


**Figure 4.11:** Data points for three different metamaterial families for a cubic material in 3D. The values for the shear modulus are omitted. A set of six structures for every family is visualized. The position of the structure is marked by the corresponding number.

**Test Setup and Method** We used an Instron E3000 frame with a 5 kN load cell for the material test. For the 2D structures, we performed tensile tests using a 10 cm pneumatic grip (see inset). We first characterized the base material using dog-bone shaped structures. To measure the tensile strength of the microstructures, we created samples consisting of a grid of  $7 \times 15$  unit cells ( $56 \text{ mm} \times 120 \text{ mm}$ ), to reduce boundary effects. After clamping, the samples were slightly pre-stretched ( $< 1 \text{ MPa}$ ) and tested with a constant displacement rate of  $50 \text{ mm}/\text{min}$ . The three-dimensional structures were tested in a compression test, using a displacement rate of  $5 \text{ mm}/\text{min}$ . The compressive properties of the base material were measured on a cylindrical sample, and the microstructure samples used a grid of  $7 \times 7 \times 6$  cells ( $56 \text{ mm} \times 56 \text{ mm} \times 48 \text{ mm}$ , images in Table 4.1).



**Tensile Test Results** In Figure 4.12 we compare the measured tangent modulus vs. strain for each generated structure with a measurement of the base material scaled by the relative Young's modulus predicted by simulation. The results indicate a good fit for small strains, and for most structures even for larger strains up to 0.1. The softest structure (9.9% of the base material's Young's modulus) deviates slightly from the scaled base material for larger strains, showing a nearly linear behavior as compared to the nonlinear behavior of the base material. This difference is likely due to the more pronounced rotations in sparse structures.



**Figure 4.12:** Tensile test results for a number of microstructures. Top: Test results for synthesized microstructures with a Young’s modulus of 9.9%, 24.6% and 50.1% of the base material. The scaled curve of the base material is shown for reference. Bottom: Test results for interpolated microstructures with orthotropic material behavior. The computed Young’s moduli were 20.5% and 50.6% of the base material’s Young’s modulus.

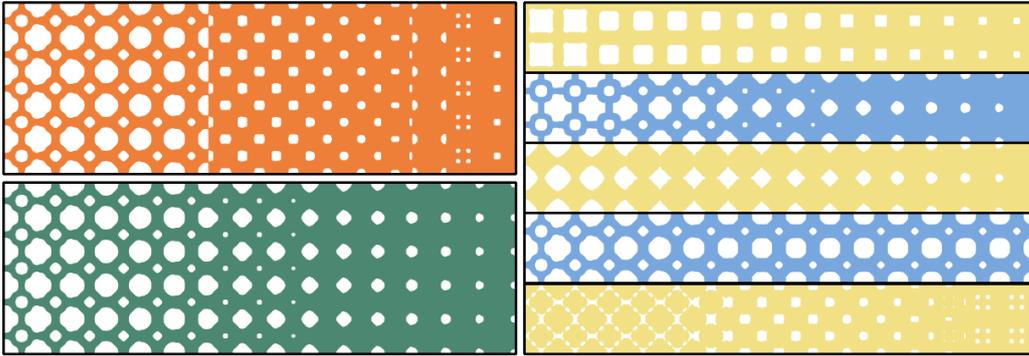


	Structure A	Structure B	Structure C
Simulation	19.0%	8.6%	12.8%
Measurement	21.5%	6.8%	12.2%

**Table 4.1:** The predicted and measured Young’s moduli for each of the three structures measured in a compression test.

**Compression Test Results** We determined the Young’s modulus of the base material and three generated structures by fitting a line to the linear region of the stress–strain measurement during a loading phase. Table 4.1 shows that the predicted and measured relative Young’s modulus match reasonably well. The stress–strain plots as well as the linear fits can be found in Appendix A.1.2.

**Material Gradient** We tested our structure synthesis on a simple material gradient example (Figure 4.13) to study how well the coarsened properties match the goal for inhomogeneous cases. We specified the relative Young’s modulus  $E$  and Poisson’s ratio  $\nu$  for a  $5 \times 15$  grid with a linear transition



**Figure 4.13:** A material gradient with a greedy tiling (top left), showing several suboptimal tile boundaries, and our optimized tiling (bottom left). Using only one of five metamaterial spaces would result in a better tiling, but worse approximation of the desired material parameters (right).

from  $E = 0.1$  and  $\nu = 0.6$  to  $E = 0.9$  and  $\nu = 0.4$ , and ran our algorithm with different settings and input. Using only structures from a single family resulted in good matching of interfaces between cells, with an average boundary dissimilarity below 5%. However, the average normalized distances between the simulated parameters of the generated structures and the desired parameters were large, with values of 0.203, 0.075, 0.177, 0.228, and 0.104 for the five metamaterial families we used for this test. By using these metamaterial families to get a set of candidate structures, we can greedily choose the best structure for each cell to achieve an average normalized distance of 0.051. However, there is no guarantee that these structures fit together, and Figure 4.13 (orange) shows that the greedy synthesis generates poorly fitting transitions, with an average boundary dissimilarity of 15.3%. Using the method described in Section 4.5, we can optimize for parameter approximation and boundary similarity at the same time. While the approximation of the desired parameters for the individual cells is slightly worse than for the greedy solution, with an average normalized distance of 0.065, the boundary dissimilarity in the resulting structure (Figure 4.13 (green)) was significantly better (4.2%).

**Boundary Forces** We presented two different approaches to represent the boundary dissimilarity during the synthesis stage. Figure 4.14 shows a situation in which simply comparing the geometry of the boundary performs worse than comparing the boundary forces under unit strain. We used numerical coarsening to determine the material parameters of these structures, as well as the parameters of a uniform mesh where each region is assigned the coarsened material parameters of the corresponding cells in the struc-



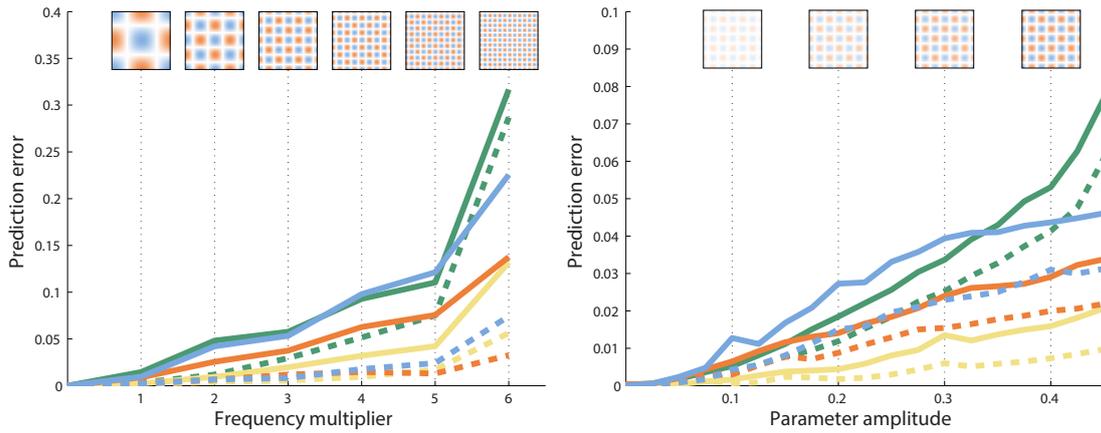
**Figure 4.14:** *Two gradients computed for the same goal parameters, using geometry (left) and boundary forces (right) to estimate the boundary dissimilarity during synthesis. The actual material parameters for the right structure deviate less from the predicted parameters (error of 0.0266, left, and 0.0137, right).*

ture. If only the geometries of the boundaries are compared, the distance between the actual parameters of the structure and the predicted parameters is 0.0266, while the error is only 0.0137 if the boundary forces are used to predict the dissimilarity of the cell boundaries.

**Effects of Heterogeneity** We explored our method’s performance for strongly heterogeneous goal parameters with two different tests, looking at the influence of the spatial frequency and amplitude of parameter changes on the prediction error. On a grid of  $12 \times 12$  cells, we synthesized target Young’s moduli in a sinusoidal pattern. We tested frequencies from a single period across the grid to 6 periods (i.e. a checkerboard), and amplitudes of parameter change ranging up to 45%, with a mean of 50% of the base material, resulting in a maximum range of 5% to 95%. We computed the difference between the homogenized material parameters of the synthesized structure and the homogenized material parameters of a uniform mesh with local material parameters set to the goals for the individual cells. Even though the assumption of infinite homogeneous tilings is violated, Figure 4.15 shows that for such heterogeneous tilings, the prediction error is below 0.05 even for very drastic parameter changes, and only above that in extreme cases where the change from maximal to minimal parameters happens in a span of less than two cells, or the amplitude is larger than 0.35.

### 4.6.3 Application Examples

**Gripper** Inspired by the field of soft robotics, we designed a simple gripper that can be actuated by air pressure (Figure 4.16). The gripper consists of two hollow tubes 16cm in length, printed with a soft material. The tubes are designed as a 2D structure with a stiff material on one half and an anisotropic material—soft along the tube and stiff along the circumferential direction—on the other half. A balloon is inserted into each tube, and increasing the



**Figure 4.15:** *The influence of a heterogeneous tiling on the approximation error, based on a  $12 \times 12$  grid with a sinusoidal distribution of Young’s moduli. The colors of the plot match the metamaterial families shown in Figure 4.9. We also show the spatial distribution of the Young’s moduli at the top of the plots (blue=soft, orange=stiff). Top: The error plotted for varying numbers of periods for the parameter distribution, with an amplitude of 45% (solid line) and 20% (dashed line). A frequency multiplier of 6 corresponds to a checkerboard distribution. Bottom: The error plotted for different amplitudes, with a frequency multiplier of 3 (solid line) and 2 (dashed line).*

pressure inside the balloons causes the tubes to bend due to the difference in stiffness. At the same time, the anisotropy of the structure prevents large changes in diameter. While this is only a very simple actuator, we believe that our method could be an important step towards a design tool for printable soft robots.

**Bunny, Teddy, and Armadillo** For the three-dimensional case, we tested our pipeline on two models (Bunny, 13 cm high; Teddy, 15 cm) with spatially varying Young’s moduli, created with an interactive material design tool [Xu et al., 2015a]. The models were subdivided into cells with 8 mm side length, and the Young’s moduli averaged for each cell. The metamaterial space used to populate these cells contained a single family of 21 microstructures. For synthesis, we chose the nearest neighbor in the database for each Young’s modulus. To keep the shape of the models, the individual voxels of each structure were set to void if they lay outside of the model. While this might lead to disconnected components in the reconstruction, these can easily be removed. We created a third model (Armadillo, 32 cm high) by manually painting the desired Young’s modulus distribution into a volumetric mesh, which was then used as an input to our method, using cells with 8 mm



**Figure 4.16:** *The optimized structure for the gripper (left), the generated model (middle). The fabricated result (right) can be used to grab and lift small objects.*



**Figure 4.17:** *The target Young's modulus distributions (colored), synthesized geometry (green) and fabricated model of the Bunny, Teddy and Armadillo.*

side length. We chose the parameter distribution such that the joints and the belly of the Armadillo are soft, while all other parts of the model are stiff. The structures of the cells were computed and tiled using our synthesis algorithm with the metamaterial space shown in Figure 4.11. The fabricated model can be easily actuated even though the base material is quite stiff (see Figure 4.17).

## 4.7 Conclusion

We presented a complete framework for automatically converting a given object with specified elastic material parameters into a fabricable representation that resembles the desired elastic deformation behavior. Our approach efficiently generates small-scale structures that obtain their elastic bulk properties from the shape and arrangement of the structures, significantly expanding the gamut of materials reproducible by 3D printers. Although our approach relies on an extensive precomputation phase for generating families of related structures that can be interpolated to smoothly vary the material properties, this only needs to be done once. To create an object with spatially varying elastic properties, our approach tiles the object's interior

with microstructures drawn from the database, using an efficient algorithm to select compatible structures for neighboring cells.

**Limitations and Future Work** Our method targets output devices that can 3D print at high resolution, and that allow easy removal of support material. In practice, we found selective laser sintering the most convenient process because the part is surrounded by unsintered powder and therefore does not require support structures. Removal of the unsintered powder from the structures can be easily achieved with compressed air. Other technologies, such as fused deposition modeling, allow printing overhangs without support structures only up to a certain angle. For future work, an interesting avenue could be incorporating these constraints into the optimization of material structures, spanning a space of metamaterials that are printable without support on these machines.

In our current work we focused on linear elasticity and small strain deformations. While this already allows us to create objects with tailored deformation properties, extensions into the nonlinear material behavior offer even more expressiveness, and could also incorporate interesting structures with buckling behavior.

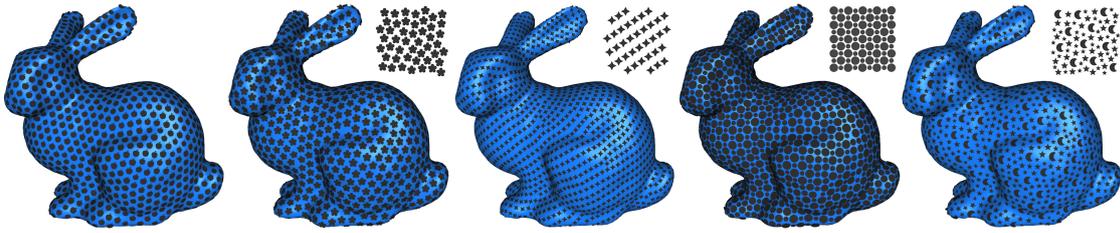
Finally, we do not explicitly treat the boundaries of the object, and obtain the boundary by simply intersecting the geometry of the structures with the shape of the object. An interesting next step would be wrapping the object with a surface for aesthetic reasons, but also taking the surface's effect on the deformation into account.

# Designing Structurally-Sound Surfaces with Decorative Patterns

Shell structures exhibit a unique combination of form and function. Ranging from curved pavilions to furniture and household items, their design very directly integrates both aesthetics and stability. Apart from overall shape, one particularly appealing way to impart style onto a shell structure is by removing material in decorative patterns. However, uninformed or too aggressive removal of material can easily degrade the structural performance of the shell and lead to failure, making manual design of such structures difficult. Automatic methods to create structurally-sound objects, on the other hand, rarely consider the aesthetic aspect of the design process, and the end result is hard to control.

In this chapter, we try to bridge this discrepancy by proposing a design tool that combines aesthetics and stability in a synergistic way. We introduce the concept of *stencils*—parametrized masks that cut out material from an underlying solid surface. Inspired by recent work on discrete element distribution [Ma et al., 2011; Roveri et al., 2015], we propose an example-based method for describing stencil patterns with custom shapes and arrangements. We formulate pattern creation as an energy minimization problem, allowing us to simultaneously optimize with respect to both aesthetic and structural goals. Using this approach, designers can quickly create structurally-sound shells with a broad range of decorative patterns, ranging from regular and homogeneous distributions to irregular and heterogeneous patterns.

The technical core of our method is an optimization algorithm that automat-



**Figure 5.1:** We show that our method can handle different types of textures. From left to right: (1) A distribution of apples using only our packing objective. (2) An exemplar with irregularly placed flowers creates a packing-like distribution, and offers the possibility to fill in missing stencils. (3) Using an exemplar with stencils placed in stripes, we can create a distribution with an anisotropic appearance. (4) An exemplar with alternating, regularly placed circles results in a highly structured output. (5) A heterogeneous exemplar can be used to create a more diverse result.

ically determines stencil parameters in order to comply as much as possible with both pattern and stability objectives. Evaluating the stability objective and its derivatives requires solving for the equilibrium state of the corresponding patterned shell, and computing its derivatives with respect to the stencil parameters. When using meshes that conform to the boundaries of the stencil, solving for static equilibrium can easily take several minutes. In order to afford interactive design iterations, we lay aside computationally-expensive approaches based on conforming discretizations and remeshing operations. Inspired by material interpolation schemes in topology optimization [Bendsøe and Sigmund, 1999], we instead perform all computations on the original surface mesh and scale the elastic energy of its (membrane and bending) elements according to the extent to which they are covered by stencils. Though approximate, this approach allows us to use a comparatively coarse mesh with constant topology, which is essential for reducing the time spent on parameter optimization to an acceptable level during design iterations. Once a satisfying design has been found, the user can verify its validity through simulation on a high-resolution conforming mesh.

The formulation that we propose in this work is general with respect to both aesthetic and structural goals, and it provides manifold possibilities for the user to guide the aesthetic appeal of the design.

## 5.1 Overview

We propose a computational design tool for physical surfaces with decorative cutouts, fulfilling both aesthetic and structural objectives. Starting from a mesh representing the input surface, the user first creates a decorative pattern that defines the aesthetic objective of the design (Section 5.2). This decorative pattern consists of a collection of stencils that are distributed across the surface according to a texture objective in the form of an energy function. We explore two options in this work: a packing scheme that places stencils on the surface according to a quasi-blue noise random distribution; and an example-based distribution scheme based on the principle of discrete element textures [2011]. We also provide a set of tools that allow the user to further control the distributions in terms of local scaling and orientation fields.

The user then defines a force distribution on the surface that is a conservative estimate of the expected loads. Furthermore, structural goals are defined to either enforce a stability criterion in terms of a fixed threshold on the maximum per-element energy density, or to minimize the compliance of the structure for a fixed amount of material. The aesthetic and structural goals provided by the user give rise to objectives that are jointly minimized in order to obtain optimal stencil parameters (Section 5.3).

## 5.2 Stencil Patterns

We present decorative patterns on curved surfaces as collections of stencils. In the following, we first lay down the representation of stencils, then proceed to texture objectives and additional art-direction mechanisms that allow users to create a variety of stencil distributions.

### 5.2.1 Stencil Representation & Projection

**Representation** Stencils are discrete elements that live on a 3D triangle mesh  $\mathcal{M}$  corresponding to the input surface. Each stencil is defined by a 2D mesh  $\mathcal{S}_i$  describing a reference shape that is transformed to its 3D counterpart  $\mathcal{S}_i$  according to the stencil’s position  $\mathbf{q}_i$ , its normal  $\mathbf{n}_i$ , an orientation vector  $\mathbf{t}_i$ , and a scaling parameter  $s_i$ . While all of these variables are required to fully determine the stencil’s transformation, not all of them are actual degrees of freedom. In particular, the position  $\mathbf{q}_i$  of the stencil has to lie on the triangle mesh of the input surface, which we achieve through projection.

Furthermore, the normal  $\mathbf{n}_i$  is determined using Phong interpolation of the vertex normal at the corresponding surface location  $\mathbf{q}_i$ . Finally, the orientation vector  $\mathbf{t}_i$  is obtained by rotating a vector from an underlying orientation field (computed with `libigl` [Jacobson et al., 2016] using the approach described in [Bommes et al., 2009]) around the normal  $\mathbf{n}_i$ . The effective degrees of freedom  $\mathbf{p}_i$  of a stencil are thus defined by five parameters corresponding to its position, scale, and scalar rotation value. We concatenate the parameters of all stencils into a vector  $\mathbf{p}$ .

**Projection** In order to compute the geometry of the stenciled mesh, we have to project the transformed stencil meshes  $\mathcal{S}_i$  onto the underlying surface mesh  $\mathcal{M}$ . For each stencil, we first orthogonally project the corresponding region of  $\mathcal{M}$  onto the plane defined by its position and normal. We then scale the distance between each vertex and the center such that Euclidean distances are preserved, i.e., the projection is equidistant. Although less accurate than geodesic distances, we found that the simpler Euclidean distance yields high-quality distributions, provided that the underlying mesh is sufficiently smooth and stencils are comparatively dense. Once projected into a two-dimensional subspace, overlaps between  $\mathcal{S}_i$  and  $\mathcal{M}$  as well as corresponding cuts for  $\mathcal{M}$  can be computed efficiently.

### 5.2.2 Pattern Objectives

We pursue a variational approach and cast pattern generation as an energy minimization problem with pattern objectives  $P(\mathbf{p})$  that measures the quality of a given stencil distribution. Our method is flexible with respect to pattern objectives and we consider two alternatives in the following.

**Collision Objective** In order for a distribution to be admissible, we require that no pair of stencils may overlap. To this end, we introduce a simple repulsive force model that resolves collisions between stencils, which we detect using their bounding circles. If an interference between stencil  $i$  and  $j$  is found, we set up an objective that penalizes overlap according to

$$O_{ij}^{\text{coll}} = ((1 + s)(r_i + r_j) - d_{ij})^2, \quad (5.1)$$

where  $r_i$  and  $r_j$  are the bounding circle radii of the stencils,  $d_{ij} = \|\mathbf{q}_i - \mathbf{q}_j\|$  is the Euclidean distance between the stencils' centers, and  $s$  is a safety factor enforcing a lower bound on inter-stencil distance such that the resulting

structure can be fabricated without running into problems with minimum feature sizes.

We use a similar approach to prevent stencils from moving off the mesh when we use a model with open boundaries. If a stencil  $i$  collides with the boundary, we introduce the objective

$$O_i^{\text{bcoll}} = ((1 + s)r_i - d_{ik})^2, \quad (5.2)$$

where  $d_{ik}$  is the distance between the stencil center and its nearest point on the boundary.

**Packing** A simple yet visually-pleasing way of generating element distributions is through random sampling. Our packing objective aims to create uniform stencil distributions by maximizing the distance between neighboring stencils. For any pair  $i$  and  $j$  of stencils with a stencil center distance  $d_{ij}$  smaller than a given neighborhood size  $\epsilon$ , we add an objective term that repels stencils according to

$$O_{ij}^{\text{rep}} = \frac{w_{ij}}{d_{ij}^2}, \quad (5.3)$$

where  $w_{ij} = (\epsilon\sqrt{2\pi})^{-1}e^{-d_{ij}^2/(2\epsilon^2)}$  is a weighting factor depending on the inter-stencil distance  $d_{ij}$ . For our examples, we set  $\epsilon$  to 25% of the length of the object. In order to encourage dense distributions, we add a simple growth objective

$$O_i^{\text{growth}} = -s_i. \quad (5.4)$$

The packing objective is then defined as

$$P^{\text{pack}} = \sum_{i,j} (O_{ij}^{\text{rep}} + O_{ij}^{\text{coll}}) + \sum_i O_i^{\text{growth}}. \quad (5.5)$$

**Discrete Element Textures** The packing objective can be used to create visually pleasing distributions as shown, e.g., in Figure 5.1, (1). However, while the user can control the shape of the stencil, the distribution is homogeneous and its structure cannot be controlled. In order to provide more artistic control over the resulting distributions, we turn to an example-based approach inspired by *Discrete Element Textures* (DET) [Ma et al., 2011]. DET was developed to synthesize and improve the distribution of points with arbitrary attributes. It works by matching points and their neighborhoods in an output distribution to neighborhoods in a user-provided exemplar, and

then computes updates that aim to improve the match between output and input. In our case, we approximate each stencil as a single point, and encode the scale and type of the stencil as attributes. Letting  $\mathcal{N}(i)$  denote the neighborhood of stencil  $i$ , its corresponding energy is defined as

$$O_i^{\text{DET}} = \sum_{j \in \mathcal{N}(i)} \|(\mathbf{q}_i - \hat{\mathbf{q}}_j) - \mathbf{T}(\mathbf{q}'_i - \mathbf{q}'_j)\|^2 + \|r_j - r'_j\|^2 + \zeta(1 - \delta_{t_j t'_j}), \quad (5.6)$$

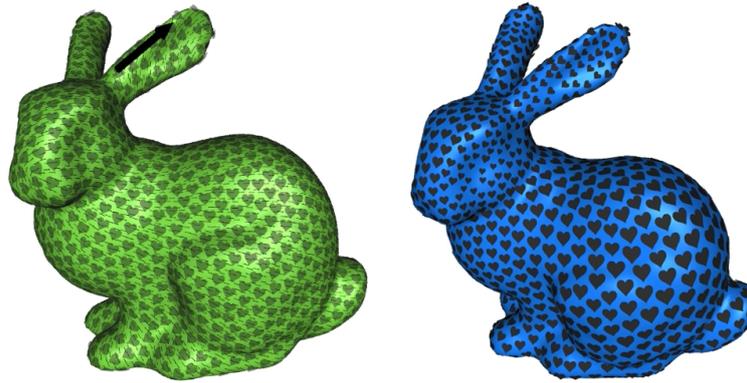
where  $\hat{\mathbf{q}}_j$  is the distance-preserving projection of the position of stencil  $j$  into the tangent plane of stencil  $i$ ,  $\mathbf{q}'_i$  and  $\mathbf{q}'_j$  are the positions matched in the exemplar for both stencil,  $\mathbf{T}$  is the transformation from the exemplar to the tangent plane of stencil  $i$ , and  $r_j$  and  $r'_j$  are the scales associated with stencil  $j$  and its match, respectively. Furthermore,  $t_j$  and  $t'_j$  are the types of stencil  $j$  and its match, and  $\delta_{t_j t'_j}$  is the Kronecker delta indicating whether  $t_j$  and  $t'_j$  are identical, such that the expression  $\zeta(1 - \delta_{t_j t'_j})$  adds a penalty of  $\zeta$  if the stencil types do not match. Similar to the packing objective, we augment the DET objective by a term that explicitly penalizes collisions.

