

# Reinforcement of General Shell Structures

FRANCISCA GIL-URETA\*, New York University  
 NICO PIETRONI, University of Technology Sydney  
 DENIS ZORIN, New York University

We introduce an efficient method for designing shell reinforcements of minimal weight. Inspired by classical Michell trusses, we create a reinforcement layout whose members are aligned with optimal stress directions, then optimize their shape minimizing the volume while keeping stresses bounded.

We exploit two predominant techniques for reinforcing shells: adding ribs aligned with stress directions and using thicker walls on regions of high stress. Most previous work can generate either only ribs or only variable-thickness walls. However, in the general case, neither approach by itself will provide optimal solutions.

By using a more precise volume model, our method is capable of producing optimized structures with the full range of qualitative behaviors: from ribs to walls and smoothly transitioning in between. Our method includes new algorithms for determining the layout of reinforcement structure elements, and an efficient algorithm to optimize their shape, minimizing a non-linear non-convex functional at a fraction of the cost and with better optimality compared to standard solvers.

We demonstrate the optimization results for a variety of shapes and the improvements it yields in the strength of 3D-printed objects.

CCS Concepts: • **Computing methodologies** → **Shape modeling**; *Mesh geometry models*; • **Mathematics of computing** → *Mathematical optimization*; • **Applied computing** → *Computer-aided design*;

Additional Key Words and Phrases: Topology optimization, design, simulation

## ACM Reference format:

Francisca Gil-Ureta, Nico Pietroni, and Denis Zorin. 2020. Reinforcement of General Shell Structures. *ACM Trans. Graph.* 39, 5, Article 153 (June 2020), 19 pages.

<https://doi.org/10.1145/3375677>

## 1 INTRODUCTION

In structural design, shells are considered to be one of the most efficient structures because they can be simultaneously lightweight and robust. Shells are common in additive fabrication applications because using shells instead of solids reduces the cost of material and decreases the fabrication time.

\*Work done prior to Amazon involvement of the author and does not reflect views of the Amazon company.

This work was supported in part by the NSF grant OAC-1835712, the NSF award OIA-1937043, and a gift from Adobe Research.

Authors' addresses: F. Gil-Ureta and D. Zorin, New York University, 60 5th Ave, 5th floor, New York, NY 10011; emails: {gilureta, dzorin}@cs.nyu.edu; N. Pietroni, University of Technology Sydney, Faculty of Engineering and IT, Building 11, 81 Broadway Ultimo NSW 2007; email: nico.pietroni@uts.edu.au.

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

© 2020 Association for Computing Machinery.

0730-0301/2020/06-ART153 \$15.00

<https://doi.org/10.1145/3375677>

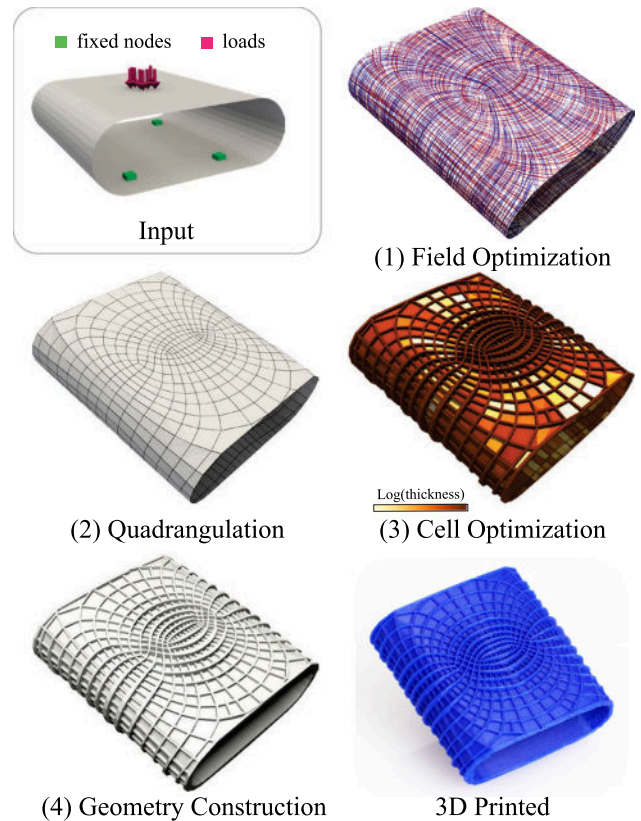


Fig. 1. Stages of our pipeline for generating a shell reinforcement structure.

An optimally-shaped shell can carry its load relying only on tensile/compression forces, with no bending involved, which is very efficient in terms of the required material. These types of shells are commonly found in architecture (domes). However, the shape of the shell may be determined by considerations other than its load-carrying properties. For example, the shape of the airplane is determined by its aerodynamics; the shape of the car body both by the aerodynamics and aesthetics; the top of a table or a shelf is flat, as objects need to be placed on it; the artistic intent primarily determines the shape of a lamp or a statue.

Shell structures with shapes fixed by considerations other than loading are often reinforced by additional means, most commonly increasing thickness in critical areas or adding ribs (Figure 2). Formally, a common optimization goal for a shell reinforcement structure is to *minimize the weight of the added material while keeping the maximal stress of the structure bounded*. The first ensures the structure remains lightweight, while the second prevents structural failure.



Fig. 2. Examples of shell structures showing variable thickness walls and ribs: (a) The Armadillo Vault, Biennale d'architecture 2016, has thicknesses varying from 5cm to 15cm (image by Jean-Pierre Dalbéra [2016]); (b) Mactan Cebu International Airport, Philippines (image by Denzeldoorn [2019]); (c) glass bottle of variable thickness (image by Dmitry Makeev [2006]); (d) plastic skateboard weighs 1.75kg and supports 100kg (image by MthwGlm [2020]); (e) plastic pails (image by Davie Bicker [2019]); (f) plastic laundry basket (image by Dodger67 [2013]).

This problem has been well studied for two dimensions in the limit of low volumes. In 2D, the optimal layout forms a pattern of orthogonal lines (Hencky-Prandtl net) and, for a given layout, the minimum weight structure has all members fully stressed. These structures form, in the limit of low volumes, classical Michell structures and can be obtained by solving a convex optimization problem. It is also well understood for pure bending problems for plates, i.e., the special case of flat shells with loads orthogonal to the surface; in this case, it also reduces to a different convex problem.

The situation is far more complex for the reinforcement of shells embedded in three dimensions. For these shells, the weight-optimal structure may be locally either beam-like, forming *ribs* aligned with stress directions, or membrane-like, forming variable-thickness walls with no perforation [Sigmund et al. 2016]. The first case typically corresponds to bending-dominated regions while the second to areas dominated by in-plane forces. The optimal local structure is determined by the surface shape, the supports, and the loads.

In this general case, the problem is no longer convex and cannot be optimally solved either by methods that assume that the result is only a variable thickness shell or by Michell-truss type methods.

In this article, we propose a novel efficient computational method for constructing optimized reinforcement structures for shells, naturally producing a full range of behaviors spanning the space between variable-thickness shell and rib-type reinforcement. Our approach can be applied to reinforcing any types of free-form manufactured objects: 3D printed plastic or metal shapes, as well as structures produced by casting.

We partition the problem into three steps (Figure 1): (1) determine the field of (approximately) optimal stress directions; (2) construct the skeleton of the reinforcement structure that follows these directions, forming polygonal (predominantly) quad cells aligned with the field; (3) optimize how material is distributed inside the cells.

#### Main contributions.

- For computing the field of optimal stress directions, we developed a generalization of Hencky-Prandtl nets, which takes bending into account and can still be solved by minimizing a convex energy.

- For material distribution optimization, we use a low-parametric structure model for cells to efficiently optimize the distribution of the material. As the global optimization problem is *fundamentally non-convex* (we discuss the reasons on Section 3), to solve it, we propose an efficient global/local method, which shows stable and fast convergence behavior.

We validate our approach by optimizing several 3D shapes. This evaluation shows that our method can handle shells with arbitrary curvature, and successfully transitions between membrane- and bending-dominated regions, obtaining the expected optimal sub-structures. We demonstrate that, by optimizing jointly for bending and compression/tension dominated regimes, we obtain lighter structures than previous work.

## 2 RELATED WORK

Our work builds on the ideas from classical structure design for 2D elasticity and plates, with the key ones originating the work of Michell [1904].

We complement these fundamental ideas with quadrangulation techniques, which can be reinterpreted as a way to transition from an infinite continuum of field-aligned beams to its discretization. We use a variation of Bommers et al. [2009], but any conforming, field-aligned method could be used [Aigerman and Lipman 2015; Campen et al. 2015; Ebke et al. 2016; Kälberer et al. 2007; Myles et al. 2014]. The optimization method for computing the optimal strain field can be viewed as a specialized cross-field optimization method. Similarly to the recently proposed method of Knöppel et al. [2015], it has the advantage of being convex.

We refer to Vaxman et al. [2016] for a complete overview of the related work on field design and to Bommers et al. [2013] for quadrangulation.

*Shell Optimization.* The closest works to ours are the recent works Kilian et al. [2017] and Li et al. [2017], with which we share a number of ideas. The former describes an elegant connection between curvature and Michell trusses and optimizes the surface shape so principal stress and curvature directions coincide. Only membrane forces are considered, and the volume approximation they use is valid for narrow beams (see Section 3). Similar to our work, Li et al. [2017] keeps the shell surface fixed. This work considers a network of ribs, aligned with stress lines, and minimizes their volume; similarly to Kilian et al. [2017], this work also uses a narrow-beam approximation for the volume, and always produces a thin-beam structure. The cross-section shape of individual beams is optimized, which produces additional weight reduction. We discuss differences to these works in more detail in Section 8.

On the other extreme, Zhao et al. [2017] considers the optimization of variable shell thickness, while keeping the topology fixed, which is suboptimal for bending reinforcement. Our work aims to bridge the gap between these extremes.

The approach of Pietroni et al. [2015] aligns a network of beams to an input stress field. Another recent related work, Jiang et al. [2017] considers structures made out of beams with a small number of distinct cross-sections. Both methods are suitable for architectural design; instead, we focus on applications, like 3D printing, which allow greater flexibility of structures.

*Structural Optimization.* The literature on structural optimization is quite extensive, and there is no chance that we can do justice to all of it. The main types of approaches found in the literature include topology optimization methods (homogenization, Solid Isotropic Material with Penalization (SIMP), or Evolutionary Structural Optimization (ESO)-based), analytic methods for optimal structures directly based on Michell-type theories, and methods based on shape derivatives (using an explicit or implicit evolving surface representation). Important books, which include reviews of many other works are Rozvany [1976], Allaire [2002], Bendsoe and Sigmund [2004], as well as recent reviews, Munk et al. [2015] and Sigmund and Maute [2013].

The most prevalent methods in topology optimization of structures are based on SIMP-type methods (see Bendsoe and Sigmund [2004] for a review), which relax the problem to optimizing a density over a domain, which is then converted to a structure by thresholding. This approach has many advantages, including simplicity of implementation [Sigmund 2001], connection to homogenization theory, flexibility in integrating functionals, and ease of scalable implementation [Aage et al. 2015; Wu et al. 2016]. Nevertheless, the result will typically depend on initialization: for the complex topology to emerge, the domain needs to be discretized using a fine grid. The parameters of the result (e.g., the sparsity of the structure, or minimal thickness) need to be controlled indirectly through algorithm parameters. Finally, the result is a voxelized structure, which then needs to be converted in some way to a form more suitable for manufacturing. In comparison, our method directly produces solutions based on a *globally optimal* field (in low-volume limit) and a beam skeleton for the optimized structure, which can be directly adjusted by the user in a variety of ways (e.g., converted to a spline-based Computer Aided Design (CAD) model if desired). We compare in more detail in Section 8.

Ground structure methods are among the oldest methods for optimizing the topology of truss/beam structures. These methods start with a structure consisting of a large number of redundant beams and optimize it to determine the cross-sections, which automatically eliminates some of the beams. Recent examples of applying these type of methods include Sokół [2011], and Zegard and Paulino [2014, 2015]. Compared to our approach, ground structure methods have to restrict the directions of beams to a small set, which affects both optimality and flexibility of the design. The larger the initial set of beams, the closer they may approximate the optimal result. In computer graphics, the ground structure method was used early in Smith et al. [2002] for truss structure design. Panetta et al. [2015] used a version of a ground structure method to obtain initial topologies for computing microstructures with prescribed material properties, followed by shape optimization.

To a great extent, our work was inspired by the beautiful structures explored in the literature on analytic or semi-analytic structure design, e.g., Rozvany [2012], which includes many examples of exact problem solutions, such as Hencky-Prandtl nets.

Our goal is to use this type of idea in the general setting of curved surface domains, taking advantage of the optimality criteria and insights into the structure of the solutions. A concise exposition of the theory underlying Michell-type optimal layouts can be found in Strang and Kohn [1983]. We note that the application of Michell-type structures in 3D is only appropriate for certain

problem settings: e.g., Sigmund et al. [2016] observed that, with no lower-bound constraints on shell thickness, variable thickness shells are likely to emerge as a solution. While their analysis is limited to compliance minimization, our experiments show it also applies to weight minimization (see Figures 23–25).

Shape-derivative based optimization techniques (e.g., Allaire and Jouve [2008]) can obtain very good results when one needs to improve an existing design, by evolving the shape to a local minimum. However, while level-set methods of this type allow for topology changes, the result does vary considerably depending on the starting point. In contrast, our goal is to obtain a starting point that is close to the global optimum, as long as the desired structure has a relatively low volume.

More recently, homogenization-based topology optimization has been used to create high-resolution *manufacturable* structures [Geoffroy-Donders et al. 2017, 2018; Groen and Sigmund 2017]. Several fields corresponding to structure parameters are optimized for a coarse mesh and later projected into a high-resolution mesh creating complex microstructures. This approach was applied to generate structures for 3D volumes [Geoffroy-Donders et al. 2018]; applying this technique to thin shells would require using a very fine resolution, as the dimension in the normal direction of the shell needs to be resolved, making the approach computationally expensive. It may also create impossible to manufacture fine-scale structures with void/solid alteration along the normal direction of the shell.

*Digital fabrication.* The works closest to ours in this domain are Li and Chen [2010], Tam et al. [2015], and Tam [2015]. These methods are based on constructing structures from stress lines on surfaces, which, while different from the optimal fields we compute, are often a close approximation. The overall pipeline of the method of Li and Chen [2010] is similar to ours: this work starts with a field and constructs trusses following the field by tracing lines from supports to loads. The method is limited to planar elasticity and demonstrated only for relatively simple structures.

Tam et al. [2015] uses Fused Deposition Modeling (FDM) to add material directly along the principal stress lines, on 2.5D surfaces. The main problems they solve are stress line generation and selection. They minimize strain energy subject to a maximal total print length (i.e., max material) and consistent maximum spacing between lines.

In 3D, Arora et al. [2019] builds volumetric Michell trusses by creating a stress-aligned 3D texture parametrization and extracting a truss structure from it. The stress field corresponds to the deformation of the initial mesh filled with solid material, not taking into account the redistribution of stress resulting from removing material. The cross-section area of all trusses are assumed to be the same; as a consequence, the volume of the resulting structure is sub-optimal.

Applying topology optimization (SIMP and ground structure methods) to 3D printing applications is discussed in Zegard and Paulino [2016].

### 3 MOTIVATING EXAMPLES

To motivate our method, we start with simple examples of qualitatively different behavior of optimized structures. With these exam-

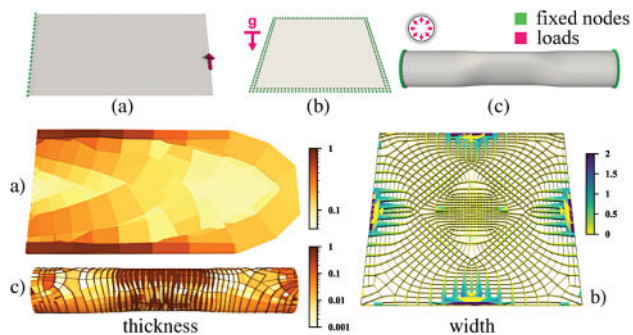


Fig. 3. (a) Optimal structure for a standard cantilever test case with variable thickness [Sigmund et al. 2016]; (b) bending plate optimization, subject to uniform vertical loads, resulting in a rib structure qualitatively consistent with analytic results; (c) an optimized pipe structure, subject to internal pressure, exhibiting a mix of behaviors.




Fig. 4. Left: beam loading and parameters; Right: surface approximated by two intersecting beams loaded in plane.

ples, we demonstrate that, in general, for in-plane loading, using a thicker surface is significantly more efficient than using narrow protruding ribs, and this solution cannot be obtained when using the convex volume approximations used in previous work.

The two key behaviors of optimal shell-like structures, observed in special-case analytic solutions and topology optimization [Sigmund et al. 2016], are the formation of discrete narrow *ribs* protruding from the surface in bending-dominated cases (most forces are perpendicular to the surface), and relatively smooth variation in shell thickness in the pure tensile/compression case (in-plane forces), as shown in Figure 3.

These two behaviors are *not* observed in the simplified models of beam networks approximating a surface: in a typical network, beams do not expand in the direction parallel to the surface, to merge into a variable-thickness shell optimal in such cases. It turns out that this is due to the qualitatively inaccurate volume computation with the volume of the beam network approximated by the sum of individual beam volumes. We now consider two simple examples showing why this is the case.

First, we consider optimization of a single horizontal beam of width  $w$  and height  $h$ , clamped at one end, and loaded at the other, at an angle  $\alpha$  to the beam direction (Figure 4, left). This example clarifies optimal behavior when there is stress in only one direction, both for bending ( $\alpha = \pi/2$ ) and tension ( $\alpha = 0$ ). The second example involves two intersecting beams (Figure 4, right). It is a simple model for a piece of a surface where there is stress in two directions.

For a single beam with an end load ( $F \cos \alpha$ ,  $F \sin \alpha$ ), the maximal stress is proportional to  $F \cos \alpha / (wh) + F \sin \alpha / (h^2 w)$ , the




Fig. 5. Loaded shells dominated by bending and membrane forces and a mix of these. Support nodes are highlighted in green while loads are shown in red. Loads on (a) and (b) include, respectively, external and internal pressure.

sum of the membrane stress along the beam and the bending stress. We optimize the beam volume  $V(w, h) = whl$  keeping this stress constant  $\sigma_0$ . Eliminating  $w$  using the constraint, we get  $V = (lF/\sigma_0)(\cos \alpha + \sin \alpha/y)$ , where  $y = h/l$ .

When  $\sin \alpha \neq 0$ , the optimal solution maximizes the relative thickness  $y$ . In contrast, when forces act along the beam ( $\alpha = 0$ ), the volume value is fixed by the constraint on stress and the choice of  $y$  makes no difference. We conclude that, for single beams, the solution can always be taken to be “thick and narrow” beams (we refer to them as *ribs*), with maximal possible thickness.

The situation is more complicated for surfaces. In the second example, we approximate the surface locally by beams aligned with the perpendicular stress directions (Figure 4, right). For simplicity, we assume the forces, widths  $w$ , lengths  $l$ , and thicknesses  $h$  to be the same for both beams. If we view the intersecting beams individually and approximate the total volume as  $V = 2whl = 2wyl^2$ , then the reasoning above applies to each beam: for in-plane forces ( $\alpha = 0$ ), we get  $wy = F$ , and the optimal volume is  $V = 2Fl^2$ , independent of the choice of  $w$ . However, even a small bending component will prioritize maximal  $h$  solutions, so both beams will be thick and narrow.

Considering beams in separation ignores the fact that *the intersection area of the beams is counted twice*: this part of material is performing “double work,” supporting loads in two directions along two beams. The correct combined volume of two beams is given by

$$V'(w, y) = 2wyl^2 - w^2yl,$$

assuming the same beam width and thickness for both. As before, for in-plane forces ( $\alpha = 0$ ), we have the constraint  $wy = F$ . Replacing  $y$ , the functional  $V'$  can be expressed then as  $V'(w) = Fl(2l - w)$ . From this expression, it is clear the minimum volume is obtained maximizing *width*, as opposed to *thickness*. For a maximal width  $l$ , the optimal volume is  $Fl^2$ . In comparison, if we use a large relative thickness  $y$ , then optimal  $w = F/y \approx \epsilon$  is close to zero, and the volume  $V'(\epsilon)$  is close to  $2Fl^2$ , two times higher than optimal  $V'(l)$ .

More generally, a combination of bending and membrane forces are required to keep an arbitrary shape structurally stable (Figure 5). In this case, given two intersecting beams with an out-of-plane load in addition to in-plane, there is an optimal tradeoff between  $w$  and  $h$  minimizing the volume.

If we further impose constraints on maximal and minimal surface thickness, even in membrane-dominated areas, ribs would form because the optimal solid shell there would be too thin. A

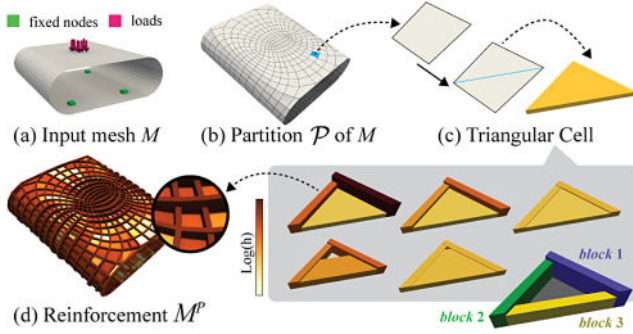


Fig. 6. Perforated variable thickness shell. (a) Input triangle mesh  $M$  with boundary conditions; (b) partition  $\mathcal{P}$  of the input into polygonal faces; (c) each triangular subcell comprises three trapezoidal blocks; (d) extruded cells forming variable-thickness shell structure  $M^P$ .

general optimization method should be able to smooth between grillage-like structures for bending dominated areas and “thin and wide” structures for the rest of the surface.

We conclude, from these examples, that to reinforce a shell in a manner close to optimal for arbitrary loads and shell shape, both solid variable-thickness and rib-like structures may be required in different areas of the surface, and for these to emerge, in a beam-based optimization problem, a non-convex volume function accounting for beam intersection areas has to be used.

## 4 PROBLEM FORMULATION

We start with a description of the variable-thickness perforated surface shell structure that we use to model surface reinforcement, and the optimization problem we aim to solve.

### 4.1 Parameters for reinforced shell structure

Our input is an initial shell  $M$  of thickness  $h_s \geq 0$  (constant per triangle) represented by a triangle mesh, with a vector of external forces  $f$  applied to its vertices, and a set of fixed vertices (supports).

We aim to compute a reinforcement structure, added to the initial shell, which we call a *perforated shell of variable thickness*  $M^P$  (Figure 6).  $M^P$  consists of a partition  $\mathcal{P}$  of the input surface into polygonal faces (typically quads), corresponding to 3D cells, and an extruded shape for each cell, consisting of blocks as described below. The blocks for sequences of cells may form rib-like structures, if the blocks are tall and narrow, or can fill the cells completely, which corresponds to the variable-thickness solid shell case.

Given an approximate user target for cell size, our goal is to optimize the edge orientation of the cell boundaries, and the thicknesses and widths of blocks forming each cell to minimize the weight while maintaining an upper bound on stresses (calculated using a beam model).  $M^P$  can be viewed as the reinforcement structure for  $M$ .

To simplify our problem, we split all polygonal cells into triangular subcells. We refer to the additional edges inserted in this way as *diagonals*. We treat these in a special way in the optimization and, in the end, ensure that the triangular sub-cells can be merged back into the original polygonal cell (Section 7).

*Cell geometry parametrization.* For each edge of a face of  $M^P$ , we introduce two parameters, width  $w_i$  and thickness  $h_i$ . With each




Fig. 7. Geometric parametrization of triangular cell. Left: perspective view of the three blocks; Center: top view showing widths; Right: triangle sides  $l_i$  and respective triangle heights  $a_i$ .

edge, we associate a hexahedron *block* constructed by creating a strip inside the face at a distance  $w_i$  from the edge and extruding the resulting trapezoid along the triangle normal by  $h_i$  (Figure 7). We model each block as a beam including tension/compression and bending forces. While this is a relatively coarse approximation of the shape, it allows obtaining an approximation of the solution robustly and quickly. These results can be further refined by shape optimization methods (e.g., Panetta et al. [2015]).

*Volume discretization.* We consider triangular cells with sides  $l_i$ ,  $i = 1, 2, 3$ , and blocks with rectangular cross-sections of width  $w_i$  and thickness  $h_i$  along the edge  $l_i$  (Figure 7). The simplest approximation of the volume, which is typically used in low-volume truss models, when applied to our system, would yield simply  $\sum_i w_i h_i l_i$ . However, as discussed in Section 3, this approximation of the volume results in highly sub-optimal results for shells regions dominated by membrane forces.

A more precise approximation is the volume of the extrusion of three trapezoidal regions by different heights. Denote the normalized width as  $y_i = w_i/a_i$ , where  $a_i$  is the height of the triangle with base on side  $l_i$ , and the triangle area as  $A = \frac{1}{2} a_i l_i$  for any  $i$ . Given a permutation  $(i, j, k)$  of  $(1, 2, 3)$  for which  $h_i \geq h_j \geq h_k$ , the volume is given by the following simple expression:

$$V(y, h) = A \left( (2 - y_i) y_i h_i + (2 - 2y_i - y_j) y_j h_j + (2 - 2y_i - 2y_j - y_k) y_k h_k \right). \quad (1)$$

This volume can be written as  $V(y, h) = \max_{(i,j,k)} V_{ijk}(y, h)$ , where  $V_{ijk}(y, h)$  is given by Equation (1) for arbitrary permutation  $(i, j, k)$ . This expression for the volume is useful for the optimization method in Section 7.

The out-of-plane heights  $h_i$  can be constrained not to exceed a user-defined value  $h_{max}$ , and the normalized widths  $y_i$  are constrained so that the trapezoidal areas do not overlap:

$$y_1 + y_2 + y_3 \leq 1, \quad y_i \geq 0, \quad 0 \leq h_i \leq h_{max}, \quad \text{for } i = 1, 2, 3. \quad (2)$$

### 4.2 Elastic deformation discretization

We model the perforated shell structure as a *beam network*: for each interior edge, there are two beams, corresponding to the blocks of incident cells along the edge.