As can be seen in Figure 5.1 (2–5), example-based stencil distributions allow users to create a wide range of decorative patterns with distinct aesthetic appeals.

### 5.2.3 Initialization

An adequate initial distribution of stencils is important to ensure that our optimization approach (Section 5.3) will converge to a desirable solution. We found that especially for the DET objective, a good initialization is crucial, and moreover, copying patches from the input exemplar to the output, as suggested in the original paper [Ma et al., 2011], rarely led to good results in our examples.

We instead opted for an incremental initialization strategy similar to [Ijiri et al., 2008]. We start by randomly selecting a point in the input exemplar, and then copy the neighborhood patch around this point onto a random location in the output domain. We then compute the DET matching from the output domain to the input exemplar, with the constraint that we do not match any patches in the input exemplar that are at the boundary of the domain. This is a one-way matching that does not include a penalty for stencils in the input neighborhood that are not matched to the output domain. However, we can use the neighborhood in the exemplar to estimate which stencils we could add to the output domain, by determining the stencils that are not matched. Adding unmatched stencils into the output domains, given that the overlap



**Figure 5.2:** *Users can control the orientation field through sketching (left) and locally adjust the size of the stencils using a brush tool (right).*

with existing stencils is not significantly larger than our threshold, allows us to expand the region covered by stencils. By iteratively applying this insertion scheme, we grow an initial stencil distribution that ultimately covers the whole model.

The nature of our packing objective simplifies the initialization in cases where only this objective is activated. In such a case, we simply use a Poisson disk sampling strategy.

#### 5.2.4 User Control

Apart from the shape of the stencils and their arrangement, we also provide additional tools that allow the user to control the aesthetics of the resulting pattern through sizing and orientation fields. In order to control the orientation of the stencils, we augment the computation of the orientation field by an objective that asks the resulting field to locally align with user-defined sketch directions (see Figure 5.2, left).

Another way of stylizing the pattern is to prescribe an inhomogeneous scaling field across the surface (see Figure 5.2, right). We provide a simple brush tool that allows the user to paint the desired stencil sizes directly onto the surface. For the DET objective, these sizing values are used to scale the exemplar when computing the matching energy in Equation (5.6), whereas for the packing objective, the scale of the stencil is simply set to the sizing value at the corresponding location.

## 5.3 Structural Optimization

Apart from aesthetic goals, a central requirement when designing functional surfaces is to satisfy structural constraints, i.e., conditions relating to the stability of the design. While the criteria that can be used to quantify stability are manifold, we focus on two common variants of structural optimization based on energy-density constraints and compliance minimization. For both of these approaches, we must compute *a*) the equilibrium state of the structure, *b*) the change in equilibrium state induced by a change in parameters, and *c*) the parameter values that lead to a desired equilibrium state. We explain these steps in the following.

### 5.3.1 Simulation

Following standard practice in graphics, we opt for a computational model that combines constant strain triangles (CST) for in-plane deformations and discrete shell elements for bending. Both types of elements rely on a discrete surface representation in terms of a triangle mesh, comprising  $m$  faces  $\mathcal{T}_i$  and  $n$  vertices  $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^T$ , with  $\mathbf{x}_j \in \mathbb{R}^3$ . For the membrane part, we use a simple St. Venant-Kirchhoff material whose strain energy density  $\Psi$  is defined as

$$\Psi = \frac{\mu}{2} \text{tr}(\mathbf{E}^t \mathbf{E}) + \lambda \text{tr}(\mathbf{E})^2, \quad (5.7)$$

where  $\mu$  and  $\lambda$  are (thickness-dependent) material parameters,  $\mathbf{E} = \frac{1}{2}(\mathbf{F}^t \mathbf{F} - \mathbf{I})$  is the  $2 \times 2$  Green strain tensor, and  $\mathbf{F} \in \mathbb{R}^{3 \times 2}$  is the deformation gradient; refer to, e.g., Skouras et al. [2012] and the textbook by Bonet and Wood [1997] for details. Using CST elements for discretization, the strain energy density is constant across the element and the membrane energy simply follows as  $W_i^{\text{memb}} = \int_{\Omega_i} \Psi = A_i \Psi$ , where  $\Omega_i$  is the parameter domain of element  $i$  and  $A_i$  its undeformed area.

Complementing the CST elements for membrane deformations, bending elements are formed by pairs of edge-adjacent triangles. The corresponding energy is defined as

$$W_{ij}^{\text{bend}} = k_b (\theta(\mathbf{x}) - \bar{\theta})^2 \frac{A_{ij}}{h_{ij}}, \quad (5.8)$$

where  $\theta$  and  $\bar{\theta}$  are the dihedral angles in the deformed and undeformed configuration, respectively,  $h_{ij}$  is a geometry factor, and  $A_{ij} = \frac{1}{3}(A_i + A_j)$  (see [Grinspun et al., 2003]). Note that  $k_b$  is a dependent coefficient that is computed from the material parameters  $\lambda$  and  $\mu$  as well as the thickness of the element according to [Garg et al., 2007].

### 5.3.2 Extension to Stenciled Shells

Perhaps the most obvious way to apply this computational model to the case of thin shells structured with decorative cutouts would be to use a triangle meshes that conforms to the boundaries of the stencils. However, this approach would require a high-resolution mesh that would drastically increase the time needed to compute equilibrium state—and it would be prohibitively expensive when used at the core of our optimization algorithm. We therefore turn away from conforming meshes and instead maintain a comparatively coarse mesh, each of whose triangles  $\mathcal{T}_i$  we endow with an additional fill-ratio variable  $\alpha_i$  corresponding to the degree to which the element is overlapped by stencils. In order to compute the overlap between a given triangle and a stencil, we first project the triangle into the tangent space of the stencil. We then compute the intersection polygon using the Sutherland-Hodgman clipping algorithm [Sutherland and Hodgman, 1974]. The ratio between the area of the overlap and the area of the projected element is then subtracted from the triangle’s fill ratio.

Given fill ratios for all triangles, we compute the strain energy for membrane elements by integrating the density only over the part of the elements that is not covered by stencils. As the deformation per CST element is constant, computing the energy of a stenciled element amounts to a simple scaling operation, i.e.,

$$\hat{W}_{ij}^{\text{memb}}(\mathbf{x}, \alpha) = \int_{\hat{\Omega}_i} \Psi = \alpha_i A_i W_i^{\text{memb}}(\mathbf{x}) \quad (5.9)$$

where  $\hat{\Omega}_i$  is the parameter domain of the stenciled element and  $\hat{A}_i = \int_{\hat{\Omega}_i} 1 = \alpha_i A_i$  its area. With similar reasoning, we define the energy of a stenciled bending element as the corresponding energy of the solid element scaled by the geometric mean of its two triangles,

$$\hat{W}_{ij}^{\text{bend}}(\mathbf{x}) = \sqrt{\alpha_i \alpha_j} W_{ij}^{\text{bend}}(\mathbf{x}). \quad (5.10)$$

Using the geometric mean ensures that the energy vanishes if one of the triangles is completely cut out. The total elastic energy of the stenciled surface is obtained as the sum of element-wise energies, i.e.,

$$\hat{W}^{\text{el}} = \sum_i \hat{W}_i^{\text{memb}} + \sum_{(i,j) \in \mathcal{H}} \hat{W}_{ij}^{\text{bend}}, \quad (5.11)$$

where  $\mathcal{H}$  is the set of edge-adjacent triangle pairs. We avoid ill-conditioned elastic stiffness matrices by enforcing a minimum value of  $10^{-4}$  for all fill ratios  $\alpha_i$ .

Finally, in order for the surface to be in equilibrium, the sum of internal forces  $\mathbf{f}^{\text{el}} = -\frac{\partial \hat{W}^{\text{el}}}{\mathbf{x}}$  and externally-applied forces  $\mathbf{f}^{\text{ext}}$  has to vanish in all nodes, i.e.,

$$\mathbf{f}_i(\mathbf{x}, \boldsymbol{\alpha}) = \mathbf{f}_i^{\text{el}}(\mathbf{x}, \boldsymbol{\alpha}) + \mathbf{f}_i^{\text{ext}} = \mathbf{0} \quad \forall i. \quad (5.12)$$

### 5.3.3 Optimization

We assume that the structural goal of a design can be quantified in terms of an objective function  $S(\alpha, \mathbf{x}(\alpha))$  with explicit dependence on both position and material fill ratios. As detailed above, the latter depend explicitly on the stencil distribution, i.e.,  $\alpha = \alpha(\mathbf{p})$ . In order to improve the structural objective, we consider its gradient with respect to the stencil parameters,

$$\frac{dS(\alpha, \mathbf{x}(\alpha))}{d\alpha} = \frac{\partial S(\alpha, \mathbf{x}(\alpha))}{\partial \alpha} + \frac{\partial S(\alpha, \mathbf{x}(\alpha))}{\partial \mathbf{x}} \frac{d\mathbf{x}}{d\alpha}. \quad (5.13)$$

It is clear from the above expression that computing the gradient requires the map between positions  $\mathbf{x}$  and material fill ratios  $\alpha$ , which is given by Equation (5.12): a combination of material fill ratios and deformed positions is admissible if and only if it corresponds to an equilibrium state, i.e.,  $\mathbf{f}(\mathbf{x}, \boldsymbol{\alpha}) = \mathbf{0}$ . Consequently, through the implicit function theorem (see Section 3.4), for an admissible change in material fill ratios it must hold

$$\frac{d\mathbf{f}}{d\alpha} = \frac{\partial \mathbf{f}}{\partial \alpha} + \frac{\partial \mathbf{f}}{\partial \mathbf{x}} \frac{d\mathbf{x}}{d\alpha} = \mathbf{0}, \quad (5.14)$$

and therefore

$$\frac{d\mathbf{x}}{d\alpha} = -\frac{\partial \mathbf{f}}{\partial \mathbf{x}}^{-1} \frac{\partial \mathbf{f}}{\partial \alpha}. \quad (5.15)$$

We can thus compute the derivative of  $\mathbf{x}$  with respect to  $\alpha$  by solving a system of linear equations, whose matrix is given by the Hessian of the elastic energy of the surface. Finally, since we ultimately solve for stencil parameters, we compute  $\frac{dS}{d\mathbf{p}}$  by applying the chain rule to Equation (5.13) and use finite differences to numerically approximate the required derivatives  $\frac{\partial \alpha}{\partial \mathbf{p}}$ .

The formulation presented above provides flexibility for different structural objectives  $S$  and we consider three examples in the following.

**Energy-Density Objective** A natural structural goal for a design is to ask that a given expected load should not lead to failure. Structural failure is typically indicated by a stability criterion that, depending on the type of

material, depends on deformation, stress, or energy density. We opt for a criterion based on per-element energy density

$$\mathcal{W}_i = \frac{1}{\hat{A}_i} \hat{W}_i^{\text{memb}} + \sum_j \frac{\hat{A}_i}{3\hat{A}_{ij}^2} \hat{W}_{ij}^{\text{bend}}, \quad (5.16)$$

with  $\hat{A}_{ij} = \frac{1}{3}(\hat{A}_i + \hat{A}_j)$ , and define a corresponding structural goal as

$$S^{\text{stab}}(\alpha, \mathbf{x}(\alpha)) = \frac{1}{m} \sum_i \begin{cases} \frac{1}{2}(\mathcal{W}_i(\alpha, \mathbf{x}(\alpha)) - \beta)^2, & \text{if } \mathcal{W}_i > \beta \\ 0, & \text{else,} \end{cases} \quad (5.17)$$

where  $\beta$  is a threshold value indicating the energy density beyond which structural failure is likely to occur.

In many cases, the materials used in a design will dictate specific threshold values that must not be exceeded. However, another way of optimizing the stability of a design is to minimize the maximum per-element energy density. We implement this strategy as an iterative scheme, each of whose steps update the current value of  $\beta$  to the average computed from the top  $y\%$  of the elements with the highest energy density, where  $y$  can be chosen depending on the use case. See Sec. Section 5.4 for an example.

**Compliance Objective** Instead of adding material in order to push the stability criterion below a given threshold, an alternative approach is to seek a distribution of a given amount of material in order to maximize the stiffness of a structure or, equivalently, minimize its compliance with respect to given loads. This *conventional* way of topology optimization typically aims to minimize the work done by externally-applied forces, but is often expressed by an equivalent formulation based on the total internal energy of a given structure. We define a corresponding structural objective as

$$S^{\text{comp}}(\alpha, \mathbf{x}(\alpha)) = \sum_{i=0}^m A_i \mathcal{W}_i(\alpha, \mathbf{x}(\alpha)), \quad (5.18)$$

which is complemented by an additional constraint  $C^{\text{area}}(\alpha) = \sum_i A_i - \sum_i \hat{A}_i(\alpha) = \text{const.}$ , requiring that the total area of all stencils be constant. In order to incorporate this constraint into our optimization method, we project the gradient of the combined (structural and aesthetic) objective onto the space of admissible directions, i.e., orthogonal to the constraint gradient. The constraint is enforced during every update of the stencil parameters.

### 5.3.4 Numerical Solution

With both the stability and pattern objectives in place, we compute optimal stencil parameters through minimization. Thanks to the equilibrium conditions (Equation (5.12)), the positions  $\mathbf{x}$  are implicit functions of the material fill ratios  $\alpha$ , which in turn are explicit functions of the stencil parameters  $\mathbf{p}$ . We thus define the joint objective as a function of the stencil parameters,

$$J(\mathbf{p}) = w_S S(\mathbf{x}(\mathbf{p}), \alpha(\mathbf{p})) + w_P P(\mathbf{p}), \quad (5.19)$$

where  $w_S$  and  $w_P$  are scaling parameters.

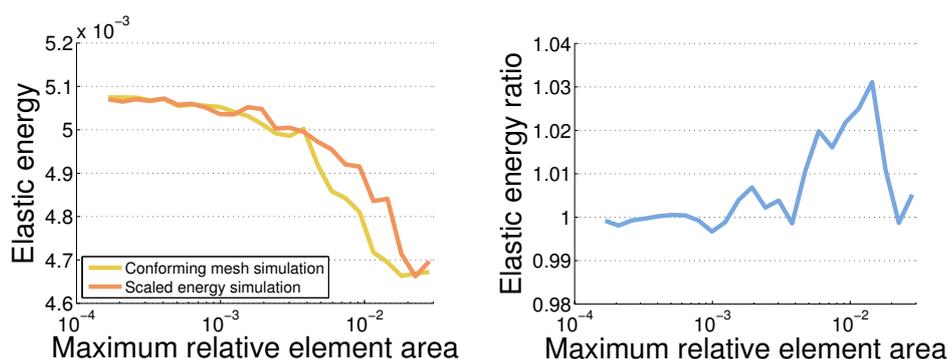
In order to minimize  $J$ , we use L-BFGS-B [Byrd et al., 1995], a Quasi-Newton method that combines analytically-computed gradients with a limited-memory BFGS approximation of the Hessian and bound constraints. We additionally employ line search in order to ensure monotonic decrease in the objective. This approach allows us to optimize for all continuous stencil parameters at the same time. For the stencil type, which is the only discrete stencil parameter, we use the update scheme proposed in [Ma et al., 2011] that selects the stencil type based on a majority vote among the stencil's neighbors. Note that in order to use line search correctly, we need to ensure that our simulation is always in equilibrium, i.e., we need to update the simulation for any change in stencil parameters.

## 5.4 Results

### 5.4.1 Validation

**Simulation with Scaled Energy Densities** Our elastic simulation using scaled energy densities is inspired by simulation approaches from topology optimization [Bendsøe and Sigmund, 1999], where the scaling of the energy density of each element is a function of its fill ratio. We check the validity of our simulation, which extends this approach to hinge elements that include two fill ratios, by ascertaining that for increasing mesh resolutions, the simulation converges to a regular simulation using a conforming mesh. Figure 5.3 shows that this convergence can indeed be observed, and moreover, that even for coarse resolutions, our approach provides an adequate approximation of a simulation on a conforming mesh.

**Physical Validation** We validate that our optimization improves the stability of an object by performing a load test on a stenciled plate, shown in



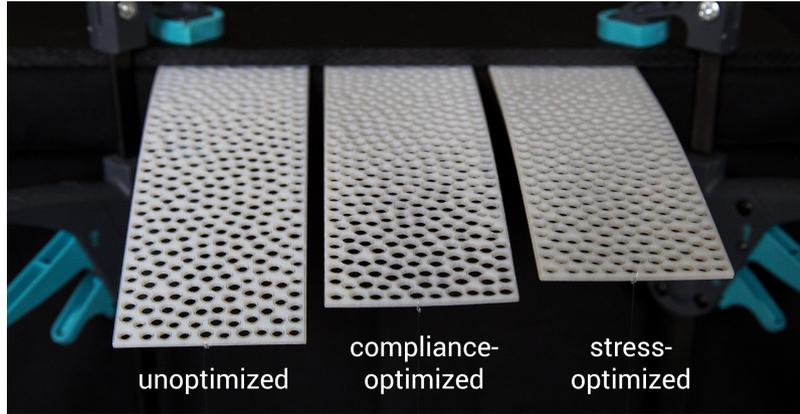
**Figure 5.3:** The convergence of elastic energy for simulations with our scaled energy density approach and simulations with a conforming mesh, for different mesh resolutions, specified by the maximum area of a mesh element. The simulation is performed on a thin, square, horizontally oriented, unit-sized sheet with a single circular stencil, with fixed outer boundaries and deformed by self-weight. Left: The elastic energy of both simulations. Right: The ratio between the elastic energies.

Figure 5.4. For the design and experiment, we fix one end of the plates, and apply a load on the center of the other end. The plate with the pattern-optimized stencil distribution shows the largest deflection, and the energy- and compliance-optimized stencil distributions show improved performance. The compliance-optimized distribution uses a fixed material constraint in this experiment, and the limited amount of material redistribution possible without any stencils overlapping explains the slightly better performance of the energy-optimized distribution.

**Coarse Simulation vs. High-resolution Simulation** For all our results that we fabricated, we compare our coarse simulation to a high-resolution conforming mesh. Figures 5.5, 5.8, 5.7, and 5.6 show that we capture the global energy density distribution well, while small local stress peaks along the boundaries of stencils might exceed the maximum energy density computed by our simulation.

## 5.4.2 Physical Results

We designed a number of results for different applications, which were 3D printed using selective laser sintering with a polyamide material (PA12). We show both the *pattern-optimized* results, which only use the pattern objective for the optimization, as well as the *stability-optimized* results that use both



**Figure 5.4:** Comparison of the deflection of plates with circular stencils, before and after running the stability optimization. The plates were fixed on one side, and a weight was attached to the opposite side.

pattern and stability objectives. For all stability-optimized examples, we set the energy density threshold  $\beta$  to a fraction of the maximum energy density  $\mathcal{W}^{max}$  of the non-stenciled mesh to ensure that we create a meaningful scenario for our optimization. We show the statistics and performance of our examples in Table 5.1, and parameters in Table 5.2.

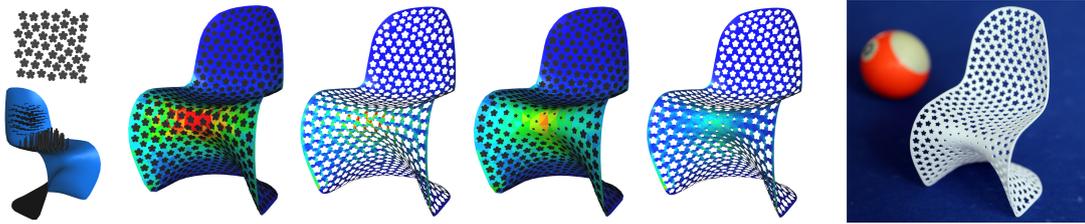
Model	Objective	#s	#v	#v <sub>sim</sub>	t [s]
Chair	Pattern-only	436	1927	520	4.1
	Energy density				12.1
Table	Pattern-only	462	1418	581	2.0
	Energy density				21.2
Bowl	Pattern-only	185	1529	772	0.8
	Energy density				16.6
Pavilion	Pattern-only	310	937	381	4.2
	Energy density				43.4
	Energy density (adaptive)				150.8
	Compliance				5.1

**Table 5.1:** Performance statistics. The table shows the number of stencils, the number of vertices for the underlying mesh and the simulation mesh, and the runtime  $t$  for all our fabricated examples. The runtime is measured from initial distribution to convergence (for the pattern-only objective) or from a pattern-optimized solution to a stability-optimized solution (for all other objectives).

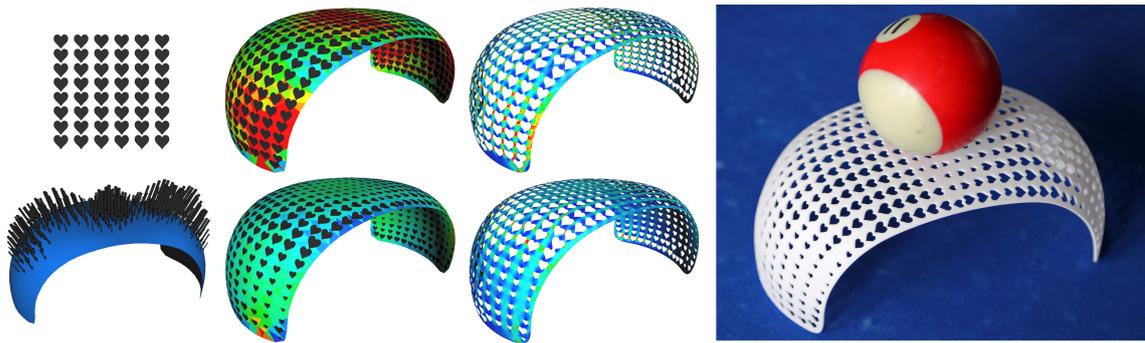
**Chair and Table** The chair example in Figure 5.5 and the table example in Figure 5.6 we both designed by fixing the bottom vertices of the model, and then applying a load on the seat and back rest for the chair, and on the table top for the table.