*Notation.* The beam network consists of a set of beams  $\mathcal{E}$  that are joined together at nodes (Figure 8). For a node  $i$ ,  $N(i)$  is the set of indices of nodes connected to it, and the vector  $e_{ij}$  connecting nodes  $i$  and  $j$  corresponds to the edge  $e_{ij}$ . For each edge,  $\hat{e}_{ij}$  denotes the unit vector along  $e_{ij}$ . The edge  $e_{ij}$  has length  $l_{ij}$ . We assume that all cells are made of uniform material with  $E$  as Young’s modulus. We use  $\sigma[ij]$  and  $\varepsilon[ij]$  to denote one-dimensional (membrane




Fig. 8. Beam bending discretization.

or bending) stresses and strains of beam connecting vertices  $i$  and  $j$  in a given cell. From now on, we use  $(\cdot)^m$  and  $(\cdot)^b$  to denote membrane and bending terms, respectively.

*Beam linear elasticity discretization.* The membrane strain along each beam is the scalar  $\varepsilon^m[ij] = (\mathbf{u}_j - \mathbf{u}_i) \cdot \hat{\mathbf{e}}_{ij}/l_{ij}$ , same for both blocks at the edge  $e_{ij}$  where  $\mathbf{u}_i$  is the displacement of a vertex  $i$ . It is related to the stress by  $\varepsilon^m[ij] = \sigma^m[ij]/E$ .

For our problem, bending discretization is critical. We use a pure displacement-based beam bending approximation, but a more standard beam element could be used. Our beams are *clamped* to a freely rotating plane at each vertex, i.e., they preserve the angle between the beam and the (freely moving) normal to that plane. Each beam is also connected to the original input shell, resulting in lower torsion. We use a simple discrete beam model for bending, where we neglect the torsion, and define the bending strain  $\varepsilon_{ij}$ , i.e., the change in the curvature of a beam, as

$$\varepsilon^b[ij] = \hat{\mathbf{e}}_{ij}^T (\Delta \hat{\mathbf{n}}_j - \Delta \hat{\mathbf{n}}_i) / l_{ij}, \quad (3)$$

where  $\hat{\mathbf{n}}_i$  and  $\hat{\mathbf{n}}_j$  are the normals meeting at nodes  $i$  and  $j$ , and  $\Delta \hat{\mathbf{n}}$  denotes a linearized change of the normal. The normal change, in turn, is expressed in terms of the displacements  $u_{i\ell}$ ,  $\ell \in N(i)$  of the incident vertices.

This leads to the expression for a scalar bending strain on beams:

$$\varepsilon[ij]^b = (D_{ij}^b)^T u \quad (4)$$

with the expressions for  $D_{ij}^b$  given in Appendix D.

At a distance  $z$  from the middle surface of a beam, the strain is given by  $\varepsilon^m + z\varepsilon^b$ , where we omit the beam index. Based on the standard Bernoulli beam assumptions, after integration over  $z \in [-h/2, h/2]$ , the total beam energy can be expressed as

$$\frac{1}{2} (w l h u^T (D^m) (D^m)^T u + \frac{1}{6} w h^3 l u^T (D^b) (D^b)^T u),$$

which leads to the following expression for the stiffness matrix of the beam system:

$$K^B = w l h (D^m) (D^m)^T + \frac{1}{6} w h^3 l (D^b) (D^b)^T. \quad (5)$$

Given the expression for strain  $\varepsilon^m + z\varepsilon^b$ , clearly both strain and stress are maximal on one of the surfaces, i.e., for  $z = h$  for a given beam. This leads to the following stress constraint:

$$\left| (D_{ij}^m + h_{ij} D_{ij}^b)^T u \right| < \sigma_0. \quad (6)$$

The stiffness of the *reinforced* system combines the stiffness of the beams and shell. We form the shell stiffness matrix  $K^S$  using the elastic discretization described in Section 6. The combined stiffness matrix  $K$  is given by  $K = K^B + K^S$ .

*Optimization problem.* We use index  $c$  for triangular cells. Let  $w$  be the vector of width parameters of all cells,  $y$  the vector of corresponding normalized widths  $w_i/a_i$ ,  $h$  the vector of all thickness parameters,  $H$  the diagonal matrix with thicknesses on the diagonal,  $u$  the displacements,  $f$  the forces, and  $\mathbf{1}$  be the vector of all ones.

We now formulate our optimization problem:

$$\begin{aligned} \min_{h, y, \mathcal{P}} \sum_{c \in \text{cells}} V(h_c, y_c) \quad \text{s.t.} \quad & K(h, y, \mathcal{P}) u = f, \\ |(D^m \pm D^b H)^T u| \leq \sigma_0 \mathbf{1}, \quad & 0 \leq h_i^c \leq h_{\max}, \quad y_i^c \geq 0, \\ y_1^c + y_2^c + y_3^c \leq 1, \quad & \text{for all cells } c, i = 1, 2, 3, \end{aligned} \quad (7)$$

where the absolute value in the stress constraint is taken element-wise, and the minimum is taken over all partitions  $\mathcal{P}$  of  $M$ . The cell volume  $V$  is defined in Equation (1), the stress constraint comes from Equation (6), and the last three constraints correspond to Equation (2). Additionally, we enforce the same thickness on two sides of all cell diagonals.

As optimization over all possible partitions into cells is an intractable problem, combining combinatorial and continuous aspects, we use heuristics to decide the partition first using *beam continuum approximation* (Section 6). Once  $\mathcal{P}$  is fixed, we optimize with respect to  $w$  and  $h$  only.

While the above formula volume differs only by a seemingly simple quadratic term from the simplest approximation, this completely changes the behavior of the problems, and, in particular, the behavior of the solvers. The problem no longer reduces to convex by a change of variables as it is the case for the simplest formula (cf., e.g., Hemp [1973]), and a different type of solver needs to be applied. In our experiments, commonly used general purpose non-convex solvers converge very slowly and often fail to make progress. Our solution is described in Section 7.

## 5 OVERVIEW OF THE APPROACH

Our pipeline for solving the optimization problem consists of the steps listed below.

- (1) **Field optimization.** Compute a per-triangle cross-field on the surface using stress-based optimization (Section 6). This field corresponds to an idealized system of densely spaced thin beams (*beam continuum*) with directions chosen to minimize weight. The problem is formulated in terms of displacements and the desired cross-field is obtained from the symmetric strain tensors. This requires solving a convex optimization problem with inequality constraints.
- (2) **Quadrangulation.** Create a quad-dominant mesh aligned to this cross-field with a user-controlled spacing of edges. This step is done using a version of mixed-integer quadrangulation [Bommes et al. 2009], although any quad layout method with field alignment can be used instead (Section 7). The faces of the mesh will correspond to the *cells* of the perforated shell  $M^{\mathcal{P}}$ .
- (3) **Cell optimization.** Optimize shape parameters of the cells (Section 7). We introduce a substructure for each cell, with a small number of control parameters defining its shape (widths  $w$  and thicknesses  $h$  of rectangular beams

along each side). We derive the optimal material distribution by efficiently solving a nonlinear, non-convex problem minimizing the total volume of all cells with respect to  $w$  and  $h$ , while keeping stresses below a user-defined maximum. To make the problem tractable, we defined an efficient local-global optimization method.

- (4) **Final geometry construction.** Finally, we derive the final geometry of  $M^P$  according to optimized widths and thickness. We obtain a triangle mesh by performing a sequence of Boolean operations between meshes representing the beams. The final watertight and manifold mesh can be directly used for 3D printing or decomposed into elements for FEM analysis.

In the following sections, we describe the details of the steps in the order in which they are applied to produce the final result; however, the key step is the third one (*cell optimization*), as it has the most impact on the optimality of the result, as demonstrated in the evaluation (Section 8).

## 6 WEIGHT-MINIMIZING FIELD OPTIMIZATION

In this section, we describe our method for constructing a field of directions on the surface, which approximates optimal directions for weight minimization with bounded stress. The cells in our construction will be aligned with these directions.

The key idea is to solve a version of the beam weight minimization in the limit case. We assume that there is a *continuum* of infinitely thin and infinitely close beams in two orthogonal directions forming the surface. The density of the beams and their orientations are the optimization variables. The idea of using this type of continua goes back to Michell [1904]; in contrast to the standard Michell continua formulations, which take only tension into account, we use both tension and bending. We describe first the classical theory of Michell continua, leading to a convex problem, and then extend it to the case of shells with bending forces, preserving convexity.

### 6.1 Michell continua

Here, we briefly review the classical solution, following Strang and Kohn [1983]. The best directions are known to be the principal stress directions of the *optimized* structure. (These are *not* the same as the stress direction on the original input shell, although these fields are often close.)

The force balance for a plate or a shell with no bending is given by the standard equations in terms of in-plane stress tensor  $\sigma$ , strain  $\varepsilon$ , possibly varying elasticity tensor  $E(\mathbf{p})$ , and external force density  $\mathbf{f}$ :

$$\operatorname{div} \sigma = \mathbf{f}, \quad \sigma = E(\mathbf{p}) : \varepsilon; \quad \varepsilon = \frac{1}{2}(\nabla \mathbf{u} + \nabla^T \mathbf{u}), \quad (8)$$

where  $A : B = \sum A_{ijkl} B_{kl}$  for a 4-tensor and a 2-tensor.

A Michell continuum is an idealization of a beam network, characterized, at every point, by beam densities  $\rho_1$  and  $\rho_2$  in two directions (Figure 9). In other words, how many beams cross a unit-length segment along one of the coordinate directions. In the limit of small thicknesses, the total fraction of a small area covered by trusses at a point  $\mathbf{p}$  is  $\rho_1(\mathbf{p}) + \rho_2(\mathbf{p})$ . The total volume of the trusses




Fig. 9. Michell continua example. In the limit of small thicknesses  $d$ , every point is characterized by two orthogonal beam densities,  $\rho_1$  and  $\rho_2$ .

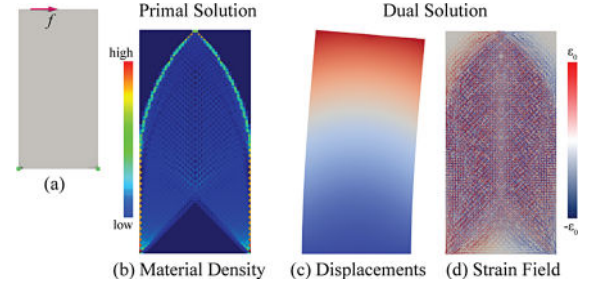


Fig. 10. Field optimization: (a) Design domain  $\Omega$  and boundary conditions; (b) Primal solution, squared densities  $\sqrt{\rho_1 + \rho_2}$ ; (c,d) Dual solutions, displacements  $\mathbf{u}$  and strain  $\varepsilon$  eigen-values and -vectors.

in the continuum can be defined as  $\int_{\Omega} \rho_1 + \rho_2 dA$ . Note that this approximation of area covered by trusses suffers from the same flaw pointed out in Section 3.

The optimal trusses have to be oriented along stress directions, and be critically stressed, i.e., all (non-averaged) stresses on the trusses are equal to maximal stress  $\sigma_0$ . This leads to the relationship between  $\rho_i$  and corresponding averaged principal stress:  $\lambda_i(\sigma) = \rho_i \sigma_0$ ,  $i = 1, 2$ , where  $\lambda_i(\cdot)$  denotes the  $i$ -th singular value.

Then, we obtain the following optimization problem for the volume, formulated entirely in terms of stresses:

$$\text{minimize } \int_{\Omega} (|\lambda_1(\sigma)| + |\lambda_2(\sigma)|) dA, \quad \text{subject to } \operatorname{div} \sigma = \mathbf{f}. \quad (9)$$

This problem is known to be *convex* [Strang and Kohn 1983] (although it is more difficult than the linear programming formulation for a truss network). Note that principal stress directions are not fixed and are determined by the optimization. We use these directions as the field for orienting beams in  $M_p$ .

The Problem (9) has a simple dual (Appendix B) of the form

$$\text{maximize } \int_{\Omega} \mathbf{f}^T \mathbf{u} dA \quad \text{subject to } |\lambda_i(\varepsilon)| \leq \varepsilon_0, \quad i = 1, 2, \quad (10)$$

where  $\varepsilon$  is the strain of the optimal solution (Figure 10). The dual problem is significantly easier to deal with in the case of continua.

We note that here we neglect the overlap volume of trusses discussed in Section 3; while it could be included as  $-\rho_1 \rho_2$  term, this would immediately make the problem non-convex, and the benefit of more precise field optimization in terms of volume reduction is minor (Section 8).

### 6.2 Continuum optimization with bending

Next, we generalize Problem (10) to include bending forces.

If the thickness of the shell remains fixed, one can add bending to the functional with relative ease without changing the convexity of the problem. We set the shell thickness in this case to half of the maximal allowable thickness; while the resulting field is sub-optimal, as we show experimentally in Section 8, inaccuracy in




Fig. 11. Vertical strain distribution in a shell.

the beam direction has less effect on the overall weight reduction, compared to width/thickness optimization of beams.

We make the standard assumption of planar stress for the shell, i.e., no stresses are active in the direction perpendicular to the shell surface. The strain at a distance  $z$  from the midline of the shell is given by (Figure 11)

$$\varepsilon(z) = \varepsilon^m + z\varepsilon^b,$$

where  $\varepsilon^b$  is the *bending strain tensor*, equal to the linearization of the change in the shape operator  $\nabla \hat{\mathbf{n}}$ .

Consistently with the Michell continuum, we seek to minimize the total weight of a beam continuum, bounding the stress everywhere by  $\sigma_0$ . We observe that the eigenvalues of a  $2 \times 2$  symmetric matrix  $A$ , using the substitutions  $a = (A_{11} + A_{22})/2$ ,  $b = (A_{11} - A_{22})/2$ , and  $c = A_{12}$ , are of the form  $a \pm \sqrt{b^2 + c^2}$ . Note that these are, respectively, convex and concave functions of the argument, and, therefore, reach their maxima (respectively, minima) on the boundary of the shell, for  $z = \pm h/2$ . For this reason, it is sufficient to bound eigenvalues of stress (or strain) only for  $z = h/2$  and  $z = -h/2$ , to guarantee the bounds elsewhere.

In the case of bending, the dual problem formulated in terms of displacements has a simpler form relative to the primal problem:

$$\text{maximize } \int_{\Omega} \mathbf{f}^T \mathbf{u} dA \text{ subject to } |\lambda_i(\varepsilon^m \pm \frac{h}{2}\varepsilon^b)| \leq \varepsilon_0, i = 1, 2. \quad (11)$$

Note that now there are two sets of constraints, corresponding to two surfaces of the shell.

The strain tensors  $\varepsilon(h/2)$  and  $\varepsilon(-h/2)$  can be interpreted as defining cross-fields on the surface. We use the angular average of these fields to align the cells.

### 6.3 Field discretization and optimization

Finally, we describe a discretization of the convex problem defined above, and how to use the resulting field to build a mesh.

As a first step, we solve a discrete version of Problem (11), which yields displacements  $\mathbf{u}$  at vertices. From these displacements, we compute the per-triangle strain field eigenvectors, forming a cross-field on the surface, i.e., an assignment of four unit vectors, aligned with perpendicular principal strain directions, to each triangle.

*Discretization of the optimization problem.* The optimization Problem (11) has a relatively simple discretization that can be readily plugged into a cone program solver, e.g., MOSEK [ApS 2015].

We assume that the surface is given as triangle mesh,  $M = (V, E, F)$ , and the same notation is used for the edge vectors and vertices that we used for beam networks.

The variables in the problem are displacements, which we discretize using standard piecewise-linear functions on the surface, with the vector of unknowns  $u$  (we use non-bold letters for high-dimensional vectors including all components of corresponding three-dimensional quantities).




Fig. 12. Discretization of the bending field.

The two quantities that need to be discretized are membrane and bending strains; we define these per triangle.

If  $\mathbf{e}_{ij}$  are the vectors along triangle edges, for a triangle  $T$ , we have the following expression for the strain, computed as  $\frac{1}{2}(\nabla u + \nabla^T u)$ :

$$\varepsilon_T^m = \frac{1}{4A_T} \left( \sum_{i=1,2,3} \mathbf{e}_{jk}^\perp \mathbf{u}_i^T + \mathbf{u} \mathbf{e}_{jk}^T \right), \quad (12)$$

where  $A_T$  is the triangle area,  $j$  is the vertex after  $i$  in CCW order,  $k$  is before  $i$ , and  $\mathbf{e}_{jk}^\perp$  is the vector in the triangle plane perpendicular to the side.

If, by abuse of notation, we use  $\varepsilon_T^m$  to denote the vector of three distinct components of the strain tensor, we can write the expression in the form  $B_T^m u$ , where  $u$  is the vector of all displacement degrees of freedom.

To discretize the bending strain, we use the triangle-based approximation of the shape operator, following the overall idea in Oñate et al. [1994] and Grinspun et al. [2006], using vertices of the edge-adjacent triangles to compute the changes of the average normals at edge midpoints, and computing the gradient of the normal using the formulas of Equation (12).

This leads to the following expressions for the bending strain on triangles, in which we neglect the triangle deformation: in the deformation modes for which the bending strain is high (i.e., if the curvature changes a lot, in-plane deformations are small). We use the following formula from Grinspun et al. [2006], Figure 12:

$$\varepsilon_T^b = \sum_{i=1,2,3} \frac{\theta_i}{2Al_i} \mathbf{e}_i^\perp (\mathbf{e}_i^\perp)^T, \quad (13)$$

where  $\theta_i$  are linearized *changes* in the angles between normals of adjacent triangles.

Similarly to  $\varepsilon^m$ , we can write  $\varepsilon^b = B^b u$ . Then, the discrete problem takes the form

$$\text{maximize } f^T u \text{ subject to, for all } T, |\lambda_i(B_T^m \pm hB_T^b)u| \leq \varepsilon_0, i=1,2. \quad (14)$$

where  $f$ , similarly to the beam case, denotes the vector of per-vertex forces, and  $u$  is the vector of all vertex displacements. We use eigenvalue formulas defined in Appendix A for Equation (19) to convert the problem to a convex cone problem, which we solve using the MOSEK solver [ApS 2015].

*Detecting field zones.* While the output of the previous step defines a tensor for each triangle, not all of these are meaningful. In some cases (if the triangle is not deformed at all, or deformed negligibly), the strain is zero. More generally, some points may have isotropic strains of the form  $kI$ , where  $k$  is a nonzero constant, for which all vectors are eigenvectors, so the cross-field is not defined uniquely on this triangle. For general fields, such points are




Fig. 13. Zones of an optimal strain field. (1) Two orthogonal directions, (2) no preferred direction, (3) one direction, (4) no beams. The crossfield directions are well-defined only for regions 1 and 3.

usually isolated. However, for the fields corresponding to the solution of the problem we are considering, the situation is different. There are four possible regimes (see, e.g., Strang and Kohn [1983]). Specifically, the possibilities include

- (1)  $\lambda_i(\varepsilon) = -\lambda_j(\varepsilon) = \varepsilon_0$ , principal strains are critical and have opposite directions; this corresponds to well-defined two orthogonal beam families;
- (2)  $\lambda_i(\varepsilon) = \lambda_j(\varepsilon) = \pm\varepsilon_0$ , principal strain are critical and have the same direction; in this case, stresses (which are dual variables to the inequality constraints) are large but beam directions are not well defined;
- (3)  $|\lambda_i(\varepsilon)| < |\lambda_j(\varepsilon)| = \varepsilon_0$ , only one strain is critical; this corresponds to a single family of beams.
- (4)  $|\lambda_{i,j}(\varepsilon)| < \varepsilon_0$ , in this case, stresses are both zero, which means there are no beams in this area.

In cases 1 and 3, the crossfield directions are well-defined (purple zones on Figure 13). In cases 2 and 4, these are either not defined or are not relevant, due to the absence of structure. For this reason, for our construction, we use only regions 1 and 3, which we detect by requiring at least one eigenvalue to be close to  $\varepsilon_0$ , and the difference of eigenvalues to be more than a constant.

We call the resulting field *salient*. The situation is essentially identical to the cross-fields used for constructing quadrangulations; typically, a field aligned with curvature directions is used as a starting point, and only in areas where the difference of two principal curvatures is high.

*Completing the field.* To complete the field on the whole surface, which is needed for a complete structure, we use cross-field constrained optimization procedure of Bommès et al. [2009]. In this algorithm, the cross-field is encoded by a per-triangle angle, with respect to a reference direction  $\beta_i$  in each triangle. The angles on salient triangles are fixed. On the remaining triangles, these are found by a greedy solve of a mixed-integer problem minimizing the energy

$$E = \sum_{\text{edges}(ij)} \left( \beta_i - \beta_j + k_{ij} \frac{\pi}{2} + \kappa_{ij} \right)^2,$$

where the summation is over all dual edges connecting triangles  $i$  and  $j$ ,  $\kappa_{ij}$  is the angle between reference directions in triangles, and  $k_{ij}$  is an integer unknown accounting for the fact that cross-field values represented by angles  $\beta + k\pi/2$  are the same.

In the resulting field, defined by the angles  $\beta_i$  for all triangles, and integers  $k_{ij}$  for all edges, one can easily compute per-vertex

field index and identify field singularities, which become irregular (valence different from four) vertices of the quad mesh at the next step. We refer to Bommès et al. [2009] for details of the index computation.

#### 6.4 Construction of the quad-dominant network

In general, there may be no optimal beam spacing (in the low-volume limit, the finer the structure, the lower the optimal volume can be for a given stress). For this reason, the beam spacing is defined by a user-specified parameter  $H$ . The most direct approach for constructing a quad mesh aligned with a field would be to trace it. However, while it was shown [Myles et al. 2014; Ray and Sokolov 2014] that this approach can be implemented robustly, in general, it requires T-joints (i.e., beams joining other beams in the middle), and it is, in general, hard to ensure uniform spacing over the whole mesh. We choose a more conservative approach based on constructing a conforming quadrangulation, without T-joints, using a version of the mixed-integer quadrangulation algorithm [Bommès et al. 2009] at this step.

While the method does not guarantee perfect alignment of the parametric lines to the input field, it minimizes the deviation in least-squares sense. We refer to Bommès et al. [2009] for further details.

### 7 OPTIMIZING CELL GEOMETRY

Given an input quad-dominant mesh, split into triangles, we aim at deriving the optimal width and thickness for each edge. As we previously stated, to obtain a structure close to optimal, the edges need to follow the directions we derived in the field optimization step. While our material distribution optimization method works for any mesh, the further the edges deviate from their optimum directions, the greater would be the total weight. The main idea of our algorithm is to use a *local-global* iteration, solving per-triangle *concave* problems for each cell, for which the solution is guaranteed to be on the boundary on the constraint domain. This yields a rapidly converging efficient algorithm.

#### 7.1 Optimization algorithm

We introduce a domain-decomposition-style algorithm for solving the Problem (7). We observe that in the optimization Problem (7), all constraints except  $Ku = f$  are localized, i.e., each constraint uses variables related to one cell. Moreover, the functional itself is a sum of local volume terms  $V^c(w^c, h^c)$ .  $Ku = f$  expresses the equation of force balance, i.e., that the sum of beam forces at each node is equal to the external force at this node. Our approach is to fix individual beam forces to their values for current values of geometry parameters  $w$  and  $h$ , and then solve for an update to  $w$  and  $h$  as a set of local volume-minimization problems with variables  $w^c, h^c$ , replacing the global constraint  $Ku = f$  with local constraints requiring that individual beam forces remain the same.

We start with an outline and then elaborate on how the local step optimization problems are solved.

Initially, we assign sufficiently high values to  $w$  and  $h$ , to ensure that max stress constraints are satisfied.

- **Global step.** The global step is just the standard solution of the elastic equilibrium problem, for fixed cell parameters:

Solve  $Ku = f$ , for fixed  $K$  defined by  $w$  and  $h$ . Compute beam forces as described below.

– **Local step.** The local step is the key part of our algorithm. Recall that an important feature of optimal structures is that they are *critically stressed*, i.e., the maximal stress on any element is equal to the maximum possible  $\sigma_0$ . The reasons for this are straightforward: if the stress on an element is below  $\sigma_0$ , one is free to remove some material, increasing the stress on the remaining part, and decreasing the weight.

This motivates our approach. For the local step, we keep the displacements computed at the global step fixed and solve for widths and thicknesses that would result in maximal stresses on blocks reaching the critical value  $\sigma_0$  for given displacements. Each system has six unknowns, with three constraints on stresses.

*Block forces.* To formulate our local algorithm, we introduce *block forces*, for individual blocks of each cell; we determine  $w^c$  and  $h^c$  for each cell by minimizing the cell volume, while keeping the block *critically stressed*, i.e., with maximal stress  $\sigma_0$  and block forces constant.

Let  $K^{loc}$  be the element stiffness matrix corresponding to a block  $B$ . The vector of forces corresponding to a beam is the vector  $\nabla_u E^{loc}$ , i.e., the derivative of the block energy  $E^{loc} = \frac{1}{2}u^T K^{loc}u$  with respect to displacements. Most forces in this vector are zero.

We use lower-case, non-bold  $d^m$  for the column vector of  $(D^m)^T$ , and  $d^b$  for column vector of  $(D^b)^T$ , corresponding to the stresses on block  $B$ . After some rearrangement, the force vector  $f^{loc} = K^{loc}u$  due to elastic forces produced by a block is

$$f^{loc} = K^{loc}u = Ewhl \left( (d^m)^T u d^m + \frac{1}{6}h^2 (d^b)^T u d^b \right).$$

Note that this equation implies that  $f$  is in the span of vectors  $d^m$  and  $d^b$ .

Let  $\tilde{d}^m, \tilde{d}^b$  be the dual pair of vectors to  $d^m, d^b$ . Let the magnitudes of membrane and bending stresses be  $|E(d^m)^T u| = \sigma^m$ ,  $|E(d^b)^T u| = \sigma^b$ ; by taking dot products of both sides with the dual vectors  $\tilde{d}^m, \tilde{d}^b$ , we arrive at the equation (dropping beam/cell subscripts):

$$wh\sigma^m = |f^m/l| = g^m, \quad \frac{1}{6}wh^3\sigma^b = |f^b/l| = g^b. \quad (15)$$

*Local optimization problem.* The stress in the block, under our assumptions, reaches its maximal value at the top or bottom, where its magnitude is equal to  $\sigma^m + h\sigma^b$ . This leads to the critical stress constraint  $\sigma^m + h\sigma^b = \sigma_0$ . Using expressions for  $g^m$  and  $g^b$  above, which we keep fixed at the local step, this is equivalent to

$$\frac{g^m}{h} + \frac{6g^b}{h^3} = \sigma_0 a y,$$

where we have switched to the variable  $y = w/a$  introduced in Equation (7), where  $a$  is the corresponding cell triangle height. Without loss of generality, we assume  $\sigma_0$  to be 1, which can be achieved by scaling all forces. The complete local problem in

variables  $h_i^c, w_i^c, i = 1, 2, 3$  is:

$$\begin{aligned} \min_{y^c, h^c} V(y^c, h^c) \quad \text{s.t.} \quad & \frac{g_i^m}{h_i^c} + \frac{6g_i^b}{(h_i^c)^3} = \sigma_0 y_i^c a_i \\ & 0 \leq h_i^c \leq h_{max}, \quad y_i^c \geq 0, \quad \text{for } i = 1, 2, 3 \\ & y_1^c + y_2^c + y_3^c \leq 1, \end{aligned} \quad (16)$$

where the volume  $V$  is given by Equation (1), and the last three constraints by Equation (2).