Model	$w_S$	$w_P$	$\beta$	$s$	$r_N$
Chair	$10^3$	1	$0.7\mathcal{W}^{max}$	0.2	0.3
Table	$10^4$	1	$0.5\mathcal{W}^{max}$	0	0.3
Bowl	200	1	$0.8\mathcal{W}^{max}$	0.1	0.3
Pavilion	500	1	$0.2\mathcal{W}^{max}$	0.2	0.3

**Table 5.2:** Parameters used for our examples: Scaling parameters  $w_S$  and  $w_P$ , energy density threshold  $\beta$  as a function of the maximum energy density  $\mathcal{W}^{max}$ , collision safety factor  $s$ , and the neighborhood radius  $r_N$  relative to the dimensions of the input exemplar.

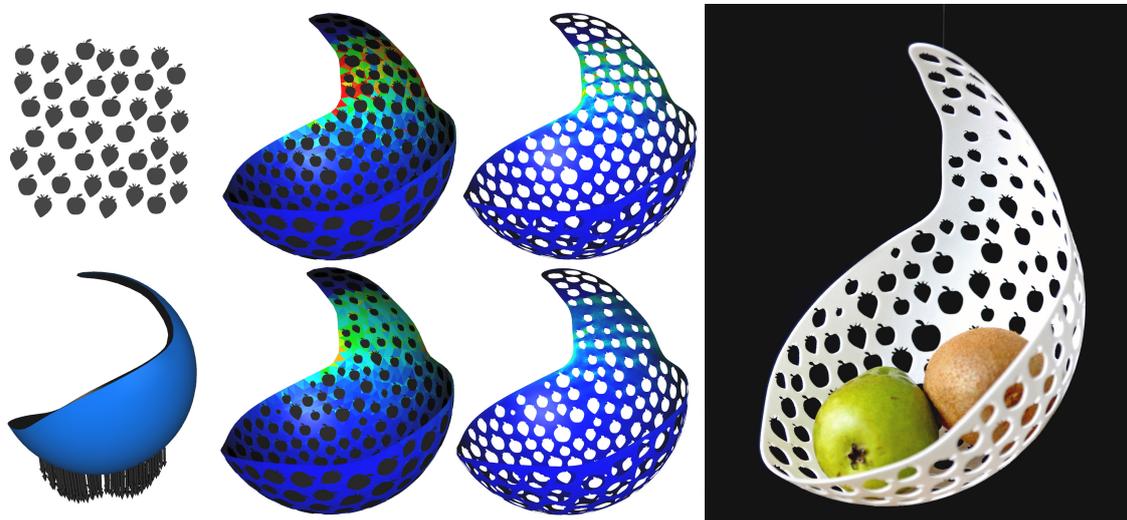


**Figure 5.5:** Comparison of energy densities between our coarse simulation mesh (second and fourth picture) and a high-resolution conforming mesh (third and fifth picture) for the texture-optimized chair (left pair) and the stability-optimized chair (right pair). Energy densities around the threshold  $\beta$  are shown in green, with higher densities in red and lower densities in blue. The pattern exemplar and load case are given on the left. The stability-optimized chair has been 3D printed (right).



**Figure 5.6:** Comparison of energy densities between our coarse simulation mesh (center left) and a high-resolution conforming mesh (center right) for the pattern-optimized table (top) and the stability-optimized chair (bottom). Energy densities around the threshold  $\beta$  are shown in green, with higher densities in red and lower densities in blue. The pattern exemplar and load case are given on the left. The stability-optimized table holds a billiard ball (right).

**Fruit Bowl** We also tested our method on alternative boundary conditions. The fruit bowl in Figure 5.7 has been designed to hang by a wire attached to its top, and the load case is a force pulling on the bottom of the bowl.

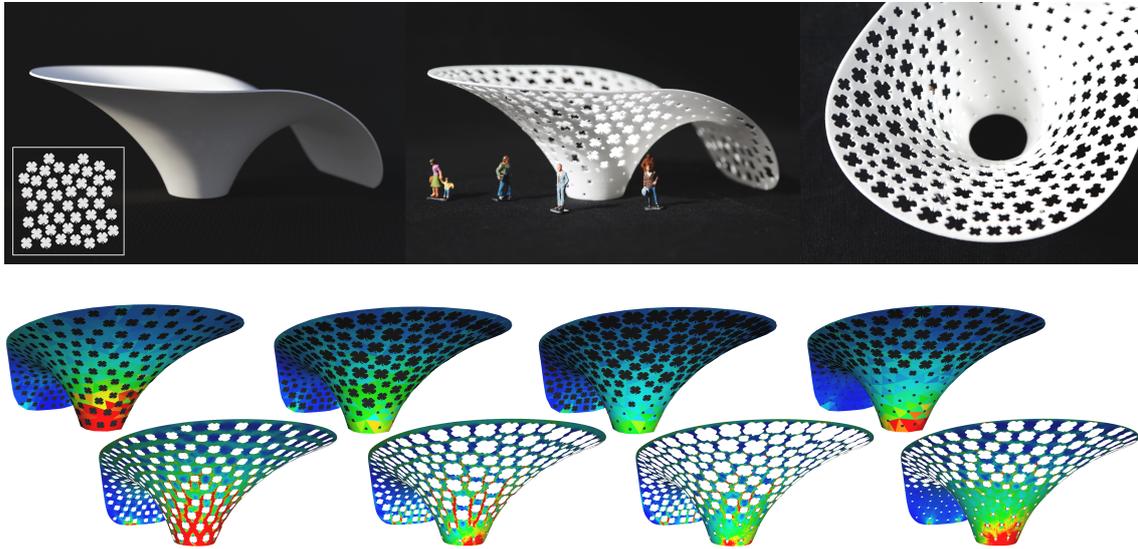


**Figure 5.7:** *The pattern exemplar and load case of the bowl example (left), and the final fabricated result (right). We compare the energy densities of our coarse simulation mesh (center left) with a high-resolution simulation on a conforming mesh (center right) for the pattern-optimized (top) and stability-optimized bowl (bottom). Energy densities around the threshold  $\beta$  are shown in green, with higher densities in red and lower densities in blue.*

**Pavilion** Figure 5.8 show the result of applying our method to an object under self-weight. In this case, the stencil distribution not only locally influences the stiffness, but also changes the load, and can have a global effect on the stability. Since the optimization considers the interplay between stencil parameters and the physical simulation, it automatically removes material from the top part of the pavilion in order to lighten the load on the base of the structure.

For this case of an object under self-weight, our approach of setting a fixed threshold for the energy density might be too restrictive. Rather, we would like to perform a minimization instead of a thresholding, with the assumption that any reduction in energy density will be beneficial. As an exercise, we performed an adaptive optimization: At the beginning of an optimization, we set the threshold to the median energy density. Whenever we reach an optimal solution, we check whether the median energy density changed. If so, we continue the optimization with the new threshold. Figure 5.8 also

shows the result of this adaptive optimization, which leads to more extreme changes to the stencils than the regular optimization.



**Figure 5.8:** *Top: The base shape and pattern exemplar for the pavilion model (left) and the final optimized result (center, right). Bottom: Comparison of energy densities for the pavilion model. The top row shows our coarse simulation mesh, the bottom row a high-resolution simulation on a conforming mesh. We show four different scenarios: The pattern-optimized pavilion (left), the energy density-optimized pavilion with a single energy density threshold (center left) and an adaptive energy density threshold (center right), and the compliance-optimized pavilion with a fixed material constraint (right). Energy densities around the threshold  $\beta$  are shown in green, with higher densities in red and lower densities in blue.*

## 5.5 Conclusion

We introduced a novel method to design stenciled surfaces that combines aesthetics and stability. We have shown that we can efficiently predict regions of failure and automatically optimize a stencil distribution to create stable objects with high visual fidelity.

### 5.5.1 Limitations and Future Work

While we can predict regions of failure, we cannot guarantee that our best solution will be able to avoid these, as either the load case or the texture objective could prevent a feasible solution. In such cases, we can provide

feedback to the user that would ideally lead to adjustments in the load case or texture objective parameters, though including more sophisticated ways to decouple the appearance and stability objectives is an interesting direction for future work.

Our fill ratio-based simulation approach allows us to efficiently approximate the deformation behavior of a stenciled model. While we have shown that with increasing resolution, the simulation result converges to the solution obtained from a conforming mesh, for practical resolutions, our approach might not be able to capture any highly anisotropic deformation behavior introduced by stencils with very thin features. XFEM [Kaufmann et al., 2009] or a homogenization method [Kharevych et al., 2009] could be used to achieve a better approximation in these cases, at the cost of computational efficiency.

All computations of positions and distances of stencils are performed in Euclidean space, and while this has not led to any problems in the examples we showed, our projection approach cannot guarantee injectivity, which might lead to unsatisfactory results in the case of folded meshes. Incorporating geodesic distances could prevent artifacts in such situations, though a sufficiently fast and accurate computation of geodesic distances is necessary to preserve our current level of efficiency.

Our current stencil representation is sufficient to generate a variety of interesting appearances. However, the stencils are not required to be rigid objects, and incorporating advanced methods to modify the stencils would increase the expressiveness of our method.

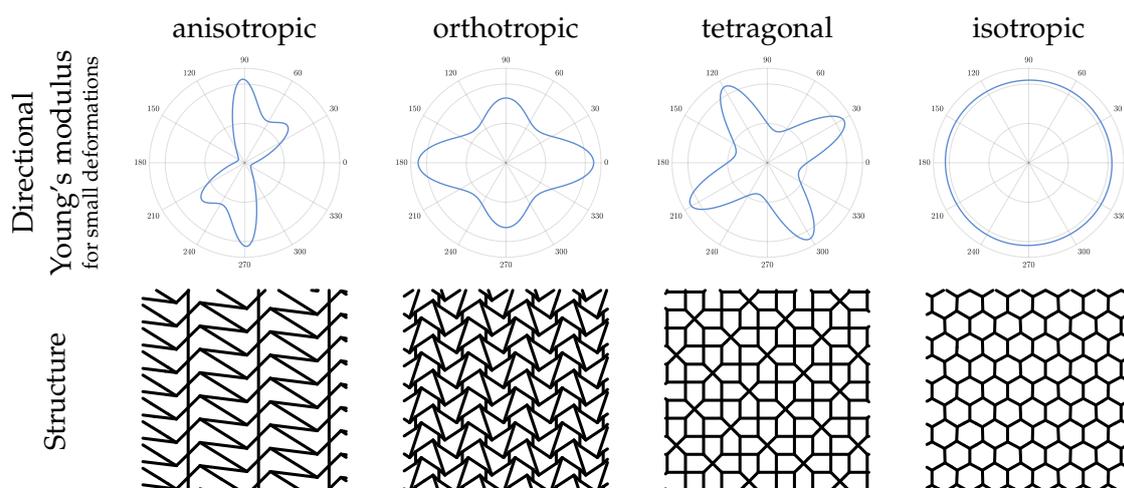
Finally, our general optimization-based approach allows us to easily integrate additional pattern and stability objectives, e.g., considering geometric features for the stencil distributions or minimizing displacement instead of energy density.

# Mechanical Characterization of Structured Sheet Materials

From simple shapes, *tessellations* can create structures of fascinating complexity and aesthetic appeal. But the diversity created by these patterns is not limited to appearance. Mechanical complexity can arise from simple shapes just as well as geometric complexity. Indeed, physical representations of these patterns—which we call *structured sheet materials*—exhibit a wide range of macromechanical properties mediated only by geometry and topology. Unlike appearance, however, the mechanical properties of structured sheet materials are often far from obvious. Exploring and understanding the connection between the geometry and mechanics of structured sheet materials is the main objective of this work. Unlike previous work that has so far focused on volumetric materials [Bickel et al., 2010; Schumacher et al., 2015; Panetta et al., 2015; Martínez et al., 2016], these structured sheets are planar networks of thin elastic rods that can stretch and *bend*.

In this chapter, we present a mechanical characterization approach that captures the macromechanical properties of these structures in a concise and intuitive way so as to allow for an informed exploration of the design space. To this end, we propose a homogenization approach that maps the mechanical behavior of a meso-scale network of elastic rods to a macro-scale thin plate model, making deliberate choices about which properties to capture and communicate. We compute a large set of deformations for each rod network and optimize for the tensors describing the anisotropic stretching and bending behavior of the sheet such that the thin plate model best approx-

## Mechanical Characterization of Structured Sheet Materials



**Figure 6.1:** *Two-dimensional linear materials can be classified into four symmetry groups: anisotropic, orthotropic, tetragonal, and isotropic. The various material symmetries lead to characteristic Young's modulus profiles and can be related to the structural symmetries of the underlying tiling.*

imates the large-scale properties of the rod network. The rather extensive material data of the thin plate model is then distilled into a compact visual representation in terms of directional Young's moduli, Poisson's ratios, bending stiffnesses, and corresponding measures of nonlinearity.

We focus on isohedral tilings, a particular class of polygonal periodic patterns in which all tiles are congruent to a single polygon. We show that this geometrically rich space of patterns offers an equally varied range of mechanical properties, characterized by strong anisotropy and nonlinearity for both stretching and bending deformations. The space of isohedral tilings is partitioned into a set of distinct families, each of which admits certain tileability-preserving transformations on the tile shape. Parameterizing these shapes within a given family yields a continuous sub-space of aesthetically pleasing patterns of the same topology but potentially vastly different mechanical properties.

### 6.1 Structured Sheet Mechanics

Structured sheets offer a rich space of heterogeneous materials with direction-dependent and nonlinear resistance to membrane and bending deformations. Our goal is to establish a formal but intuitive language for describing the mechanical properties of these materials, and to provide insight in the connection between their structure and mechanics.

In order for this analysis to be meaningful, we must be able to predict the deformation of structured sheets under imposed loads or boundary conditions using simulation (Section 6.1.2). To investigate their macromechanical behavior, we turn to numerical homogenization (Section 6.1.4) and condense the mechanics of the heterogeneous networks, simulated at their native mesoscopic level, into the familiar framework of linear elasticity. This process gives rise to elasticity tensors that characterize the direction-dependent stress-strain behavior around a given state of deformation. Since this representation does not readily lend itself to interpretation, we convert these elasticity tensors into radial plots of Young's modulus, Poisson's ratio, and bending stiffness (Section 6.1.5).

### 6.1.1 Phenomenology

In order to make an informed decision on the computational models to use in our analysis, we must first determine the salient characteristics that govern the deformation behavior of structured sheet materials. We are particularly interested in how deformations at the native, mesoscopic scale (from 1–5mm) of the material give rise to macroscopic deformation effects (from 5–10cm). To this end, we experimentally investigated the behavior of network structures subjected to different boundary conditions that create in-plane and out-of-plane deformations. We briefly summarize the central qualitative findings here and provide extensive, quantitative analysis in Section 6.3 and Appendix A.2.5.

**Membrane** Structured sheet materials are networks of interconnected rods. When imposing an external deformation onto the network, the individual rods will either stretch, bend, or twist—rods do not compress but buckle out of plane instead. For a given imposed deformation, the ratio between bent and stretched rods is an indicator of the stiffness of the network for that specific deformation. Since the resistance to bending is much smaller than the resistance to stretching, *bending-dominated* networks will be softer than *stretching-dominated* networks. For a given structure, this property, and therefore the overall stiffness, depends on the direction of the imposed deformation.

Beyond a certain maximum deformation, however, all rods will eventually be stretched and the structure will transition into a stiff regime. These two observations indicate that structured sheet materials are characterized by anisotropic and nonlinear in-plane stiffness.

**Bending** Similar to their in-plane deformation behavior, the stiffness response of structured sheets to bending is strongly anisotropic, but nonlinearities are generally less pronounced. It is insightful to draw the analogy to the simpler case of thin sheets made of homogeneous material, whose bending resistance depends only on their thickness and the mechanical properties of the bulk material. But unlike sheets made of homogeneous material, and due to the heterogeneity of structured sheets, the correlation between membrane and bending stiffness is only weak, making it necessary to treat bending separately from stretching to capture a complete picture of the mechanical properties.

With their salient characteristics spelled out, we can now proceed to the question of which computational models to use for structured sheets at the mesoscopic and macroscopic scales, respectively.

### **6.1.2 Mesoscopic Model**

To accurately model the mechanics of our networks at their native mesoscopic scale, their structure needs to be captured in sufficient geometric detail. As a good compromise between accuracy and efficiency, we opt to model structured sheets as networks of thin Kirchhoff rods.

Our implementation is based on the discrete elastic rod model [Bergou et al., 2010; Kaldor et al., 2010]. We use the extension to networks described by Zehnder et al. [2016] to model connections with more than two rod segments: we place a rigid body frame at each such connection, and define the bending and twisting energies at the connection not between two neighboring segments, but between a segment and the rigid body frame. The frame is added to the simulation variables, and will rotate to minimize the energy. The discrete rod model can be applied almost directly to our case, except for one important modification: due to the complex structure of the networks, stretching the material in a given direction will also induce compressions in some rods. In physical reality, these compressions immediately resolve into out-of-plane buckling, but under the perfectly planar and symmetric conditions in simulation, compressions will persist and give rise to unstable equilibrium configurations. Besides the associated numerical difficulties, persistent compressions lead to an overall much stiffer deformation response than what is observed in reality. To avoid such parasitic stiffening, we use a full three-dimensional model even for deformations that are, on the macroscopic level, in-plane, and slightly perturb the initial state into the normal direction, forcing compressions to resolve into bending. Thanks to this modification, the rod model shows very good agreement with physical experiments, both

in terms of force-deformation behavior and local deformations observed in the structures; see Section 6.3.1.

### 6.1.3 Macroscopic Model

On the macroscopic level, the most salient characteristic of structured sheets is arguably their direction-dependent stiffness response to stretching and bending deformation. Together with their thin nature and planar rest state, this deformation behavior suggests a macromechanical model based on the theory of anisotropic Kirchhoff plates (see, e.g., [Hwu, 2010, Section 1.4]), which postulates a strain energy density of the form

$$W(\boldsymbol{\epsilon}, \boldsymbol{\kappa}) = \frac{1}{2} \boldsymbol{\epsilon} : \mathbf{C} : \boldsymbol{\epsilon} + \frac{1}{2} \boldsymbol{\kappa} : \mathbf{B} : \boldsymbol{\kappa} = W^M + W^B, \quad (6.1)$$

where  $\boldsymbol{\epsilon}$ ,  $\boldsymbol{\kappa}$  are membrane and bending strains, respectively, and  $\mathbf{C}$ ,  $\mathbf{B}$  are corresponding material stiffness tensors. The strains and stresses are rank-2 tensors that can be represented as symmetric 2x2-matrices;  $\mathbf{C}$  and  $\mathbf{B}$  are symmetric rank-4 tensors with 16 entries. Due to their symmetry, only 6 entries can be independent, even in the case of complete anisotropy. We can see that this model is an extension of the two-dimensional linear model (see Section 3.2), with an additional term for bending deformations. Differentiating the energy density with respect to membrane and bending strains gives rise to membrane and bending stresses,

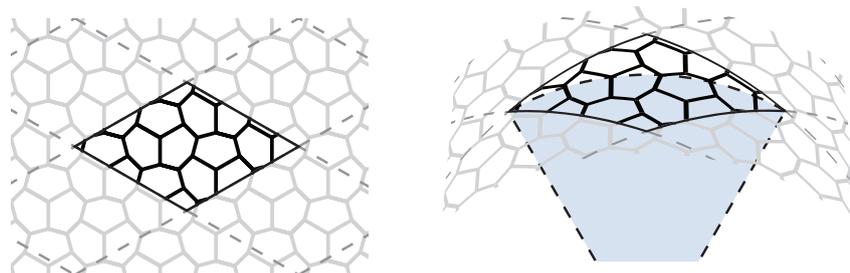
$$\begin{aligned} \boldsymbol{\sigma} &= \mathbf{C} : \boldsymbol{\epsilon} \\ \mathbf{M} &= \mathbf{B} : \boldsymbol{\kappa}. \end{aligned} \quad (6.2)$$

The above expressions underline the linear nature of this material model—we will describe extensions to account for nonlinearities in Section 6.1.4—and the decoupling of the membrane and bending energies. While this model only approximates the full nonlinear material behavior, it is sufficient to capture the salient material properties that we are interested in.

With the mesoscopic and macromechanical models defined, we can now proceed to the mapping between the two via homogenization.

### 6.1.4 Homogenization

The essential idea of homogenization is to subject a tileable *unit cell* of material to various boundary conditions that create different states of deforma-



**Figure 6.2:** *Periodic boundary conditions for in-plane deformations emulate a flat tiling of patches (left). For curved configurations, the tiling is mapped to a cylinder, where adjacent patches are related through a rotation and translation (right). Additional details can be found in Section 3.3.3 and Appendix A.2.1.*

tion. These boundary conditions have to preserve tileability on the mesoscopic scale while at the same time leading to deformations that are easily quantified on the macroscopic scale. Figure 6.2 shows the general concept of these boundary conditions: for flat configurations, they emulate a tiling of a single patch of the structure in the plane, while curved configurations are possible if we introduce a rotational component to the tiling. We refer to Section 3.3.3 for the periodic in-plane boundary conditions. They follow the same principle as the boundary conditions used in Chapter 4, but while we prescribed a specific stress to the microstructures, we apply a *uniaxial strain* in this case. For the curved periodic boundary conditions, we refer to Appendix A.2.1 for a detailed description.

When using these periodic boundary conditions, special care has to be taken to ensure that the simulated structure is not too small to capture all deformation properties of the full structure. Specifically, the buckling modes of a structure can often extend over two of the smallest tileable unit cells. To capture these buckling modes properly, we always apply our homogenization on a  $2 \times 2$  patch of the smallest tileable unit cells.

While classic plate theories generally establish a shared set of parameters that connect the membrane and bending behavior—both are fully described by a single Young’s modulus and Poisson’s ratio in the case of a homogeneous isotropic plate—this connection has been found to disappear in structured materials [O’Donnell and Langer, 1962]. Having confirmed these findings in our own physical experiments, we perform homogenization independently for membrane and bending properties.

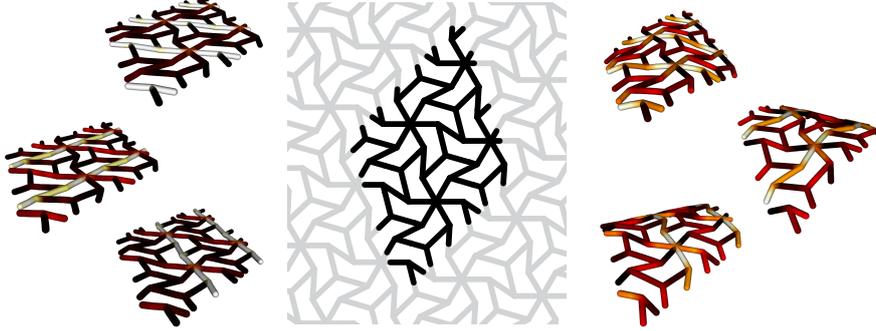
**Membrane Material Tensor Fitting** Given a structured sheet material, we seek to find elasticity parameters for the macroscopic model that best approximate the true stress-strain behavior of the mesoscopic model. To this end, we must first determine a set of *test deformations* and corresponding boundary conditions that will be used to probe the material in simulation. For the membrane part, a natural choice are uniaxial tension states, for which the strain in a specific direction is prescribed while the perpendicular direction is only subject to periodic boundary conditions (see Figure 6.3). Uniaxial deformations generally describe the material behavior sufficiently well, but we found that some structures require more information to fit a macroscopic material tensor. Whenever such a case is detected (from the rank deficiency of the fitting matrix), we augment the test set with an additional biaxial deformation that prescribes a uniform stretch in all directions.