By eliminating the stresses, we arrive at a *single constraint per block* relating  $w$  and  $h$ , which we express as follows:

$$y = (6g^b z^2 + g^m z)/a, \quad (17)$$

where  $z = h^{-1}$  is a new variable we introduce to simplify the expressions. This allows us to eliminate all variables  $w_i^c$  from the local optimization problem, leaving only three variables  $z_i^c$ , with constraints  $0 \leq z_{min} \leq z_i^c, i = 1, 2, 3$ , and  $z_{min} = (h_{max})^{-1}$ .

We say that a cell is *filled* if the equality  $y_1^c + y_2^c + y_3^c = 1$  is satisfied, i.e., the blocks completely fill the cell.

Without loss of generality, we assume that for the solution,  $z_1 \leq z_2 \leq z_3$ ; in practice, six problems corresponding to six permutations of  $(1, 2, 3)$  need to be solved and minimal solution picked.

**PROPOSITION 1.** *The function  $V(z^c)$  is a concave function of  $z_i$ . As a consequence, its minima are reached on the boundary of the constraint domain; specifically, it is reached at one of the five types of configurations:*

- (1) all three blocks have maximal thickness:  $z_i^c = z_{min}, i = 1, 2, 3$ ;
- (2) the cell is filled, i.e.,  $y_1^c + y_2^c + y_3^c = 1$ , and no inequality constraint reaches equality;
- (3) the cell is filled, and two thicker beams have equal thickness,  $z_1^c = z_2^c$ ;
- (4) the cell is filled, and two thinner beams have equal thickness,  $z_2^c = z_3^c$ ;
- (5) the cell is filled, and the thickest beam has maximal thickness  $z_1^c = z_{min}$ .

In the first case, the solution is completely determined. In the second case, there are four possibilities: no inequality constraint is active (a 2-variable unconstrained optimization problem needs to be solved, e.g., parametrized by  $z_1^c, z_2^c$ ); the other three cases define one-parametric families of solutions, and one-dimensional unconstrained optimization needs to be performed to find exact values, as we explain below. These families can be parametrized by, e.g.,  $z_3^c = (h_3^c)^{-1}$ , with the values of the remaining  $z_i^c$  and  $y_i^c$  determined from the active constraints. The proposition is proved in Appendix C.

This behavior of  $V$  is in stark contrast to the low-volume formulation ignoring common areas of beam-like parts of the structure: one can see that in three cases out of four, it creates a completely filled cell.

*Solving the optimization problem.* Proposition 1 leads to an efficient algorithm for the local step.

Observe that the constraint  $y_1^c + y_2^c + y_3^c = 1$  has the form

$$\sum_i g_i^m z_i^c + 6g_i^b (z_i^c)^2 = 1, \quad (18)$$




Fig. 14. Inflation: (a) Each block is meshed independently, leading to a visible artifact; (b) An entire stream of edges is meshed to produce a smoother result.

i.e., it is quadratic in  $z_i$ . This allows us to reduce the problem to a set of unconstrained optimization problems in one or two variables.

- (1) Compute  $g_i^b, g_i^m, i = 1, 2, 3$ , from Equation (15) for current displacements.
- (2) Evaluate  $V(z)$ , for the case 1 solution with  $z_i^c = z_{min}$ .
- (3) For each permutation of (1, 2, 3), solve three one-dimensional optimization problems, minimizing  $V(z^c)$ , for each of the cases 3–5, and the two-dimensional problem for case 2 of Proposition 1. In each case 3–5, substitute the active constraint for  $z_i^c$  into Equation (18), yielding a quadratic equation in two remaining free variables, one of which is  $z_3^c$ . Solve it to express the other variable in terms of  $z_3^c$ , and solve a one-dimensional optimization problem for  $V(z_3^c)$ . This yields a set of *solution candidates*; the minimal solution is guaranteed to belong to it.
- (4) Pick the minimal solution from the set of solutions obtained for all possible permutations and cases on the previous step.
- (5) Update  $w_i^c$  using formula  $w_i^c = (6g_i^b z_i^c + g_i^m z_i^c)$ , and recompute the global stiffness matrix  $K$ .

The convergence behavior of the method is considered in Section 8.

*Handling polygonal cells and postprocessing.* There are three factors not considered in the solution method above: (1) possible inconsistency of thickness values across diagonal edges inside triangulated polygonal cells; (2) the coherence of block widths and thicknesses along the edge lines of the quad mesh, approximating the optimized stress lines (Figure 14). While jumps in thickness/width along these lines do not affect the stresses in our simplified model; in practice, these are likely to lead to localized stress concentrations close to jumps, and they are aesthetically objectionable. (3) Stress values may slightly exceed the maximal stress after a final global step.

We experimentally observe that many of the candidate solutions have close values, especially in areas with no predominant stress direction.

We address (1) primarily in the process of optimization, at step 4, we pick a minimal candidate solution with lowest block thickness on the diagonal, which may not be the most optimal one, as long as it does not deviate above a threshold. Once the optimization is complete, for each subcell of a polygonal cell, we increase the lowest block thickness to the maximal minimal thickness over the whole cell. In addition, for each cell, we store a number of candidate solutions with the smallest volumes.

We address (2) in a post-processing step using stored candidate solutions: for each non-diagonal edge of a cell, we find its continuation edges along a quad mesh edge line in both directions, and choose the candidate solution closest in width to the average of the previous and next edge widths along the edge line.

To address (3), we find all blocks with stress exceeding  $\sigma_0 = 1$ , and increase their thickness and width, while maintaining constraints, to decrease the stress to the bound. This process is repeated iteratively until convergence. We note that all additional steps are designed to ensure that the final result satisfies stress constraints: in all cases, we never decrease the amount of material in cells, so while the resulting solution may be suboptimal, it always satisfies stress constraints. In practice, the effect of these alterations on the resulting weight is small.

Alternatively, to alleviate these issues, one could optimize for the width and thickness per-vertex instead of per-edge. To be more precise, we would require two widths/thicknesses per-vertex since two beams crossing at a vertex may be different. While this approach is, in principle, possible, in practice, it is much more expensive—for vertex-based cells, the volume function is not likely to be concave, so a general nonconvex solver needs to be used. Our experiments showed that as the number of variables and constraints increases, a general solver is likely to get stuck in local minima (see Section 8, Figure 21).

*Inflation.* Once we have performed the weight optimization, we have one value of thickness and one value of width for each half-edge of the optimized polygonal mesh. If we simply extrude the solid block that matches thickness and width for each half edge cell, we end up in the situation illustrated in Figure 14(a), where there are visible discontinuities between adjacent blocks, causing possible structural discontinuities. Instead, for each continuous stream of quad edges, we generate a unique solid block that interpolates thickness and width along its length (Figure 14(b)). In detail, given a sequence of aligned edges, we first derive a thickness value for each vertex by averaging the thickness from its adjacent half edges. Similarly, we interpolate widths, but this time, we derive two different values for each vertex, one for each side of the sequence. We then define a tangent vector for each vertex as the cross product between its normal and its direction along the edge sequence (obtained by averaging the direction of the two incident edges). Then, having defined a proper reference frame, a thickness, and a width for each vertex, we have all the information to extrude a proper volumetric block. We perform a Boolean operation to merge all the blocks together to a manifold watertight mesh using the approach of Zhou et al. [2016].

## 8 EVALUATION

*Topology optimization.* To validate our approach vs. a general-purpose topology optimization method, we solve a similar problem with topology optimization code [Aage et al. 2015] by restricting the volume of the material to a cylinder of *fixed* small thickness  $h$ , and choosing the volume grid resolution to be half of the cylinder thickness, leaving little room for shell shape variation; with this grid resolution, Aage et al. [2015] can only generate beams (or walls) of thickness  $h$  and  $h/2$ . To make the comparison fair, we configure our optimization to have a minimum thickness close



Fig. 15. (a) Cylinder beam structure obtained using topology optimization with a target volume of 0.375. The pixelated effect is caused by the voxel grid; (b) structure obtained using our method with a maximum thickness of 0.05, minimum thickness of 0.04, no support shell ( $h_s = 0.00$ ).

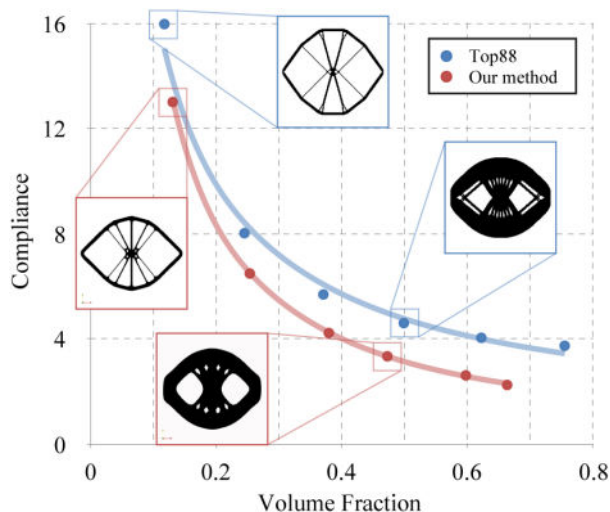


Fig. 16. Volume fraction vs. compliance energy for a standard example, a double truss, SIMP topology optimization vs. our method.

to the maximum thickness (near-constant thickness). Default parameters were used in Aage et al. [2015]. We observe that for small target volume fractions, as expected, Aage et al. [2015] generates substructures similar to the beam structures we construct (Figure 15). For large volume fractions, both methods result in variable-thickness shells. However, due to the limitations on the thickness mentioned above (large minimum thickness), the thickness variation is minimal. We note that topology optimization requires high resolution (with multiple cells fitting in the thickness direction of the shell) to achieve more variation in shell-thickness.

Figure 16 shows the results of SIMP topology optimization vs. our method, with compliance as a function of the volume fraction.

While our algorithm does not have a target volume fraction, we try to match these using different load magnitudes. To measure compliance for our results, we use HyperWorks finite element analysis (FEA), with a unit load. For SIMP, we use the code provided by Andreassen et al. [2011] with a square grid of  $200 \times 200$  pixels, and the compliance reported by this software. The main reason not to use HyperWorks for SIMP is that it would require triangulating the result and smoothing its sharp corners, so the result will be inaccurate. Instead, we verified that the compliance reported by both systems were the same for various shapes and load cases. In the plot, we observe similar behaviors for both methods. We note that in our case, we need to choose the beam spacing parameter: when this parameter is chosen to be too coarse, the performance deteriorates.

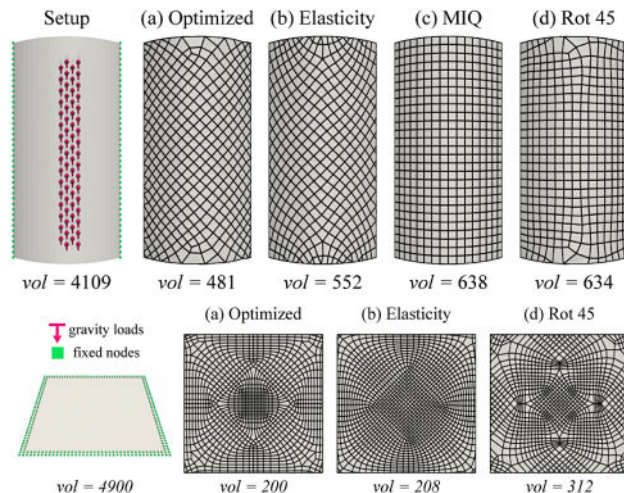


Fig. 17. Comparison of volumes obtained using different fields for the shell structure. Images (a–d) show the quad-mesh obtained with the different fields: (a) optimized stress directions from solving Equation (14); (b) stress directions from solving elasticity; (c) MIQ field constructed using the smoothing method of Bommers et al. [2009]; and (d) constructed rotating (a) by 45 degrees.




Fig. 18. Effects of the choice of diagonal directions: on the left, the original diagonals computed by our method, choosing the shorter diagonal; on the right, the diagonals of the first boundary ring are preserved and internal ones are flipped.

*The role of beam direction.* In Figure 17, we explore the dependence of the role of beam direction in structure optimality, by comparing structure volume for several fields, in addition to our optimized field (Equation (14)). As a “worst-case” baseline, we use the cross-field at a maximal distance from the original field (d); as one can see, the field makes a significant difference when it is very far away from the optimal directions, i.e., the choice of directions matters. Similarly, the curvature field (c), unrelated to stress directions, produces relatively high values of the volume. On the other hand, the difference between the optimized field and the stress field (a and b), while present, is not large in most cases. This experimentally justifies using, e.g., Li et al. [2017], the elasticity stress field instead of the optimized field; however, the additional expense of using an optimized field is minimal, so there is effectively no penalty for this improvement.

To evaluate the effect of the choice of diagonals in conversion from a quad mesh to a triangle mesh, we compare the optimized volumes for the shorter and longer diagonal choices. For the pipe example (Figure 18, left), the volumes are equal to 14.6% and 16.9% of the original volume, respectively. Choosing shorter diagonals

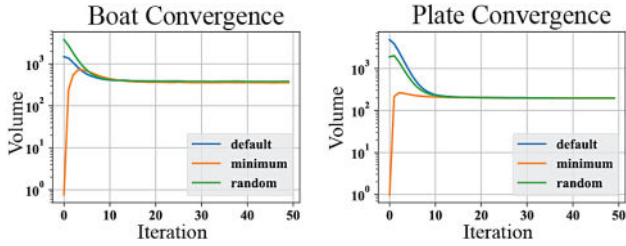


Fig. 19. Left: Volume convergence for our method, for the boat model, using different starting points. Right: similar plots for the bending plate.

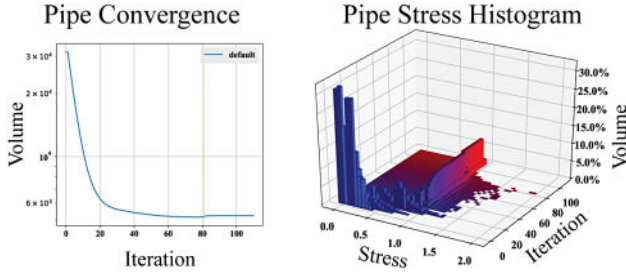


Fig. 20. Convergence plots for the Pipe model (Figure 29(d)). On the left, the volume vs. iteration count; On the right, the evolution of the histogram of the stress where volume percentage is measured with respect to the initial volume. At iteration 81, we start the post-process that ensure stresses are below  $\sigma_0 = 1$ .

(our default) results in better-shaped triangles, which we have observed to require less material to keep the stress low. If instead of flipping all diagonals, we flip only interior diagonals, using the shortest-diagonal on the ring near the boundary, we obtain a volume of 15.9% (Figure 18, right).

*Convergence and dependence on the starting point.* We have observed that our algorithm almost invariably converges in several iterations, and yields the expected behavior of solutions in the extreme cases (ribs in the case of bending-dominant shells, and wide-and-thin beams in high-tension/low-bending areas). The plots in Figure 19 show the volumes at each iteration for a basic and more complex problem. Note that sometimes the optimal value is approached from below: the local step overshoots the volume reduction and the stress exceeds the maximal allowed level. Nevertheless, the method recovers reliably.

The plots on Figure 20 show the volume and stress histogram vs. iteration count for a complex model. Using our default initialization, all stresses start with values below  $\sigma_0$ , by construction, and during the optimization, they converge to values near  $\sigma_0$ . As mentioned on Section 7.1, during the optimization, some stress values may exceed  $\sigma_0$ . These are resolved during the post-process (after iteration 80 for this example), without significantly increasing the total volume.

*Comparison to other optimization methods.* We also compare our method to two general-purpose constrained optimization methods, Sequential Least Squares Programming (SLSQP) [Kraft 1988] and Ipopt, a barrier interior point method [Byrd et al. 2000]. In this setup, we used an approximation of the beams volume that is smoother (it uses average instead of *max*) and simple lower and

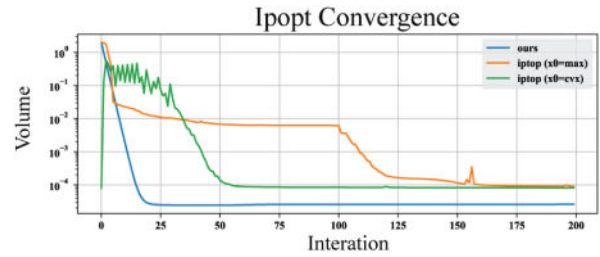


Fig. 21. Volume vs. iteration step for a small cantilever example (408 beams) using Ipopt vs. our method.




Fig. 22. In a bending dominated structure, changing the maximal allowed thickness from small to large causes the results to change from a solid plate to a rib-like structure.




Fig. 23. In a membrane dominated structure, changing the thickness lower bound from small to large causes the results to change from a solid plate to a rib-like structure. This example does not use a support shell ( $h_s = 0.00$ ).

upper bounds on the width and thickness. As a starting point, we have used the solution of the convex problem with volume ignoring the overlaps. Somewhat surprisingly, these methods were not able to change this initial solution by much after a few hundred iterations; although moving in the right direction, in terms of values, it may differ by a factor up to two from the optimal solution.

While the SLSQP solutions exhibited oscillatory behavior, alternating decreasing the functional with decreasing constraints violations, the interior point method solutions mostly stayed close to the initial values. Figure 21 shows comparative results with Ipopt for a small cantilever test case. We observe that Ipopt converges to a volume almost an order of magnitude larger. When Ipopt is initialized with the solution of the convex problem with simplified volume functional (green), the optimization fails to find a better solution. Using a different initialization (orange: maximum width and thickness) does not provide better results. In contrast, our method converges to a much better solution in only a few iterations.

*Effects of constraints on structure parameters.* Figure 22 shows the effect of increasing maximal allowable thickness in a bending scenario, with the structure moving from fully solid, to a structure of narrow but tall beams (if the bound is increased to infinity, in principle, any bending force can be realized by a zero-volume infinitely thin rib).

Similarly, Figure 23 shows the effect of increasing *minimal* thickness in a *tensile/compression* scenario. In this case, tall beams

appear when increasing the lower bound because widths decrease to compensate for the increase of thickness. The usage of beams in tensile/compression scenarios, while common, is often a consequence of constraints on minimum element size and are sub-optimal in the absence of these constraints. On the torsion cylinder experiment, the effect of using a minimal thickness of 0.04 vs. 0.00 increases the volume by 60%.

*Comparisons to related work.* A direct apples-to-apples comparison of different methods for optimization for shells is extremely difficult for a number of reasons, most importantly, because of different models used (e.g., all approaches use a variation of a beam model, and the specific choice has a significant impact on stress estimation). Another important reason includes a strong dependence of the results on the maximal width/thickness/aspect ratio constraints. We summarize qualitative differences from the most closely related works—Kilian et al. [2017] and Li et al. [2017].

The key difference from Kilian et al. [2017] is that it changes the shape of the surface, focuses on the case when bending forces can be neglected, and uses the narrow-beam volume approximation. The approach by Kilian et al. [2017] might be appropriate for the targeted type of architectural applications, combining surface shape design with structural constraints; but in many cases, the surface cannot be modified arbitrarily. Similarly, form-finding methods for structural shape optimization [Veenendaal and Block 2012] apply a minimal modification to the input shape such that the final form can transfer their loads purely through axial or in-plane forces. In the optimal case, no out-of-plane stress is present in the final shape, which may require a significant change to the surface, e.g., for gravity loads, resulting shapes need to have a very specific form for bending loads to be eliminated. For most examples we consider, decreasing out-of-plane forces, for given loads, without a drastic change on the input shape is impossible, e.g., it would be impossible to keep the surface of the shelf or the bottom of a boat flat.

In comparison, we deal with a shell of fixed surface shape, take bending forces into account, and use a more precise volume approximation, which leads to a non-convex problem.

The method of Li et al. [2017] is closer to ours: they construct a quadrangulation following a field to determine beam directions, and also take bending into account in the elasticity model. However, rather than using an optimized stress field as we do, they use the stress field of the non-optimized shell. Although the fields are typically close, in some cases, the optimized field provides an advantage (Figure 17).

More importantly, Li et al. [2017] uses a narrow-beam approximation for the volume, which has a major effect on the efficiency of the reinforced structure. To measure this effect, we compare the results of our optimization using different volume formulas as functional (Figure 24). The results show that using the simpler volume approximation can increase the volume used by up to 44%.

As explained on Section 3, the convex volume approximation biases the solution against using variable-thickness shells. To demonstrate the effects of this, we compare the stiffness of two reinforcements, obtained using different methods, with the same material cost. The shell and rib-reinforcement structure (Figure 25(b)) were provided by the authors of Li et al. [2017], while the second




Fig. 24. Optimization results using different volume formula: (a) exact volume and (b) approximate convex volume.




Fig. 25. Finite Element Analysis of reinforced structures: (a) variable thickness reinforcement generated by our system; (b) rib-reinforcement from Li et al. [2017]; (c–d) simulated displacements on reinforced structures.

type of reinforcement (Figure 25(a)) was generated with our algorithm.

To match the setup of Li et al. [2017], we set a minimum width for the beams. In contrast to that work, for similar loads, our optimization results in a lower-weight variable-thickness wall without protruding ribs, due to our use of the exact volume formulation (Section 3).

To compare the results fairly, with no dependence on the beam model used, we post-process both results using TetWild meshing software [Hu et al. 2018] to create tetrahedral meshes of similar density, and then use Finite Element Analysis to solve for the displacements with the loads and supports used in optimization (we use the PolyFem implementation [Schneider et al. 2019]). We impose zero-displacements on the nodes at floor level, a uniform vertical load on the rest of the structure, a Young’s modulus  $E = 210,000$  and Poisson’s ratio  $\nu = 0.3$ . The resulting displacements are shown in Figure 25(c) and (d). Our reinforced structure shows a maximum displacement of 67.31, while the rib structure of Li et al. [2017] shows a higher maximum of 97.16. Thus, under the same boundary conditions, our result achieves 1.44 times higher stiffness compared to a rib-reinforced structure of the same weight.

This comparison does not include the I-beam profile optimization of beam cross-sections introduced by Li et al. [2017], as I-beam shapes were not included in the geometry provided by the authors, and relevant code is no longer available. However, in the case of dominant tensile/compression loads along the beam, this technique is not likely to yield significant improvement, as these forces depend primarily on the cross-section area, not its shape.

A minor difference with Li et al. [2017] concerns the smoothing of the final result. The approach of Li et al. [2017] is to add an




Fig. 26. (a) Shelf of uniform thickness; (b) optimized shelf of same weight; (c) model of optimized shelf.

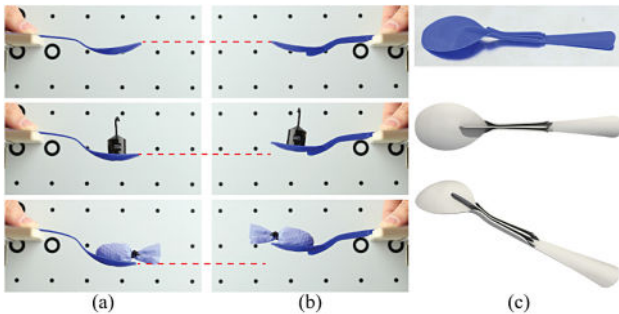


Fig. 27. (a) Spoon of uniform thickness; (b) optimized spoon of same weight; (c) model of optimized spoon.




Fig. 28. (a) Leaf of uniform thickness; (b) optimized leaf of same weight; (c) we use a digital scale to measure the pulling force; (d) model of optimized leaf.

additional smoothing term in the function that optimizes the material distribution. This term smooths the volume along sequences of consecutive aligned edges, classified into curved, circular, and tree-like rib elements. Our smoothing procedure is a part of the post-optimization geometric process along consecutive aligned edges, that is easier to control (see Figure 14(b)). As a future improvement, we plan to include a global smoothing term in the energy formulation.

*Physical Experiments.* We have printed several simple structures to validate our optimization experimentally (Figures 26–28). Due to the highly approximate nature of the physical model used, we did not attempt a quantitative match to the simulated values,




Fig. 29. Examples of consumer product objects. For each one, images show loads and initial stress distribution, quad-rangulation, cell optimization colored by thickness, and final geometry.

but we did closely match the values of the printed models. The comparison in all cases is between an optimized model and a uniform-thickness model of the same weight. Observed displacements in all cases differed by a factor more than 3, suggesting a similar difference in stress.

The optimization of the shelf was done using distributed loads (Figure 30, last row). However, when the shelf bends, the physical load is applied in only two separate regions. Using properly distributed loads would, most likely, only increase the difference in performance, as the structure was optimized for this case.

Finally, Figures 29 and 30 show a set of optimized shells obtained for a variety of shapes using our method (see Table 1 and 2 for details); in all cases, we have preserved a minimal width/thickness beam to indicate the mesh edges, but the load is carried by a relatively small number of beams; we found that ribs tend to appear, even in tension areas, unless the thickness bound is set to very low values. This is consistent with the observation that with no thickness bound, the load carried by bending forces is maximized, if the goal is to reduce the volume.

## 9 CONCLUSIONS AND FUTURE WORK

In this article, we have described a way to approximate and efficiently solve the problem of minimizing the weight of a support structure for a shell. Our work focuses on choosing a simple, yet sufficiently expressive, geometric description of the support structure, and an optimization algorithm capable of optimizing it. The proposed method separates the construction into three stages, with the first stage optimizing the field the beam directions must follow and creating a corresponding quad-dominant mesh, the second stage creating a cell structure with optimized shape parameters, and the third stage creating an actual realization.




Fig. 30. Examples of structures obtained for a variety of shapes and loads. For each one, images show loads and initial stress distribution, quadrangulation, cell optimization (colored by thickness in logarithmic scale), and final geometry.