Once the set of test deformations is defined, we perform meso-scale simulations and compute the macroscopic strains and stresses using the approach described in Appendix A.2.2. The resulting set of  $N$  stress-strain pairs are then used to fit a homogenized *compliance tensor* in a least squares sense,

$$\mathbf{S}^H = \operatorname{argmin}_{\mathbf{S}} \sum_{i=1}^N \frac{1}{\|\epsilon_i\|_F^2} \|\mathbf{S} : \sigma_i - \epsilon_i\|_F^2, \quad (6.3)$$

where  $\epsilon_i$  and  $\sigma_i$  are the strain and stress of the  $i$ -th test deformation, respectively, and  $\|\cdot\|_F$  is the Frobenius norm. The solution to Equation (6.3) is a homogenized material compliance tensor  $\mathbf{S}^H$ , and the corresponding homogenized material stiffness tensor can be computed through the symmetric inverse,  $\mathbf{C}^H = (\mathbf{S}^H)^{-1}$ . One could also fit the homogenized material stiffness tensor  $\mathbf{C}^H$  directly; however, we found that since the significant anisotropy of certain structures leads to extreme stresses, measuring the error in strain space is more robust.

**Bending Material Tensor Fitting** The bending behavior of a material is characterized by its bending moment response to applied curvatures. Similar to the in-plane stiffness, we design a set of test deformations that fully describe this behavior. Cylindrical bending is a natural choice for this purpose, since this type of deformation does not induce Gaussian curvature and hence minimizes interference with membrane deformations. Cylindrical bending alone, however, is not sufficient to uniquely determine the bending stiffness tensor, as is apparent by the rank deficiency of the fitting matrix. We therefore account for the missing data by augmenting the cylindrical test deformations by a single deformation with spherical curvature. For each of



**Figure 6.3:** Example deformations obtained by applying different loads to the boundary of a periodic patch of a structure. Left: membrane deformations corresponding to uniaxial stretching. Right: bending deformations corresponding to uniaxial (or cylindrical) curvature.

the  $M$  boundary conditions in the resulting test set, we then perform meso-scale simulations to obtain the corresponding deformed network.

While the direct analogy to the membrane case would be to compute macroscopic bending moments for the deformed networks, the curved patch boundaries complicate the direct computation of moments. We therefore use the energy density of the deformed patch in order to fit the homogenized bending stiffness

$$\mathbb{B}^H = \operatorname{argmin}_{\mathbb{B}} \sum_{i=1}^M \left( \frac{1}{2} \boldsymbol{\kappa}_i : \mathbb{B} : \boldsymbol{\kappa}_i - W_i \right)^2, \quad (6.4)$$

where  $\boldsymbol{\kappa}_i$  and  $W_i$  are the prescribed curvature tensor and energy density for the  $i$ -th test deformation, respectively.

**Determining the Test Set** The set of deformations that are used to fit the stiffness tensors influence the quality of the result. While a set of only three deformations is sufficient to fit the tensor for infinitesimal strains, the nonlinearities that appear for finite strains require that we choose a suitable set of deformations that captures the material properties in sufficient detail and will reproduce the global deformation behavior in the fitted material tensor. We found that for most structures, a small number of up to 20 deformations is sufficient for a robust fitting to a linear material tensor. For highly anisotropic and nonlinear materials, however, a larger number of deformations generally leads to improved fitting results. To this end, we use a cross-validation scheme that adaptively determines the number of test deformations needed to obtain robust homogenization results. From a given set of

$N$  test deformations, we use multiple sets of 80% as training data to fit a stiffness tensor, then measure its performance on the remaining 20% of the test set. If the error on the test data is larger than on the training data, we increase the number of test deformations to  $2N$  and repeat the process.

**Nonlinearities** The homogenized stiffness tensors provide a linearized description of a material’s direction-dependent stress response. For a given structure, the result of the fitting process is influenced by two factors: the reference state of deformation for which the tensor is computed, and the magnitude of the imposed deformation away from the reference state. The choice of reference state, not to be confused with the rest state, is important when there is nonlinear coupling between deformation modes. In principle, a comprehensive description of the nonlinear material behavior can be obtained by densely sampling the deformation space. However, besides the computational complexity of such an approach, the massive amount of high-dimensional data points generated in this way would simply be overwhelming for the user. We argue that, for the purpose of understanding and navigating the space of materials, conciseness is far more important than completeness. In order to convey a compact description of a material’s direction-dependent nonlinearity, we simply create multiple test sets with uni- and biaxial deformations of varying magnitude (0.1% and 10.0%) and curvature ( $0.1\text{m}^{-1}$  and  $5\text{m}^{-1}$ ) around the rest state. Taken together, the corresponding stiffness tensors fitted to these different test sets can then be condensed into a concise and intuitive description of nonlinear, direction-dependent material behavior as described next.

### 6.1.5 Representing Direction-Dependent Elasticity

Although the stiffness tensors provide a full picture of the material behavior around a given state of deformation, their entries are rarely used to describe a material: they are hard to correlate to quantities that can be measured directly and, consequently, provide little intuition about the material behavior. In the following, we describe how to distill these stiffness tensors into representations that lend themselves more readily to interpretation. We will briefly introduce the formulas to extract these representations from the stiffness tensors, but refer to Appendix A.2.3 for a more in-depth explanation.

**Membrane** An intuitive set of measures for the membrane behavior are the *Young’s modulus*  $E$ , describing the force per area required to stretch the material to one unit of deformation, and the *Poisson’s ratio*  $\nu$ , which is the amount

of transverse contraction for a unit extensional deformation. Taken together, Young's modulus and Poisson's ratio completely characterize the behavior of linear isotropic materials. By allowing  $E$  and  $\nu$  to vary depending on the direction  $\mathbf{d}$ , they can also be used to describe anisotropic materials.

In order to derive direction-dependent expressions for  $E$  and  $\nu$ , we use the compliance tensor  $\mathbb{S} = \mathbb{C}^{-1}$ . This tensor allows us to compute the deformation induced by the uniaxial unit stress  $\boldsymbol{\sigma}^{\mathbf{d}} = \mathbf{d}\mathbf{d}^T$ , which is the configuration in which the two material properties are measured. The directional Young's modulus can then be computed as

$$E(\mathbf{d}) = \frac{1}{(\mathbf{d}\mathbf{d}^T) : \mathbb{S} : (\mathbf{d}\mathbf{d}^T)}. \quad (6.5)$$

Similarly, we compute the directional Poisson's ratio, describing the relative compression of the material along the direction  $\mathbf{n}$  perpendicular to the stretch direction  $\mathbf{d}$ , as

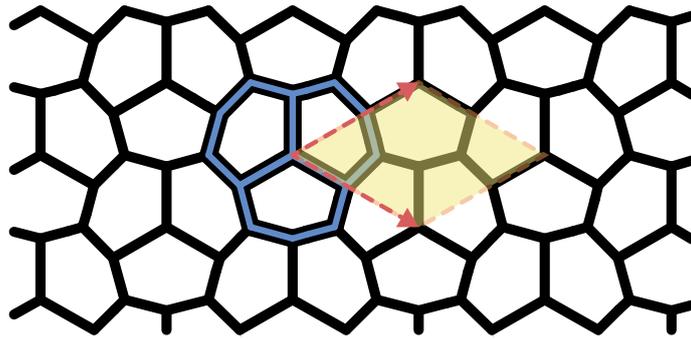
$$\nu(\mathbf{d}) = -\frac{(\mathbf{d}\mathbf{d}^T) : \mathbb{S} : (\mathbf{n}\mathbf{n}^T)}{(\mathbf{d}\mathbf{d}^T) : \mathbb{S} : (\mathbf{d}\mathbf{d}^T)}. \quad (6.6)$$

**Bending** We can follow a similar approach for the characterization of the bending behavior of a structured sheet material. However, we found that the Young's modulus equivalent of the bending stiffness matrix, which is measured by applying a uniaxial bending moment to the structure, describes a state that is hard to reproduce in a real application, since it will automatically cause in-plane deformations that generally dominate the deformation mode. Instead, we characterize the bending behavior of the material using a natural low energy state of purely cylindrical curvature, or zero Gaussian curvature. Given a direction  $\mathbf{d}$ , we compute the directional bending moment generated by a unit curvature,

$$b(\mathbf{d}) = (\mathbf{d}\mathbf{d}^T) : \mathbb{B} : (\mathbf{d}\mathbf{d}^T). \quad (6.7)$$

## 6.2 Tilings and Symmetries

Apart from periodicity constraints, the machinery introduced so far makes no assumptions on the nature of the patterns. To investigate the interplay between structure and mechanics, we will from now on focus on the class



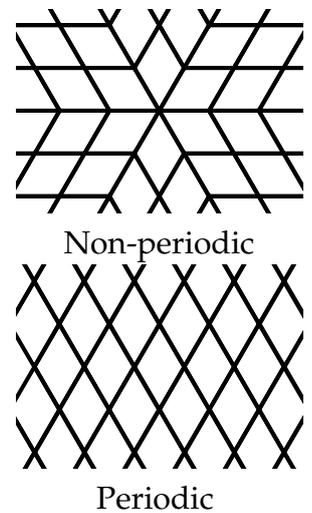
**Figure 6.4:** The translational unit (blue) of an isohedral tiling creates an infinite periodic tiling of the plane using only translations (red). The translation vectors define the minimal tileable unit cell (yellow) necessary for a simulation with periodic boundary conditions.

of isohedral tilings, which we briefly introduce below. We pay special attention to the symmetry properties of the patterns and how they relate to symmetries in the material behavior.

### 6.2.1 Isohedral Tilings

Isohedral tilings are a special class of monohedral tilings—tessellations built from a single base tile—in which only tiles transformed under the symmetry group of the tiling are allowed. In practical terms, this means that an isohedral tiling has translational periodicity (see inset figure).

From a combinatorial point of view, 93 different families of isohedral tilings can be identified [Kaplan, 2009], referred to individually as IH01 to IH93. Each family has a polygonal base tile ranging from a triangle to a hexagon, with a parameterization having zero to six degrees of freedom. Up to 12 base tiles then form a *translational unit*, the smallest structure of a family that can be tiled infinitely by only applying translations (see Figure 6.4). For our structures, we create a consistent scale between families by normalizing the area of the translational unit.



**Symmetries** Isohedral tilings exhibit a high degree of visual regularity that can be quantified in terms of symmetries. Intuitively, the individual

translational units of an isohedral tiling span a hexagonal or parallelogrammatic grid. Such a grid, and therefore the tiling, may have twofold, threefold, fourfold or sixfold rotational symmetry. Similarly, isohedral tilings can have glide reflection symmetries—a combination of reflection and translation along the reflection axis—along one, two, three, four, or six reflection axes. We can observe all of these rotational symmetries in our isohedral tilings, and, as it turns out, these are also the only rotational and reflectional symmetries that isohedral tilings can have [Grünbaum and Shephard, 1986].

Besides their visual impact, geometric symmetries in the tiling also induce symmetries in the material behavior. In order to investigate the nature of this relation, we first review material symmetries below.

### 6.2.2 Material Symmetries

The mechanical properties of a material can exhibit various symmetries. For example, any two-dimensional material rotated by  $180^\circ$  will still show the same material properties. Additional symmetries allow us to define different *symmetry classes* for materials that characterize their qualitative behavior. For a two-dimensional linear material model, there are four distinct symmetry classes: isotropic, tetragonal, orthotropic, and (fully) anisotropic materials.<sup>1</sup> See Figure 6.1 for example structures from these categories and their corresponding Young's modulus profiles. Rather than determining the symmetry class through visual inspection of the Young's modulus profiles, it can be identified from a set of five invariants of the material stiffness tensor  $\mathbf{C}$  that measure the various rotation-invariant non-symmetries of the tensor [de Saxcé and Vallée, 2013].

*Anisotropic materials* are the most general symmetry class and do not possess any additional symmetries.

*Orthotropic materials* have distinct material properties along two orthogonal directions, which introduces a reflection symmetry along these two axes. This symmetry is apparent in the directional Young's modulus plot of a material, where two axes that show a local maximum or minimum are perpendicular to each other.

*Tetragonal materials* are special types of orthotropic materials whose material properties along any pair of orthogonal directions are identical, making

---

<sup>1</sup>Unless specifically noted, we use the name of a symmetry class to refer to patterns that have that symmetry but not any greater symmetry; for example *orthotropic* means the material is orthotropic but not tetragonal or isotropic.

them invariant to rotations of  $90^\circ$ , and giving them a total of four equally-spaced axes of reflection. Due to the rotational symmetry, tetragonal materials are easily identifiable from their Young's modulus plot.

*Isotropic materials* are the most constrained type of elastic material. They show a direction-independent material behavior, and are therefore invariant under any rotation or reflection. On a Young's modulus plot, isotropic materials can be easily identified as circles.

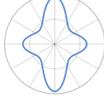
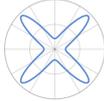
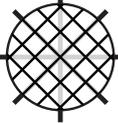
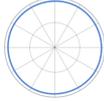
### 6.2.3 Shared Symmetries

We can now investigate the connection between the symmetries of the material properties and the geometric symmetries of our structured sheet materials, and we summarize the mapping between the symmetries in Table 6.5. Note that this mapping is one-way—while the geometry of the structure dictates the necessary material symmetries, a specific material symmetry does not require any geometric symmetry. Figure 6.10 shows an example of a structure that is isotropic, but does not possess the geometric symmetries that imply this material symmetry.

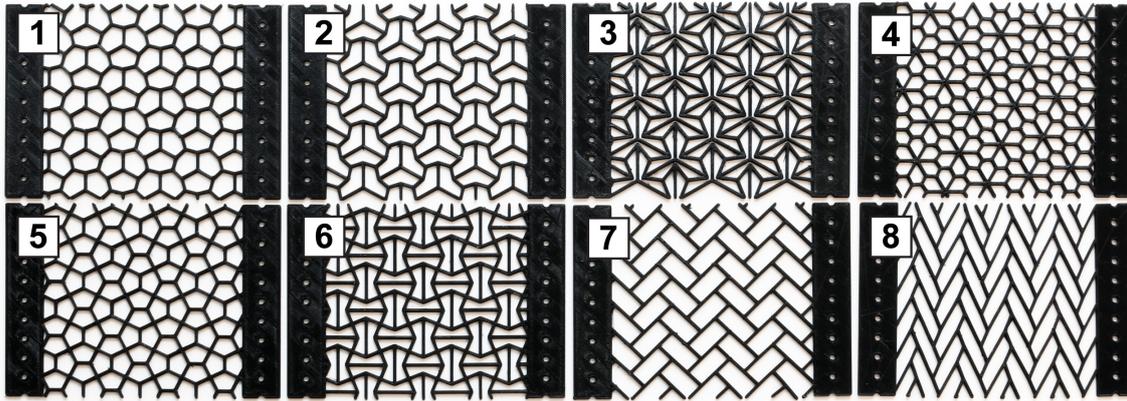
The most important factor when comparing geometric symmetries to material symmetries is the inherent rotational symmetry of the material properties: every material property is invariant to rotations of  $180^\circ$ . Tilings with a twofold rotational symmetry will therefore not add any new symmetries to the material properties, and in general, the material symmetry will be a combination of the geometric symmetry and a  $180^\circ$  rotation. A special case are structures with threefold and sixfold geometric symmetries, and three-axis and six-axis geometric glide reflection symmetries, which all show the same material response along at least three axes, or six directions. In the case of a linear material, these responses provide a complete basis of the underlying space, and since all responses are rotationally symmetric, every other response must also be rotationally symmetric, making the material *isotropic*.

## 6.3 Results

We used our method to characterize a large number of structures from the space of isohedral tilings and present our findings in this section. We first start with general observations, then highlight specific properties on selected examples. Finally, we present extensions to (inverse) material design to create tilings with a desired deformation behavior.

<i>Geometric symmetry</i>			<i>Material symmetry</i>
 Twofold rotation	$\Rightarrow$		Anisotropic
 One-axis reflection	$\Rightarrow$		Orthotropic
 Two-axis reflection			
 Fourfold rotation	$\Rightarrow$		Tetragonal
 Four-axis reflection			
 Threefold rotation	$\Rightarrow$		Isotropic
 Sixfold rotation			
 Three-axis reflection			
 Six-axis reflection			

**Figure 6.5:** The different symmetry classes for isohedral tilings (left) and the corresponding material symmetries that they induce (right).

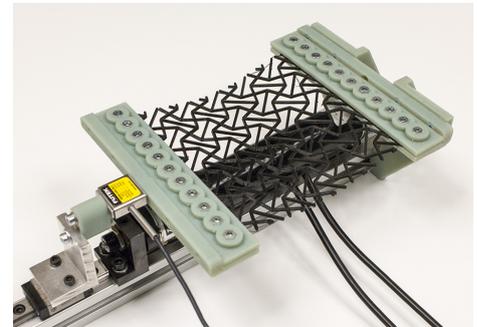


**Figure 6.6:** The structures used in the experimental validation and for the characterization in Figure 6.9, covering a wide range of visual styles and mechanical properties. Structures 1 to 3 are part of the IH07 family, structure 4 is from IH21, structures 5 and 6 from IH29, and structures 7 and 8 are part of IH02.

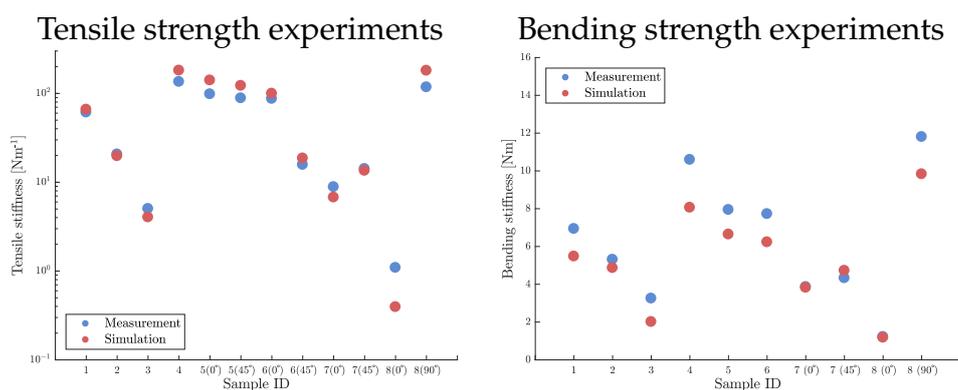
But first, we will show a validation of the rod model introduced for the mesoscopic simulation in Section 6.1.2.

### 6.3.1 Mechanical Testing

We validated our simulation model on a set of tension and bending tests, using the structures shown in Figure 6.6. The structures were printed as patches of  $100\text{mm} \times 100\text{mm} \times 1.2\text{mm}$  on an Ultimaker 2 printer with a Flex3Drive extruder, using the rubber-like NinjaFlex material. The tensile testing setup, shown in the inset figure, uses a linear actuator and a load cell to perform uniaxial tension tests, automatically capturing reaction forces of samples under tensile loads. We measured the Young's modulus of the base material to be 15.5 MPa, based on a tensile test of a *dogbone* (tensile specimen), and determined a Poisson's ratio of 0.48 based on standard values for TPU materials [Qi and Boyce, 2005]. Using these values in simulation leads to results consistent with the experimental data for a large range of structures (see Figure 6.7).

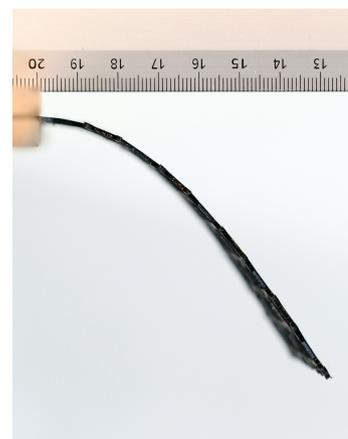


Since the bending stiffness of structured sheets is significantly lower than its in-plane stiffness, the self-weight of the material needs to be accounted for during measurements. We therefore opted to follow Pabst et al. [2008] and



**Figure 6.7:** The tensile stiffnesses (left) and bending stiffnesses (right) measured in our experiments (blue), compared to our simulation results (red), at 10% strain and  $20m^{-1}$  curvature, respectively. The full measurements can be found in Appendix A.2.5.

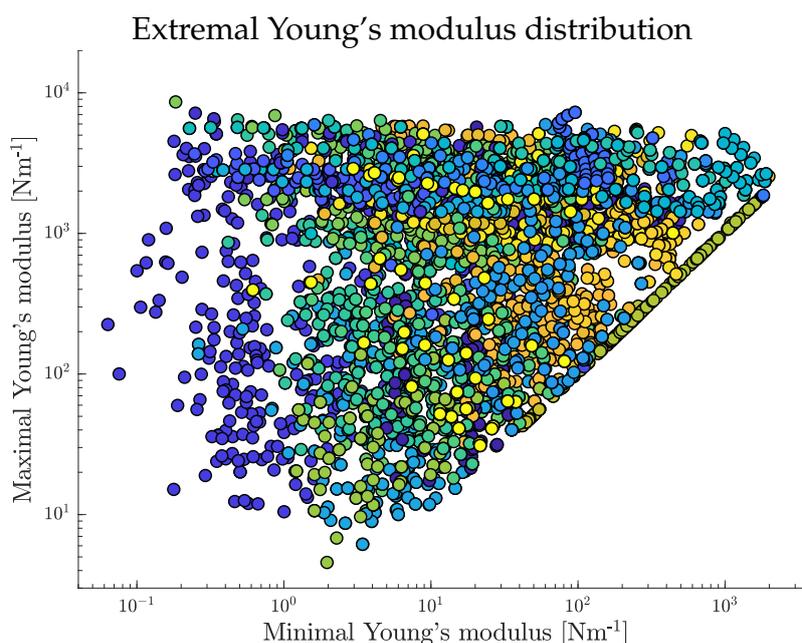
Miguel et al. [2013], measuring the curvature-moment curves of our samples using the gravity-assisted method of Clapp et al. [1990], which uses visual data and polynomial fitting to extract material parameters from a specimen (see inset). We use this method on both the printed and simulated structures. Figure 6.7 shows an overview of the resulting bending stiffnesses for small curvatures. As can be seen from the data, the deviation between measurements and simulations is larger for the bending tests than for the tensile tests. We speculate that this difference is largely due to the inaccuracies of the visual capture approach and the increased uncertainty of the fabrication process along the height dimension.



The full results for both the tensile and bending tests can be found in Appendix A.2.5.

### 6.3.2 Space of Structures

Applying our mechanical characterization approach to many structures from the space of isohedral tilings reveals the *gamut* of mechanical properties that they cover. The visualization of this gamut is, however, not a trivial task. Figure 6.8 shows a plot of the minimal and maximal directional Young's moduli of all isohedral tiling families, using a regular sampling of the tiling parameters in the interval  $[-1, 2]$  resulting in around 5500 valid



**Figure 6.8:** *The gamut of directional Young's moduli covered by isoedral tilings, plotted as the minimal and maximal directional Young's modulus of each structure of our sampling. The colors of the data points represent the isoedral tiling family the structure belongs to.*

structures. We used the material parameters presented in Section 6.3.1 on patches with an area of  $16 \text{ cm}^2$ , a rod diameter of  $1.5 \text{ mm}^2$ , and a maximum segment length of  $2.5 \text{ mm}^2$ . The data shows that our sampling covers a directional Young's modulus range of roughly five orders of magnitude, with the minimal directional Young's modulus of a structure covering four orders of magnitude, and the maximal directional Young's modulus covering three orders of magnitude. The bending stiffness shows a smaller variation, covering a range of two orders of magnitude. While this visualization is useful to determine the practical boundaries of Young's moduli we can expect from these structures, it does not lend itself to exploration, and only shows a limited view of the mechanical properties.

We found that a proper exploration of the mechanical properties of such a space is best done interactively. We organized our dataset into an interactive material browser that offers visualization of various mechanical properties for the whole dataset or individual families, as well as a detailed mechanical characterization and visualization of individual structures. In such a way, a user can easily narrow down the number of structures based on their requirements, and then interactively explore the combined aesthetic and mechanical properties of the rest.

**Change of Rod Parameters and Dimensions** The results of the mechanical characterization of isohedral tilings are specific to the rod parameters that we determined in Section 6.3.1. While the differences in geometry of the individual structures allow us to navigate within the gamut of mechanical properties, these rod parameters allow us to modify the gamut itself. The underlying mesoscopic simulation model directly specifies the parameters that influence the characterization, and therefore this gamut: the Young's modulus and Poisson's ratio of the base material, the cross sectional area of the rods, and the size of a single patch of the tiling.

The Young's modulus appears as a linear factor in the formulation of the rod energy, and allows us to linearly scale the mechanical properties connected to the stiffness, i.e., Young's modulus and bending stiffness.

The Poisson's ratio of the base material is only used in the twisting energy of the mesoscopic simulation. Since twisting has a negligible effect on the mechanical in-plane behavior in the structures we investigated, the membrane characterization is largely independent of the Poisson's ratio. The bending behavior, on the other hand, can show significant twisting of rods. Since a rod's resistance to twisting will increase with lower Poisson's ratio, the whole gamut of macroscopic bending stiffness will shift towards higher values if we decrease this parameter.

Finally, the cross sectional area and patch area are linked quantities. If we keep their ratio, but scale them, the material properties will scale by the same amount. Interesting effects emerge when this ratio changes, and we can investigate this effect by keeping the patch size fixed and increasing the rod diameter: the mesoscopic simulation model tells us that the stretch energy term of the rod model will scale linearly with the cross sectional area, while the bending and twisting energy terms scale quadratically. The change in elastic response of a structure then lies within these bounds, and the actual scaling depends on whether the structure is bending-dominated or stretching-dominated: for bending-dominated structures, it will be close to quadratic, while for stretching-dominated structures, it will be close to linear.

Since stretching-dominated structures create the stiffest elastic response, and bending-dominated structures are generally softer, this scaling effectively shrinks the gamut of properties if we increase the rod cross section. Additionally, the difference in scaling leads to a decrease in buckling and out-of-plane bending in the structures. Accordingly, a decrease in cross section has the opposite effect: it increases the gamut of material properties, and encourages more buckling and out-of-plane bending.

### 6.3.3 Analysis of Individual Structures

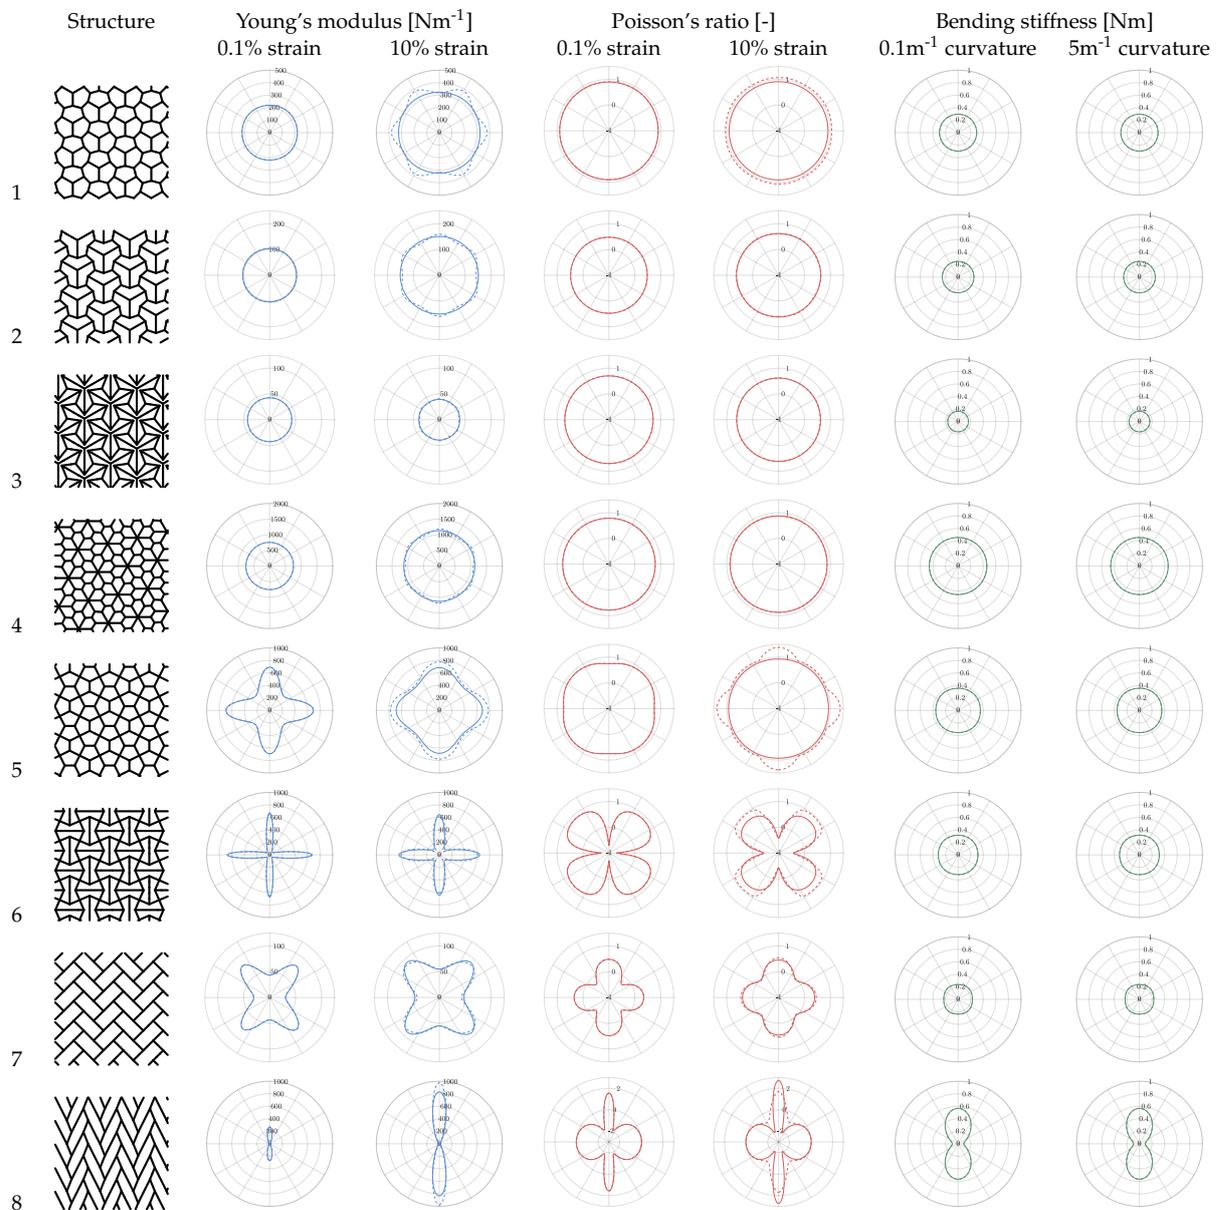
We now highlight a number of interesting properties on individual structures. In Figure 6.9, we present an extensive analysis of the eight structures already used in the validation of the simulation, shown in Figure 6.6. We chose structures from four different families (IH02, IH07, IH21, and IH29), showing the potential differences in material properties between structures from the same family and across different families. For each structure, we plot the directional Young's modulus, Poisson's ratio, and bending stiffness computed from the homogenized stiffness tensors, at two different strain magnitudes (0.1% and 10%) or curvatures ( $0.1\text{m}^{-1}$  and  $5\text{m}^{-1}$ ) to show the non-linearity of the material properties. Besides these homogenized values, we also show the ground truth simulation data, i.e., the values extracted from the same type of uniaxial or cylindrical simulation used in the homogenization. We extract the Young's modulus  $E$  and Poisson's ratio  $\nu$  along a direction  $\mathbf{d}$  from the simulation with uniaxial stretch along  $\mathbf{d}$  as  $E(\mathbf{d}) = \mathbf{d}^T \boldsymbol{\sigma}_{macro} \mathbf{d}$  and  $\nu(\mathbf{d}) = \mathbf{n}^T \boldsymbol{\epsilon}_{macro} \mathbf{n}$ , where we use the macroscopic strain  $\boldsymbol{\epsilon}_{macro}$  and stress  $\boldsymbol{\sigma}_{macro}$  described in Appendix A.2.2, and the direction  $\mathbf{n}$  normal to  $\mathbf{d}$ . The directional bending stiffness is computed as the energy density divided by the squared curvature.

**Isotropy & Anisotropy** Structures 1 to 4 in Figure 6.9 show a very good fit for an isotropic material behavior for small strains, predicted by their three-fold rotational symmetry (see Section 6.2.2). The variations in Young's moduli between the structures are large, and especially structures 1 to 3 demonstrate that even for a single family, we can expect significant variations in the elastic responses.

The remaining structures display different types of anisotropic material behavior. Structures 5 and 6, which belong to the same family, show orthotropic material properties. The directional Young's moduli of structure 6 vary by an order of magnitude depending on the direction, and its Poisson's ratio switches from positive to negative for stretching along the main axes.

Structures 7 and 8, which also belong to the same family, show a transition from a tetragonal material to an orthotropic material whose Young's modulus along the horizontal direction is two orders of magnitude smaller than along the vertical direction. This change can be explained by the structure transitioning from a bending-dominated elastic response to a stretching-dominated elastic response along the vertical axis, creating a stiffer mechan-

## Mechanical Characterization of Structured Sheet Materials



**Figure 6.9:** Simulation results for all structures listed in Figure 6.6. We show the directional Young's modulus and Poisson's ratio at 0.1% and 10% strain, visualizing the homogenized result (solid line) as well as the simulated value (dashed line). The last two columns show the directional bending stiffness for a curvature of  $0.1\text{m}^{-1}$  and  $5\text{m}^{-1}$ . Only for larger strains is there a difference between the homogenized result and the full simulation.

ical response. Conversely, the response along the horizontal axis becomes more bending-dominated for structure 8, making it softer.

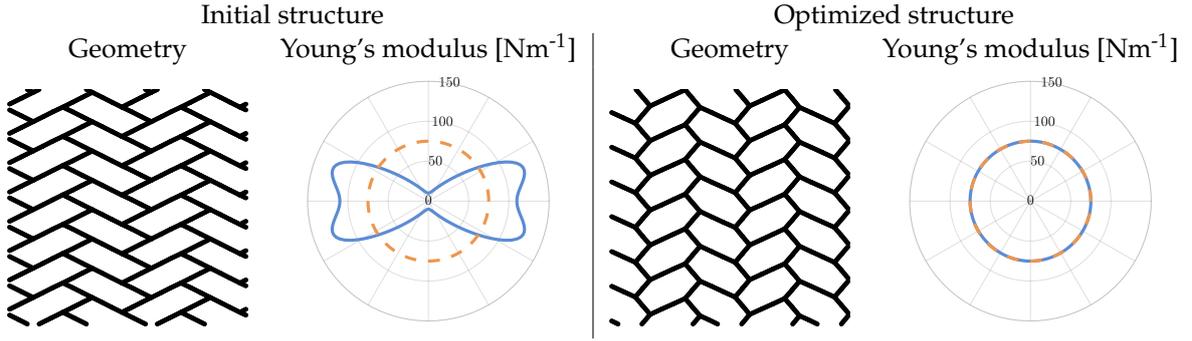
The bending-dominated or stretching-dominated nature of the elastic response plays an important role in the mechanical properties of these structures. For example, structure 1 and 5 are visually similar, but the directional Young's modulus of structure 5 is up to three times higher. This stems purely from the fact that the geometry of structure 5 contains straighter paths, creating a more stretching-dominated response along these directions.

**Nonlinearity** The transition from a bending-dominated deformation to a stretching-dominated deformation can not only be observed when we change the geometry of a structure, but also if we increase the strain on a structure. Indeed, most nonlinear effects we can observe in the structures are due to this transition. For example, when the zigzag pattern in structure 8 in Figure 6.9 unfolds, the induced stretch creates a significantly stiffer elastic response. An analogy can be drawn to the *stiffness percolation* in random materials (see, e.g., [Wilhelm and Frey, 2003]), i.e., the point at which a material forms a path that is aligned with a deformation, significantly increasing its stiffness.

Structures 1, 2 and 4 all exhibit stiffness percolation, shown by the increase in stiffness for larger strains. Additionally, while they show an isotropic behavior for small strains, they develop an increasingly anisotropic material response for larger strains, indicated by the six bumps in the directional Young's modulus that align with the symmetry directions of the structure. While the linear approximation cannot capture this anisotropy, it still achieves a good fit for the average Young's modulus.

A different type of nonlinear effect can be observed in structure 3. While it shows the same transition to an anisotropic material response as structures 1, 2 and 4, the average stiffness decreases slightly for larger strains. This decrease can be attributed to the out-of-plane deformations that the structure experiences, allowing it to avoid stiffness percolation for the applied strains.

**Bending** The bending response of our structures shows less pronounced anisotropic behavior than their membrane resistance. Even for structure 8, whose Young's moduli span two orders of magnitude, the ratio between the stiffest and softest bending direction is around 5. Additionally, the bending does not exhibit any stiffness percolation, making the bending response more linear, with little change in stiffness between different curvatures.



**Figure 6.10:** Starting from a structure from IH09 (left) with an orthotropic material behavior (blue), and a desired isotropic target material behavior (orange), our structure optimization is able to find a similar structure with the desired material behavior (right).

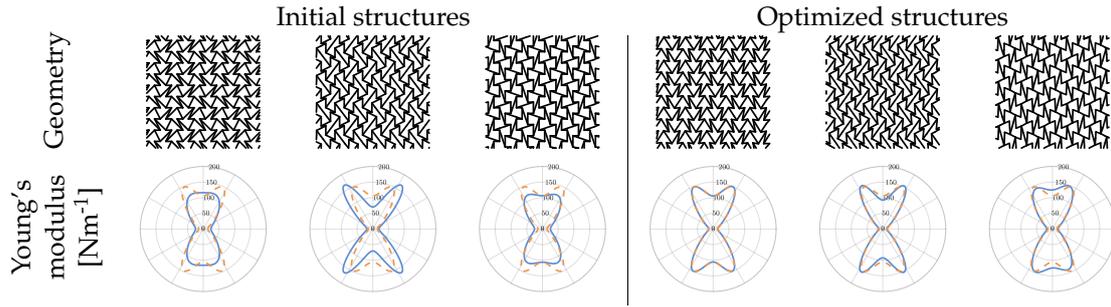
### 6.3.4 Structure Optimization

So far, we have set structure parameters and observed the mechanical behavior that they induce. A natural extension to this forward exploration is to search for structure parameters that lead to desired mechanical properties. As a proof-of-concept for such an inverse design approach, we optimize for the directional Young's modulus of a structure.

**Objective** We define an objective based on Equation (6.5) by measuring the difference in the directional Young's modulus of our structure and a target function. The directional Young's modulus  $E(\mathbf{d})$  is sampled along a set of  $n$  directions  $\mathbf{d}(\phi_i)$  with equidistant angles  $\phi_i$  between 0 and  $\pi$ . Using the structure parameters  $\mathbf{p}$  and an additional rotation parameter  $\alpha$  as variables, the optimization objective then becomes

$$O(\mathbf{p}, \alpha) = \frac{1}{2n} \sum_{i=1}^n \left( \frac{E(\mathbf{d}(\phi_i); \mathbf{p}, \alpha)}{E^{goal}(\mathbf{d}(\phi_i))} - 1 \right)^2. \quad (6.8)$$

**Optimization** A simple gradient-based optimization approach is already sufficient to create structures with desired elasticity profiles. We opt for a gradient descent approach with basic backtracking line search. While the derivative of a homogenized material stiffness tensor with respect to the structure parameters can be computed from a series of chain rules, the small number of structure parameters (up to six) allows for an efficient computation of the derivative through finite differences.



**Figure 6.11:** The results of a structure optimization in the family IH06 for an orthotropic Young's modulus target. The Young's modulus for each initial (left) and optimized structure (right) are shown in blue, along with the target in orange.

Figure 6.10 shows the result of optimizing for the isotropy of a structure. Given an initial pattern and a target Young's modulus, the optimization finds a similar structure with isotropic material behavior, even though the isotropy cannot be inferred from the geometric symmetries.

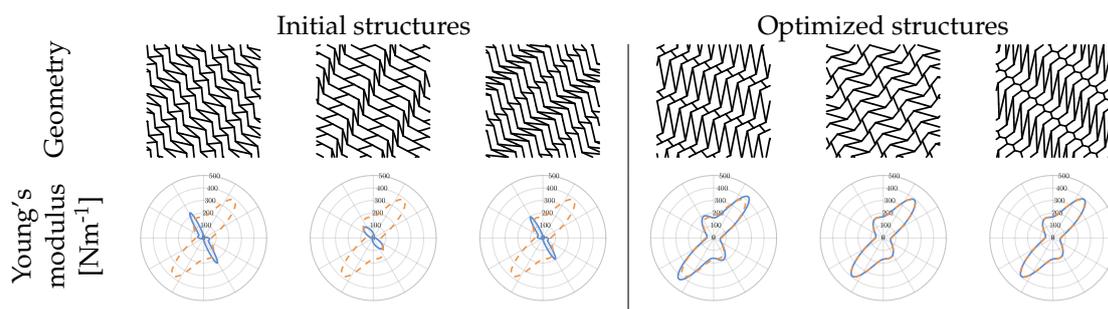
We improve our inverse design approach by running multiple optimizations with different starting points, chosen from our dataset with around 5500 structures. This allows us to discard undesired local minima, and gives us the possibility to provide several results with similar behavior but (potentially) different aesthetics for the user to choose from; see examples shown in Figure 6.11.

**Target Specification** While our optimization approach handles arbitrary Young's modulus targets, defining a physically meaningful Young's modulus distribution without any feedback is a difficult task. We help users by letting them draw a set of points onto the Young's modulus profile and then fit a material compliance tensor  $\mathbb{S}^{goal}$  to those points. The fitting uses a linear least squares approach based on Eq. (6.5):

$$\mathbb{S}^{goal} = \operatorname{argmin}_{\mathbb{S}} \sum_i \left( E_i((\mathbf{d}_i \mathbf{d}_i^T) : \mathbb{S} : (\mathbf{d}_i \mathbf{d}_i^T)) - 1 \right)^2 \quad (6.9)$$

for a set of goal directions  $\mathbf{d}_i$  and directional Young's moduli  $E_i$ .

**Optimization Results** Figure 6.11 shows an optimization result for a target Young's modulus profile corresponding to an orthotropic material. While the initial structures selected from the dataset already match the target fairly well, the optimization can significantly improve on this initial match.



**Figure 6.12:** The results of a structure optimization in the family IH04 for an anisotropic Young's modulus target. The Young's modulus for each initial (left) and optimized structure (right) are shown in blue, along with the target in orange.

Figure 6.12 shows results for an anisotropic target profile. Here, the initial structures are much further from an optimal result. Nonetheless, our optimization finds structures with the desired mechanical properties by significantly altering, and rotating, the initial structures.

## 6.4 Conclusions

We presented a method to characterize the mechanical properties of structured sheet materials, providing an intuitive approach to analyze and explore their deformation behavior. Using numerical homogenization, we condense the complex deformation behavior of structured sheets, simulated at their native level, into a macromechanical Kirchhoff plate model. Our method thus captures the salient mechanical properties of structured sheet materials, including anisotropic stretching and bending resistance as well as nonlinearities for larger strains.

We applied our mechanical characterization method to isohedral tilings, which provide an extensive space of visually pleasing structures with interesting mechanical properties. We investigated the connection between mechanical and geometric symmetries, and determined the conditions which lead to orthotropic, tetragonal, or isotropic materials.

Our method offers an intuitive approach to explore materials, suitable for applications that require a joint design approach that combines aesthetics and material properties.

### 6.4.1 Limitations and Future Work

Our current approach characterizes flat structures, but especially for applications in architecture, structured materials with curved rest shapes are of interest. Investigating the influence of rest curvature on the mechanical properties of the structure might offer insights on how to create optimal synergies between shape and structure.

The linear material model we use for the characterization offers an intuitive way to visualize material properties. However, even though it can also characterize nonlinear material behavior, the model is not accurate for larger strains, and it does not model the interaction between bending and stretching. We consider this trade-off between complexity and conciseness suitable for material characterization, but applications that would use the macromechanical model for simulation and (shape) optimization might require a more accurate and extensive material description.

Many applications of structured materials can benefit from the potential to create *spatially-varying* or *aperiodic* structures, be it for mechanical or purely aesthetic purposes. The required irregular tiling can be created by choosing a specific type of pattern or by relaxing the geometric constraints on isohedral tilings. However, the mechanical characterization of such a tiling is not trivial, and offers many possibilities for future work.

*Mechanical Characterization of Structured Sheet Materials*

## Conclusion

In this thesis, we have addressed three design problems through the use of computational design methods.

We have shown that we can use a data-driven approach to efficiently synthesize microstructures with desired elastic properties. Such an approach requires an appropriate set of precomputed data points, and we have introduced *metamaterial families* that allow us to interpolate between structures and material properties. Equally important as selecting the right microstructure is how to create spatially varying material properties through a tiling of different microstructures. Such an assembly of structures introduces an error in the material behavior that depends on the compatibility of neighboring structures. We have investigated the effect of such a spatially-varying tiling, and shown that unless there is a very abrupt change in material properties, the transition between different material properties will likely not induce a large error.

We proposed a synergistic design approach for shell structures with decorative cutouts that combines structural and aesthetic objectives. With such an approach, a designer who wants to create visually appealing structures can get feedback about the structural properties of the design, and choose to automatically optimize the pattern for structural objectives, reducing the time spent on checking the stability after the design phase and thus having to iterate on the design without proper feedback. While we showed that discrete element textures are a suitable description for the aesthetic objective, our approach is universal in that it can support any objective that translates aesthetics into an energy-based formulation.

## Conclusion

We explored the mechanical properties of space-filling tilings built from simple geometric shapes, and discovered that they can exhibit a diverse set of interesting elastic properties that match their visual appeal. To communicate mechanical properties in an intuitive way that is suitable for the exploration of this large space of structures, we introduced a homogenization approach that captures the salient features of networks of rods that are built from these tilings. We also showed that such a mechanical characterization of a space of structures with a compact parameterization is suitable for an inverse design of structured sheet materials, where structure parameters are automatically optimized to reach desired mechanical properties.