This makes our approach particularly flexible and allows us to integrate a variety of additional user inputs and constraints, e.g., by modifying the field to change the truss directions, or by adding beams in the second stage to support a connection to a separate

object. It is also very efficient, with optimization converging in a few iterations, and quite consistently.

While we use a highly simplified model for cell mechanics, the overall approach admits replacing this model with a more

Table 1. Statistics for the 3D Models in Figures 30 and 1

	$d(mm)$	$ F $	$h$	$h^s$	$k$	$s$	$t_{\text{solve}}(s)$
Aquadom	85.60	7,742	0.5	0.005	100	0.2	72.330
Basket	127.56	10,943	2.0	0.06	30	0.5	72.513
Beetle	169.01	7,558	1.2	0.012	50	0.5	35.999
Boat	92.60	9,084	1.0	0.05	100	0.2	33.067
Botanic	42.59	2,152	0.5	0.01	100	0.5	35.586
Bowl	117.77	4,779	2.0	0.0	32	0.2	12.354
Bucket	150.96	7,860	2.1	0.21	50	0.5	45.744
Bunny	106.00	7,471	2.0	0.9	100	0.25	22.291
Duct	460.58	6,932	1.0	0.01	100	0.2	69.593
Leaf	99.24	2,980	2.0	0.9	46	0.5	1.954
Neumunster	111.92	6,072	2.0	1.2	63	0.2	6.959
Pipe	159.48	10,409	2.0	0.0	82	0.2	80.226
Shelf	88.54	2,400	4.0	0.5	60	0.5	1.663
Spoon	133.11	6,528	4.0	0.9	29	0.5	4.129
Stevia	90.68	4,840	2.0	0.9	50	0.5	23.941
Vase	69.64	4,798	1.0	0.45	100	0.5	21.817

The Bounding Box Diagonal  $d$  (mm), the Number of Triangles  $|F|$  in the Input Mesh, the Maximum Thickness  $h$  (mm) of the Beams, the Constant Shell Thickness  $h^s$  (mm), the Number of Iterations  $k$  and Step-size  $s$  of the Optimization Solve, and the Time  $t_{\text{solve}}$  it Took to Complete All Iterations.

Table 2. Performance for the 3D Models in Figure 30

	$V^s(mm^3)$	$V^b(mm^3)$	$h_{eq}(mm)$	$c_{eq}/c$
Aquadom	16.989	90.335	0.032	1.601
Basket	1,207.584	3,392.022	0.229	7.983
Beetle	170.128	1,287.227	0.103	5.739
Boat	189.997	418.539	0.160	7.099
Botanic	9.134	80.533	0.098	1.266
Bowl	0.000	4,083.710	0.410	4.437
Bucket	105.195	58.274	0.326	26.924
Bunny	7,270.428	1,134.806	1.040	1.229
Duct	1,075.645	19,777.817	0.194	2.105
Leaf	2,304.430	659.348	1.158	1.105
Neumunster	44,378.721	3,902.169	9.139	1.076
Pipe	0.000	45,833.029	2.551	1.161
Shelf	1,176.000	1,386.550	1.090	7.252
Spoon	1,544.960	326.733	1.090	18.621
Stevia	8,329.664	2,568.409	1.178	1.198
Vase	3,505.556	492.212	1.026	1.953

We show the ratio  $c_{eq}/c$  between the compliance  $c$  of our results and the compliance  $c_{eq}$  of a shell of constant thickness of the same weight. The fixed shell volume  $V^s$ , the support structure volume  $V^b$ , the thickness  $h_{eq}$  of the equivalent-weight constant thickness shell.

advanced finite element formulation. In this case, the local step would likely require numerical optimization; however, as long as the model for a cell stays low-parametric, one is likely to be able to solve it efficiently.

*Future Work.* In the future, we will study the effect of alternating optimizing the cell mesh  $\mathcal{P}$ , which defines the layout of the beams, and optimizing beam parameters. While our evaluations on the role of beam directions (Figure 17) suggest only an incremental benefit, it might help regions with large stress concentrations. Multiple loading scenarios can be optimized for simultaneously,

following the approach of Sokół and Rozvany [2016]), i.e., computing stresses for each load case and constraining them to be bounded.

*Limitations.* Our work has two main limitations: first, the formulation for the field optimization still has restrictive low-volume and fixed-thickness assumptions. Based on our evaluation of field direction sensitivity of the final design, we do not view this as a significant limitation. The second, more significant, limitation is the highly simplified model we used. If exact results are needed, this model can be used to quickly obtain an initial result, which can then be refined using a more advanced mechanical description of cells and shape optimization.

## APPENDIX

### A CONVEXITY OF TRUSS CONTINUUM OPTIMIZATION

Expressing the functional in terms of entries of the stress matrix, using  $a = (\sigma_{11} + \sigma_{22})/2$ ,  $b = (\sigma_{11} - \sigma_{22})/2$ , and  $c = \sigma_{12}$ , we obtain

$$|\lambda_1(\sigma)| + |\lambda_2(\sigma)| = 2 \max(a, \sqrt{b^2 + c^2}), \quad (19)$$

which verifies convexity of the energy.

### B DUAL OF THE CONTINUUM SHELL PROBLEM

This derivation follows Strang and Kohn [1983], which we include here for completeness.

$$L(\sigma, \mathbf{u}) = \frac{E}{\sigma_0} \int_{\Omega} (|\lambda_1(\sigma)| + |\lambda_2(\sigma)|) dA + \int_{\Omega} \mathbf{u}^T \text{div} \sigma dA - \int_{\Omega} \mathbf{f}^T \mu dA.$$

Integrating by parts, and assuming either free boundary  $\partial_n \mathbf{u} = 0$  or fixed boundary  $\mathbf{u} = 0$ , we obtain

$$L(\sigma, \mathbf{u}) = \frac{E}{\sigma_0} \int_{\Omega} (|\lambda_1(\sigma)| + |\lambda_2(\sigma)|) dA - \int_{\Omega} \varepsilon(\mathbf{u}) : \sigma dA - \int_{\Omega} \mathbf{f}^T \mathbf{u} dA.$$

To obtain the dual, we need to minimize over all possible  $\sigma$  in a coordinate system aligned with  $\varepsilon$  the expression

$$\frac{E}{\sigma_0} (|\lambda_1(\sigma)| + |\lambda_2(\sigma)|) - \lambda_1(\varepsilon) \sigma_{11} - \lambda_2(\varepsilon) \sigma_{22}.$$

The minimum of this expression is  $-\infty$ , if  $|\lambda_i(\varepsilon)| > 1$ , for either  $i = 1$  or  $i = 2$ ; otherwise, it is zero. This leads to the dual Problem (10).

### C PROOF OF THE PROPERTIES OF $V(Z)$ .

We drop the cell superscript  $c$  in this proof. We assume that  $z_1 \leq z_2 \leq z_3$ . By the constraints of the problem,  $y_1 + y_2 + y_3 \leq 1$ ,  $0 < z_{\min} \leq z_i$ .

$V(y, h)$  in Equation (1) is a quadratic function of  $y_i$  and linear in  $h_i$ ; once we express it in terms of  $z_i$ , to eliminate the constraint between  $h_i$  and  $z_i$ , it becomes a cubic function of  $z_i$ , with Hessian

$$H(z) = \begin{bmatrix} A & -k_1 g_2^b & -k_1 g_3^b \\ -k_1 g_2^b & B & -k_2 g_3^b \\ -k_1 g_3^b & -k_2 g_3^b & -1/2 g_3^b (g_3^m + 3k_3) \end{bmatrix},$$

where  $k_i = g_i^m + 2g_i^b z_i = dw_i/dz_i$ ,  $B = -1/2 g_2^b (g_2^m + 3k_2 + 2g_3^m + 2k_3)$  and  $A = -1/2 g_1^b (g_1^m + 3k_1 + 2g_2^m + 2k_2 + 2g_3^m + 2k_3)$ .

A direct evaluation shows that  $v^T H(z) v$  evaluated for  $v = [1, 0, 0]$  is negative for positive  $z_i$  and  $g_i^b \neq 0$ . Similar is true for  $v = [0, 1, 0]$  and  $v = [0, 0, 1]$ . We conclude if  $g_i^b \neq 0$  for some  $i$ , any critical point in the interior of the domain is a maximum or a saddle, so there are no minima in the interior. If all  $g_i^b$  are zero, the volume is a linear function of  $z_i$ , and the optimum is also on the boundary.

The constraint  $y_1 + y_2 + y_3 \leq 1$  defines, if  $g_i^b \neq 0$  for some  $i$  a quadratic surface in  $z_i$  coordinates; and the constraints  $z_1 \leq z_2$ ,  $z_2 \leq z_3$ ,  $z_1 \geq z_{min}$  define three halfspaces. Their intersection is a tetrahedron with a curved face on the ellipsoid. The faces of the feasible domain correspond to one of the inequality constraints becoming an equality. Similarly, by a direct calculation, substituting either  $z_1 = z_{min}$ ,  $z_2 = z_1$ , or  $z_2 = z_3$  into  $V(z)$  and computing  $2 \times 2$  Hessians, we observe that these are not positive definite; hence, there can be no solution on faces corresponding to the linear constraint. Finally, the same fact holds for pairs of constraints that define edges of the constraint domain, e.g.,  $z_1 = z_{min}$  and  $z_2 = z_1$ . We conclude that either the solution is on the face corresponding to the constraint  $y_1 + y_2 + y_3 = 1$  or at the only vertex not on this face of the feasible domain,  $z_i = z_{min}$ ,  $i = 1, 2, 3$  (case 1 of the proposition). If the minimum satisfies  $y_1 + y_2 + y_3 = 1$ , it is either in the interior (case 2), or on one of the three edges of the boundary of that face (cases 3–5).

## D DISCRETE BENDING STRAIN FOR BEAMS

The common normal  $\mathbf{n}_i$  at  $i$  is defined as

$$\mathbf{n}_i = \sum_{\ell \in N(i)} \mathbf{e}_{i\ell} \times \mathbf{e}_{i \text{ next}(\ell)},$$

where  $\text{next}(\ell)$  is the edge following  $\ell$  in CCW order around  $i$ .

At each vertex, we define a normal  $\mathbf{n}_i$ , as the average of cross-products of pairs of incident edges. For an edge  $\mathbf{e}_{i\ell}$ , incident at a vertex  $i$ , let  $\mathbf{e}_{im}$  be the next CCW edge, and  $\mathbf{e}_{ip}$  be the previous edge. We define

$$\mathbf{q}_{lm} = \mathbf{e}_{i\ell} \times \mathbf{e}_{im}$$

the (unscaled) normal to the triangle formed by  $i, l, m$ . Then,  $\mathbf{n}_i = \sum_{l, m \in N(i)} \mathbf{q}_{lm}$ , and  $\hat{\mathbf{n}}_i = \mathbf{n}_i / |\mathbf{n}_i|$ .

We express the bending strain in the form  $D_{ij}^b u$ , where  $u$  is the vector of all vertex displacements. Define  $\Delta \mathbf{e}_{ij} = \mathbf{u}_j - \mathbf{u}_i$ . Then,  $\Delta \mathbf{q}_{lm} = \Delta \mathbf{e}_{i\ell} \times \mathbf{e}_{im} + \mathbf{e}_{i\ell} \times \Delta \mathbf{e}_{im}$ . The part of  $\Delta \mathbf{q}_{lm}$  affecting  $\Delta \hat{\mathbf{n}}_i$  is the part perpendicular to  $\hat{\mathbf{n}}_i$ :

$\Delta q_{lm}^\perp = (I - \hat{\mathbf{n}}_i \hat{\mathbf{n}}_i^T) \Delta \mathbf{q}_{lm} = P_i \Delta \mathbf{q}_{lm}$ , where  $P_i$  is the matrix  $(I - \hat{\mathbf{n}}_i \hat{\mathbf{n}}_i^T)$ . By substitution into the expression for  $n$ , we get

$$\Delta \hat{\mathbf{n}}_i = \frac{1}{|\mathbf{n}_i|} P_i \sum_{l \in N(i)} g_l \times (\mathbf{u}_l - \mathbf{u}_i),$$

where  $g_l$  is defined as follows:

- $g_l = \mathbf{e}_{ip} - \mathbf{e}_{im}$  if both  $p, m$  are in  $N(i)$ , i.e.,  $\mathbf{e}_{i\ell}$  is not on the boundary.
- $g_l = \mathbf{e}_{ip}$  if  $m \notin N(i)$ .
- $g_l = -\mathbf{e}_{im}$  if  $p \notin N(i)$ .

For a vector  $\mathbf{a} = [a_x, a_y, a_z]$ , let  $R(a)$  be the infinitesimal rotation matrix about  $\mathbf{a}$ : 
$$\begin{pmatrix} 0 & -a_z & a_y \\ a_z & 0 & -a_x \\ -a_y & a_x & 0 \end{pmatrix}.$$

Then,

$$\Delta \hat{\mathbf{n}}_i = \sum_{l \in N(i)} M_l^i (\mathbf{u}_l - \mathbf{u}_i),$$

where  $M_l^i = \frac{1}{|\mathbf{n}_i|} P_i R(\mathbf{g}_l)$  is a  $3 \times 3$  matrix.

$$\varepsilon^b [ij] = h \left( \sum_{l \in N(j)} (\mathbf{d}_{ji}^l)^T (\mathbf{u}_l - \mathbf{u}_j) + \sum_{l \in N(i)} (\mathbf{d}_{ij}^l)^T (\mathbf{u}_l - \mathbf{u}_i) \right),$$

where  $\mathbf{d}_{ij}^l = -(M_l^i)^T \hat{\mathbf{e}}_{ij} / l_{ij}$ , a vector of length 3. From this expression, we can immediately obtain  $D_{ij}$ .