While these three approaches only scratched the surface of the possibilities of structured materials, they present clear directions and building blocks for any computational design approach of such materials.

In the presence of a large design space, splitting this space into a hierarchy through the use of microstructures and homogenization allows for more efficient solutions to computational design problems, and makes large problems manageable. It also facilitates data-driven methods that use large sets of precomputed structures to synthesize the desired properties. For such methods, we have seen that the use of *structure families* instead of individual structures improves the results of interpolations and the tiling of spatially varying structures by ensuring that structures with similar properties will have similar geometries. We have also seen that we can leverage the general ambiguity in the design space—different structures can have the same mechanical properties—to automatically create optimized tilings, where we choose the right structures from a set of candidates, or we can expose the ambiguity to the user and present multiple results to choose from.

Apart from using hierarchical approaches to tackle a complex design space, we can simplify the space by choosing the appropriate *parameterization*. The microstructure optimization approach in Chapter 4 showed us that, while the voxel-based parameterization offers a very general description of spatial structures and allows for various types of solutions, we have to make sure that these solutions fulfill manufacturing constraints, and it is easy to arrive at undesirable solutions due to all the regularization that is part of the optimization. Likewise, the simulation of such a general solution space requires a substantial amount of processing power. When compared to the method presented in Chapter 6, we see that even if we have very compact representation of a space of structures, we can still cover a broad range of mechanical properties simply *because* we can remove many undesirable or infeasible solutions by choosing the right parameterization. This can not only make the simulation of such structures much easier—we can use specialized ap-

proaches like rod simulations—the smaller number of design variables also helps with solving the inverse design problem.

We have also seen that, when aesthetics and functionality are both important parts of the design of an object, a computational design approach that offers interactive feedback and allows some level of user control greatly improves the design process. For such an approach, simulation methods that can quickly create such feedback are crucial, and we are required to find the right balance between speed and accuracy. As we have seen in Chapter 5, even a relatively coarse simulation can provide valuable feedback during the design process, even though it would not pass the requirements of a full structural analysis.

### 7.1 Future Work

Structured materials offer a wealth of possibilities for future work. While we have presented a useful set of methods to leverage the control over the geometry of objects and structures that new manufacturing techniques provide us, we have just scratched the surface of all the possibilities that this manufacturing advance gives us, and is likely to give us in the future.

The material models that we use to simulate structured material are often tailored to the application, and simplified to the point where we achieve the best trade-off between accuracy and efficiency. Even in the realm of elasticity, if we want to control more complex material behavior such as instabilities, wave propagation properties, or just highly nonlinear deformations, the material models will have to grow with this complexity, making the inverse design of such structures all the more difficult. And especially when we pass into the realm of discontinuous properties such as buckling, we have to find more advanced optimization approaches or simulation methods that allow us to optimize structures for these properties.

Besides elasticity, there is a host of other material properties that can be tuned through structures. While some of them have a similar complexity as elasticity, for example thermal conductivity, properties like the airflow through a structure require complex simulation methods for a prediction. Many of these additional properties have been explored through manual design, but computational design approaches promise more control over the material properties, and allow for an easier translation from functional descriptions to material structures, even for non-expert users.

In practical terms, using a sufficiently restricted subset of structures with a suitable parameterization, such as the one presented in Chapter 6 or in

## *Conclusion*

[Panetta et al., 2015], often helps to efficiently find structures with the desired material properties. Especially as fabrication processes become more precise and provide higher resolutions, the design space grows to sizes where we can no longer subdivide it into blocks and solve complex inverse design problems for every one of them. Application-specific structures and parameterizations can have a large benefit in such cases, though finding the right subset of structures and their parameterization is not trivial. Ideally, this subspace could also encode fabrication constraints that would guarantee that every instance of it will always be a valid result.

Finally, we presented only one method that has a full pipeline where a user can go from a functional description to a final result, and two methods that can create structured materials with specific mechanical properties. While this allowed us to focus all progress on this part, it is useless without the full pipeline. And while there are methods that compute distributions of elastic properties, translating a functional description of an object into such a distribution that is suitable for microstructures is still an open problem.

# Appendix

## A.1 Microstructure Optimization

### A.1.1 Numerical Coarsening

We use a homogenization method to describe the coarse-scale behavior of a microstructure, and as the basis of our microstructure optimization. Such a method computes material parameters for a homogeneous material that approximates a structure. In the following, we summarize the *Numerical Coarsening* approach [Kharevych et al., 2009] and highlight differences due to its application to microstructures.

**Harmonic Displacements** To describe the deformation behavior of the microstructure, a set of representative displacements have to be computed for different load cases. These *harmonic displacements*  $\mathbf{h}_{ab}$  (see Figure A.1 for an illustration) are defined as the solution to the following boundary value problem:

$$\begin{aligned}\nabla \cdot \boldsymbol{\sigma}(\mathbf{h}_{ab}) &= 0 && \text{inside } \Omega \\ \boldsymbol{\sigma}(\mathbf{h}_{ab}) \cdot \mathbf{n} &= \frac{1}{2}(\mathbf{e}_a \mathbf{e}_b^T + \mathbf{e}_b \mathbf{e}_a^T) \cdot \mathbf{n} && \text{on } \partial\Omega.\end{aligned}\tag{A.1}$$

Here,  $\mathbf{e}_a$  is the unit vector along the  $a$ -th coordinate direction,  $\frac{1}{2}(\mathbf{e}_a \mathbf{e}_b^T + \mathbf{e}_b \mathbf{e}_a^T)$  describes the tractions on the surface  $\partial\Omega$  of the object domain  $\Omega$ , and  $\mathbf{n}$  is the surface normal. For tiled structures, this surface is the boundary of the cell.

## Appendix

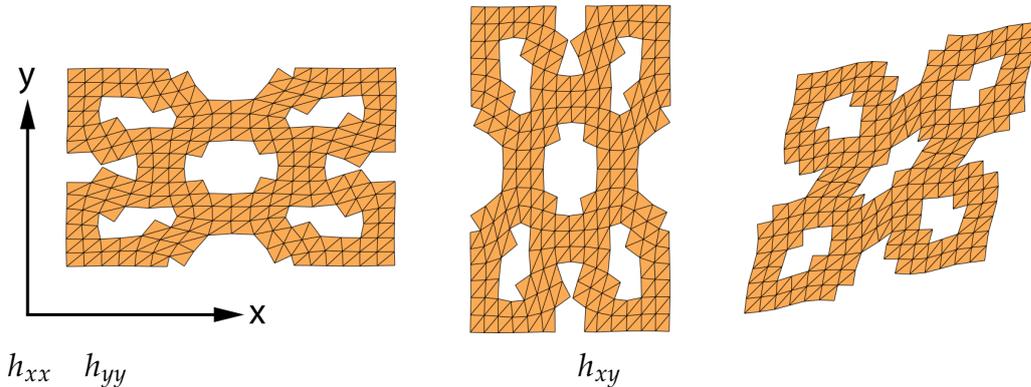
Considering symmetries, there are 3 and 6 distinct harmonic displacements in 2D and 3D, respectively. From these displacements, a 4-th order deformation tensor  $\mathbb{G}$  can be defined per element:

$$\mathbb{G}_{klab} = (\epsilon(\mathbf{h}_{ab}))_{kl}. \quad (\text{A.2})$$

This tensor contains the Cauchy strain for every displacement, and by considering the elasticity equation  $W = \epsilon : \mathbb{C} : \epsilon$  as a bilinear equation, the term  $\mathbb{G}^T : \mathbb{C} : \mathbb{G}$  describes the energy density for any pair of harmonic displacements.

**Coarsening** The homogenized material stiffness tensor can then be computed from the deformation behavior of the microstructure. The deformation is first transferred from the harmonic displacements of the microstructure to a coarse mesh consisting of only a single voxel of the size of the cell. For the case where the corners of the cell correspond to vertices of the fine mesh, this simply means transferring the displacements at the corner. For the general case, the displacement is transferred by computing a distance-weighted interpolation of a set of nearest neighbors in the fine mesh, while adhering to the periodic boundary conditions. After the deformation has been transferred to the coarse mesh, a single coarse-scale deformation tensor  $\mathbb{G}$  can be defined in a manner similar to Equation (A.2). The coarsened material stiffness tensor for the coarse mesh is then obtained analytically as

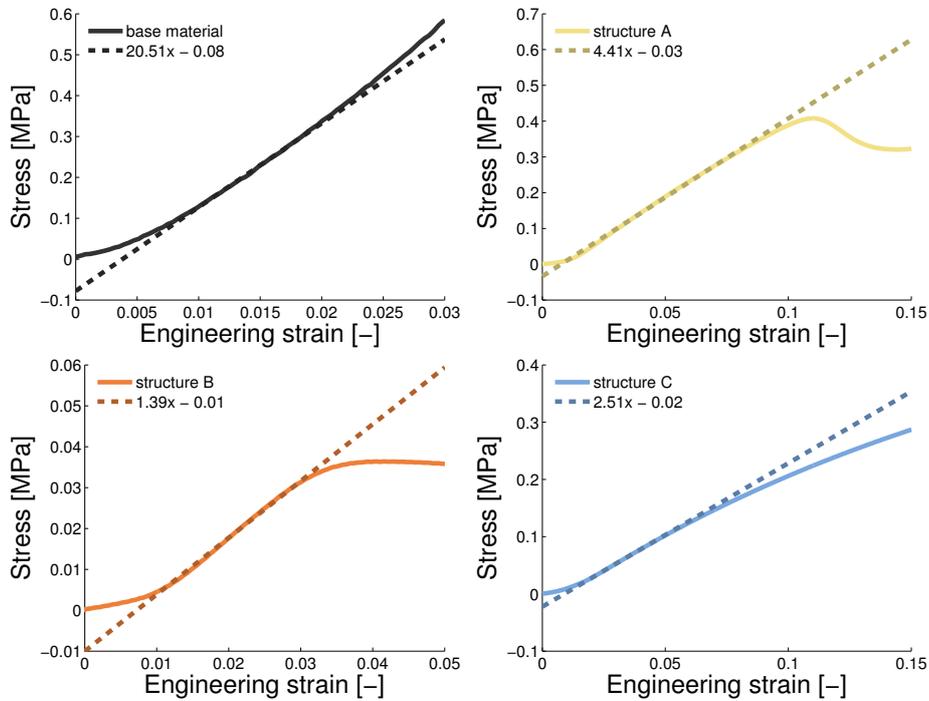
$$\mathbb{C} = \mathbb{G}^{-T} : \left( \sum_{i=1}^k \frac{V_i}{V} \mathbb{G}_i^T : \mathbb{C}_i : \mathbb{G}_i \right) : \mathbb{G}^{-1}, \quad (\text{A.3})$$



**Figure A.1:** Harmonic displacements of a microstructure cell in 2D.

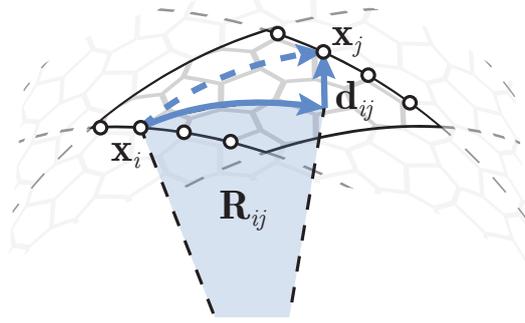
where in the 2D (3D) case  $V_i$  is the area (volume) of element  $i$  in the fine mesh, and  $V$  is the area (volume) of the entire cell. This approach differs slightly from Kharevych et al. [2009], where a coarsened material stiffness tensor is computed for every element in a coarse tetrahedral mesh. Due to the periodic boundary condition for microstructures, the strain in the coarse mesh is uniform, and a single material stiffness tensor can be computed for the whole cell. Note that we use the Voigt compressed matrix representation to express all tensors. This is especially important when computing  $\mathbf{G}^{-1}$ , which can be computed as a simple matrix inverse instead of a more complex symmetric tensor inverse.

### A.1.2 Compression Test Data



**Figure A.2:** The stress–strain measurements for the base material and the three structures tested in the compression test. The tangents of the linear part of the curve describes the Young’s modulus of the structure.

Figure A.2 shows the data from the compression tests of the base material and three synthesized structure. We determined the Young’s modulus of the structures by fitting a linear polynomial to the linear part of the stress–strain curve.



**Figure A.3:** For periodic boundaries in curved configurations, we split the periodic boundary vertex relationship for planar configuration into a translation  $\mathbf{d}_{ij}$  and rotation  $\mathbf{R}_{ij}$ .

## A.2 Structured Sheet Materials

### A.2.1 Curved Periodic Boundary Conditions

We incorporate the tileability of a structure patch by defining a set of *periodic boundary conditions* that impose constraints on vertices on opposite boundaries. These boundary conditions guarantee that the patch behaves as if it were surrounded by identical copies of itself, without having to fix the position of vertices, which would artificially stiffen the structure. For planar configurations, we refer to Section 3.3.3. Here, we extend this formulation for curved deformations (see Figure A.3).

We introduce an additional rotation  $\mathbf{R}_{ij}$  to the regular relationship between periodic boundary vertices (see Equation (3.16)). We then describe the relationship for any vertex pair  $(\mathbf{x}_i, \mathbf{x}_j)$  on opposite boundaries as

$$\mathbf{x}_j = \mathbf{R}_{ij}\mathbf{x}_i + \mathbf{d}_{ij} \quad (\text{A.4})$$

The rotation  $\mathbf{R}_{ij}$  has to be chosen such that the macroscopic curvature of the patch matches the target curvature. For cylindrical curvature configurations with cylinder direction  $\mathbf{v}$  and curvature  $\kappa_C$ , we use the difference in rest state positions,  $\Delta\mathbf{X}_{ij} = \mathbf{X}_j - \mathbf{X}_i$ , to determine the rotational and translational component:  $\mathbf{R}_{ij}$  is a rotation of  $\kappa_C\|\Delta\mathbf{X}_{ij} - \mathbf{v}\mathbf{v}^T\Delta\mathbf{X}_{ij}\|$  around  $\mathbf{v}$ , and  $\mathbf{d}_{ij} = \mathbf{v}\mathbf{v}^T\Delta\mathbf{X}_{ij}$ .

In the case of the spherical curvature  $\kappa_S$ , we only use rotations to define the relationship between the vertices, and set  $\mathbf{d}_{ij} = \mathbf{0}$ . The rotation  $\mathbf{R}_{ij}$  is then

defined as the rotation of  $\kappa_S \|\Delta \mathbf{X}_{ij}\|$  around the rotation axis aligned with  $(\mathbf{X}_i + (0, 0, 1/\kappa_S)^T) \times \Delta \mathbf{X}_{ij}$ , which puts the macroscopic deformation onto a sphere with the desired curvature, but violates the strict tileability of the structure.

### A.2.2 Computing the Macroscopic In-Plane Strain and Stress

The macroscopic strain and stress of a simulation with periodic boundary conditions can conveniently be extracted from the deformations and forces at the boundaries.

In particular, we can consider two pairs of vertices  $(\mathbf{x}_i, \mathbf{x}_j)$  and  $(\mathbf{x}_k, \mathbf{x}_l)$  on different opposing boundaries (see inset). The deformations of these vertices relative to each other directly describe the deformation of the boundary, and with that the macroscopic deformation. In terms of the macroscopic deformation gradient, this relationship is given by  $\mathbf{x}_i - \mathbf{x}_j = \mathbf{F}_{macro}(\mathbf{X}_i - \mathbf{X}_j)$ , where  $\mathbf{X}_i$  is the undeformed position of  $\mathbf{x}_i$ . By using the two pairs of vertices, we can set up a system of equations whose solution is the macroscopic deformation gradient:

$$\mathbf{F}_{macro} = [\mathbf{x}_i - \mathbf{x}_j \quad \mathbf{x}_k - \mathbf{x}_l] [\mathbf{X}_i - \mathbf{X}_j \quad \mathbf{X}_k - \mathbf{X}_l]^{-1} \quad (\text{A.5})$$

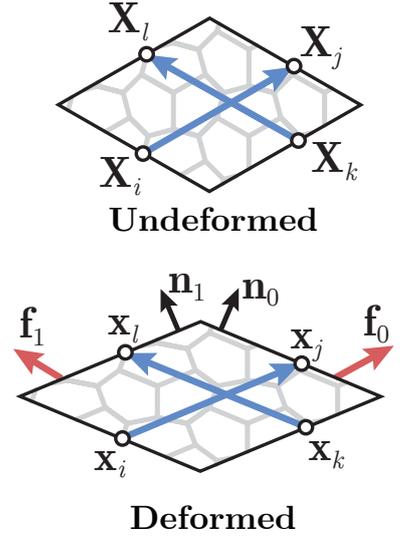
From this expression, we then compute the macroscopic Cauchy strain tensor as  $\boldsymbol{\epsilon}_{macro} = \frac{1}{2}(\mathbf{F}_{macro} + \mathbf{F}_{macro}^T) - \mathbf{I}$ .

Similarly, we can compute the macroscopic stress by looking at the forces at the boundary of the simulation domain. For two non-opposing boundaries, we compute the forces per unit length  $\mathbf{f}_0$  and  $\mathbf{f}_1$ . Then, using the boundary normals  $\mathbf{n}_0$  and  $\mathbf{n}_1$ , which are perpendicular to  $\mathbf{x}_l - \mathbf{x}_k$  and  $\mathbf{x}_j - \mathbf{x}_i$ , respectively, we derive a system of equations based on the stress equation  $\mathbf{f} = \boldsymbol{\sigma}_{macro} \mathbf{n}$ , and obtain

$$\boldsymbol{\sigma}_{macro} = [\mathbf{f}_0 \quad \mathbf{f}_1] [\mathbf{n}_0 \quad \mathbf{n}_1]^{-1}. \quad (\text{A.6})$$

### A.2.3 Direction-Dependent Elasticity Measures

We use three elasticity parameters, the Young's modulus, Poisson's ratio, and bending stiffness, as a way to visualize the direction-dependent mate-



rial behavior of structured sheet materials in an intuitive way. While the formulas to compute these properties from the stiffness tensors are well-known [Böhlke and Brüggemann, 2001], we think that a more thorough explanation of their derivation is important for the understanding of direction-dependent material behavior.

## A.2.4 Membrane

We characterize the membrane behavior of a material using the Young's modulus and Poisson's ratio. Intuitively, these properties describe how much the material resists stretch along a direction, and how much the material contracts perpendicular to this direction. Both of these properties are measured in a *uniaxial stress* configuration, where the material is stretched along a given direction, and any stress perpendicular to this direction is resolved by deformation.

While the in-plane material behavior is most often described using the stiffness tensor  $\mathbf{C}$ , which maps strains to stresses, the case of uniaxial stress is more easily covered by the *compliance tensor*  $\mathbf{S} = \mathbf{C}^{-1}$ , its symmetric inverse, mapping stresses to strains.

**Young's modulus** We can define a *uniaxial unit stress* along a given direction  $\mathbf{d}$  using a simple outer product,  $\boldsymbol{\sigma}^{\mathbf{d}} = \mathbf{d}\mathbf{d}^T$ . Applying this stress to the compliance tensor results in the strain  $\boldsymbol{\epsilon}^{\mathbf{d}} = \mathbf{S} : \boldsymbol{\sigma}^{\mathbf{d}}$  that is induced by this unit stress. From this strain tensor, we want to extract the deformation along the direction  $\mathbf{d}$ , which we get by again applying the tensor  $\mathbf{d}\mathbf{d}^T$  to the strain, resulting in the expression  $(\mathbf{d}\mathbf{d}^T) : \boldsymbol{\epsilon}^{\mathbf{d}}$ . The ratio between the applied stress and the induced deformation then defines the Young's modulus, and since we used the unit stress, we arrive at the formula

$$E(\mathbf{d}) = \frac{1}{(\mathbf{d}\mathbf{d}^T) : \mathbf{S} : (\mathbf{d}\mathbf{d}^T)}. \quad (\text{A.7})$$

**Poisson's ratio** The computation of the Poisson's ratio is based on the same strain response  $\boldsymbol{\epsilon}^{\mathbf{d}}$  to the uniaxial unit stress  $\boldsymbol{\sigma}^{\mathbf{d}}$ , but we now compare two different deformations instead of a deformation and stress. Additionally to the deformation along the direction  $\mathbf{d}$ , we therefore also need to extract the deformation along the direction  $\mathbf{n}$  that is perpendicular to  $\mathbf{d}$ , which we get by applying the tensor  $\mathbf{n}\mathbf{n}^T$  to  $\boldsymbol{\epsilon}^{\mathbf{d}}$ . The negative ratio between these two deformations defines the Poisson's ratio, resulting in the expression

$$\nu(\mathbf{d}) = -\frac{(\mathbf{d}\mathbf{d}^T) : \mathbb{S} : (\mathbf{n}\mathbf{n}^T)}{(\mathbf{d}\mathbf{d}^T) : \mathbb{S} : (\mathbf{d}\mathbf{d}^T)}. \quad (\text{A.8})$$

## Bending

We compute the bending properties based on a simpler approach. Instead of applying a uniaxial unit moment, we measure the bending stiffness on a deformation with purely cylindrical curvature directly on the bending stiffness tensor  $\mathbb{B}$ . Given a direction  $\mathbf{d}$ , we apply the unit curvature tensor  $\boldsymbol{\kappa}^{\mathbf{d}} = \mathbf{d}\mathbf{d}^T$  to  $\mathbb{B}$ , which results in the bending moment  $\mathbf{M}^{\mathbf{d}} = \mathbb{B} : \boldsymbol{\kappa}^{\mathbf{d}}$ . From this bending moment, we can, similar to the membrane case, extract the directional bending moment by again applying the tensor  $\mathbf{d}\mathbf{d}^T$  to it. This then gives us the formula to compute the bending stiffness as

$$b(\mathbf{d}) = (\mathbf{d}\mathbf{d}^T) : \mathbb{B} : (\mathbf{d}\mathbf{d}^T). \quad (\text{A.9})$$

### A.2.5 Test Results

#### Tensile Tests

We present the full plots of our tensile test results in Figure A.4. For the anisotropic structures 5, 6, 7, and 8, we performed additional measurements on a rotated sample to capture the direction dependence of the tensile strength.

#### Bending Tests

We present the full plots of our bending test results in Figure A.5. For structure 7 and 8, we performed additional measurements on a rotated sample to investigate their anisotropy.