## ACKNOWLEDGMENTS

We thank Chelsea Tymms for help with printing and photography, and the anonymous reviewers for their useful comments.

## REFERENCES

- Niels Aage, Erik Andreassen, and Boyan Stefanov Lazarov. 2015. Topology optimization using PETS: An easy-to-use, fully parallel, open source topology optimization framework. *Structural and Multidisciplinary Optimization* 51, 3 (2015), 565–572.
- Noam Aigerman and Yaron Lipman. 2015. Orbifold tutte embeddings. *ACM Transactions on Graphics* 34, 6 (2015), 190:1–190:12.
- Grégoire Allaire. 2002. *Shape Optimization by the Homogenization Method*. Number v. 146 in Applied Mathematical Sciences. Springer.
- Grégoire Allaire and François Jouve. 2008. Minimum stress optimal design with the level set method. *Engineering Analysis with Boundary Elements* 32, 11 (2008), 909–918.
- Erik Andreassen, Anders Clausen, Mattias Schevenels, Boyan S. Lazarov, and Ole Sigmund. 2011. Efficient topology optimization in MATLAB using 88 lines of code. *Structural and Multidisciplinary Optimization* 43, 1 (2011), 1–16.
- MOSEK ApS. 2015. *The MOSEK Optimization Toolbox for MATLAB Manual. Version 7.1 (Revision 28)*.
- Rahul Arora, Alec Jacobson, Timothy R. Langlois, Yijiang Huang, Caitlin Mueller, Wojciech Matusik, Ariel Shamir, Karan Singh, and David I.W. Levin. 2019. Volumetric Michell trusses for parametric design & fabrication. In *Proceedings of the ACM Symposium on Computational Fabrication*. ACM, 6.
- Martin P. Bendsoe and Ole Sigmund. 2004. *Topology Optimization: Theory, Methods, and Applications*. Springer Berlin.
- David Bicker. 2019. Plastic Pails. Retrieved from <https://pixabay.com/photos/buckets-wash-pail-clean-collect-4554791/>.
- David Bommes, Bruno Lévy, Nico Pietroni, Enrico Puppo, Claudio Silva, Marco Tarini, and Denis Zorin. 2013. Quad-mesh generation and processing: A survey. In *Computer Graphics Forum*, Vol. 32. Wiley Online Library, 51–76.
- David Bommes, Henrik Zimmer, and Leif Kobbelt. 2009. Mixed-integer quadrangulation. *ACM Transactions on Graphics (TOG)* 28, 3 (2009), 77.
- Richard H. Byrd, Jean Charles Gilbert, and Jorge Nocedal. 2000. A trust region method based on interior point techniques for nonlinear programming. *Mathematical Programming* 89, 1 (2000), 149–185.
- Marcel Campen, David Bommes, and Leif Kobbelt. 2015. Quantized global parametrization. *JTOG* 34, 6 (2015), 192.
- Jean-Pierre Dalbéra. 2016. The Armadillo Vault (Biennale d’architecture 2016, Venise). Retrieved from <https://www.flickr.com/photos/dalbera/30245801805/>.
- Denzeldoor. 2019. Nederlands: Internationale Terminal. Retrieved from <https://commons.wikimedia.org/wiki/File:EC0B3F5B-AD67.jpg>.
- Dodger67. 2013. A Blue pvc Injection Moulded Laundry Basket. Retrieved from [https://commons.wikimedia.org/wiki/File:Plastic\\_laundry\\_basket.jpg](https://commons.wikimedia.org/wiki/File:Plastic_laundry_basket.jpg).
- Hans-Christian Ebke, Patrick Schmidt, Marcel Campen, and Leif Kobbelt. 2016. Interactively controlled quad remeshing of high resolution 3D models. *ACM Transactions on Graphics* 35, 6 (2016), 218:1–218:13.
- Perle Geoffroy-Donders, Grégoire Allaire, Julien Cortial, and Olivier Pantz. 2017. Optimization of oriented and parametric cellular structures by the homogenization method. In *World Congress of Structural and Multidisciplinary Optimisation*. Springer, 767–778.
- Perle Geoffroy-Donders, Grégoire Allaire, and Olivier Pantz. 2020. 3-D topology optimization of modulated and oriented periodic microstructures by the homogenization method. *Journal of Computational Physics* 401 (2020), 108994.

- Eitan Grinspun, Yotam Gingold, Jason Reisman, and Denis Zorin. 2006. Computing discrete shape operators on general meshes. In *Computer Graphics Forum*, Vol. 25. Wiley Online Library, 547–556.
- Jeroen P. Groen and Ole Sigmund. 2018. Homogenization-based topology optimization for high-resolution manufacturable microstructures. *International Journal for Numerical Methods in Engineering* 113, 8 (2018), 1148–1163.
- W. S. Hemp. 1973. *Optimum Structures*. Oxford, Clarendon Press.
- Yixin Hu, Qingnan Zhou, Xifeng Gao, Alec Jacobson, Denis Zorin, and Daniele Panozzo. 2018. Tetrahedral meshing in the wild. *ACM Transactions on Graphics* 37, 4, Article 60 (July 2018), 14 pages. DOI: <https://doi.org/10.1145/3197517.3201353>
- Caigui Jiang, Chengcheng Tang, Hans-Peter Seidel, and Peter Wonka. 2017. Design and volume optimization of space structures. *ACM Transactions on Graphics* 36, 4 (2017), 1–14.
- Felix Kälberer, Matthias Nieser, and Konrad Polthier. 2007. QuadCover-surface parameterization using branched coverings. In *Computer Graphics Forum*, Vol. 26. Wiley Online Library, 375–384.
- Martin Kilian, Davide Pellis, Johannes Wallner, and Helmut Pottmann. 2017. Material-minimizing forms and structures. *ACM Transactions on Graphics (TOG)* 36, 6 (2017), 173.
- Felix Knöppel, Keenan Crane, Ulrich Pinkall, and Peter Schröder. 2015. Stripe patterns on surfaces. *J-TOG* 34, 4 (2015), 39:1–39:11.
- Dieter Kraft. 1988. A software package for sequential quadratic programming. *Forschungsbericht- Deutsche Forschungs- und Versuchsanstalt für Luft- und Raumfahrt* (1988).
- Wei Li, Anzong Zheng, Lihua You, Xiaosong Yang, Jianjun Zhang, and Ligang Liu. 2017. Rib-reinforced shell structure. In *Computer Graphics Forum*, Vol. 36. Wiley Online Library, 15–27.
- Yongqiang Li and Yong Chen. 2010. Beam structure optimization for additive manufacturing based on principal stress lines. In *Solid Freeform Fabrication Proceedings*. 666–678.
- Dmitry Makeev. 2006. Bottle. Colored Glass. Retrieved from [https://commons.wikimedia.org/wiki/File:Bottle\\_Colored\\_glass\\_img\\_01.jpg](https://commons.wikimedia.org/wiki/File:Bottle_Colored_glass_img_01.jpg)
- A. G. M. Michell. 1904. The limits of economy of material in frame-structures. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 8, 47 (1904), 589–597.
- MthwGlm. 2020. The Bright Light 22” from Penny Skateboards 2020 Classics Collection. Retrieved from [https://commons.wikimedia.org/wiki/File:PCAMP0004\\_Classics\\_Tile\\_BrightLight.png](https://commons.wikimedia.org/wiki/File:PCAMP0004_Classics_Tile_BrightLight.png).
- David J. Munk, Gareth A. Vio, and Grant P. Steven. 2015. Topology and shape optimization methods using evolutionary algorithms: A review. *Structural and Multidisciplinary Optimization* 52, 3 (2015), 613–631.
- Ashish Myles, Nico Pietroni, and Denis Zorin. 2014. Robust field-aligned global parametrization. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 135.
- E. Oñate, F. Zárte, and F. Flores. 1994. A simple triangular element for thick and thin plate and shell analysis. *International Journal for Numerical Methods in Engineering* 37 (1994), 2569.
- Julian Panetta, Qingnan Zhou, Luigi Malomo, Nico Pietroni, Paolo Cignoni, and Denis Zorin. 2015. Elastic textures for additive fabrication. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 135.
- Nico Pietroni, Davide Tonelli, Enrico Puppo, Maurizio Froli, Roberto Scopigno, and Paolo Cignoni. 2015. Statics aware grid shells. In *Computer Graphics Forum*, Vol. 34. Wiley Online Library, 627–641.
- Nicolas Ray and Dmitry Sokolov. 2014. Robust polylines tracing for n-symmetry direction field on triangulated surfaces. *ACM Transactions on Graphics (TOG)* 33, 3 (2014), 30.
- George I.N. Rozvany. 1976. *Optimal Design of Flexural Systems: Beams, Grillages, Slabs, Plates and Shells*. Elsevier.
- George I.N. Rozvany. 2012. *Structural Design Via Optimality Criteria: The Prager Approach to Structural Optimization*. Vol. 8. Springer Science & Business Media.
- Teseo Schneider, Jérémie Dumas, Xifeng Gao, Denis Zorin, and Daniele Panozzo. 2019. Polyfem. Retrieved from <https://polyfem.github.io/>.
- Ole Sigmund. 2001. A 99 line topology optimization code written in Matlab. *Structural and Multidisciplinary Optimization* 21, 2 (2001), 120–127.
- Ole Sigmund, Niels Aage, and Erik Andreassen. 2016. On the (non-) optimality of Michell structures. *Structural and Multidisciplinary Optimization* 54, 2 (2016), 361–373.
- Ole Sigmund and Kurt Maute. 2013. Topology optimization approaches. *Structural and Multidisciplinary Optimization* 48, 6 (2013), 1031–1055.
- Jeffrey Smith, Jessica Hodgins, Irving Oppenheim, and Andrew Witkin. 2002. Creating models of truss structures with optimization. *ACM Transactions on Graphics* 21, 3 (July 2002), 295–301.
- Tomasz Sokół. 2011. A 99 line code for discretized Michell truss optimization written in Mathematica. *Structural and Multidisciplinary Optimization* 43, 2 (2011), 181–190.
- Tomasz Sokół and George I. N. Rozvany. 2016. A new adaptive ground structure method for multi-load spatial Michell structures. *Advances in Mechanics: Theoretical, Computational and Interdisciplinary Issues* (2016), 525–528.
- Gilbert Strang and Robert V. Kohn. 1983. Hencky-prandtl nets and constrained Michell trusses. *Computer Methods in Applied Mechanics and Engineering* 36, 2 (1983), 207–222.
- Kam-Ming Mark Tam. 2015. *Principal stress line computation for discrete topology design*. Ph.D. Dissertation. Massachusetts Institute of Technology.
- Kam-Ming Mark Tam, James R. Coleman, Nicholas W. Fine, and Caitlin T. Mueller. 2015. Stress line additive manufacturing (SLAM) for 2.5-D shells. *Proceedings of International Symposium on Shell and Spatial Structures*.
- Amir Vaxman, Marcel Campen, Olga Diamanti, Daniele Panozzo, David Bommes, Klaus Hildebrandt, and Mirela Ben-Chen. 2016. Directional field synthesis, design, and processing. In *Computer Graphics Forum*, Vol. 35. Wiley Online Library, 545–572.
- D. Veenendaal and P. Block. 2012. An overview and comparison of structural form finding methods for general networks. *International Journal of Solids and Structures* 49, 26 (2012), 3741–3753. DOI: <https://doi.org/10.1016/j.ijsolstr.2012.08.008>
- Jun Wu, Christian Dick, and Rüdiger Westermann. 2016. A system for high-resolution topology optimization. *IEEE Transactions on Visualization and Computer Graphics* 22, 3 (March 2016), 1195–1208.
- Tomás Zegard and Glaucio H. Paulino. 2014. GRAND—Ground structure based topology optimization for arbitrary 2D domains using MATLAB. *Structural and Multidisciplinary Optimization* 50, 5 (2014), 861–882.
- Tomás Zegard and Glaucio H. Paulino. 2015. GRAND3—Ground structure based topology optimization for arbitrary 3D domains using MATLAB. *Structural and Multidisciplinary Optimization* 52, 6 (2015), 1161–1184.
- Tomás Zegard and Glaucio H. Paulino. 2016. Bridging topology optimization and additive manufacturing. *Structural and Multidisciplinary Optimization* 53, 1 (2016), 175–192.
- Haiming Zhao, Weiwei Xu, Kun Zhou, Yin Yang, Xiaogang Jin, and Hongzhi Wu. 2017. Stress-constrained thickness optimization for shell object fabrication. In *Computer Graphics Forum*, Vol. 36. Wiley Online Library, 368–380.
- Qingnan Zhou, Eitan Grinspun, Denis Zorin, and Alec Jacobson. 2016. Mesh arrangements for solid geometry. *ACM Transactions on Graphics* 35, 4, Article 39 (July 2016), 15 pages. DOI: <https://doi.org/10.1145/2897824.2925901>

Received July 2019; revised February 2020; accepted March 2020