Stress–strain curves

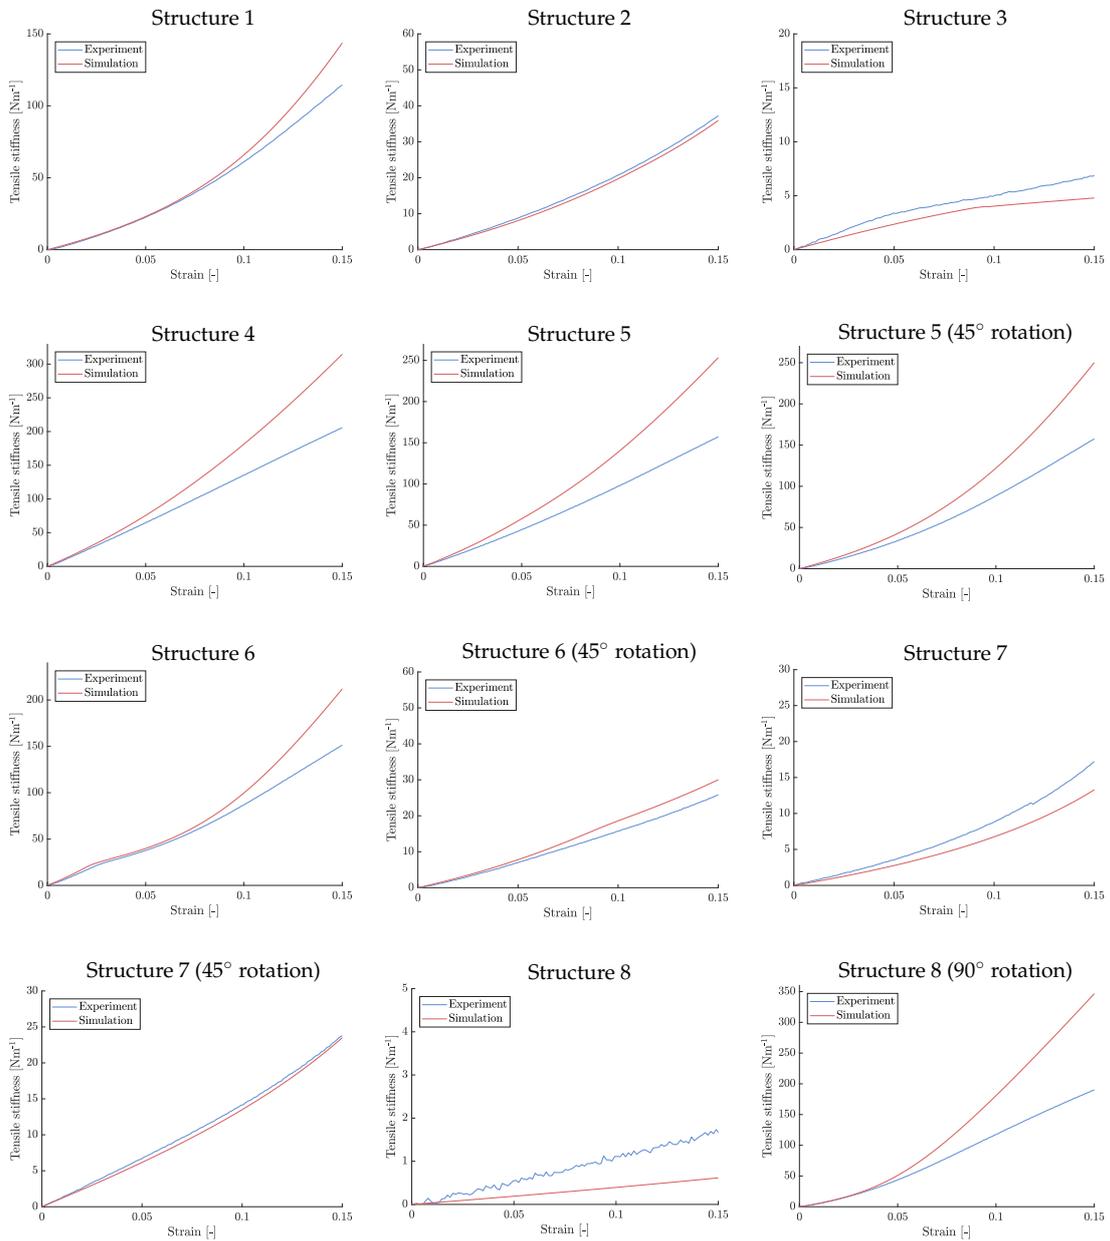
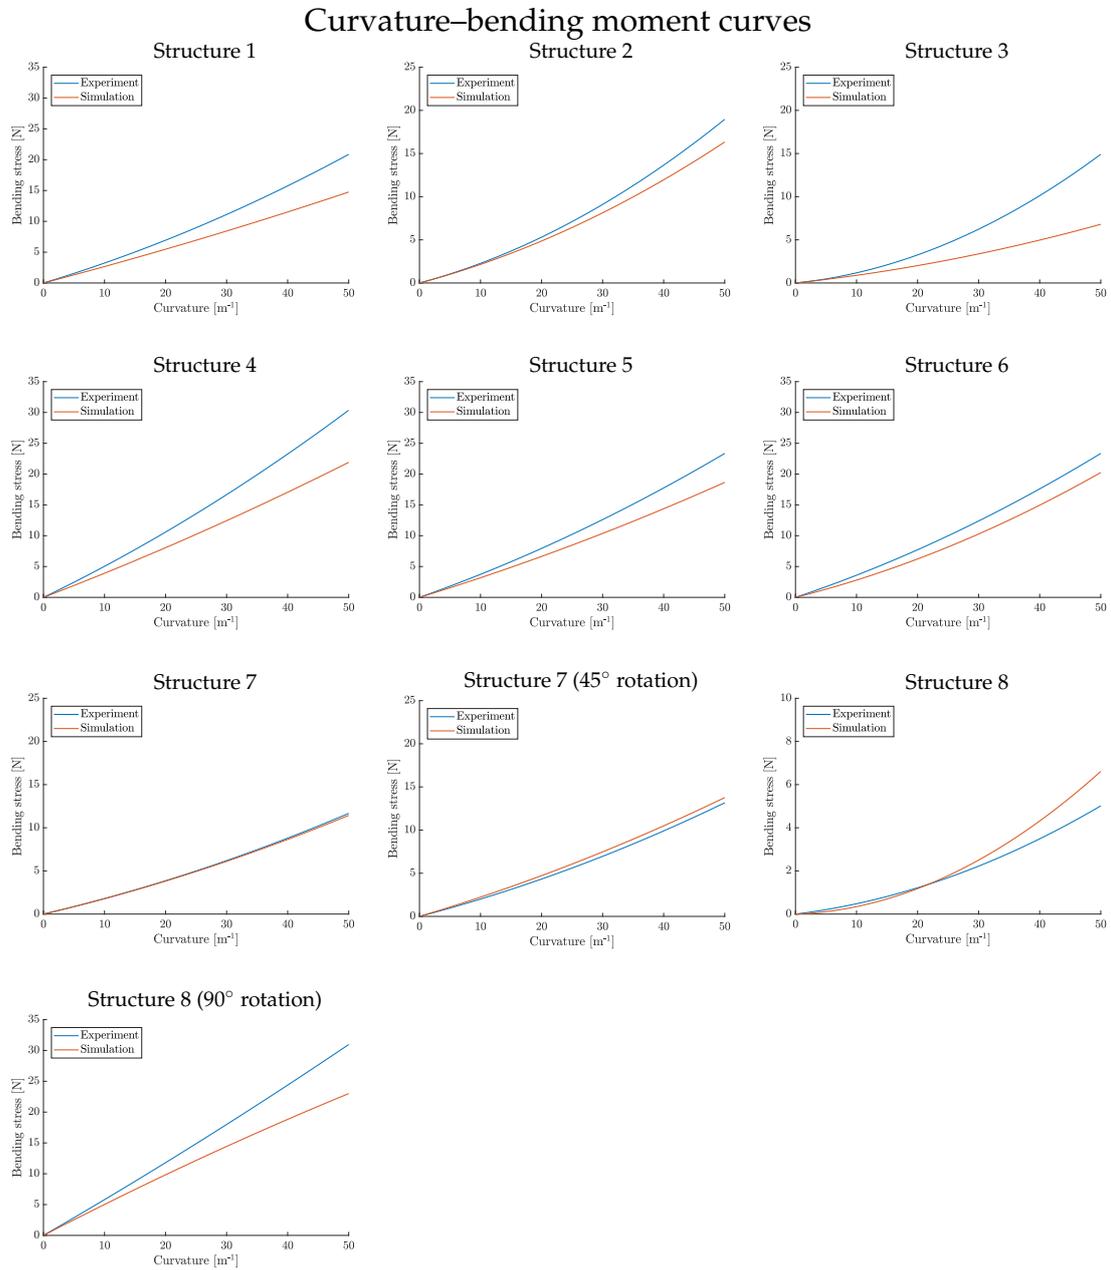


Figure A.4: Tensile test results and comparison to our simulation.



**Figure A.5:** *Bending test results and comparison to our simulation.*

## *Appendix*

## References

- [Akleman et al., 2009] Ergun Akleman, Jianer Chen, Qing Xing, and Jonathan L. Gross. Cyclic Plain-weaving on Polygonal Mesh Surfaces with Graph Rotation Systems. *ACM Trans. Graph.*, 28(3):78:1–78:8, 2009.
- [Ashby et al., 2000] Michael F. Ashby, Anthony G. Evans, Norman A. Fleck, Lorna J. Gibson, John W. Hutchinson, and Haydn N. G. Wadley. *Metal Foams: A Design Guide*. Elsevier Science, 2000.
- [Ashby, 2005] Michael F. Ashby. The properties of foams and lattices. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 364(1838):15–30, 2005.
- [Attard and Grima, 2012] Daphne Attard and Joseph N. Grima. A three-dimensional rotating rigid units network exhibiting negative Poisson’s ratios. *physica status solidi (b)*, 249(7):1330–1338, 2012.
- [Babae et al., 2013] Sahab Babae, Jongmin Shim, James C. Weaver, Nikita Patel, and Katia Bertoldi. 3D Soft Metamaterials with Negative Poisson’s Ratio. *Advanced Materials*, 25(36):5044–5049, 2013.
- [Bächer et al., 2014] Moritz Bächer, Emily Whiting, Bernd Bickel, and Olga Sorkine-Hornung. Spin-it: Optimizing Moment of Inertia for Spinnable Objects. *ACM Trans. Graph.*, 33(4):96:1–96:10, 2014.
- [Bathe, 2006] Klaus-Jürgen Bathe. *Finite Element Procedures*. Prentice Hall, 2006.
- [Becker and Teschner, 2007] Markus Becker and Matthias Teschner. Robust and Efficient Estimation of Elasticity Parameters using the linear Finite Element

## References

- Method. In *Proceedings of the Conference on Simulation and Visualization (SimVis)*, pages 15–28, 2007.
- [Bendsøe and Sigmund, 1999] Martin P. Bendsøe and Ole Sigmund. Material Interpolation Schemes in Topology Optimization. *Arch. Applied Mech.*, 69(9):635–654, 1999.
- [Bendsøe and Sigmund, 2003] Martin Philip Bendsøe and Ole Sigmund. *Topology Optimization: Theory, Methods, and Applications*. Springer Berlin Heidelberg, 2003.
- [Bergou et al., 2008] Miklós Bergou, Max Wardetzky, Stephen Robinson, Basile Audoly, and Eitan Grinspun. Discrete Elastic Rods. *ACM Trans. Graph.*, 27(3), 2008.
- [Bergou et al., 2010] Miklós Bergou, Basile Audoly, Etienne Vouga, Max Wardetzky, and Eitan Grinspun. Discrete Viscous Threads. *ACM Trans. Graph.*, 29(4), 2010.
- [Bern et al., 2017] James M. Bern, Kai-Hung Chang, and Stelian Coros. Interactive Design of Animated Plushies. *ACM Trans. Graph.*, 36(4):80:1–80:11, 2017.
- [Bertoldi et al., 2017] Katia Bertoldi, Vincenzo Vitelli, Johan Christensen, and Martin van Hecke. Flexible mechanical metamaterials. *Nature Reviews*, 2:17066, 2017.
- [Bickel et al., 2009] Bernd Bickel, Moritz Bächer, Miguel A. Otaduy, Wojciech Matusik, Hanspeter Pfister, and Markus Gross. Capture and Modeling of Non-linear Heterogeneous Soft Tissue. *ACM Trans. Graph.*, 28(3):89:1–89:9, 2009.
- [Bickel et al., 2010] Bernd Bickel, Moritz Bächer, Miguel A. Otaduy, Hyunho Richard Lee, Hanspeter Pfister, Markus Gross, and Wojciech Matusik. Design and Fabrication of Materials with Desired Deformation Behavior. *ACM Trans. Graph.*, 29(3):63:1–63:10, 2010.
- [Böhlke and Brüggemann, 2001] Thomas Böhlke and C. Brüggemann. Graphical representation of the generalized Hooke’s law. *Technische Mechanik*, 21(2):145–158, 2001.
- [Bommes et al., 2009] David Bommes, Henrik Zimmer, and Leif Kobbelt. Mixed-integer Quadrangulation. *ACM Trans. Graph.*, 28(3):77:1–77:10, 2009.
- [Bonet and Wood, 1997] Javier Bonet and Richard D. Wood. *Nonlinear Continuum Mechanics for Finite Element Analysis*. Cambridge Univ. Press, 1997.
- [Buannic and Cartraud, 2001] Natacha Buannic and Patrice Cartraud. Higher-order effective modeling of periodic heterogeneous beams. i. asymptotic ex-

- pansion method. *International Journal of Solids and Structures*, 38(40):7139–7161, 2001.
- [Byrd et al., 1995] Richard H. Byrd, Peihuang Lu, Jorge Nocedal, and Ciyou Zhu. A Limited Memory Algorithm for Bound Constrained Optimization. *SIAM Journal on Scientific Computing*, 16(5):1190–1208, 1995.
- [Cartraud and Messenger, 2006] Patrice Cartraud and Tanguy Messenger. Computational homogenization of periodic beam-like structures. *International Journal of Solids and Structures*, 43(3):686–696, 2006.
- [Chen et al., 2013] Desai Chen, David I. W. Levin, Piotr Didyk, Pitchaya Sitthiamorn, and Wojciech Matusik. Spec2Fab: A Reducer-tuner Model for Translating Specifications to 3D Prints. *ACM Trans. Graph.*, 32(4):135:1–135:10, 2013.
- [Chen et al., 2015] Desai Chen, David I. W. Levin, Shinjiro Sueda, and Wojciech Matusik. Data-driven Finite Elements for Geometry and Material Design. *ACM Trans. Graph.*, 34(4):74:1–74:10, 2015.
- [Chen et al., 2016] Weikai Chen, Xialong Zhang, Shiqing Xin, Yang Xia, Sylvain Lefebvre, and Wenping Wang. Synthesis of Filigrees for Digital Fabrication. *ACM Trans. Graph.*, 35(4), 2016.
- [Chen et al., 2017] Desai Chen, David I. W. Levin, Wojciech Matusik, and Danny M. Kaufman. Dynamics-aware Numerical Coarsening for Fabrication Design. *ACM Trans. Graph.*, 36(4):84:1–84:15, 2017.
- [Chen et al., 2018] Jiong Chen, Hujun Bao, Tianyu Wang, Mathieu Desbrun, and Jin Huang. Numerical Coarsening using Discontinuous Shape Functions . *ACM Trans. Graph.*, 37(4), 2018.
- [Cioranescu and Donato, 1999] Doina Cioranescu and Patrizia Donato. *An Introduction to Homogenization*. Oxford lecture series in mathematics and its applications. Oxford University Press, 1999.
- [Clapp et al., 1990] Timothy G. Clapp, Hong Peng, Tushar K. Ghosh, and Jeffrey W. Eischen. Indirect measurement of the moment-curvature relationship for fabrics. *Textile Research Journal*, 60(9):525–533, 1990.
- [Coelho et al., 2008] Pedro G. Coelho, Paulo R. Fernandes, Jose M. Guedes, and Hélder C. Rodrigues. A hierarchical model for concurrent material and topology optimisation of three-dimensional structures. *Structural and Multidisciplinary Optimization*, 35(2):107–115, 2008.
- [Coenen et al., 2010] Erica W. C. Coenen, Varvara G. Kouznetsova, and Marc G. D. Geers. Computational homogenization for heterogeneous thin sheets. *International Journal for Numerical Methods in Engineering*, 83(8-9):1180–1205, 2010.

## References

- [Coulais et al., 2016] Corentin Coulais, Eyal Teomy, Koen de Reus, Yair Shokef, and Martin van Hecke. Combinatorial design of textured mechanical metamaterials. *Nature*, 535:529—532, 2016.
- [de Saxcé and Vallée, 2013] Géry de Saxcé and Claude Vallée. Invariant Measures of the Lack of Symmetry with Respect to the Symmetry Groups of 2D Elasticity Tensors. *Journal of Elasticity*, 111(1):21–39, 2013.
- [Deaton and Grandhi, 2014] Joshua D. Deaton and Ramana V. Grandhi. A Survey of Structural and Multidisciplinary Continuum Topology Optimization: Post 2000. *Struct. Multidiscip. Optim.*, 49(1):1–38, 2014.
- [Derbinsky et al., 2013] Nate Derbinsky, José Bento, Veit Elser, and Jonathan S Yedidia. An improved three-weight message-passing algorithm. *arXiv:1305.1961*, 2013.
- [Dong et al., 2010] Yue Dong, Jiaping Wang, Fabio Pellacini, Xin Tong, and Baining Guo. Fabricating Spatially-varying Subsurface Scattering. *ACM Trans. Graph.*, 29(4):62:1–62:10, 2010.
- [Dumas et al., 2015] Jérémie Dumas, An Lu, Sylvain Lefebvre, Jun Wu, and Christian Dick. By-Example Synthesis of Structurally Sound Patterns. *ACM Trans. Graph.*, 34(4):137:1–137:12, 2015.
- [Fleck et al., 2010] Norman A. Fleck, Vikram S. Deshpande, and Michael F. Ashby. Micro-architected materials: past, present and future. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 466(2121):2495–2516, 2010.
- [Florijn et al., 2014] Bastiaan Florijn, Corentin Coulais, and Martin van Hecke. Programmable Mechanical Metamaterials. *Physical Review Letters*, 113:175503, 2014.
- [François et al., 2017] Marc L. M. François, Letian Chen, and Michel Coret. Elasticity and symmetry of triangular lattice materials. *International Journal of Solids and Structures*, 129(Supplement C):18–27, 2017.
- [Gaillac et al., 2016] Romain Gaillac, Pluton Pullumbi, and François-Xavier Coudert. ELATE: an open-source online application for analysis and visualization of elastic tensors. *Journal of Physics: Condensed Matter*, 28(27):275201, 2016.
- [Garg et al., 2007] Akash Garg, Eitan Grinspun, Max Wardetzky, and Denis Zorin. Cubic Shells. In *Proc. of the ACM SIGGRAPH/Eurographics Symp. on Computer Animation '07*, pages 91–98, 2007.

- [Garg et al., 2014] Akash Garg, Andrew O. Sageman-Furnas, Bailin Deng, Yonghao Yue, Eitan Grinspun, Mark Pauly, and Max Wardetzky. Wire Mesh Design. *ACM Trans. Graph.*, 33(4):66:1–66:12, 2014.
- [Geers et al., 2007] Marc G. D. Geers, Erica W. C. Coenen, and Varvara G. Kouznetsova. Multi-scale computational homogenization of structured thin sheets. *Modelling and Simulation in Materials Science and Engineering*, 15(4):S393–S404, 2007.
- [Gibson and Ashby, 1997] Lorna J. Gibson and Michael F. Ashby. *Cellular Solids: Structure and Properties*. Cambridge Solid State Science Series. Cambridge University Press, 2nd edition, 1997.
- [Gingold et al., 2004] Yotam Gingold, Adrian Secord, Jefferson Y. Han, Eitan Grinspun, and Denis Zorin. A discrete model for inelastic deformation of thin shells. Technical report, Courant Institute of Mathematical Sciences, New York University, 2004.
- [Grima and Evans, 2000] Joseph N. Grima and Kenneth E. Evans. Auxetic behavior from rotating squares. *Journal of Materials Science Letters*, 19(17):1563–1565, 2000.
- [Grinspun et al., 2003] Eitan Grinspun, Anil Hirani, Mathieu Desbrun, and Peter Schröder. Discrete Shells. In *Proc. of the ACM SIGGRAPH/Eurographics Symp. on Computer Animation '03*, pages 62–67, 2003.
- [Grünbaum and Shephard, 1986] Branko Grünbaum and Geoffrey Colin Shephard. *Tilings and Patterns*. W. H. Freeman & Co., New York, NY, USA, 1986.
- [Guseinov et al., 2017] Ruslan Guseinov, Eder Miguel, and Bernd Bickel. Curve-Ups: Shaping Objects from Flat Plates with Tension-Actuated Curvature. *ACM Trans. Graph.*, 36(4), 2017.
- [Hahn et al., 2006] Mariah S. Hahn, Jordan S. Miller, and Jennifer L. West. Three-Dimensional Biochemical and Biomechanical Patterning of Hydrogels for Guiding Cell Behavior. *Advanced Materials*, 18(20):2679–2684, 2006.
- [Hašan et al., 2010] Miloš Hašan, Martin Fuchs, Wojciech Matusik, Hanspeter Pfister, and Szymon Rusinkiewicz. Physical Reproduction of Materials with Specified Subsurface Scattering. *ACM Trans. Graph.*, 29(4):61:1–61:10, 2010.
- [Hashin and Shtrikman, 1963] Zvi Hashin and Shmuel Shtrikman. A variational approach to the theory of the elastic behaviour of multiphase materials. *Journal of the Mechanics and Physics of Solids*, 11(2):127–140, 1963.

## References

- [Hassani and Hinton, 1998] Behrooz Hassani and Ernest Hinton. A review of homogenization and topology optimization i-homogenization theory for media with periodic structure. *Computers & Structures*, 69(6):707–717, 1998.
- [Hassani et al., 2013] Behrooz Hassani, Seyed Mehdi Tavakkoli, and Hossein Ghasemnejad. Simultaneous Shape and Topology Optimization of Shell Structures. *Struct. Multidiscip. Optim.*, 48(1):221–233, 2013.
- [Hiller and Lipson, 2012] Jonathan Hiller and Hod Lipson. Automatic design and manufacture of soft robots. *IEEE Transactions on Robotics*, 28(2):457–466, 2012.
- [Huang et al., 2011] Xiaodong Huang, Arash Radman, and Yi Min Xie. Topological design of microstructures of cellular materials for maximum bulk or shear modulus. *Computational Materials Science*, 50(6):1861–1870, 2011.
- [Hughes et al., 2010] Tom P. Hughes, Arnaud Marmier, and Ken E. Evans. Auxetic frameworks inspired by cubic crystals. *International Journal of Solids and Structures*, 47(11–12):1469–1476, 2010.
- [Hutchinson and Fleck, 2006] Robert G. Hutchinson and Norman A. Fleck. The structural performance of the periodic truss. *Journal of the Mechanics and Physics of Solids*, 54(4):756–782, 2006.
- [Hwu, 2010] Chyanbin Hwu. *Anisotropic Elastic Plates*. Springer US, 2010.
- [Igarashi et al., 2012] Yuki Igarashi, Takeo Igarashi, and Jun Mitani. Beady: Interactive Beadwork Design and Construction. *ACM Trans. Graph.*, 31(4):49:1–49:9, 2012.
- [Ijiri et al., 2008] Takashi Ijiri, Radomír Měch, Takeo Igarashi, and Gavin Miller. An Example-based Procedural System for Element Arrangement. *Comput. Graph. Forum*, 27(2):429–436, 2008.
- [Jacobson et al., 2016] Alec Jacobson, Daniele Panozzo, et al. libigl: A simple C++ geometry processing library, 2016. <http://libigl.github.io/libigl/>.
- [Jiang et al., 2015] Caigui Jiang, Chengcheng Tang, Amir Vaxman, Peter Wonka, and Helmut Pottmann. Polyhedral Patterns. *ACM Trans. Graph.*, 34(6):172:1–172:12, 2015.
- [Kaldor et al., 2010] Jonathan M. Kaldor, Doug L. James, and Steve Marschner. Efficient Yarn-based Cloth with Adaptive Contact Linearization. *ACM Trans. Graph.*, 29(4):105:1–105:10, 2010.
- [Kaplan and Salesin, 2000] Craig S. Kaplan and David H. Salesin. Escherization. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '00*, pages 499–510, 2000.

- [Kaplan and Salesin, 2004] Craig S. Kaplan and David H. Salesin. Islamic Star Patterns in Absolute Geometry. *ACM Trans. Graph.*, 23(2):97–119, 2004.
- [Kaplan, 2009] Craig S. Kaplan. *Introductory Tiling Theory for Computer Graphics*. Synthesis Lectures on Computer Graphics and Animation. Morgan & Claypool Publishers, 2009.
- [Kaufmann et al., 2009] Peter Kaufmann, Sebastian Martin, Mario Botsch, Eitan Grinspun, and Markus Gross. Enrichment Textures for Detailed Cutting of Shells. *ACM Trans. Graph.*, 28(3):50:1–50:10, 2009.
- [Kharevych et al., 2009] Lily Kharevych, Patrick Mullen, Houman Owhadi, and Mathieu Desbrun. Numerical Coarsening of Inhomogeneous Elastic Materials. *ACM Trans. Graph.*, 28(3):51:1–51:8, 2009.
- [Kochmann and Bertoldi, 2017] Dennis Kochmann and Katia Bertoldi. Exploiting microstructural instabilities in solids and structures: From metamaterials to structural transitions. *Applied Mechanics Reviews*, 69:050801, 2017.
- [Konaković et al., 2016] Mina Konaković, Keenan Crane, Bailin Deng, Sofien Bouaziz, Daniel Piker, and Mark Pauly. Beyond developable: Computational design and fabrication with auxetic materials. *ACM Trans. Graph.*, 35(4):89:1–89:11, 2016.
- [Lakes, 1987] Roderic Lakes. Foam Structures with a Negative Poisson’s Ratio. *Science*, 235(4792):1038–1040, 1987.
- [Lan et al., 2013] Yanxiang Lan, Yue Dong, Fabio Pellacini, and Xin Tong. Bi-scale appearance fabrication. *ACM Trans. Graph.*, 32(4):145:1–145:12, 2013.
- [Landes et al., 2013] Pierre-Edouard Landes, Bruno Galerne, and Thomas Hurtut. A Shape-Aware Model for Discrete Texture Synthesis. *Comput. Graph. Forum*, 32(4):67–76, 2013.
- [Langlois et al., 2016] Timothy Langlois, Ariel Shamir, Daniel Dror, Wojciech Matusik, and David I. W. Levin. Stochastic Structural Analysis for Context-aware Design and Fabrication. *ACM Trans. Graph.*, 35(6):226:1–226:13, 2016.
- [Lee et al., 2012] Jae-Hwang Lee, Jonathan P. Singer, and Edwin L. Thomas. Micro-/Nanostructured Mechanical Metamaterials. *Advanced Materials*, 24(36):4782–4810, 2012.
- [Li and Barbič, 2015] Yijing Li and Jernej Barbič. Stable Anisotropic Materials. *IEEE Trans. on Visualization and Computer Graphics*, 21(10):1129–1137, 2015.
- [Livesu et al., 2017] Marco Livesu, Stefano Ellero, Jonàs Martínez, Sylvain Lefebvre, and Marco Attene. From 3d models to 3d prints: an overview of the processing pipeline. *Comput. Graph. Forum*, 36(2):537–564, 2017.

## References

- [Lu et al., 2014] Lin Lu, Andrei Sharf, Haisen Zhao, Yuan Wei, Qingnan Fan, Xuelin Chen, Yann Savoye, Changhe Tu, Daniel Cohen-Or, and Baoquan Chen. Build-to-last: Strength to Weight 3D Printed Objects. *ACM Trans. Graph.*, 33(4):97:1–97:10, 2014.
- [Ma et al., 2011] Chongyang Ma, Li-Yi Wei, and Xin Tong. Discrete Element Textures. *ACM Trans. Graph.*, 30(4):62:1–62:10, 2011.
- [Ma et al., 2017] Li-Ke Ma, Yizhong Zhang, Yang Liu, Kun Zhou, and Xin Tong. Computational Design and Fabrication of Soft Pneumatic Objects with Desired Deformations. *ACM Trans. Graph.*, 36(6):239:1–239:12, 2017.
- [Marmier et al., 2010] Arnaud Marmier, Zoe A.D. Lethbridge, Richard I. Walton, Christopher W. Smith, Stephen C. Parker, and Kenneth E. Evans. ElAM: A computer program for the analysis and representation of anisotropic elastic properties. *Computer Physics Communications*, 181(12):2102–2115, 2010.
- [Martínez et al., 2015] Jonàs Martínez, Jérémie Dumas, Sylvain Lefebvre, and Li-Yi Wei. Structure and Appearance Optimization for Controllable Shape Design. *ACM Trans. Graph.*, 34(6):229:1–229:11, 2015.
- [Martínez et al., 2016] Jonàs Martínez, Jérémie Dumas, and Sylvain Lefebvre. Procedural Voronoi Foams for Additive Manufacturing. *ACM Trans. Graph.*, 35(4):44:1–44:12, 2016.
- [Martínez et al., 2017] Jonàs Martínez, Haichuan Song, Jérémie Dumas, and Sylvain Lefebvre. Orthotropic K-nearest Foams for Additive Manufacturing. *ACM Trans. Graph.*, 36(4):121:1–121:12, 2017.
- [Martínez et al., 2018] Jonàs Martínez, Samuel Hornus, Haichuan Song, and Sylvain Lefebvre. Polyhedral Voronoi diagrams for additive manufacturing. *ACM Trans. Graph.*, 37(4):15, 2018.
- [Megaro et al., 2017] Vittorio Megaro, Jonas Zehnder, Moritz Bächer, Stelian Coros, Markus Gross, and Bernhard Thomaszewski. A Computational Design Tool for Compliant Mechanisms. *ACM Trans. Graph.*, 36(4):82:1–82:12, 2017.
- [Meza et al., 2015] Lucas R. Meza, Alex J. Zelhofer, Nigel Clarke, Arturo J. Mateos, Dennis M. Kochmann, and Julia R. Greer. Resilient 3D hierarchical architected metamaterials. *Proceedings of the National Academy of Sciences*, 112(37):11502–11507, 2015.
- [Michel et al., 1999] Jean Claude Michel, Hervé Moulinec, and Pierre Suquet. Effective properties of composite materials with periodic microstructure: a computational approach. *Computer Methods in Applied Mechanics and Engineering*, 172(1):109–143, 1999.

- [Miehe and Koch, 2002] Christian Miehe and Andreas Koch. Computational micro-to-macro transitions of discretized microstructures undergoing small strains. *Archive of Applied Mechanics*, 72(4):300–317, 2002.
- [Miguel et al., 2012] Eder Miguel, Derek Bradley, Bernhard Thomaszewski, Bernd Bickel, Wojciech Matusik, Miguel A. Otaduy, and Steve Marschner. Data-Driven Estimation of Cloth Simulation Models. *Comput. Graph. Forum*, 31(2):519–528, 2012.
- [Miguel et al., 2013] Eder Miguel, Rasmus Tamstorf, Derek Bradley, Sara C. Schwartzman, Bernhard Thomaszewski, Bernd Bickel, Wojciech Matusik, Steve Marschner, and Miguel A. Otaduy. Modeling and Estimation of Internal Friction in Cloth. *ACM Trans. Graph.*, 32(6):212:1–212:10, 2013.
- [Mori and Igarashi, 2007] Yuki Mori and Takeo Igarashi. Plushie: An Interactive Design System for Plush Toys. *ACM Trans. Graph.*, 26(3), 2007.
- [Mullin et al., 2007] Tom Mullin, Stéphanie Deschanel, Katia Bertoldi, and Mary C. Boyce. Pattern Transformation Triggered by Deformation. *Physical Review Letters*, 99:084301, 2007.
- [Musialski et al., 2015] Przemyslaw Musialski, Thomas Auzinger, Michael Birsak, Michael Wimmer, and Leif Kobbelt. Reduced-order Shape Optimization Using Offset Surfaces. *ACM Trans. Graph.*, 34(4):102:1–102:9, 2015.
- [Musialski et al., 2016] Przemyslaw Musialski, Christian Hafner, Florian Rist, Michael Birsak, Michael Wimmer, and Leif Kobbelt. Non-linear Shape Optimization Using Local Subspace Projections. *ACM Trans. Graph.*, 35(4):87:1–87:13, 2016.
- [Nealen et al., 2006] Andrew Nealen, Matthias Müller, Richard Keiser, Eddy Boxerman, and Mark Carlson. Physically Based Deformable Models in Computer Graphics. *Comput. Graph. Forum*, 25(4):809–836, 2006.
- [Nealen, 2004] Andrew Nealen. An As-Short-As-Possible Introduction to the Least Squares, Weighted Least Squares and Moving Least Squares Methods for Scattered Data Approximation and Interpolation. Technical report, TU Darmstadt, 2004.
- [Nesme et al., 2009] Matthieu Nesme, Paul G. Kry, Lenka Jeřábková, and François Faure. Preserving Topology and Elasticity for Embedded Deformable Models. *ACM Trans. Graph.*, 28(3):52:1–52:9, 2009.
- [O’Donnell and Langer, 1962] William J. O’Donnell and Bernard F. Langer. Design of Perforated Plates. *ASME Journal of Engineering for Industry*, 84(3):307–319, 1962.

## References

- [Ostromoukhov, 2007] Victor Ostromoukhov. Sampling with Polyominoes. *ACM Trans. Graph.*, 26(3), 2007.
- [Overvelde et al., 2017] Johannes T. B. Overvelde, James C. Weaver, Chuck Hoberman, and Katia Bertoldi. Rational design of reconfigurable prismatic architected materials. *Nature*, 541:347–352, 2017.
- [Pabst et al., 2008] Simon Pabst, Sybille Krzywinski, Andrea Schenk, and Bernhard Thomaszewski. Seams and Bending in Cloth Simulation. In *Workshop in Virtual Reality Interactions and Physical Simulation "VRIPHYS" (2008)*, 2008.
- [Panetta et al., 2015] Julian Panetta, Qingnan Zhou, Luigi Malomo, Nico Pietroni, Pablo Cignoni, and Denis Zorin. Elastic Textures for Additive Fabrication. *ACM Trans. Graph.*, 34(4):135:1–135:12, 2015.
- [Panetta et al., 2017] Julian Panetta, Abtin Rahimian, and Denis Zorin. Worst-case Stress Relief for Microstructures. *ACM Trans. Graph.*, 36(4):122:1–122:16, 2017.
- [Pérez et al., 2015] Jesús Pérez, Bernhard Thomaszewski, Stelian Coros, Bernd Bickel, José A. Canabal, Robert Sumner, and Miguel A. Otaduy. Design and Fabrication of Flexible Rod Meshes. *ACM Trans. Graph.*, 34(4):138:1–138:12, 2015.
- [Pérez et al., 2017] Jesús Pérez, Miguel A. Otaduy, and Bernhard Thomaszewski. Computational Design and Automated Fabrication of Kirchhoff-plateau Surfaces. *ACM Trans. Graph.*, 36(4):62:1–62:12, 2017.
- [Picu, 2011] Catalin R. Picu. Mechanics of random fiber networks—a review. *Soft Matter*, 7:6768–6785, 2011.
- [Pietroni et al., 2015] Nico Pietroni, Davide Tonelli, Enrico Puppo, Maurizio Froli, Roberto Scopigno, and Paolo Cignoni. Statics Aware Grid Shells. *Comput. Graph. Forum*, 34(2):627–641, 2015.
- [Prévost et al., 2013] Romain Prévost, Emily Whiting, Sylvain Lefebvre, and Olga Sorkine-Hornung. Make It Stand: Balancing Shapes for 3D Fabrication. *ACM Trans. Graph.*, 32(4):81:1–81:10, 2013.
- [Prévost et al., 2016] Romain Prévost, Moritz Bächer, Wojciech Jarosz, and Olga Sorkine-Hornung. Balancing 3D Models with Movable Masses. In *Proceedings of the Conference on Vision, Modeling and Visualization, VMV '16*, pages 9–16, 2016.
- [Qi and Boyce, 2005] Hang J. Qi and Mary C. Boyce. Stress-strain behavior of thermoplastic polyurethanes. *Mechanics of Materials*, 37(8):817–839, 2005.
- [Qu et al., 2017] Jingyuan Qu, Muamer Kadic, Andreas Naber, and Martin Wegener. Micro-Structured Two-Component 3D Metamaterials with Negative

- Thermal-Expansion Coefficient from Positive Constituents. *Scientific Reports*, 7(40643), 2017.
- [Rafsanjani and Pasini, 2016] Ahmad Rafsanjani and Damiano Pasini. Bistable auxetic mechanical metamaterials inspired by ancient geometric motifs. *Extreme Mechanics Letters*, 9:291–296, 2016.
- [Reddy, 2004] Junuthula Narasimha Reddy. *Mechanics of laminated composite plates and shells: theory and analysis*. CRC press, 2004.
- [Reiner et al., 2014] Tim Reiner, Nathan Carr, Radomir Mech, Ondrej Stava, Carsten Dachsbacher, and Gavin Miller. Dual-Color Mixing for Fused Deposition Modeling Printers. *Comput. Graph. Forum*, 33(2):479–486, 2014.
- [Rodrigues et al., 2002] Hélder Rodrigues, Jose M. Guedes, and Martin P. Bendsøe. Hierarchical optimization of material and structure. *Structural and Multidisciplinary Optimization*, 24(1):1–10, 2002.
- [Rouiller et al., 2013] Olivier Rouiller, Bernd Bickel, Jan Kautz, Wojciech Matusik, and Marc Alexa. 3D-Printing Spatially Varying BRDFs. *IEEE Computer Graphics and Applications*, 33(6):48–57, 2013.
- [Roveri et al., 2015] Riccardo Roveri, A. Cengiz Öztireli, Sebastian Martin, Barbara Solenthaler, and Markus Gross. Example based repetitive structure synthesis. *Proc. of Symp. on Geometry Processing*, 34(5):39–52, 2015.
- [Schaedler et al., 2011] Tobias A. Schaedler, Alan J. Jacobsen, Anna Torrents, Adam E. Sorensen, Jie Lian, Julia R. Greer, Lorenzo Valdevit, and William B. Carter. Ultralight Metallic Microlattices. *Science*, 334(6058):962–965, 2011.
- [Schüller et al., 2016] Christian Schüller, Daniele Panozzo, Anselm Grundhöfer, Henning Zimmer, Evgeni Sorkine, and Olga Sorkine-Hornung. Computational Thermoforming. *ACM Trans. Graph.*, 35(4):43:1–43:9, 2016.
- [Schumacher et al., 2015] Christian Schumacher, Bernd Bickel, Jan Rys, Steve Marschner, Chiara Daraio, and Markus Gross. Microstructures to Control Elasticity in 3D Printing. *ACM Trans. Graph.*, 34(4):136:1–136:13, 2015.
- [Schumacher et al., 2016] Christian Schumacher, Bernhard Thomaszewski, and Markus Gross. Stenciling: Designing Structurally-Sound Surfaces with Decorative Patterns. *Comput. Graph. Forum*, 35(5):101–110, 2016.
- [Schumacher et al., 2018] Christian Schumacher, Steve Marschner, Markus Gross, and Bernhard Thomaszewski. Mechanical Characterization of Structured Sheet Materials. *ACM Trans. Graph.*, 37(4):148:1–148:15, 2018.

## References

- [Sifakis and Barbic, 2012] Eftychios Sifakis and Jernej Barbic. FEM Simulation of 3D Deformable Solids: A Practitioner’s Guide to Theory, Discretization and Model Reduction. In *ACM SIGGRAPH 2012 Courses*, pages 20:1–20:50, 2012.
- [Sigmund and Petersson, 1998] Ole Sigmund and Joakim Petersson. Numerical instabilities in topology optimization: A survey on procedures dealing with checkerboards, mesh-dependencies and local minima. *Structural optimization*, 16(1):68–75, 1998.
- [Sigmund, 2009] Ole Sigmund. Systematic Design of Metamaterials by Topology Optimization. In *IUTAM Symposium on Modelling Nanomaterials and Nanosystems*, pages 151–159, 2009.
- [Skouras et al., 2012] Mélina Skouras, Bernhard Thomaszewski, Bernd Bickel, and Markus Gross. Computational Design of Rubber Balloons. *Comput. Graph. Forum*, 31(2):835–844, 2012.
- [Skouras et al., 2013] Mélina Skouras, Bernhard Thomaszewski, Stelian Coros, Bernd Bickel, and Markus Gross. Computational Design of Actuated Deformable Characters. *ACM Trans. Graph.*, 32(4):82:1–82:10, 2013.
- [Skouras et al., 2014] Mélina Skouras, Bernhard Thomaszewski, Peter Kaufmann, Akash Garg, Bernd Bickel, Eitan Grinspun, and Markus Gross. Designing Inflatable Structures. *ACM Trans. Graph.*, 33(4):63:1–63:10, 2014.
- [Skouras et al., 2015] Mélina Skouras, Stelian Coros, Eitan Grinspun, and Bernhard Thomaszewski. Interactive Surface Design with Interlocking Elements. *ACM Trans. Graph.*, 34(6):224:1–224:7, 2015.
- [Smit et al., 1998] Robert J. M. Smit, W. A. Marcel Brekelmans, and Han E. H. Meijer. Prediction of the mechanical behavior of nonlinear heterogeneous systems by multi-level finite element modeling. *Computer Methods in Applied Mechanics and Engineering*, 155(1):181–192, 1998.
- [Stava et al., 2012] Ondrej Stava, Juraj Vanek, Bedrich Benes, Nathan Carr, and Radomír Měch. Stress Relief: Improving Structural Strength of 3D Printable Objects. *ACM Trans. Graph.*, 31(4):48:1–48:11, 2012.
- [Sullivan, 1896] Louis H. Sullivan. The Tall Office Building Artistically Considered. *Lippincott’s Magazine*, pages 403–409, 1896.
- [Sutherland and Hodgman, 1974] Ivan E. Sutherland and Gary W. Hodgman. Reentrant Polygon Clipping. *Communications of the ACM*, 17(1):32–42, 1974.
- [Tang et al., 2014] Chengcheng Tang, Xiang Sun, Alexandra Gomes, Johannes Wallner, and Helmut Pottmann. Form-finding with Polyhedral Meshes Made Simple. *ACM Trans. Graph.*, 33(4):70:1–70:9, 2014.

- [Vidimče et al., 2013] Kiril Vidimče, Szu-Po Wang, Jonathan Ragan-Kelley, and Wojciech Matusik. OpenFab: A Programmable Pipeline for Multi-Material Fabrication. *ACM Trans. Graph.*, 32:136:1–136:12, 2013.
- [Wang et al., 2011] Huamin Wang, James F. O’Brien, and Ravi Ramamoorthi. Data-driven elastic models for cloth: Modeling and measurement. *ACM Trans. Graph.*, 30(4):71:1–71:12, 2011.
- [Wang et al., 2013] Weiming Wang, Tuanfeng Y. Wang, Zhouwang Yang, Ligang Liu, Xin Tong, Weihua Tong, Jiansong Deng, Falai Chen, and Xiuping Liu. Cost-effective printing of 3D objects with skin-frame structures. *ACM Trans. Graph.*, 32(6):177:1–177:10, 2013.
- [Wegst et al., 2015] Ulrike G. K. Wegst, Hao Bai, Eduardo Saiz, Antoni P. Tomasia, and Robert O. Ritchie. Bioinspired structural materials. *Nature Materials*, 14(1):23, 2015.
- [Wei et al., 2009] Li-Yi Wei, Sylvain Lefebvre, Vivek Kwatra, and Greg Turk. State of the Art in Example-based Texture Synthesis. In *Eurographics 2009, State of the Art Report, EG-STAR*, 2009.
- [Weyrich et al., 2009] Tim Weyrich, Pieter Peers, Wojciech Matusik, and Szymon Rusinkiewicz. Fabricating Microgeometry for Custom Surface Reflectance. *ACM Trans. Graph.*, 28(3):32:1–32:6, 2009.
- [Wilhelm and Frey, 2003] Jan Wilhelm and Erwin Frey. Elasticity of Stiff Polymer Networks. *Phys. Rev. Lett.*, 91:108103, 2003.
- [Wright, 1953] Frank Lloyd Wright. *The future of architecture*. Mentor book. Horizon Press, 1953.
- [Wu et al., 2016] Jun Wu, Christian Dick, and Rüdiger Westermann. A System for High-Resolution Topology Optimization. *IEEE Trans. on Visualization and Computer Graphics*, 22(3):1195–1208, 2016.
- [Xu et al., 2015a] Hongyi Xu, Yijing Li, Yong Chen, and Jernej Barbič. Interactive Material Design Using Model Reduction. *ACM Trans. Graph.*, 34(2):18:1–18:14, 2015.
- [Xu et al., 2015b] Hongyi Xu, Funshing Sin, Yufeng Zhu, and Jernej Barbič. Non-linear Material Design Using Principal Stretches. *ACM Trans. Graph.*, 34(4):75:1–75:11, 2015.
- [Xu et al., 2017] Lizhi Xu, Terry C. Shyu, and Nicholas A. Kotov. Origami and Kirigami Nanocomposites. *ACS Nano*, 11(8):7587–7599, 2017.

## References

- [Zehnder et al., 2016] Jonas Zehnder, Stelian Coros, and Bernhard Thomaszewski. Designing Structurally-sound Ornamental Curve Networks. *ACM Trans. Graph.*, 35(4):99:1–99:10, 2016.
- [Zehnder et al., 2017] Jonas Zehnder, Espen Knoop, Moritz Bächer, and Bernhard Thomaszewski. Metasilicone: Design and fabrication of composite silicone with desired mechanical properties. *ACM Trans. Graph.*, 36(6):240:1–240:13, 2017.
- [Zheng et al., 2014] Xiaoyu Zheng, Howon Lee, Todd H. Weisgraber, Maxim Shusteff, Joshua DeOtte, Eric B. Duoss, Joshua D. Kuntz, Monika M. Biener, Qi Ge, Julie A. Jackson, Sergei O. Kucheyev, Nicholas X. Fang, and Christopher M. Spadaccini. Ultralight, ultrastiff mechanical metamaterials. *Science*, 344(6190):1373–1377, 2014.
- [Zhou and Li, 2008] Shiwei Zhou and Qing Li. Design of graded two-phase microstructures for tailored elasticity gradients. *Journal of Materials Science*, 43(15):5157–5167, 2008.
- [Zhou et al., 2006] Kun Zhou, Xin Huang, Xi Wang, Yiyong Tong, Mathieu Desbrun, Baining Guo, and Heung-Yeung Shum. Mesh Quilting for Geometric Texture Synthesis. *ACM Trans. Graph.*, 25(3):690–697, 2006.
- [Zhou et al., 2013] Qingnan Zhou, Julian Panetta, and Denis Zorin. Worst-case Structural Analysis. *ACM Trans. Graph.*, 32(4):137:1–137:12, 2013.
- [Zhu et al., 2017] Bo Zhu, Mélina Skouras, Desai Chen, and Wojciech Matusik. Two-Scale Topology Optimization with Microstructures. *ACM Trans. Graph.*, 36(4), 2017.